Ambulance Assignment for Medical Emergencies

João Silva

Instituto Superior Técnico, Lisboa, Portugal

November 2021

Abstract

In the medical world, one of the purposes of ambulances is that of providing aid to emergency situations and transporting people to a facility where they can get further medical attention. However, there is a possibility that there are too many emergencies for the ambulances available in a certain region or that too many ambulances are made available resulting in resource wastage. Finding a balance between the availability of resources, coverage of all medical emergencies and minimizing response times has been shown to be a difficult task.

In this document, we analyze the literature on the Emergency Medical Services (EMS) ambulance location problem, addressing several models presented in previous research work addressing the EMS subject. Some define and formulate models with greater insight while other research works present innovative techniques

Additionally, we define a mathematical model that is able to represent any situation on a defined period of time with multiple emergency occurrences with several vehicles that can provide aid to said occurrence. We do this using Multi-Objective Combinatorial Optimization (MOCO) to tackle the ambulance location problem while focusing on minimizing the overall number of resources used, as well as minimizing the response times. We apply this model to real data retrieved from three different districts in Portugal, in various time periods. We then solve the instances created by our model in these scenarios and inferred conclusions and possible improvements to the Emergency Medical System in those situations.

Keywords: Multi-Objective Combinatorial Optimization, Ambulance Location Problem, Emergency Medical System

1. Introduction

An Emergency Medical System (EMS) can be defined as a system that aims to provide urgent treatment or stabilization in medical emergencies. Structurally, a control facility serves as an integrating part of this system, serving the purpose of receiving incoming emergency calls, which can be placed by any person, and then assigning emergency vehicles to the required location, depending on the number of people in need of immediate medical assistance and/or on the seriousness of the emergency at hand. It is important that this service is provided in the least amount of time possible in order to provide medical care to the person in need in the fastest possible way. Once the patient's situation is stabilized and he is transferred to the established health facility, the ambulance that was dispatched becomes available once more as soon as it returns to its base, where it can be given a new task, and all the procedures have been completed for the ambulance to be available again.

In the context of EMS, the vehicle location problem consists of locating the vehicles in some potential service sites in order to reduce the delay of covering emergency service demands [9]. However, in real-world scenarios, there is the need to keep the solutions feasible, while still granting a satisfiable level of optimality. This work focuses on analysing proposed solutions to the vehicle location problem, as the basis to then formulate a model to apply to existing data from the Instituto Nacional de Emergência Médica (INEM) and try to see if there is a more efficient way of providing aid to emergency situations.

This document is organized as follows. In Section 2, the fundamental concepts of the problem we are addressing are explained, and a description of the preparation we performed on our data is presented. In Section 3 we introduce our model in detail, explaining the incremental steps that led to our final model. Section 4 depicts the results of our investigation regarding the application of the created model. Section 5 concludes the document and presents possible enhancements or ideas to continue the study of this subject.

2. Background

The EMS system has an optimization problem regarding the allocation of the emergency vehicles (ambulances) which has been tackled in a number of different ways throughout the years, all of them ultimately aiming at increasing the efficiency of resource usage. The motivation for these works come from the fact that EMS's exist with the purpose of not only assigning ambulances to emergencies when these happen, but also conveying a distribution of these ambulances in a way that allows them to maximize the area covered, as it was first described, by Church & ReVelle [6], in their work that considered a fixed size fleet of ambulances. After this work, there have been an enormous number of authors dwelling upon the intricacies of the ambulance location, relocation and assignment problem, some of them suggest the usage of a dynamic approach, in which each ambulance is able to communicate with other ambulances (multi-agent approach) or with an ambulance coordinator (centralized approach) in order to decide where to go after it has been dispatched to an emergency and therefore maximize the coverage of each zone after each emergency has been dealt with [10, 8, 5]. This dynamic approach is presented with many computational or scalability problems when applied to either large fleets of ambulances or extensive land coverage. This means that dynamic approaches may not be achievable in the expected amount of time, effectively harming the end goal of reducing the overall response time to emergencies.

Since the allocation of emergency vehicles is a very complex problem that involves a great deal of variables and different scenarios, it is important to specify the context for each proposed solution. The details of several of relevant works are present in an very complete and interesting set of tables present in the work of Bélanger [3] in which the author divides and classifies the different works based on the specific details they cover. These tables cover variables such as the number of different types of ambulances available, a list of covering and standby site constraints and the objectives of the work developed in the literature. Different types of ambulances define an interesting alteration to the standard scenario, for example, a certain type of ambulance is more effective in heart disease related emergencies, and the system should, therefore, prioritize assigning these ambulances to emergencies of that nature.

