




INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Increasing the Responsiveness of Firefighter Services by Relocating Base Stations in Amsterdam

Pieter L. van den Berg, Guido A. G. Legemaate, Rob D. van der Mei

To cite this article:

Pieter L. van den Berg, Guido A. G. Legemaate, Rob D. van der Mei (2017) Increasing the Responsiveness of Firefighter Services by Relocating Base Stations in Amsterdam. INFORMS Journal on Applied Analytics 47(4):352-361. <https://doi.org/10.1287/inte.2017.0897>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Increasing the Responsiveness of Firefighter Services by Relocating Base Stations in Amsterdam

Pieter L. van den Berg,^a Guido A. G. Legemaate,^b Rob D. van der Mei^c

^a Rotterdam School of Management, Erasmus University, Rotterdam 3062 PA, Netherlands; ^b Brandweer Amsterdam-Amstelland, Amsterdam 1101 CK, Netherlands; ^c Centrum Wiskunde en Informatica, Amsterdam 1098 XG, Netherlands

Contact: vandenberg@rsm.nl,  <http://orcid.org/0000-0002-8550-6769> (PLvdB); g.legemaate@brandweeraa.nl (GAGL); mei@cw.nl (RDvdM)

Received: December 9, 2015

Revised: June 14, 2016; December 19, 2016; January 12, 2017

Accepted: March 24, 2017

Published Online in Articles in Advance: May 19, 2017

<https://doi.org/10.1287/inte.2017.0897>

Copyright: © 2017 INFORMS

Abstract. In life-threatening situations where every second counts, the timely presence of firefighter services can make the difference between survival and death. Motivated by this, in collaboration with Fire Department Amsterdam-Amstelland in the Netherlands, we developed a mathematical programming model for determining the optimal locations of the vehicle base stations, and for optimally distributing firefighter vehicle types over the base stations. The model is driven by practical considerations. It (1) allows for fixing any subset of existing base locations that cannot be relocated (e.g., for historical reasons); (2) includes multiple vehicle types, each of which may have a type-dependent response-time target; and (3) includes crews that consist of arbitrary mixtures of professional (i.e., career) and volunteer firefighters. Extensive analysis of a large data set obtained from the Fire Department Amsterdam-Amstelland demonstrates: (1) that a reduction of over 50 percent in the fraction of firefighter late arrivals can be realized by relocating only three of the current 19 base locations; and (2) that adding new base locations to improve performance is unnecessary: optimization of the locations of the current base stations is as effective, and saves money. The results show an enormous potential for substantially reducing the fraction of late arrivals of firefighter services, with little investment in relocating a small number of stations.

History: This paper was refereed.

Funding: This research was financed in part by Technology Foundation STW [Contract 11986], which the authors gratefully acknowledge.

Keywords: firefighters • location model • integer linear programming

In 1874, the Dutch capital of Amsterdam was the first city in the Netherlands with a professional firefighter service. With 144 people and nine fire stations covering 30 square kilometers, it ensured fire-protection safety for approximately 285,000 inhabitants. Today, it is the regionally organized Amsterdam-Amstelland fire department, which has 1,150 people and 19 fire stations covering 354 square kilometers, and is responsible for over 1,000,000 inhabitants (Brandweer 2011). Over time, new requirements and means for fire-protection safety emerged; thus, the questions of how many fire stations were needed for a given coverage and where they should be located had to be answered numerous times.

In this study, we introduce a model to determine optimal locations for firefighter vehicles and an

optimal distribution of the vehicles over the selected locations. As bases arose within a growing city, the locations were selected to match the needs at the time. This local approach does not necessarily result in bases that are appropriate for the current shape of the city. We consider a greenfield scenario (i.e., a scenario in which we do not consider the current base structure) to determine the optimal configuration if we could select the bases from scratch. This gives insights that can be considered in long-term decision making. For practical purposes, however, it is more reasonable to consider only small changes in the configuration. To that end, we evaluate the impact of changing only a few of the base locations.

For the model, we have two main, but potentially conflicting, requirements. It must be relatively easy to

implement; conversely, it must capture the dynamics of the fire department. For example, it should be able to deal with multiple types of vehicles with different response-time targets. A given vehicle type might have different response-time targets within a region. Additionally, because most fire departments, including Amsterdam-Amstelland, have both a voluntary and a professional crew, the model must be able to incorporate both.

