Decision support for operational ambulance control

T. Andersson
Department of Science and Technology, Linköping University, Sweden

Abstract

Operational ambulance control naturally includes the dispatching of ambulances to calls, but equally important, and an even more complicated task, is the need to maintain a sufficient level of preparedness in the area of responsibility. Preparedness is a measure of the ability to serve potential patients with ambulances in a swift and efficient manner. Here, a number of enhancements to the existing information systems are presented, including way of calculating and visualising the preparedness in a geographical information system (GIS). Furthermore, a simulation module and real-time decision support tools for automatic dispatching and dynamic ambulance relocation, based on mathematical modelling and heuristic solution algorithms, are developed.
Keywords: ambulance logistics, optimisation, preparedness, decision support tools, geographical information systems.

1 Introduction

SOS Alarm [11] is the company in Sweden 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 central of which there is one in each county (administrative district) in Sweden. One of the services offered by SOS Alarm is called *ambulance logistics*, and the customers are the county councils. An ambulance can be requested by the public or by the health care. A SOS operator receives the request and prioritises it according to three degrees:

- Prio 1: Urgent, life threatening symptoms.
- Prio 2: Urgent, not life threatening symptoms.
- Prio 3: Non-urgent calls.

Management Information Systems, C. A. Brebbia (Editor)

A Prio 3 call is often a transport to or from the hospital and the patient's home, or between hospitals. These requests are often made by the health care but can also be made by the public at a special ambulance order number.

Orders received more than a day before they are to be executed can be used to create a transportation schedule. Today, the planning mostly consists of checking that it is possible to perform the pre-ordered transports and still have enough ambulances left to cover for incoming urgent calls, i.e. without compromising the *preparedness* in the county.

Calls that are received the same day as they are to be executed may be urgent or non-urgent. A SOS operator decides which is the case when the call is prioritised. An ambulance controller (that may or may not be the same person as the operator) then has to decide which ambulance to assign to the incoming call. This decision depends on a number of different parameters; e.g. the location of the ambulances, their equipment and also the preparedness in the area. Preparedness refers to the ability to get one or more ambulances to a potential call site, i.e. where the patient is located, within a reasonable time. At a certain location, the preparedness depends on how many ambulances that can reach the location within a certain time frame, as well as their travel time. Of course, the preparedness also depends on the expected need for ambulances in the area; e.g. in a densely populated area it may be necessary to have more than one ambulance close by in order to get an adequate preparedness.

To a Prio 1 call, the closest ambulance (or ambulances if more than one is needed) is always assigned. This assignment may however cause a drop in the preparedness, probably for the zones close to the ambulance that was assigned to the new call. This might make it necessary for the controller to relocate another ambulance to the affected areas, in order to cover for the ambulance that is busy. For Prio 2 and 3 calls, the controller may choose to assign another ambulance than the closest one in order to preserve the preparedness, if this can be done without compromising the safety of the patient.

Keeping an adequate preparedness is one of the most complex tasks for the ambulance controller. It requires knowledge of where call sites are likely to appear and of how fast the ambulances can travel through different parts of the county, as well as knowledge of where the ambulances currently are located and if they are available. Today, most ambulances in Sweden have GPS receivers and transmit their position and status to the SOS central. However, to know where ambulances might be needed in the future, and how fast they can get there requires experience.

2 Information systems in the SOS central

The main computer system in a SOS Central is called CoordCom, originally developed by Ericsson [6] in 1987. CoordCom supports communication by radio, telephone and data, and is also used to book, edit and store information about ambulances and calls. Information about a new call might, for example, be called in through the public switched telephone network, in which case CoordCom automatically checks, and enters the address from where the call

Management Information Systems, C. A. Brebbia (Editor)
© 2004 WIT Press, www.witpress.com, ISBN 1-85312-728-0

originated. The SOS operator prioritises the call and enters additional information into CoordCom, and CoordCom indicates for the ambulance controllers that there is a new call to be served. The ambulance controller assigns an ambulance to the call, and the system sends a data message, using the Mobitex network [10], to the ambulance personnel. When the ambulance is on its way to the call site, the ambulance personnel acknowledges this by sending a data message back to the central. The Mobitex network is also used to transmit the position of the ambulances, which is received using the global positioning system (GPS), at regular intervals. A new system that will replace CoordCom is being developed by Ericsson, and will be taken into service in the beginning of 2005.