However, other approaches have been taken in order to tackle this problem, namely some deterministic models, probabilistic models [4], and some more recent approaches that use the heuristics, for example, the tabu search heuristic[7]. Discussing all the papers in depth would make this work too extensive and difficult to follow, hence, added to the more concrete analysis of a few examples, the tables in the next pages will provide an overview on which papers cover what topics and which techniques are used.

The works mentioned thus far describe a summarized state of the art of the EMS location problem. There have been various ways of formulating and solving the problem throughout the years. However, for simplicity, we selected and analyzed those which we deemed more important in order the understand how the state of the art has evolved.

There have also been studies and reviews on the effectiveness, advantages and disadvantages of several approaches and techniques used to try to solve the EMS ambulance location or relocation problem. These studies are particularly helpful as an introduction to the matter at hand because they are often very generic in their descriptions in order to fit most models that are referenced in it and also because they allow the reader to become aware of which papers address what specific topics before actually reading them, effectively serving as a filtering method for whoever needs to search for works in this area of study[2, 1].

The database used, courtesy of the Instituto Nacional de Estatistica (INEM) has a variety of different fields that track a number of different variables regarding each emergency, vehicle, station and possibly other fields that were not used for the purposes of this study.

From an analytical point of view, we wanted to retrieve two important measures: the average distance that each vehicle traveled to get to an emergency site and the average time between each dispatch of the same vehicle. The first problem we came across had to do with missing values and incorrectly formatted records, and this problem occurred in the process of reviewing both measures. Since there is no way to infer or calculate the missing values, we discarded any records that were either.

This initial data cleaning allowed us to perform the first proper analysis of the data we were working with and start to draw conclusions from the results of that same analysis. However, in the average distance measure, we found values that were inconsistent with what we were expecting. More specifically, and given the fact that we were analysing data from continental Portugal, in which the longest straight line that connects two parts of the country is just below six hundred (600) kilometers long, we found records that read distances over four thousand (4000) kilometers. After analysing these records, we discovered that they all the same emergency coordinates corresponding to the point at zero degrees latitude and zero degrees longitude (0°N 0°E), as seen in Fig.1. Because of this we decided to not take into account the records that used this point as their emergency site.

Even though we had dealt with these erroneous

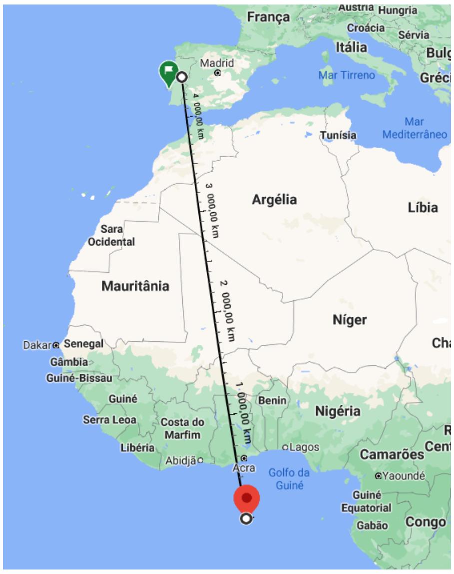

Figure 1: Example of emergency being recorded at position (0°N 0°E)

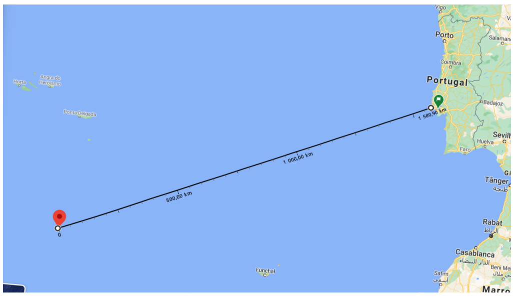

Figure 2: Example of emergency in the Madeira Archipelago

records, we still had a number of emergencies occurring at more than fifteen hundred kilometers (1500) which were still out of context. After further investigation, we found out that these records referred to marine emergencies in the Atlantic Ocean. These records still appear on this database because the Portuguese territorial sea area is about fifty one thousand squared kilometers (50957), sixteen thousand (16460) of those belong to the continental portion of Portugal and the rest belonging to the Açores and Madeira archipelagos, as seen in Fig.2. Since we are only dealing with terrestrial emergencies in this study, we imposed a coordinates limit that ensured that the emergency sites were inside continental Portugal.