As a location problem, the placement of fire stations and, more generally, the planning of emergency services, have been extensively studied by scholars. Since the 1970s, models have been developed. Some focus on coverage requirements or on maximizing coverage (Ingolfsson et al. 2008 and Chong et al. 2016); others on response times (Dzator and Dzator 2013) and (or) survival probabilities (Erkut et al. 2008, McLay and Mayorga 2010). Most of these studies focus on the location of ambulances. Ahmadi-Javid et al. (2017) presented an overview of healthcare facility location models. Aringhieri et al. (2017) and Reuter-Oppermann et al. (2017) give a broader overview of planning problems that arise in emergency medical service systems. Hogg (1968) introduced one of the first firefighter-specific models. In this model, the set of locations of fire stations is determined to minimize both the losses from fire and the cost of providing the service. Building on this and on the detailed study of Toregas et al. (1971), Plane and Hendrick (1977) use response time as a standard for coverage and apply a location set-covering problem (LSCP) optimization model. The amount of resources needed (i.e., firemen or fire apparatus) is included as a decision rule (Swersey 1982), expanding the work on square root laws for fire engine response distances (Kolesar and Blum 1973). Later, Batta and Mannur (1990) also optimize the number of resources sent to an emergency; however, they consider only one type of fire apparatus. The model developed in Andersson and Särndqvist (2007) allows for combinations of multiple resource types and multiple event types. In 1974, Church and ReVelle (1974) introduced the maximal coverage location problem (MCLP), which Schilling et al. (1980) apply to fire-protection location decisions. Murray and Tong (2009) and Chevalier et al. (2012) use variants of the MCLP, and further improve the model by including a risk-modeling approach to estimate demand.

In this study, we determine the optimal number of base locations, their optimal geographical location, and the optimal distribution of vehicles of various types over the selected locations based on historical demand. The model provides sufficient flexibility to capture the dynamic aspects of fire departments. For example, we permit different response-time targets for different demand points and vehicle types. We can model risk factors by assigning weights to different parts of the regions. Additionally, the model can distinguish between voluntary and professional crews. This distinction is typical for fire departments, but is uncommon in other emergency response services. We maximize the coverage and minimize the number of locations that are used to provide that coverage by penalizing the use of a potential location. In addition to the insights obtained from a greenfield scenario, our practical contribution is in the analysis of small changes to the vehicle configuration, which has already significantly improved responsiveness.

First, we introduce the initial version of the model, in which we allow multiple types of vehicles. In this model, we consider only professional firefighter crews. Then, we discuss the results of this model for the Amsterdam-Amstelland region and the data we use in the case study. Subsequently, we introduce a model extension that includes different types of crews. The main application of this extension is the distinction between volunteer and professional firefighters, which resulted in different response times. This section also includes the results of this extended model. Finally, we present the main conclusions and recommendations.

Model Description

As is common in location models, we divide the region into a set of N demand points from which fire emergency calls can arise. We also define a set of M locations, any of which can be used as a potential base location. Typically, M is a subset of the demand points. Because the model considers multiple types of vehicles that are used to serve different types of calls, we define a set of vehicle types, denoted by K . For each combination of a demand point i and a vehicle type k , we have a response-time target. This target gives the maximum allowed response time for a vehicle of type k to reach demand point i to cover the demand point. Our model differs from other location models, because we allow

the target to deviate for each demand point and each vehicle type. Thus, we can set a different target for each part of the region. To compute the coverage, we also need the travel times between potential base locations and demand points, and the pre-trip delay. The pre-trip delay is the elapsed time before the vehicle begins to drive to the scene of the emergency. It consists of two parts: (1) triage and dispatch, and (2) chute time. The triage and dispatch time is the time spent within the call center to assess the importance of the call and assign a vehicle. The chute time is the time from the assignment of a call until a vehicle departs from the base. For now, we assume that this delay is fixed. In our computations, the travel times are assumed to be independent of the vehicle type; however, the model can handle vehicle-type-dependent travel times.

The objective of the model is to maximize the sum of the coverage that is provided for the different vehicle types. A demand point is said to be covered if the response time of the closest vehicle of the appropriate type is, at most, the given response-time target. Because the call intensity is typically not distributed evenly throughout a region, we add weights to the demand points. These weights should indicate the importance of covering demand point i with vehicle type k . We commonly use the expected number of calls for the weights; however, we could include other risk measures. The model does not consider backup coverage in the objective.

In our model, we include a fixed number of available vehicles of each type. To avoid the situation in which too many bases are opened, we add a penalty on the number of opened bases. This penalty defines a trade-off between the additional coverage of a new base and the cost of opening the base. The value of the penalty should be interpreted as the number of previously uncovered calls that a new base should cover to make opening the new base beneficial. In practice, however, determining this value may be difficult. In such a case, we can exclude this part from the objective function and add a constraint on the number of bases to open. We can compute the optimal coverage for multiple numbers of bases, thus making the choice easier for the decision maker. The decision maker's choice implies a penalty value that can be used in future calculations.

Note that we do not include backup coverage in the models. This is in contrast to many other studies in the literature, especially those dealing with ambulance modeling. In these studies, the backup coverage is often included by considering ambulance unavailability (Daskin 1983, Hogan and ReVelle 1986, Gendreau et al. 1997). Since the call volumes for fire departments are significantly lower than for ambulance providers, overlapping calls are less common. They occur, but not at an alarming rate. Data analysis we conducted showed that the Amsterdam-Amstelland fire department handles an average of 52 calls per day with its 19 base locations and 34 vehicles. Furthermore, we do not include all vehicles currently in use in our computations; thus, these other vehicles can be used to provide backup coverage. In 2015, we found 572 occurrences of overlapping calls that share the closest base station. In 242 of these cases, the involved base station was one of the stations in which one of the additional vehicles is located. With an average of 19,000 calls per year, less than two percent of the calls overlapped. As a result of our focus on single coverage, our model will never locate two vehicles of the same type at the same base station. For computational convenience, we add this as a constraint in the model.

