Linköping University Post Print

Decision support tools for ambulance dispatch and relocation

Tobias Andersson Granberg and Peter Värbrand

N.B.: When citing this work, cite the original article.

This is a post-peer-review, pre-copyedit version of an article published in Journal of the Operational Research Society. The definitive publisher-authenticated version:

Tobias Andersson Granberg and Peter Värbrand, Decision support tools for ambulance dispatch and relocation, 2007, Journal of the Operational Research Society, (58), 195-201.

is available online at: <http://dx.doi.org/10.1057/palgrave.jors.2602174>

Copyright: Palgrave Macmillan

<http://www.palgrave-journals.com/pal/index.html>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-17063>

DECISION SUPPORT TOOLS FOR AMBULANCE DISPATCH AND RELOCATION

TOBIAS ANDERSSON PETER VÄRBRAND

Division of Communications and Transport Systems
Linköping University, ITN
SE-601 74 Norrköping, Sweden
{toban,petva}@itn.liu.se

In this paper, the development of decision support tools for dynamic ambulance relocation and automatic ambulance dispatching is described. The ambulance dispatch problem is to choose which ambulance to send to a patient. The dynamic ambulance relocation problem occurs in the operational control of ambulances. The objective is to find new locations for some of the ambulances, to increase the preparedness in the area of responsibility. Preparedness is a way of evaluating the ability to serve potential patients with ambulances now and in the future. Computational tests using a simulation model show that the tools are beneficial in reducing the waiting periods for the patients.

1. Introduction

The two most important ambulance logistics services are the medical treatment of patients and the transportation of patients. An ambulance call typically starts with a triage, where medically qualified personnel determine the urgency of the call. This is followed by an ambulance assignment, where an ambulance dispatcher decides which ambulance to send to the call site (where the patient is located). The time from when the call has been received until the ambulance personnel have reached the patient is called the *waiting period*. The waiting period is more commonly referred to as the response time, but as it has been shown that many different definitions of response time are used (Moeller, 2004), waiting period is used here instead. The waiting period is the time from when the call has reached the emergency operator until the ambulance personnel have reached the patient. After the medical treatment, if this is necessary, the ambulance will transport the patient to a health care facility. Not all ambulance calls are urgent; non-urgent transportations can be ordered several days in advance, making it possible to perform some sort of transportation planning. The core of ambulance logistics is however to plan and control the emergency medical services, i.e. take care of the urgent calls. An ambulance dispatcher is commonly found in an emergency centre, where distress calls are received and from where the resources are controlled. The main tasks for the dispatcher are to assign ambulances to incoming calls, and to ensure that there are units available to serve future calls. The dispatchers are also expected to support the ambulance personnel with route guidance and possibly medical advice.

In Sweden, the foremost ambulance logistics provider is called SOS Alarm AB. They are responsible for receiving all calls to the national emergency number, 112, and also for controlling all ambulance movements. The operations are run from a SOS centre, of which there is one in each county (administrative region) in Sweden. This paper describes a number of applications that can support the ambulance dispatchers in the SOS centres. The

development of these decision support tools has been done as a part of OPAL – Optimized Ambulance Logistics, which is a joint project between SOS Alarm and Linköping University. From an operations research perspective, the contributions in the ambulance logistics area have mainly focused on reducing the waiting periods by trying to find optimal locations for ambulance stations. The first models dealing with the emergency station location problem surfaced in the 1970s (e.g. Toregas, 1971 and Church *et al.* 1974), and new models and algorithms keep appearing, as do surveys of the area (see e.g. Brotcorne *et al.* 2003 and Goldberg, 2004).

In order to evaluate the value of a set of ambulance station locations, it is possible to use simulation. One early simulation model was used for evaluating possible improvements in ambulance service is described in Savas (1969), and a more recent simulation study is described in Henderson *et al.* (2004). Furthermore, the hypercube model (Larson, 1974), and later extensions of this work, can be used to evaluate a solution from a location model.

The contributions in operational ambulance control are much more sparse. It is however possible to identify two problems that have received some attention; the ambulance dispatch problem and the ambulance relocation problem.