The preparations for calculating the average dispatch time measure were not as extensive as the ones for the average distance. In order to get the records for this measure we first separated and organised the records according to the vehicle identifier and we arranged these same records in a chronological order, which made it trivial to calculate the consequent dispatch times and calculate the average dispatch time for all the vehicles. After getting this initial value, we tried to get more detailed data, namely isolating each one of the eighteen districts in Portugal and after that getting data by vehicle type and emergency priority levels.

3. Implementation

In this section we will go over the several models we developed at different stages of the elaboration of this document, explaining our thought process behind each change we made along the way, as well as providing insight as to what implications those changes had in the way we treated data and had to adapt our process.

After the initial analysis and preparation of the data retrieved from the INEM database, the first step was to create a very simple, traceable model with which we could check for implementation mistakes as well as start testing very small examples in the solver. These first examples were handmade and not derived from data. Our main goals with this approach were testing specific cases and how the solver would react to them while also optimizing the generation of an instance that represents each problem mathematically.

The ambulance location model can be defined as a graph $G = (V, E)$ where $V = N \cup M$, $N = (v_1, ..., v_n)$ and $M = v_{n+1}, ..., v_{n+m}$ being two vertex sets representing, respectfully, emergency sites and standby vehicles with and associated position. Additionally, E is an edge set where each edge $\{(v_i, v_j) : v_i, v_j \in V, i < j\}$ is associated with a travel time or distance t_{ij}. The variable x_{ij} will take the value 1 if and only if, the the vehicle j provides aid to the emergency site i, and the value 0 otherwise.

For simplicity purposes, we designed a simple example with a given set of emergency sites, $N = A, B, C, D$, as well as a set of standby vehicles $M = 1, 2, 3$ spread randomly on an example map, represented in Fig. 3. This map does not represent any specific zone, it merely serves as an example.

In this map, we implemented a baseline model. It is meant to be a blueprint for Static Models, where we only considered edges that represent connections between emergency sites and standby vehicles. The simplified graph for this implementation of a static model is represented in Fig. 4.

The objective this type of models aims to achieve is that of minimizing the number of vehicles and the overall travel time in a way that all emergencies are accounted for. For simplicity, we will consider that a vehicle is unavailable for a period of time correspondent to $2t_{ij}$ in these static models. These time periods represent the minimum amount of time that a vehicle takes to answer an emergency call and

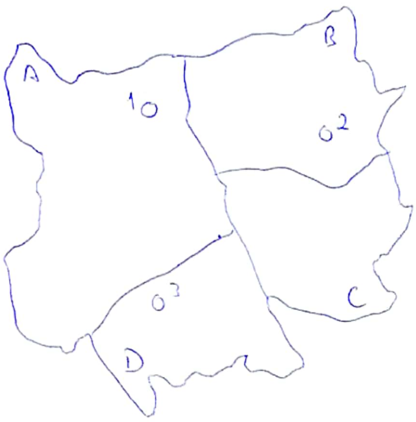

Figure 3: Example Map

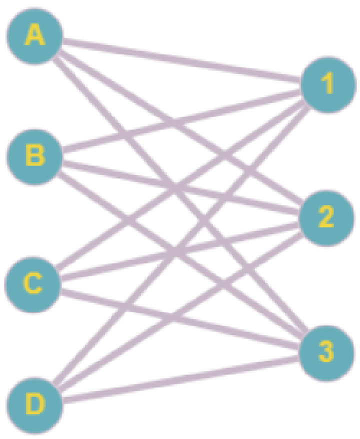

Figure 4: Static Model Simplified Graph

$$\min \sum_{j=1}^{M} x_{ij} t_{ij} \tag{1}$$

$$\text{subject to:} \sum x_{ij} \geq 1, \qquad i = 1, ..., n, \tag{2}$$

$$x_{ij} \in \{0, 1\}, \quad i = 1, ..., n, j = 1, ..., m. \tag{3}$$

Figure 5: Vehicle usage minimizing Model definition

$$\min \sum_{j=1}^{M} y_j \tag{4}$$

$$\text{subject to:} \sum y_j \geq n, \qquad j = 1, ..., m, \tag{5}$$

$$y_j \in \{0, 1\}, \qquad j = 1, ..., m. \tag{6}$$

Figure 6: Vehicle usage minimizing Model definition

go back to a standby site.