We formulate the model as an integer linear programming (ILP) problem. This type of problem can generally not be solved in polynomial time. However, commercial solvers like CPLEX (ILOG 2009) can provide an optimal solution in reasonable time for instances with realistic sizes. Therefore, in this paper, we use CPLEX to solve the problem. The appendix includes a complete mathematical description of the model.

Results

Data Description

To apply the model to the Amsterdam region in the Netherlands, we must determine the appropriate data. In close cooperation with the Amsterdam-Amstelland fire department, we defined a set of 2,643 demand points, which corresponds to the sections currently in use. This includes some areas of neighboring regions at the borders. We assume that a base can be located in most demand points that are part of the region. We exclude demand points in neighboring regions and demand points that contain only highways. This gives

us 2,223 potential base locations. The fire department provided us with the travel times between potential base locations and demand points; these times are based on estimated travel times on the road network between each location. For every demand point, we calculated the travel time to every other demand point by using the fastest route from the road closest to the centroid of the start polygon to the center of the ending polygon, using the given road network. The driving speed on the roads is based on guidelines set by the national firefighters' organization for a fire truck driving with lights and sirens.

In our analysis, we include the four most common types of vehicles used at Dutch fire departments: fire apparatus (FA), aerial apparatus (AA), rescue apparatus (RA), and marine rescue units (MR). The number of available vehicles of each type is 22, 9, 3, and 2, respectively. The current configuration has 19 bases. Since the objective does not benefit from backup coverage, three of 22 FA vehicles do not contribute to the coverage in the current situation. Therefore, we do include these three vehicles in our computations. Thus, the number of FA vehicles in the computation is 19, and the other three vehicles can be used as backup in the rare occurrence of overlapping calls. For each demand point and for each vehicle type, we must define the weight. We take the absolute number of calls per vehicle that occurred between 2011 and 2015. In this period, 93,975 calls were registered. The number of calls per vehicle type is 70,022, 20,433, 1,842, and 1,678, respectively.

Since different vehicle types are used to serve different kinds of calls, a vehicle type may also have different response-time targets. Dutch law states a response-time target for both FA and AA vehicles. These targets depend on the type and function of a building, and vary between 6 and 10 minutes for FA vehicles, and between 6 and 15 minutes for AA vehicles. Based on these requirements, we set a response-time target for each demand point for FA and AA vehicles. For RA and MR vehicles, no requirements are set by law. For these vehicle types, we set the response time requirement to 15 minutes for all demand points. This differs in one aspect from the current practice at fire departments in the Netherlands. Currently, coverage targets

Table 1. This Table Gives the Available Number of Vehicles, the Number of Calls, and the Response-Time Targets for Four Vehicle Types

Vehicle type	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit
No. of vehicles	19	9	3	2
No. of calls	70,022	20,433	1,842	1,678
Minimum target	6	6	15	15
Average target	7.98	14.68	15	15
Maximum target	10	15	15	15

are set for each building, rather than each area. Therefore, one building could have a five-minute response-time target, while a neighboring building has a 10-minute target. We define the response-time target for each demand point; thus, this will deviate from the real response-time target for some buildings. However, the Dutch government will soon change its coverage definition to adopt an area-based target. Therefore, we changed our model to accommodate this new approach. Since only travel times are known and the response time also includes the time spent before the actual driving (of the fire truck) begins, we add three minutes to the travel time to correct for the delay prior to the trip. Table 1 provides a summary of the data.

Results

Given the model and the data described above, we can compute the optimal set of base locations and the optimal distribution of the available vehicles over the selected bases. Finding an appropriate penalty value for the number of opened bases is difficult; therefore, we will fix the maximum number of bases and compute the corresponding maximum coverage. The decision maker can determine the number of bases to open. This choice defines a penalty value that can be used in future computations, ensuring that we do not have to compute all instances again.

Table 2 shows the results for a maximum number of bases, which vary from 9 to 19. If no restriction is imposed on the number of bases, 28 bases will be opened and coverage of 98.18 percent will be achieved. Currently, 19 bases are in use in the region we are considering. The last two rows show the coverage achieved by the current set of bases. The second-to-last row gives the coverage after optimally distributing the vehicles

Table 2. In the Greenfield Scenario, the Coverage, Shown as a Fraction of the Total Number of Calls, for the Various Vehicle Types Increases as the Number of Bases Increases