The most common and natural dispatch rule is to send the closest unit, since a general objective is to minimise the response times. However, this rule is not always optimal (Carter *et al.* 1972). Consider a case where two units, A and B, have equally large areas of responsibility, but A's area has a significantly higher call frequency. In this case, the mean response time will decrease if B is allowed to respond to some of the calls for which A is the closest unit. This result can be generalised for cases involving more than two units (Cunningham-Green *et al.* 1988), and it may be better to send unit C to take A's call when A already is busy, than to send the closer unit B (Repede *et al.* 1994). This is done if the call frequency in B's primary district is higher than in C's. In Weintraub *et al.* (1999) a dispatch system for vehicles servicing the electrical system in Santiago de Chile is described. Vehicles travel from call site to call site, and the dispatcher tries to maintain an adequate preparedness for quickly servicing high priority calls when deciding which unit that should be assigned to each call.

The demand for ambulances commonly varies with time, and some efforts to compensate for this by matching the amount of resources to the demand have been made. Even without changing the number of ambulances, it is possible to plan relocations of the existing fleet to better match the demand, if it changes in the area of responsibility during the day for example. (Carson *et al.* 1990) present a way of locating just one ambulance at a campus, where its position is changed several times during a 24 hours period, to compensate for population changes. To find pre-planned relocations, a location problem can be solved for each interesting time period. For a specific time period, one set of location points is used, and in the shift to a new period, ambulances have to be relocated to a new set of location points ((Repede *et al.* 1994).

In practice, it is maybe more common to dynamically relocate ambulances to cover for busy units. A dynamic relocation algorithm for fire companies is developed in Kolesar *et al.* (1974), and a call for relocation is triggered when some part of the city is not covered by any unit. In Gendreau *et al.* (2001) a tabu search heuristic for the dynamic relocation of ambulances is described, and a similar model for physician cars is presented in Gendreau *et al.* (2005). Both models maximise the coverage of the area.

In this paper, new algorithms for the ambulance dispatch and the dynamic ambulance relocation problems are presented. The assumptions made in the dispatch algorithm is similar to the work in Weintraub *et al.* (1999) in that the closest unit is not always sent to a new call, but adapted for the pick-up and delivery nature of ambulance calls rather than the repair problem that is studied in Weintraub *et al.* (1999). The relocation algorithm is dynamic, i.e.

the problem is solved when there is a lack of ambulances somewhere in the area. This is a similar problem to the one studied in Kolesar *et al.* (1974) and in Gendreau *et al.* (2001). The major difference is that it is possible to relocate an ambulance to any zone in the area, not just to vacant stations, as in Kolesar *et al.* (1974). It is also possible to run the algorithm on a common PC, i.e. no expensive hardware is needed, and still obtain solutions within a few seconds. Another novel aspect of the algorithms is that both of them utilise a new quantitative preparedness measure.

2. Preparedness

In ambulance logistics, preparedness has been used as a qualitative measure for a long time, but two people does not always mean the same thing when using the word. Also, two ambulance controllers may have different opinions on what can be considered good or bad preparedness, depending on their experience, their risk aversion and their personality. For example, one controller may think that less than twenty available ambulances in the county means that the preparedness is low, while another controller thinks that twenty ambulances are more than enough for an adequate preparedness.

In order to find a quantifiable measure for preparedness, we first divide the area of consideration into a number of zones. To each zone j , a weight c_j is associated, which mirrors the demand for ambulances in the zone. The weight can for example be proportional to the number of calls served in the zone during a specific time period, or to the number of people currently resident in the zone. The preparedness in zone j can then be calculated as:

$$p_j = \frac{1}{c_j} \sum_{l=1}^{L_j} \frac{\gamma^l}{t_j^l} \quad (1)$$

where L_j is the number of ambulances that contribute to the preparedness in zone j , t_j^l is the travel time for ambulance l to zone j , γ^l is the contribution factor for ambulance l and the following properties hold:

$$t_j^1 \leq t_j^2 \leq \dots \leq t_j^{L_j} \quad (2)$$

$$\gamma^1 > \gamma^2 > \dots > \gamma^{L_j} \quad (3)$$

Thus, the preparedness is calculated by letting the L_j closest ambulances to zone j contribute to the preparedness with an impact that is decreasing as the travel time to the zone increases.