In order to better understand each goal, we are first going to formulate two separate model definitions that represent single objective models. In the first model, represented in Fig. 5, the goal is to minimize the travel time throughout all operations within a time frame. For simplicity, we consider only the parcel t_{ij} associated with each variable x_{ij}.

After this, we defined a model in which the goal was to minimize the number of vehicles being used. In order to do that we will define a variable y_j which will take the value of 1 if and only if a vehicle j is used to provide aid to an emergency. A reduction in the number of vehicles used can mean that that area in particular has a more vehicles than it should, and these can be relocated to areas that are struggling with more emergencies. Using this information, we can also cut our vehicle fleet if we realise that there are vehicles that are never used in the long term. This model definition is represented in Fig.6.

Since these two objectives are conflicting, meaning that a bigger number of vehicles used will result in a smaller value for the distance covered and vice-versa. Our goal is then to find a Pareto front, or a set of optimal solutions that minimizes both the number of vehicles used and the distance covered by those vehicles.

After both these models are defined, the Multi-Objective problem definition consists of joining both the single objective formulations into a single one, represented in Fig. 7.

As a result, we had to make several changes to

$$\min \sum_i \sum_j x_{ij} t_{ij} \tag{7}$$

$$\min \sum_{j=1}^{M} y_j \tag{8}$$

$$\text{subject to:} \sum x_{ij} \geq n, \quad i = 1, ..., n, j = 1, ..., m, \tag{9}$$

$$\sum y_j \leq 1, \quad j = 1, ..., m, \tag{10}$$

$$x_{ij} \in \{0, 1\}, \quad i = 1, ..., n, j = 1, ..., m. \tag{11}$$

$$y_j \in \{0, 1\}, \quad j = 1, ..., m. \tag{12}$$

Figure 7: Double Objective Model definition

both the way the data was being inserted in the programmatic generation file, as well as the generation process itself, since it was not accounting for a number of cases we had not initially considered, for example, the case where two vehicles were being used in subsequent time periods and therefore were unavailable when they were needed in these limited models where we did not have more vehicles than emergencies, causing the solver to deem this example as unsatisfiable.

With the initial model we were able to test some examples for small time periods where each vehicle could not be called to two distinct emergencies. However, in a real scenario, it is important that the vehicle can be dispatched to another emergency as soon as it is available again. In order to apply this concept to our model we divided a shift sized time frame into smaller fractions. This alteration meant that in our model we had to alter our set of constraints so that there was a variable for each vehicle for every one of the smaller time frames, represented by u_{ij}. Additionally, we also added a constraint that made sure that a vehicle could not be used in two consecutive time frames, effectively giving each vehicle at least one small time frame where it is unavailable between each dispatch. These time frames might differ in size depending on the situation and the value has to come from an analysis to the different time between dispatches in several similar real case scenarios.

In order to mathematically model this addition, we divided our large time period into a group of smaller time periods $E = E_1, ..., E_k$ where k represents the index of the period of time within a certain group E, which contains e time periods. This index k also serves the purpose of identifying a vehicle assignment u_{ijk}, which is a new variable that takes

$$\sum u_{ijk} \geq 1, \quad i = 1, ..., n, j = 1, ...m, k = 1, ..., e \tag{13}$$

$$u_{ijk} + u_{ij(k+1)} \leq 1, \quad i = 1, ...n, j = 1, ...m, k = 1, ..., e \tag{14}$$

Figure 8: Time Period Constraints

$$\sum u_{ijk} - y_j \geq 0, \quad i = 1, ...n, j = 1, ...m, k = 1, ..., e \tag{15}$$

Figure 9: New Time Minimization Expression

the value 1 if and only if the vehicle j provides aid to the emergency i in the time period k. As for the model itself, we adapted the previous version to account for the addition of time periods and we added a constraint that makes sure each vehicle can only be assign in a certain time period if it has not been assigned in the previous time period as seen in Fig.8.