No. of bases	Coverage				Total
	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit	
9	0.8549	0.9685	0.9723	0.9130	0.8830
10	0.8862	0.9817	0.9843	0.9076	0.9093
11	0.9088	0.9818	0.9843	0.9076	0.9262
12	0.9239	0.9964	0.9821	0.9088	0.9405
13	0.9370	0.9964	0.9821	0.9082	0.9503
14	0.9482	0.9965	0.9843	0.8915	0.9584
15	0.9580	0.9964	0.9832	0.8826	0.9655
16	0.9665	0.9965	0.9832	0.8826	0.9719
17	0.9706	0.9981	0.9848	0.8862	0.9754
18	0.9744	0.9975	0.9805	0.8969	0.9781
19	0.9771	0.9982	0.9875	0.9029	0.9805
Unlimited	0.9773	0.9998	0.9935	0.9368	0.9818
Current bases	0.8234	0.9606	0.9571	0.8826	0.8569
Current distribution	0.8234	0.9313	0.8768	0.8492	0.8484

Notes. The second-to-last row shows the coverage of the optimal distribution of the vehicles over the current bases. The last row gives the coverage of the current distribution of the vehicles.

over the current bases, while the last row gives the coverage of the current distribution. We see that a coverage increase of more than 0.85 percent can be achieved without changing the bases—that is, by changing only the assignment of vehicles to bases. Note that the coverage for FA vehicles does not increase, because we have sufficient FA vehicles to locate one at every base.

Table 2 further shows that a high coverage can be achieved with the current number of bases. By replacing all 19 bases, we can achieve coverage of 98.05 percent. The coverage provided by the current set of bases can be achieved with only nine optimally located bases. Note that if the number of bases is smaller than the number of vehicles of a particular type, the optimal solution does not use all vehicles, because single coverage suffices for complete coverage. We further observe that increasing the number of bases does not yield a significant coverage improvement. If no limit is set to the number of bases, 28 bases are used and the coverage increases by only 0.0013 percentage points compared to the case with 19 bases. One surprising observation is that the coverage of the RA and MR vehicles can be improved significantly. Since only three or two vehicles, respectively, are available, we would expect that the current 19 bases would contain a good tuple,

Table 3. A Significant Improvement in Coverage Can be Achieved with a Limited Number of Base Changes

No. of changes	Coverage				Total
	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit	
0	0.8234	0.9606	0.9571	0.8826	0.8569
1	0.8544	0.9921	0.9571	0.8826	0.8868
2	0.8911	0.9921	0.9571	0.8963	0.9144
3	0.9119	0.9921	0.9587	0.8886	0.9299
4	0.9229	0.9921	0.9587	0.8838	0.9380
5	0.9320	0.9910	0.9587	0.8838	0.9445
6	0.9406	0.9917	0.9843	0.8838	0.9515
7	0.9491	0.9925	0.9864	0.8880	0.9582
8	0.9566	0.9918	0.9864	0.8880	0.9636
Unlimited	0.9771	0.9982	0.9875	0.9029	0.9805

Note. The marginal improvement decreases as the number of changes increases.

or pair, of bases for these types. However, this is not the case, because we can obtain a coverage increase of 3.64 and 5.42 percent, respectively, for these types of vehicles.

Table 2 shows that significant improvement can be obtained by changing the base locations. In our calculations, we do not consider the current set of base locations. However, this is not realistic, because changing all locations would result in high costs. It is, therefore, interesting to see the improvements we can attain if we limit the number of changes. Table 3 shows the maximum coverage attainable when we allow only a limited number of bases to be changed. Note that the total number of bases remains fixed at 19.

We see that changing only one base yields a coverage increase of 2.99 percent. The next two changes also result in an increase of more than one percent. Although each change adds coverage, the marginal benefits decrease rapidly. If we do not limit the number of changes, we achieve coverage with 19 bases, as we report in Table 2. Another observation we can make is that with the first five changes, the focus is on FA and AA vehicles—the ones with highest call volume. Despite the large potential improvement for RA and MR vehicles, RA coverage does not significantly increase as a result of the first five changes. The sixth change is the first change that focuses on this vehicle type.

In contrast to many papers in the literature, we allow for response-time targets based on the demand points.

Table 4. The Coverage Obtained with Respect to the Real Response-Time Targets Is Significantly Lower When One Does Not Consider Any Differentiation in the Response-Time Targets

Target used	Estimated coverage	Real coverage
Minimum	0.7503	0.9248
Average	0.9852	0.9000
Maximum	0.9973	0.7783
Real	0.9805	0.9805

Notes. The estimated coverage gives the overall coverage with respect to the targets used in the optimization. The real coverage gives the coverage with respect to the real response-time target for each demand point.

This allows us to make region-specific adjustments in the model, for example, to distinguish the 17th-century canal ring area of the Amsterdam city center (a UNESCO World Heritage-listed site with specific fire risks) from the rest of the region, which has other types of buildings and (or) infrastructure. To quantify the impact of considering this differentiation, we compare the solutions of three instances with the results we previously present. We compute the optimal solution with respect to the minimum, average, and maximum response-time targets for all demand points, and evaluate the resulting solutions with respect to the real response-time targets. This provides insights into the importance of incorporating multiple response times into the model. For each of the four instances, Table 4 shows the coverage based on the response-time targets considered and the coverage based on the real targets. The values used for the minimum, maximum, and average targets are shown in Table 1.