The ambulance controllers also have a geographical information system (GIS) to support them in their work. The old GIS was recently replaced by the ResQMap system, developed by Carmenta [5]. ResQMap is connected to CoordCom and can graphically display the information that is stored in CoordCom, e.g. the location of calls and ambulances. Today it is not possible to issue any kind of order in the GIS, due to the inability of CoordCom to process them, but with the arrival of the new communication and call handling system, the ambulance controller will be able to perform more of the work tasks directly in ResQMap.

Although the current systems, and the forthcoming system, provide substantial support to the ambulance controllers, there are a number of desirable features that they do not currently offer, including:

- **Calculation and visualisation of preparedness.** To state the preparedness in ambulance logistics is a way of describing the current operational situation. If the preparedness is “good”, the ambulance controller feels confident that it is possible to efficiently serve a normal amount of calls in the near future. This situation might however change rapidly; suppose that there is a serious traffic accident, forcing the controller to dispatch a large number of ambulances to the call site. These ambulances will probably be the ones that are closest to the call site, making the surrounding area devoid of available ambulances. Suddenly the preparedness is “bad”, at least around the call site. Naturally, “good” and “bad” are quite vague expressions, and does not give a clear picture of the operational situation. Therefore an application capable of calculating and visualising the preparedness would be valuable.
- **Simulation.** Simulation of ambulance control would as a minimum include the generation of calls and simulated ambulance movements. This would make it possible to use the simulator as an educational tool, by letting an ambulance controller make all the decisions, such as the assignment of calls to ambulances and the dynamic relocation of ambulances. In order to use simulation to evaluate more strategic questions, e.g. fleet sizing and ambulance station location, it is

necessary to simulate the ambulance controller as well, which means finding algorithms for automatic assignment and automatic relocation.

- **Automatic dispatch**. In order to simulate an ambulance controller, a decision support tool capable of selecting which ambulance that should serve a new incoming call is needed. The tool should of course pick the most suitable ambulance, from a patient safety and a preparedness perspective. This also makes the tool useful as decision support in operational ambulance control, where it can give the controller suggestions on possible assignments.
- **Automatic relocation**. In order to keep a high level of preparedness in the area of responsibility, an ambulance controller may choose to relocate ambulances to cover zones with a low preparedness. A decision support tool that can judge when a relocation is necessary and which ambulances that should be relocated is also needed to simulate the work performed by an ambulance controller. Just as for the automatic assignment tool, automatic relocation can be used in operational ambulance control.

3 Enhancements to current systems

3.1 Calculation and visualisation of preparedness

In order to find a quantifiable measure for preparedness, we first divide the area of consideration into a number of different 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:

$$p_j = \frac{1}{c_j} \sum_{l=1}^{L_j} \frac{\gamma^l}{t_j^l} \tag{1}$$

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

$$t_j^1 \le t_j^2 \le \ldots \le t_j^{L_j} \tag{2}$$

$$\gamma^1 > \gamma^2 > \ldots > \gamma^{L_j} \tag{3}$$

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. A more thorough discussion on the preparedness measure (1), and alternative ways of calculating the preparedness, can be found in Andersson *et al* [2].

Given a certain fleet of ambulances with a known status and location, the preparedness in each of the zones can be calculated. The most obvious area of use for this is in the operational control, where the preparedness dynamically can

Management Information Systems, C. A. Brebbia (Editor)

be updated. The ambulance controller can check when the preparedness is low in a zone and do something to correct this, e.g. relocate an ambulance to cover the affected zone.

To benefit fully from the preparedness measure, it should be presented in a way that makes it easy for the ambulance controller to evaluate the situation. This can, for example, be implemented by letting different colours represent different levels of preparedness in the GIS.

3.2 Simulation

Today, the GIS is capable of, among other things, to visualise and manage information about calls and ambulances. This makes it an ideal platform for an ambulance control simulator, especially since the ambulance controllers are already used to the system interface.

In order to simulate calls, a call generator that imitates the information that is usually received from CoordCom is needed. Calls can be generated stochastically according to historically verified distributions, which is proper if the simulator is to be used for evaluating strategic decisions. A constructed sequence of calls may be a preferred input if the simulator is used as an educational tool, since this makes it possible to build scenarios to educate and test the skills of the ambulance controllers. Of course, the constructed sequence may also be a set of historical calls.

For the simulation to work, it is also necessary to supply the GIS with ambulance information, and for the ambulances to automatically find and travel a route from an origin to a destination when ordered.