Additionally, because of this alteration, we have also had to make changes to the expression we are trying to minimize that represents the usage of each vehicle. With these changes, if a vehicle is used in any period of time, the corresponding variable y_j should be equal to one. Hence we have added a constraint as seen in Fig.9

Emergencies are classified by a priority group that is associated with severity. INEM has an eight point priority rating system which can be converted into three major priority groups, due to a very negligible amount of emergencies in some of the categories in the eight point system. Depending on the priority value assigned to each emergency, we wanted to have the more commonly assigned vehicles for each of the three major priority groups be preferred when selecting the vehicle that is going to be dispatched. Therefore, we retrieved the type of each vehicle and the priority of each emergency, and imposed constraints that prevent certain sets of vehicle types to respond to certain emergencies depending on their priority.

Mathematically, this means that we will have three sets V_1, V_2, V_3 of vehicle types corresponding to the allowed vehicles to our three major priority groups $p1, p2, p3$, which will contain the priority level of each emergency p_{ij}. When building the model, we will only consider vehicles for a certain emergency if their vehicle type is contained within the set for the specific priority level of that emergency.

As for the model itself, we added a constraint that

$$u_{ijk} = 0, \qquad i \notin V_1, j \in E_1, k = 1, ..., e \qquad (16)$$

$$u_{ijk} = 0, \qquad i \notin V_2, j \in E_2, k = 1, ..., e \qquad (17)$$

$$u_{ijk} = 0, \qquad i \notin V_3, j \in E_3, k = 1, ..., e \qquad (18)$$

Figure 10: Vehicle Priority Constraint

represents the vehicle exclusivity described above, where only a certain group of vehicles can be assigned to a particular type of emergencies as seen in Fig.10.

After gradually building our model with by adding features to a base initial model, we ended up with a final model that covered all of the aspects we needed to be able to create instances for every situation we were going to be dealing with. All these different additions led us to have a model that aims to provide the best assignment for the available vehicles to the emergencies that occur in a determined amount of time, taking into account several relevant details about both the vehicles, namely their vehicle type and their availability as to not overuse a small number of vehicles, as well as the emergencies, namely their level of priority, the distance to each vehicle station, the number of vehicles needed and even taking into account district borders as to not have vehicles from different cities attend to emergencies in other cities if it is not expected.

Even though this model covers a variety of conditions, it does not scale well, meaning that the solver gets exponentially slower at coming up with optimum or even satisfiable results. Because of this, we decided to apply the model to 8 hour shifts, thus dividing the day in 3 equal parts, as a way of allowing the solver to be able to come up with solutions that can be applied in real situations in a reasonable amount of time.

4. Results

Upon having a reliable model built and tested, we started comparing the results of real case scenarios to applications of a solver algorithm to the instances we created of those same scenarios. However, after we analysed a small amount of these scenarios, and consistently getting a Pareto front of solutions that was better than the real case scenario, we wanted to try to infer some additional conclusions as well as come up with some possible solutions to help improve the efficiency of the ambulance assignment process for future situations.

Firstly, our main goal was to compare real case scenarios, analyse how many vehicles were used and how much distance they covered, and then compare these results with the Pareto front solutions from the application a solver to the instance we created under the same conditions. From this comparison, we expected to verify if the real case scenarios were sub-optimal, and try to quantify how much the solutions on our Pareto front were performing better than these scenarios. In order to do this we analysed cases with different periods of vehicle inactivity, namely thirthy minutes and one hour and different acceptable distances for a vehicle to provide aid to an emergency, namely ten and fifteen kilometers. We selected four separate days across three different years, namely the first day of February, May, August and December, from 2017, 2018 and 2019. Each of the days was divided in intervals of one, two, four and six hours, giving us a total of 46 instances created per day analysed. Cumulatively, we generated a total of 6624 instances, given that we generated all of the previously mentioned conditions in three different districts, namely Faro, Guarda and Lisboa.

Our goal with this analysis is to compare solutions between districts and seasons and try to understand whether or not we can allocate the vehicles in a different way in order to try to improve the overall performance in all of the places we are analysing.

We wanted to see how big the impact was for the overall solutions if we changed the time a vehicle becomes inactive after it provided aid to an emergency, therefore we tested using a thirty minute inactivity period and using a one hour inactivity period. We used examples with only a ten kilometer radius of activity for the vehicles of each emergency. These two values, thirty minutes and one hour, both came from the analysis of data from previous emergencies

As expected, the execution of the solver on the instances with the thirty minutes inactivity time took more time than the instances with the one hour inactivity time due to the amount of vehicles available for each emergency. This is especially true for larger examples. In Figure 12, we see an example from a two hour period in the district of Lisboa where in the first one we use a thirty minute vehicle inactivity period, represented by the green crosses, and a one hour vehicle inactivity period, represented by the red crosses. Since this is a larger example, it is the only one where the solutions differ when we change the inactivity time. If we look at Figures 13 and 14 on the other hand, the solutions found are exactly the same for periods of four hours, hence the green crosses and the red crosses are in the same positions. Although these are larger time periods, because of the difference in number of emergencies, the examples in Lisboa are far larger, and give a perspective of the impact the vehicle inactivity period has on the search of a solution.