The results show that in all three cases, we lose coverage as a consequence of not considering the variability in response-time targets in the optimization. The impact is smallest when we use the minimum response-time target for all demand points, because this considers the worst-case perspective. However, in this case, the estimation of the coverage is too conservative. In particular, when we optimize using the maximum response-time targets, we obtain poor solutions.

Extending the Crew Types in the Model

One aspect common to fire departments, which we have not yet incorporated into our model, is the distinction between voluntary firefighters and professional

firefighters. As a consequence of the low call volumes for firefighters, staffing some vehicles with a voluntary (i.e., on-call) brigade is often efficient. A voluntary crew is not located at the base, but is called only in an emergency. Clearly, this approach increases the chute time, which is part of the delay encountered before the crew begins to drive to the emergency. Staffing bases with voluntary brigades significantly reduces the cost. Thus, decision makers must make a trade-off between cost and efficiency. Therefore, we developed an adapted version of the model to assist decision makers with this decision. Note that since the model assumes that single coverage suffices, choices between voluntary and professional brigades are based solely on coverage, and not on workload.

To incorporate the different types of crews in a general way, we introduce a set of different crew types, denoted by L . In our experiments, L contains two types. For each type of crew, we define the pre-trip delay. This delay consists of a triage and dispatch time, which we assume to be independent of the type of crew, and a crew-dependent chute time. In the model, we now must determine for each vehicle whether it will be staffed with a professional or a voluntary crew. Because professional crews are more costly, the number of crews of this type has a limit, which we can implement in various ways. For example, we could introduce a crew budget and costs for crews of different types. Then, we could allow the model to determine the best way to spend the budget. However, we currently fix the total number of vehicles, and crews only have value if assigned to a vehicle; therefore, the total number of crews is fixed. We decided to fix the number of crews of each type. The appendix contains the complete formulation of our extended model.

Results of the Extended Model

We apply the extended model to the same data as the original model. As is the case in the current execution, we allow for 27 professional and six voluntary crews. Currently, all six voluntary brigades operate on an FA vehicle. For both the professional and the voluntary brigades, we must set the pre-trip delay. For the professional staff, we set this delay at three minutes, as in our previous experiments. For the voluntary staff, we set the delay at six minutes. Note that more voluntary personnel than necessary are often called. In such a case,

Table 5. In the Greenfield Scenario, with Only Nine Optimally Located Bases, We Obtain a Higher Coverage Than with the Current 19 Bases

No. of bases	Coverage				Total
	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit	
9	0.8549	0.9685	0.9723	0.9130	0.8830
10	0.8862	0.9817	0.9843	0.9076	0.9093
11	0.9088	0.9818	0.9843	0.9076	0.9262
12	0.9239	0.9964	0.9821	0.9088	0.9405
13	0.9370	0.9964	0.9821	0.9082	0.9503
14	0.9478	0.9963	0.9821	0.9082	0.9583
15	0.9580	0.9961	0.9832	0.8826	0.9654
16	0.9665	0.9961	0.9832	0.8826	0.9718
17	0.9706	0.9976	0.9848	0.8743	0.9750
18	0.9742	0.9966	0.9864	0.8743	0.9775
19	0.9768	0.9959	0.9875	0.9035	0.9798
Unlimited	0.9773	0.9988	0.9848	0.9368	0.9814
Current	0.8231	0.9590	0.9452	0.8826	0.8561

Note. In all cases reported, the vehicles can be freely distributed over the selected bases, and a maximum of 27 professional crews are used.

the crew can depart when sufficient crew members have arrived at the base. First, we consider a greenfield scenario in which we do not consider the current bases. Second, we fix the current bases and evaluate the performance of the various crew configurations. Finally, we include the current bases, but allow for a limited number of base changes.

Greenfield Scenario. In this experiment, we ignore the current bases and evaluate the maximum performance with the given vehicles, and 27 professional and six voluntary crews to change the total number of bases. We consider all cases from 9 to 19 bases, which is the current number of bases. We further consider the case in which we set no limit on the number of bases. Note that we again deviate slightly from the model formulation by not including the penalty term for each base, but adding a constraint on the maximum number of bases. Table 5 shows the results; the last row shows the performance of the best distribution of vehicles and crews over the current bases.

As with the original model, we see that there is an enormous potential to improve the responsiveness by a better distribution of bases over the region. Surprisingly, staffing six vehicles with voluntary crews does not have a significant impact on the coverage. Given the current set of bases, this reduces coverage by only 0.08 percent, compared to an all-professional crew; that

is, we compare Table 5 with Table 2.