One basic quality of (1) is that the preparedness in a zone increases if an ambulance moves closer to the zone, i.e. some t_j^l decreases. Furthermore, if the call frequency in a zone, i.e. c_j , increases, the preparedness decreases (Andersson, 2005).

3. Dispatching Support

In Sweden, the prioritisation of an ambulance call results in one of three degrees, Prio 1, 2 or 3. Prio 1 calls are the most urgent, life threatening calls, while Prio 2 are urgent but not life threatening and Prio 3 are non-urgent calls.

Sometimes it is trivial to decide which ambulance to assign to a new call, e.g. for a Prio 1 call that requires only one ambulance, the ambulance with the shortest expected travel time to the call site is always dispatched. If the call is not as urgent, an ambulance dispatcher may choose to assign an ambulance with a longer travel time, if this assignment means that the drop in preparedness will be less significant. The dispatcher may also reassign an ambulance already on its way to a call site, if the new call is more urgent.

The implementation of the preparedness measure (1) includes a list of the closest ambulances for each zone, sorted according to the expected travel time. Thus, it is easy to find the closest

ambulance to a certain zone. To check which ambulance to dispatch to a Prio 2 or 3 call, an algorithm (see Table 1) has been developed that checks all available ambulances within a certain travel time from the zone, and picks the one whose unavailability causes the least drop in the preparedness as calculated by (1).

Table 1: The ambulance dispatch algorithm

1.	Let j be the zone to where an ambulance needs to be dispatched, and $l = 1, \dots, L_j$ an ordered list of the ambulances that contribute to the preparedness in j . Let $A = \emptyset$ be the ambulance that is dispatched and let $p_{min} = 0$.
2.	IF $\text{PRIO}(j) == 1$
3.	Set $A = 1$ and dispatch A , i.e. dispatch the ambulance that is closest to zone j and therefore first in the list
4.	IF $\text{PRIO}(j) == 2$ OR 3
5.	Check the ambulances in the list, beginning with the closest: FOR $l = 1, \dots, L_j$ IF $t_j^l < T_2$ (or T_3 if $\text{PRIO}(j) == 3$)
6.	Remove ambulance l from the list of ambulances contributing to the preparedness, and recalculate the preparedness, p_i , in all zones that are affected by this action. IF $\min_{i \in N} \{p_i\} > p_{min}$
7.	$p_{min} = \min_{i \in N} \{p_i\}, A = l$
8.	Dispatch A

The algorithm in Table 1 starts, after the initialisations, by checking the priority of the call. If it is a Prio 1 call, i.e. $\text{PRIO}(j) == 1$ in Step 2, the closest ambulance is dispatched to the call. If the call is a Prio 2 or 3 call, the algorithm starts by checking if the closest ambulance, i.e. the one first in the list, can reach the zone within T_2 (or T_3) minutes. If not, the algorithm will stop and the closest ambulance will be dispatched. If it can, the ambulance will be set as unavailable, and new levels of preparedness will be calculated for all zones in NC^k , which is the set of zones that will be affected by the assignment (Step 6). The lowest level for any of the zones, p_{min} , is saved, and used as a measure on how the preparedness is affected by the dispatch. When the first ambulance in the list has been processed, the algorithm checks if ambulance number two can reach the zone within T_2 (or T_3) minutes. If it can, a new minimum preparedness level is calculated and compared to the current p_{min} . This is continued until an ambulance is too far away from the zone or until there are no ambulances left in the list.

By letting ambulances on their way to a Prio 2 or 3 call still contribute to the preparedness, it is also possible to assign these to calls that are more urgent, e.g. an ambulance on its way to a Prio 3 call, can be assigned to a new Prio 2 or Prio 1 call. In this case it is necessary to check, in Step 5, if the ambulance is already on its way to serve a call.

To ensure that the waiting periods for the less urgent calls do not grow beyond what is practically feasible, pseudo priorities are used when ambulances are reassigned. The pseudo priority for a call changes if the call has not been served within a certain time, e.g. a Prio 3

call that has not been reached by an ambulance in T_3 minutes changes pseudo priority from 3 to 2. This means that an ambulance that is on its way to serve this Prio 3 (pseudo Prio 2) call cannot be reassigned to a Prio 2 call, but still to a new Prio 1 call. It may be noted that also the real priority of a call, and thus not only the pseudo priority, may change if a patient has to wait for medical care.