Analogously, we wanted to test the impact of the radius of activity for the vehicles of each emergency,

$$\min \sum_{i} \sum_{j} u_{ijk} t_{ij} \tag{19}$$

$$\min \sum_{j=1}^{M} y_j \tag{20}$$

$$\begin{aligned}
\text{subject to:} \sum u_{ijk} &\geq n, & i = 1, ..., n, j = 1, ..., m, k = 1, ..., e \qquad (21)\\
u_{ijk} + u_{ij'(k+1)} &\leq 1, & i = 1, ...n, j = 1, ..., m, j' = 1, ..., m, k = 1, ..., e \qquad (22)\\
\sum u_{ijk} - y_j &\geq 0, & i = 1, ...n, j = 1, ..., m, k = 1, ..., e \qquad (23)\\
y_j - u_{ijk} &\geq 0, & i = 1, ...n, j = 1, ..., m, k = 1, ..., e \qquad (24)\\
u_{ijk} &= 0, & i \notin V_1, j \in E_1, k = 1, ..., e \qquad (25)\\
u_{ijk} &= 0, & i \notin V_2, j \in E_2, k = 1, ..., e \qquad (26)\\
u_{ijk} &= 0, & i \notin V_3, j \in E_3, k = 1, ..., e \qquad (27)\\
u_{ijk} &\in \{0, 1\}, & i = 1, ..., n, j = 1, ..., m. \qquad (28)\\
y_j &\in \{0, 1\}, & j = 1, ..., m. \qquad (29)
\end{aligned}$$

Figure 11: Double Objective Model definition

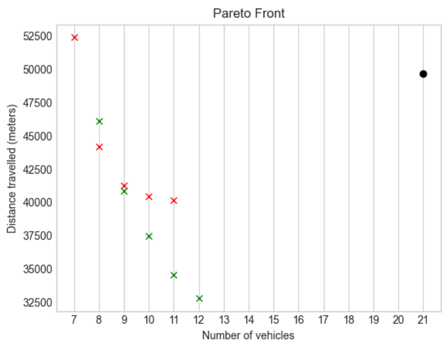

Figure 12: Comparison between inactivity times in Lisbon. The green crosses represent the thirty minute vehicle inactivity period and the red crosses represent the one hour vehicle inactivity period.

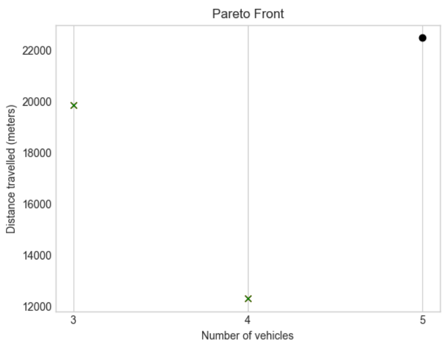

Figure 13: Comparison between inactivity times in Guarda. The green crosses represent the thirty minute vehicle inactivity period and the red crosses represent the one hour vehicle inactivity period.

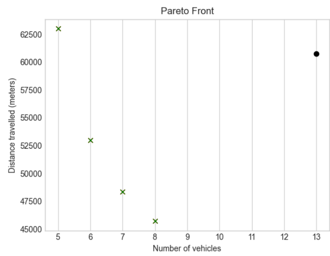

Figure 14: Comparison between inactivity times in Faro. The green crosses represent the thirty minute vehicle inactivity period and the red crosses represent the one hour vehicle inactivity period.

or in other words, how far away from the emergency a vehicle can be to be taken into account as a possible vehicle to provide aid to an emergency. When we first started this test we had three values for this distance, however, upon several tests, we decided to rule out the distance of five kilometers because a lot of emergencies were getting no available vehicles to provide aid to them. Therefore we ended up with the values of ten and fifteen kilometers and running the tests on these two values.