Fixed Set of Bases. To better understand of the impact of voluntary crews on the coverage, we evaluate the impact of different crew configurations, given the current set of bases. First, all vehicles can be optimally distributed over the bases, while fixing the voluntary crews to their base and vehicle type. Second, the crews can be switched for vehicles of the same type. We currently have only voluntary crews on FA vehicles; therefore, the coverage for the three other vehicle types will be the same as in the previous experiment, because all vehicles are staffed with professionals. Third, all crews and vehicles can be distributed freely over the existing bases. The number of crews of each type remains fixed. In the final case with a fixed set of bases, we evaluate the impact of replacing voluntary crew members by professionals and vice versa. We include all cases up to three changes.

Table 6 shows that 0.84 percent of coverage for FA vehicles can be gained by reassigning crew members over the FA vehicles. Only one vehicle is staffed differently. If we also allow the assignment of crew to different vehicle types, the overall coverage increases by 0.65 percent from 84.96 to 85.61 percent. In this case, four of the nine AA vehicles, and one of the three RA vehicles, are operated by voluntary staff. This leads to a coverage decrease of 0.16 percent for AA vehicles and 1.19 percent for RA vehicles, whereas the FA vehicle coverage increases by 0.95 percent. The last six rows of the table show that the effect of replacing voluntary staff by professionals is limited, given the current set of bases. Replacing three professional crews by volunteers reduces coverage by only 0.13 percent.

Limited Number of Base Changes. Finally, we analyze the impact of changing some of the existing bases based on the coverage provided by the solution. We fix the number of voluntary crews and the number of professional crews, but allow for redistribution of the crews over bases and vehicles types. We consider the cases of adding one, two, three, or four bases and replacing one, two, three, or four bases. The results are shown in Table 7.

Surprisingly, we observe only a slight difference between adding and replacing bases. For one and two changes, closing the same number of bases does not reduce coverage. This indicates that some of the current bases are not positioned adequately. We further

Table 6. In the Extended Model with a Fixed Set of Bases, the Coverage Can be Improved by Reconfiguring the Crew

No. of crew changes	Coverage				
	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit	Total
Crews fixed	0.8052	0.9606	0.9571	0.8826	0.8433
Crews fixed to vehicles	0.8136	0.9606	0.9571	0.8826	0.8496
Free assignment of crew	0.8231	0.9590	0.9452	0.8826	0.8561
Add three professional crews	0.8234	0.9598	0.9571	0.8826	0.8567
Add two professional crews	0.8234	0.9590	0.9571	0.8826	0.8566
Add one professional crew	0.8234	0.9590	0.9452	0.8826	0.8563
Add one voluntary crew	0.8227	0.9590	0.9452	0.8826	0.8558
Add two voluntary crews	0.8227	0.9569	0.9452	0.8826	0.8553
Add three voluntary crews	0.8220	0.9569	0.9452	0.8826	0.8548

Notes. In the base case, 27 professional and six voluntary crews are available. When changing the number of professional or voluntary crews, the total number of crews remains fixed at 33, which is the total number of vehicles.

Table 7. Based on the Results of the Extended Model, Replacing Some of the 19 Current Bases Gives a Coverage Improvement That Is Similar to Adding the Same Number of Bases

Base changes	Coverage				Total
	Fire apparatus	Aerial apparatus	Rescue apparatus	Marine rescue unit	
Add one	0.8540	0.9902	0.9452	0.8826	0.8859
Add two	0.8907	0.9902	0.9452	0.8963	0.9135
Add three	0.9119	0.9884	0.9452	0.8886	0.9288
Add four	0.9227	0.9884	0.9549	0.8886	0.9370
Change one	0.8540	0.9902	0.9452	0.8826	0.8859
Change two	0.8907	0.9902	0.9452	0.8963	0.9135
Change three	0.9117	0.9883	0.9549	0.8886	0.9288
Change four	0.9227	0.9903	0.9549	0.8617	0.9369
No changes	0.8231	0.9590	0.9452	0.8826	0.8561
Unlimited	0.9768	0.9959	0.9875	0.9035	0.9798


Note. In all cases, the vehicles can be distributed freely over the bases and 27 professional and six voluntary crews are available.

see that only a few changes lead to a significant coverage improvement. For example, three changes yield a 6.29 percent coverage increase. Figure 1 shows the movements of the bases when we change three bases.

Implementation and Benefits

On numerous occasions, we have presented the results of this study to groups within the Amsterdam-Amstelland fire department and to Firefighters Netherlands, the national organization of firefighters in the Netherlands. The results have provided insights for these organizations to consider as they make decisions

Figure 1. (Color online) We Show the Distribution of Fire Stations When We Move Three Stations



Notes. Empty circles correspond to the current bases that we replace. Fire symbols correspond to the location of existing base stations. Encircled fire symbols correspond to the locations of new bases. The black line shows the border of the region for which the Amsterdam-Amstelland fire brigade is responsible.

in the future. In the fall of 2015, the results supported a decision not to replace one of the fire stations in the region. Our results quantified the very limited loss in coverage after removal of this station.