4. Dynamic Ambulance Relocation

By colour coding the measure (1) in a geographical information system, an ambulance dispatcher can manually check where the preparedness is low, and thus where to send ambulances. What is even more useful is a tool that automatically checks the preparedness in the zones, and suggests ways to relocate ambulances in order to maintain a sufficient level.

The measure (1) can be used as a base for this kind of tool. First however, it is necessary to calibrate the measure (i.e. deciding parameter values for c and γ) and finding a lowest level of preparedness, P_{min} , that should be kept in all zones. In Andersson (2005), the measure (1) is calibrated for use in the county of Stockholm in Sweden, and a level P_{min} of 0.923 is identified.

The ambulance relocation problem occurs when one or more zones have a preparedness level less than P_{min} . The objective is then to reach the P_{min} level in all zones as quickly as possible. The preparedness is increased by relocating one or more ambulances closer to the zones that suffer from a low level of preparedness. A model, DYNAROC, that solves the dynamic ambulance relocation problem follows:

$$\min \quad z \quad (4)$$

$$z \geq \sum_{j \in N^k} \tau_j^k x_j^k \quad k = 1, \dots, A \quad (5)$$

$$\sum_{j \in N^k} x_j^k \leq 1 \quad k = 1, \dots, A \quad (6)$$

$$s.t. \quad \sum_{k=1}^A \sum_{j \in N^k} x_j^k \leq M \quad (7)$$

$$\frac{1}{c_j} \sum_{l=1}^{L_j} \frac{\gamma^l}{t_j^l(\mathbf{x})} \geq P_{min} \quad j = 1, \dots, N \quad (8)$$

$$\mathbf{x} \in \{0,1\} \quad (9)$$

The objective (4) is to minimise the variable z , which is the maximum travel time for any of the relocated ambulances, i.e. the time it will take until the preparedness is at least P_{min} in all zones, which is required in (8). Constraint (5) states that z has to be greater than or equal to any of the travel times τ_j^k , which is the time required for ambulance k to reach zone j . The variable x_j^k equals 1 if ambulance k is relocated to zone j . Each of the ambulances can be relocated to at most one zone in the set N^k (6), which is the set of zones that can be reached by ambulance k in less than R minutes. By setting R low, the set N^k and the set of feasible solutions will be smaller. This will however also decrease the number of feasible solutions to

the model, with the risk that there will not exist any solutions in some instances. R is an upper bound on the objective function variable z , which means that if R is set e.g. to 20 minutes, no ambulance will have a relocation travel time longer than 20 minutes. Constraint (7) ensures that not more than M ambulances are relocated. $t_j^l(\mathbf{x})$ in constraint (8) is a function of the variable $\mathbf{x}$, which is the vector form of x_j^k . Naturally, the travel time for the l 'th closest ambulance to zone j , i.e. t_j^l , depends on where the ambulances are located, which is decided by the values on the variable $\mathbf{x}$.

Table 2: The tree search algorithm that finds solutions to DYNAROC

1. Let the current (infeasible) solution, i.e. $x_j^k = 0 \ \forall j,k$, be the root of the tree.
2. Let j = the zone with the lowest preparedness.
3. REPEAT
4. Find the n ambulances, with the minimum travel times, that can be relocated in a way that ensures that $p_j \geq P_{min}$. Save a maximum of m zones for each of these ambulances that satisfy the conditions above. The ambulances must not have been relocated once already (earlier in the tree) and not more than $M-I$ ambulances must have been moved.
5. Each of the moves in Step 4 gives a potential solution.
FOR all new solutions
6. Check if $p_i \geq P_{min} \ \forall i = 1, \dots, N$ and check the longest travel time against the best solution found so far if this is true.
If the new solution is not feasible, create a new node and connect it to its parent solution.
7. Pick a new node and let j = the zone with the lowest preparedness in the new solution.
8. UNTIL there are no nodes left to examine, or some other stop criterion triggers