In some cases, making the radius of search bigger allows the solver to find more solutions that normally use less vehicles but increase the distance travelled by these vehicles overall. These additional solutions are usually cases in which the distance traveled is larger than the real case scenario, and therefore are not interesting when we have other solutions that can both reduce the number of vehicles used and the distance traveled. In Figure 15 we can see that the usage of fifteen kilometers adds

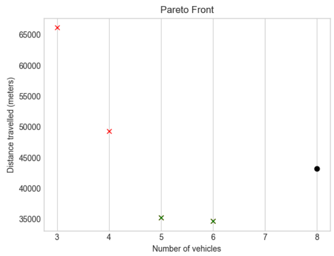

Figure 15: Comparison between distance times in Guarda. The green crosses refer to a ten kilometer radius and the red crosses refer to a fifteen kilometer radius

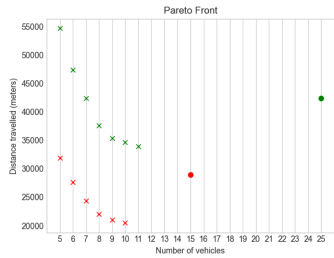

Figure 16: Seasonality comparison in Lisboa. The red figures refer to the month of August and the green figures refer to the month of January

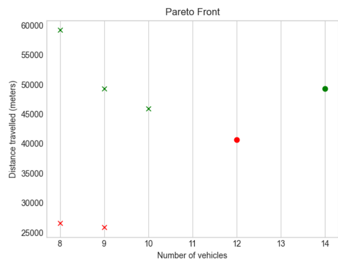

Figure 17: Seasonality comparison in Faro. The red figures refer to the month of August and the green figures refer to the month of January

two solutions, respectfully using three and four vehicles and combined distances of 49348 and 66176 kilometers, but these represent cases that we do not want to consider exactly because the two solutions that we already had from the ten kilometer radius case, respectfully using five and six vehicles and combined distances of 35239 and 34673 kilometers, reduce both the distance travelled and the number of vehicles used when compared to the real case scenario which used eight vehicle and with a combined distance of 43197 kilometers.

On top of this, using a fifteen kilometer radius also largely increases the execution time of the solver for the same scenarios, especially for larger examples. In most cases for Lisboa, the solutions found by the solver were not even close to being optimal as they were worse than the real case scenario. This happened because adding five kilometers means that there are a lot more vehicles to consider for each emergency, and since Lisboa has a larger number of emergencies than the Faro and Guarda and the complexity of this problem makes it grow exponentially fast, the solver could not provide acceptable solutions in the established amount of time of one hour.

Practically, this means that the ten kilometer range is more appropriate in most scenarios, since it grants enough vehicles to reach optimal solution and also because it reduces the universe of vehicles we would have to consider using, making the decision of which vehicle to assign less complex. We could not infer how the Emergency Medical System decided which vehicles were considered for each emergency from the data we had, however, if like us, they use a radius around the emergency, we would suggest the usage of a ten kilometer range as the standard from the three hypothesis we tested. Eventually, if the population density drops heavily there might be a need to use the fifteen kilometer range since there should be less vehicles in that area. Among our three districts, Guarda is the one with the least population density and we did not find any example time period in which there was a need for the fifteen kilometer range, so we infer this is a very unlikely scenario.

Provided that Portugal has a big affluence of tourists towards the South in the summer, we wanted to see if this had an influence on the number of occurrences when compared to other times of the year, especially between Faro and Lisboa.

The first thing to note about the examples shown in Figures 16 and 17, which represent periods of one hour and four hours from the same time period of the day in the year 2019, in Lisboa and Faro respectively, is that the seasonality is present in both districts. This means that both districts account for a consistently different number of emergencies during the month of August as opposed to the month of January. However, in Lisboa January is the month with more emergencies as opposed to Faro, where August is the month with more emergencies. We present only the results for the year 2019, because in the remaining two years of 2017 and 2018, the results were analogous.

We wanted to see if there was a possibility that these seasonal changes were not being addressed properly. In order to infer this, we looked ate the difference between the solutions in our Pareto front and the real case scenarios in both districts and we can see that in both districts, our real case scenario is closer to the optimal solutions in the month of August as opposed to the month of January. However, this difference is more evident in the district of Faro, which can mean that there are fewer vehicles available than there should be for this district at this time of year. This difference could be attenuated if some of the vehicles that are allocated to Lisboa in the month of August were reallocated temporarily to Faro, in an attempt to allow for a better response to a month where there are more emergencies at the expense of a slightly less optimal performance in that month in the district of Lisboa. In the months where the Faro has less emergencies, these vehicles would then be reallocated to Lisboa again as there will be a higher need for them there.