The model is currently being used in two other projects. In the first, it is computing the best location for two new fire stations, which will replace one fire station. Our study demonstrated the limited coverage that the current fire station provides. Second, the model is determining the best distribution of the MR vehicles, incorporating the availability of MR vehicles in neighboring regions. In a follow-up study, we intend to adapt the model to compute relocation scenarios in cases of very large fires, for which multiple fire trucks are dispatched.

Conclusions and Recommendations

The results of this study show that there is enormous potential to improve the responsiveness of the of the Amsterdam-Amstelland fire department. Even with very modest changes, the coverage can improve significantly. For example, without changing any bases, but changing only the vehicle distribution, we can achieve a coverage improvement of 0.85 percent. If we further permit reassignment of crew, without changing the number of crews of each type, we obtain an additional 1.28 percent improvement. Finally, if we allow up to three base changes, we can achieve another 3.0 percent, 5.7 percent, or 7.2 percent coverage increase. Note that for these changes, neither the number of bases nor the number of vehicles must increase. Making these changes would increase coverage from 83.48 to 92.88 percent. This corresponds to a reduction of more than 50 percent of calls for which the vehicle arrives late.

Appendix

Formulation of the Initial Model

Input

- N Set of demand locations;
- M Set of potential base locations;
- K Set of vehicle types;
- c_k Number of available vehicles of type $k \in K$;
- d_{ik} Demand of demand point $i \in N$ for calls of type $k \in K$;
- r_{ik} Response time target for demand point $i \in N$ for calls of type $k \in K$;
- t_{ij} Travel time from location $j \in M$ to demand point $i \in N$;
- τ Pre-trip delay;
- β Penalty on number of opened base locations;
- M_{ik} Set of base locations that can cover demand point $i \in N$ with vehicle type $k \in K = \{j \in M: t_{ij} + \tau \leq r_{ik}\}$.

Variables

- x_{jk} Number of type k vehicles that are located at potential base location j ;
- y_{ik} Binary variable indicating whether demand point i is covered by a vehicle of type k ;
- z_j Binary variable indicating whether at least one vehicle is located at base location j .

Model

$$\max \left\{ \sum_{i \in N} \sum_{k \in K} d_{ik} y_{ik} - \beta \sum_{j \in M} z_j \right\} \quad (A.1)$$

$$\sum_{j \in M_{ik}} x_{jk} \geq y_{ik} \quad \forall i \in N, k \in K \quad (A.2)$$

$$\sum_{j \in M} x_{jk} \leq c_k \quad \forall k \in K \quad (A.3)$$

$$x_{jk} \leq z_j \quad \forall j \in M, k \in K \quad (A.3)$$

$$y_{ik}, z_j, x_{jk} \in \{0, 1\} \quad \forall i \in N, j \in M, k \in K.$$

Formulation of the Extended Model

Additional Input

- L Set of different crew types;
- p_l Number of available crews of type $l \in L$;
- τ_l Pre-trip delay of crew of type $l \in L$;
- M_{ikl} Set of base locations that can cover demand point $i \in N$ with vehicle type $k \in K$ staffed with crew of type $l \in L = \{j \in M: t_{ij} + \tau_l \leq r_{ik}\}$.

Additional Variables

- x_{jkl} Number of type k vehicles that are located at potential base location j .

Model

$$\max \left\{ \sum_{i \in N} \sum_{k \in K} d_{ik} y_{ik} - \beta \sum_{j \in M} z_j \right\} \quad (A.4)$$

$$\sum_{l \in L} \sum_{j \in M_{ikl}} x_{jkl} \geq y_{ik} \quad \forall i \in N, k \in K \quad (A.5)$$

$$\sum_{j \in M} \sum_{l \in L} x_{jkl} \leq c_k \quad \forall k \in K \quad (A.6)$$

$$\sum_{j \in M} \sum_{k \in K} x_{jkl} \leq p_l \quad \forall l \in L \quad (A.7)$$

$$x_{jkl} \leq z_j \quad \forall j \in M, k \in K, l \in L \quad (A.7)$$

$$y_{ik}, z_j, x_{jkl} \in \{0, 1\} \quad \forall i \in N, j \in M, k \in K, l \in L. \quad (A.8)$$

References

- Ahmadi-Javid A, Seyedi P, Syam SS (2017) A survey of healthcare facility location. *Comput. Oper. Res.* 79(March):223–263.
- Andersson T, Sårdqvist S (2007) Planning for effective use of fire and rescue service resources. Accessed March 26, 2017, <http://liu.diva-portal.org/smash/get/diva2:260234/FULLTEXT01.pdf>.
- Aringhieri R, Bruni M, Khodaparasti S, van Essen J (2017) Emergency medical services and beyond: Addressing new challenges through a wide literature review. *Comput. Oper. Res.* 78(February):349–368.