Since a short computation time is of uttermost importance, a tree-search heuristic is used to solve DYNAROC. It is schematically described in Table 2 and illustrated in Figure 1. Starting with the current situation, the heuristic iteratively tries to raise the preparedness in the zone with the lowest preparedness. This is done by moving ambulances closer to this zone. Only ambulances that can be moved close enough to raise the preparedness in the zone to at least P_{min} are evaluated. Ambulances that can raise the preparedness enough are compared to each other concerning the required travel time for the relocation (Step 4). If possible, n ambulances are saved, and each of these ambulances is relocated to a maximum of m zones. The reason for not just saving a set of the best possible relocations is that these may be performed by a single ambulance, and a certain divergence in the search for solutions is desired. Every move that has raised the preparedness is checked for feasibility in Step 6. If some zone has a preparedness below P_{min} , the move has given a new infeasible solution that is saved in the search tree. In the new solutions, another ambulance is relocated to raise the preparedness. In Figure 1, n and m are both set to 3, i.e. three ambulances are marked as potential relocation candidates, and for each of these the three best relocation zones are identified. Best in this aspect means the zones to which the travel times are shortest, but where a relocation would

raise the preparedness to above the threshold value. The relocation of an ambulance gives rise to a new potential solution, which is feasible if the preparedness in all zones is above the threshold value. If the solution is not feasible, it is possible that further relocations are needed. Therefore, a new node is created and the infeasible solution is used as a start solution (a parent node) further down in the search tree.

Figure 1. Illustrating the tree-search heuristic that solves the DYNAROC model

The first level of the search tree contains (infeasible) solutions where only one ambulance has been relocated, and the second level where two ambulances have been moved. Thus, the tree will never grow more than M levels and each node can have at most $n*m$ children. If no feasible solution can be found, the search tree will have searched a total of $1+(n*m)^M$ nodes. With a feasible solution, it is however possible to cut extensively during the search process. It is never beneficial to evaluate a solution where the one of the relocation travel times are longer than the objective function for the currently best feasible solution. It is therefore possible that all nodes will be examined and the algorithm terminated because of this, but an alternative stop criterion, e.g. number of iterations or elapsed computation time should be used as well.

5. An Ambulance Operations Simulator

Simulation of the ambulance operations can be used for evaluating strategic decisions, such as where to locate ambulance stations or how large the ambulance fleet should be. It can also be a valuable tool for education and training of the ambulance dispatchers. Moreover, it can be used as a visualisation tool for information purposes towards customers, decision makers, and the public.

The two decision support tools described earlier, which give suggestions on ambulance assignments and ambulance relocations, are necessary for it to be possible to simulate the ambulance operations faster than real time. When simulating in real time, an ambulance dispatcher can decide which ambulance to send to the incoming calls, and if any relocations should be made. This would however be very time consuming if the simulation tool is to be used for evaluating strategic decisions, when it is necessary to simulate days, weeks or months of operations in order to get significant results. Thus, the ambulance dispatcher, or more accurately the decisions made by him or her, has to be simulated as well. This includes foremost dispatching decisions and relocation decisions.

Figure 2. The ambulance operations simulator

The developed simulation model is schematically described in Figure 2. The *call generator* simulates incoming calls to the emergency centre; in the current implementation, these are stochastically generated, but in the future, it should be possible to use historical sequences of calls, or to use sequences that are constructed to serve a certain, e.g. educational, purpose. The *resource simulator* handles the incoming calls and the ambulances that are used to serve them. The simulation is time based, and a time step of one minute is used, so each minute there is a certain possibility, depending on the population, that a call will be generated in a zone. When a new call is generated, the *automatic dispatch module* finds an ambulance to assign to the call. If the ambulance was already on its way to serve a call, the dispatch model finds a new ambulance to serve the old, less urgent call. If there are no available, or soon to be available, ambulances to assign to a call, it will be put in a queue. Each iteration starts with the processing of the calls on queue, to see if any ambulances have become available and are able to serve them.

The assignment of ambulance to a call may affect the preparedness level in some zones, which is why the *preparedness calculator* is used to check if the level has dropped below a certain threshold value, P_{min} . If it has, the *automatic relocation module* tries to find a relocation of one or more ambulances that will raise the preparedness in the affected zones.