5. Conclusions

Firstly, when we first set out to do this project, the main goal was to be able to create a model that could accurately represent any situation in the context of the ambulance assignment problem and use it to determine what the optimal situations would be and compare them to the data we had from INEM. This objective was achieved successfully as we created a working model that can in fact represent every situation we have idealized.

Apart from the main goal, we discussed several possible smaller possibilities for further research using the model we had created. Ultimately, we ended up conducting an analysis on three different districts of Portugal and inferring information and possible adjustments to the current way the ambulances are being allocated by the Portuguese Emergency Medical System. This analysis required an analysis of the records we had access to in a variety of different measures like number of emergencies, number of vehicles available for each emergency, emergency priority levels, specific vehicle information, as well as an extensive stage of creation of instances followed by the solving of these same instances. This allowed us to contextualize each situation and be able to critically analyse and compare the real case scenarios and the solutions we obtained from running our examples through the solver in order to come up with possible solutions for less optimal situations we encountered.

In the development process of our algorithm we always considered a given situation in which the total number of vehicles available as well as total number of emergencies were known variables. This means that we have privileged insight when looking for a solution that the operators who choose which vehicles to assign to each emergency. With this in mind, it would be interesting to develop a real-time vehicle assignment mechanism that would better emulate the situation in which the operators have to make the assignment decisions. Afterwards, it would be possible to compare both performances to the solution given by the model developed in this paper. Furthermore, that real-time vehicle assignment mechanism could then benefit from information retrieved from various solutions given by our model like identifying periods of time more prone to a large number of emergencies in a certain area, and trying to preserve more vehicles in that area, even if at the cost of a more lengthy assignment on some other emergency beforehand.

Furthermore, we have talked about using a predictive model to predict possible future emergencies and then apply our model to allocate a higher amount of vehicles in the zones we predict are going to have more emergencies in a certain time period, and leaving zone we deem to be less likely to have emergencies with a smaller amount of vehicles. This would serve as an attempt at optimizing the performance of the Emergency Medical System even further, focusing on probabilities of what will happen in the future, and not only on data from what has happened in the past.

References

[1] L. Aboueljinane, E. Sahin, and Z. Jemai. A review on simulation models applied to emergency medical service operations. *Computers & Industrial Engineering*, 66(4):734 – 750, 2013.

[2] V. Bélanger, Y. Kergosien, A. B. Ruiz, and P. Soriano. An empirical comparison of relocation strategies in real-time ambulance fleet management. *Comput. Ind. Eng.*, 94:216–229, 2016.

[3] V. Bélanger, A. B. Ruiz, and P. Soriano. Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. *Eur. J. Oper. Res.*, 272(1):1–23, 2019.

[4] P. Beraldi and M. E. Bruni. A probabilistic model applied to emergency service vehicle location. *Eur. J. Oper. Res.*, 196(1):323–331, 2009.

[5] H. Billhardt, M. Lujak, V. Sánchez-Brunete, A. Fernández, and S. Ossowski. Dynamic coordination of ambulances for emergency medical assistance services. *Knowl. Based Syst.*, 70:268–280, 2014.

[6] R. Church and C. ReVelle. The maximal covering location problem. *Papers of the Regional Science Association*, 32(1):101–118, 1974. cited By 1537.

[7] M. Gendreau, G. Laporte, and F. Semet. Solving an ambulance location model by tabu search. *Location Science*, 5(2):75 – 88, 1997.

[8] S. Ibri, M. Nourelfath, and H. Drias. A multi-agent approach for integrated emergency vehicle dispatching and covering problem. *Eng. Appl. Artif. Intell.*, 25(3):554–565, 2012.

[9] G. Laporte, F. V. Louveaux, F. Semet, and A. Thirion. Application of the double standard model for ambulance location. In J. A. Nunen, M. G. Speranza, and L. Bertazzi, editors, *Innovations in Distribution Logistics*, pages 235–249, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

[10] B. López, B. Innocenti, and D. Busquets. A multiagent system for coordinating ambulances for emergency medical services. *IEEE Intelligent Systems*, 23(5):50–57, 2008.