- Batta R, Mannur NR (1990) Covering-location models for emergency situations that require multiple response units. *Management Sci.* 36(1):16–23.
- Brandweer (2011) Historie 1874–1895: Beroepskorps. Accessed October 11, 2014, <https://www.brandweer.nl/amsterdamamstelland/het-korps/historie/historie-1874-1895?param=0>.
- Chevalier P, Thomas I, Geraets D, Goetghebeur E, Janssens O, Peeters D, Plastria F (2012) Locating fire stations: An integrated approach for Belgium. *Socio-Econom. Planning Sci.* 46(2):173–182.
- Chong KC, Henderson SG, Lewis ME (2016) The vehicle mix decision in emergency medical service systems. *Manufacturing Service Oper. Management* 18(3):347–360.
- Church RL, ReVelle CS (1974) The maximal covering location problem. *Papers Regional Sci.* 32(1):101–118.
- Daskin MS (1983) A maximum expected covering location model: Formulation, properties and heuristic solution. *Transportation Sci.* 17(1):48–70.
- Dzator M, Dzator J (2013) An effective heuristic for the p -median problem with application to ambulance location. *OPSEARCH* 50(1):60–74.
- Erkut E, Ingolfsson A, Erdoğan G (2008) Ambulance location for maximum survival. *Naval Res. Logist.* 55(1):42–58.
- Gendreau M, Laporte G, Semet F (1997) Solving an ambulance location model by tabu search. *Location Sci.* 5(2):75–8.
- Hogan K, ReVelle C (1986) Concepts and applications of backup coverage. *Management Sci.* 32(11):1434–1444.
- Hogg JM (1968) The siting of fire stations. *J. Oper. Res. Soc.* 19(3):275–287.
- ILOG (2009) IBM ILOG CPLEX optimization studio. Accessed March 26, 2017, <http://www-03.ibm.com/software/products/us/en/ibmilogcpleoptstud>.
- Ingolfsson A, Budge S, Erkut E (2008) Optimal ambulance location with random delays and travel times. *Health Care Management Sci.* 11(3):262–274.
- Kolesar P, Blum EH (1973) Square root laws for fire engine response distances. *Management Sci.* 19(12):1368–1378.
- McLay LA, Mayorga ME (2010) Evaluating emergency medical service performance measures. *Health Care Management Sci.* 13(2):124–136.
- Murray AT, Tong D (2009) GIS and spatial analysis in the media. *Appl. Geography* 29(2):250–259.
- Plane DR, Hendrick TE (1977) Mathematical programming and the location of fire companies for the Denver Fire Department. *Oper. Res.* 25(4):563–578.
- Reuter-Oppermann M, van den Berg PL, Vile JL (2017) Logistics for emergency medical service systems. *Health Systems*, ePub ahead of print February 27, <http://dx.doi.org/10.1057/s41306-017-0023-x>.
- Schilling DA, ReVelle CS, Cohon J, Elzinga DJ (1980) Some models for fire protection locational decisions. *Eur. J. Oper. Res.* 5(1):1–7.
- Swersey AJ (1982) A Markovian decision model for deciding how many fire companies to dispatch. *Management Sci.* 28(4):352–365.
- Toregas C, Swain R, ReVelle C, Bergman L (1971) The location of emergency service facilities. *Oper. Res.* 19(6):1363–1373.

Verification Letter

I. J. Stelstra, Commandant, a.i., Brandweer (Fire Department) Amsterdam-Amstelland, writes:

“Herewith I declare that the mathematical models and methods described in the paper ‘*Increasing the responsiveness of firefighter services by relocating base stations in Amsterdam*’ (co-authored by P. L. van den Berg, G. A. G. Legemaate and R. D. van der Mei) are highly valuable for optimizing the resource planning within our organization.

“The results convincingly show that—and how—significant improvements of our service quality can be realized by easily implementable re-allocation of our resources. While pro-actively re-allocating current base stations is costly and time-consuming, we recognize the benefits improved coverage provides. We have successfully integrated results from the model into our decision making process, and will continue to do so.

“Furthermore, we have identified another process which can greatly benefit from optimizations the model provides. When during a large scale incident multiple base stations are being called upon, we are now able to re-allocate remaining resources (vehicles) to better positions to regain optimal overall coverage. Results from this project are to be implemented in the Spring of 2016.”

Pieter L. van den Berg is an assistant professor at the Rotterdam School of Management at Erasmus University. He obtained his PhD in applied mathematics from the Delft University of Technology. His research focuses on the application of operations research to the logistics of emergency medical service providers and public transportation.

Guido A. G. Legemaate is an adviser working for Fire Department Amsterdam-Amstelland and collaborates with the logistics research team at Centrum Wiskunde en Informatica (CWI). He obtained a Master of Science in general health sciences from VU University Amsterdam. His research and professional activities focus on improving fire department operations using data science techniques.

Rob D. van der Mei is a full professor in operations research at VU University Amsterdam, the head of the Logistics research team, and the Industrial Liaison Officer at CWI. He is a co-founder of the spin-off company Stokhos Emergency Mathematics. His research interests include performance modeling and scalability analysis of ICT systems, logistics, grid computing, revenue management, military operations research, sensor networks, call centers, queueing theory, and applications of big data. He is widely published.