6. Computational Results

The algorithms are tested using data for the county of Stockholm in Sweden. It offers the most complex ambulance control situation in Sweden, with at most 58 ambulances and about 400 ambulance calls a day. The data that is used consists of travel times in minutes between the 1240 zones into which the county is divided. The travel times were originally collected for fire engines, but were deemed directly translatable to ambulances. Population data for each of the zones is used to calculate the weights, c_j . The population in the zones varies between 1 and 17985 and these values are divided by 20000, in order to get more convenient values to work with, which means that $1/c_j$ varies between 0.9 and 20000. γ^l is set to $1/2^{l-1}$ for $l = 1, 2, \dots, 7$, i.e. $\gamma^1 = 1$, $\gamma^2 = 0.5$, $\gamma^3 = 0.25$, $\gamma^4 = 0.125$, etc. A maximum of 7 ambulances are used to calculate the preparedness for a zone. Furthermore, a maximum relocation travel time (R) of 40 minutes is used in DYNAROC, meaning that N^k is the set of zones that can be reached by ambulance k within 40 minutes. M is set to 3, as it is reasoned that an ambulance dispatcher would be reluctant to relocate too many ambulances. n and m in the DYNAROC algorithm in Table 2 is set to 5 and 3 respectively, and a maximum of 50 iterations is performed by the DYNAROC algorithm. On a PC computer with a 1000 MHz Pentium III processor, it takes about 6 seconds to perform 50 iterations. However, the heuristic most often finishes before 50 iterations are performed. In a quick test where DYNAROC is solved 2275 times, the mean running time for the algorithm is 2.24 seconds, and the maximum running time 5.89 seconds. About 53% of the calls generated in the simulations are Prio 1 calls, 33% Prio 2 and 14 % Prio 3 calls, in accordance with historical data for the county of Stockholm. The probability that a call will appear is proportional to the population in the zone.

One proposition made in Andersson (2005), is that if the minimum level of preparedness is kept above 0.923 in the county of Stockholm, the waiting period targets should be satisfied. The targets differ depending on the priority of the call. For Prio 1 calls, the target is that 75% of all calls should be served within 10 minutes, 95% within 15 minutes and 99% within 20 minutes. For Prio 2 and 3 calls, these times are allowed to be longer. To test this, simulations of 1000 hours of operations with $P_{min} = 0.923$, are performed and compared to simulations with $P_{min} = 0.277$ and 0. 0.277 is the lowest level of preparedness in the county when all ambulances are available at their stations, and when P_{min} is zero, no relocations are performed. As can be seen in Figure 3, the number of calls served within the specified times increases as P_{min} is increased. Especially the waiting period target that 99% of all Prio 1 calls should be served within 20 minutes, benefits from the relocations. It should be noted that instant relocations are used. This means that the relocation travel time (τ_j^k in DYNAROC) for each relocated ambulance is set to a very small value, and thus the ambulances are instantly transferred to their new zones. The reason for this is that the simulation model has no way of keeping track of the ambulances while they are relocating. Thus, when the actual relocation travel times are used and an ambulance that is being relocated is assigned to call, the travel time to the new call will be overestimated.

Even though instant relocations are used, the minimum level of preparedness does not stay above the threshold value at all times; in fact as the number of calls is increased, the time p_{min} is above 0.923 drops from 96.1% at 200 calls per day to 0.22% at 800 calls a day. This is because the heuristic increasingly fails to find relocations that will improve the situation (at 800 calls a day the heuristic fails to find a solution for $P_{min} = 0.923$ in 38.4% of the cases). Here it should be noted that 800 calls a day is almost double to the normal amount, and the ambulance fleet is not large enough to handle that kind of call frequency for a longer period. Therefore it is not surprising that improving relocations are hard to find. Still, it seems reasonable to state that the waiting period targets probably will be satisfied if a minimum preparedness level of 0.923 is kept at all times.

Figure 3. The mean number of relocated ambulances per day as a function of the mean number of calls per day, when using a threshold value of 0, 0.277, and 0.923 respectively

However trying to maintain a high minimum preparedness level requires an unreasonable amount of relocations, as indicated by Table 3. Looking at the results when $P_{min} = 0.923$, in average more than one ambulance is relocated for each call when the call volumes are high, even though the heuristic often fails in finding a solution.

How often an ambulance may be relocated has to be negotiated between the ambulance personnel, or the company providing the vehicles, and the people responsible for the ambulance health care. Relocations will most probably increase the total travel distance for the ambulances, and they cause the ambulance personnel to spend more time on the road and less at the station. Therefore, it is natural that the ambulance personnel want to be compensated if the number of relocations increases.

Table 3: The average number of relocated ambulances per day

nr of calls / day	200	250	300	350	400	450	500	550	600	650	700	750	800
$P_{\min} = 0.277$	4.4	6.7	9.3	10.2	14.7	19.6	25.1	33.9	51.6	65.1	101.5	146.2	223.6
$P_{\min} = 0.923$	105.5	148.7	194.0	253.4	331.4	412.5	500.4	578.5	662.2	700.9	752.3	778.4	819.6

7. Conclusions

In this paper, decision support tools for dynamic ambulance relocation and automatic ambulance dispatching are presented. The tools utilise a measure for preparedness, which is a way of evaluating the ability to serve current and future calls anywhere in the area.

Simulations show that maintaining a high level of preparedness as calculated by the developed measure, is helpful in reaching the waiting period targets that are set by the county councils in Sweden. During the simulations, the preparedness is improved by dynamically relocating ambulances and a large amount of relocations is necessary to get significant results. The purpose of the relocations is to keep the preparedness in the area high. Using a different set of ambulance station locations is another possible way of raising the minimum preparedness level, at least initially. By changing these locations as the demand changes in the area, i.e. using pre-planned relocations, it will be easier to maintain this level, which may decrease the need for dynamic relocations. Dynamic relocations can then be used as they are in practice now, to cover for a temporary and serious lack of ambulances somewhere in the area.

Acknowledgements

This research was partly sponsored by SOS Alarm. The authors would like to thank SOS Alarm for the privilege of working with them and for the constant support and encouragement.

References

- Andersson, T. (2005). *Decision support for dynamic fleet management*. PhD thesis, Department of Science and Technology, Linköping University, Sweden.
- Brotcorne, L., Laporte, G., Semet, F. 2003. Ambulance location and relocation models. *European Journal of Operational Research*. 147, 451-463.
- Carson, Y., Batta, R. 1990. Locating an ambulance on the Amherst campus of the State University of New York at Buffalo. *Interfaces*. 20(5), 43-49.
- Carter, G., Chaiken, J., Ignall, E. 1972. Response areas for two emergency units. *Operations Research*. 20(3), 571-594.
- Church, R., ReVelle, C. 1974. The maximal covering location problem. *Papers in Regional Science*. 32(1), 101-120.
- Cunningham-Green, R., Harries, G. 1988. Nearest-neighbour rules for emergency services. *ZOR – Zeitschrift for Operations Research*. 32, 299-306.

- Gendreau, M., Laporte, G., Semet, S. 2001. A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. *Parallel Computing*. 27, 1641-1653.
- Goldberg, J. 2004. Operations research models for the deployment of emergency service vehicles. *EMS Management Journal*. 1(1), 20-39.
- Henderson, S., Mason, A. 2004. Ambulance service planning: Simulation and data visualization. In Brandeau, M., Sainfort, F., Pierskalla, W., (eds.), *Operations Research and Health Care*. Kluwer Academic Publishers, Boston, 77-102.
- Kolesar, P., Walker, W. 1974. An algorithm for the dynamic relocation of fire companies. *Operations Research*. 22(2), 249-274.
- Larson, R. 1974. A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers and Operations Research*. 1(1), 67-95.
- Larson, R. 1975. Approximating the performance of urban emergency systems. *Operations Research*. 23(5), 845-868.
- Moeller, B. 2004. Obstacles to measuring emergency medical services performance. *EMS Management Journal*. 1(2), 8-15.
- Repede, J., Bernardo, J. 1994. Developing and validating a decision support system for locating emergency medical vehicles in Louisville, Kentucky. *European Journal of Operational Research*. 75, 567-581.
- Savas, E. 1969. Simulation and cost-effectiveness analysis of New York's emergency ambulance service. *Management Science*. 15(12), B608-B627.
- Toregas, C., Swain, R., ReVelle, C., Bergman, L. 1971. The location of emergency service facilities. *Operations Research*. 19(6), 1363-1373.
- Weintraub, A., Aboud, J., Fernandez, C., Laporte, G., Ramirez, E. 1999. An emergency vehicle dispatching system for an electric utility in Chile. *Journal of the Operational Research Society*. 50, 690-696.