3.3 Automatic dispatch

Sometimes it is easy 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. However, the controller must still ensure that the ambulance carries the necessary equipment and that the ambulance personnel are qualified to handle the call. If the call is not as urgent, an ambulance operator may choose to dispatch an ambulance with a longer travel time, if this assignment means that the drop in preparedness will be less significant. The controller may also reassign an ambulance already on its way to a call site, if the new call is more urgent.

All of the decisions above have to be simulated by an automatic dispatch algorithm, for a simulator to be credible. Since the implementation of the preparedness measure (1) keeps track of which ambulances that are close to a zone, it is easy to find the closest ambulance to a certain zone. The check, if a certain ambulance may handle a certain call, can be made beforehand and stored together with the ambulance information. To check which ambulance to dispatch to a Prio 2 or 3 call, an algorithm 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).

Management Information Systems, C. A. Brebbia (Editor)

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. To ensure that the waiting times for the less urgent calls do not grow beyond what is practically feasible, pseudo priorities are used when making the assignments. 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. A similar approach is used in Weintraub *et al* [12] where requests for service vehicles used for repairing breakdowns in the electricity system move up in priority the longer they remain in the service queue. Gendreau *et al* [7] use a slightly different approach and only picks ambulances already on their way to a call if it is the only available ambulance that can reach the more urgent call site quickly, or if there is another ambulance that can cover the less urgent call within a certain time. It may be noted that the real priority of a call, and thus not only the pseudo priority, may change if a patient has to wait for medical care.

3.4 Automatic relocation

An ambulance controller may relocate an ambulance if there is reason to believe that there exists a location where the ambulance is more likely to be able to serve a new call in a shorter time. This relocation will then increase the general preparedness. In order to simulate this decision, it is assumed that relocating one or more ambulances is beneficial if the level of preparedness, as given by (1), drops below a certain threshold value, P_{min}. This gives rise to the following optimisation problem:

Minimise

{the maximum relocation time among the relocated ambulances}

subject to

{all zones must have a preparedness level of at least P_{min} after the relocation is complete}

{at most R_{max} ambulances may be relocated}

A mathematical representation, and a solution algorithm for the problem can be found in Andersson *et al* [4]. Today, the ambulance controllers solve the problem above manually, although the objective and the constraints are not quantified. In order to simulate the relocation decisions, quantifications are necessary. The decision to solve the problem is triggered by any zone having a level of preparedness below P_{min}, and the objective is to correct this situation as quickly as possible. The constraint that not more than R_{max} ambulances may be relocated is necessary to ensure that it will be possible to use the solutions to the problem in practice; no ambulance controller would approve of a solution that involved too many relocations. If more than two or three ambulances are to be relocated, it is hard for the controller to evaluate the solution, and might feel that it is not worth the trouble to execute it.

4 Evaluating the system enhancements

Once the system enhancements have been developed, they have to be calibrated, verified and evaluated. The calibration phase includes collecting necessary data and finding values for parameters. For the calculation of the preparedness, this means constructing the zones, finding proper weights and obtaining expected travel times between the zones. Each one of these tasks is a research challenge in itself, but fortunately this also means that there often has been some prior work in the area, some of which is described in Goldberg [8]. The construction of the zones and the travel time modelling often comes hand in hand if previous developed models are to be used for obtaining the travel times. One way of finding expected travel times is to use an equilibrium based traffic model; see for example Alsalloum *et al* [1], where traffic models are used to obtain travel times for ambulances. This type of model is often used for evaluating alternative modifications in the traffic infrastructure. The solution from the traffic model is a state of equilibrium where no traveller can find a shorter route from their origin to their destination. The model primarily determines the traffic flow on links in a traffic network. The travel time on a link in the model is expressed as a function of the flow on this link; typically, the function prescribes that when the flow increases the travel time increases. The travel time, as well as the traffic flow, for each link in the network in the equilibrium situation can then be extracted from the computed equilibrium solution. The traffic equilibrium can be run for different traffic situations and travel demands, making it possible to obtain travel times for, for example, different time periods during a day.

The demand for ambulances in a zone, or the probability that an ambulance will be needed in a zone, can be found in different ways. Common methods for forecasting demand in other areas include causal methods, time series models and judgemental methods. In causal methods the forecasted demand is expressed as a function of independent variables such as e.g. population and employment in the zone. Kamenetzky *et al* [9] describe the development of a causal model for estimating the demand for pre-hospital care in Southwestern Pennsylvania. Time series models are based on historical data, which is analysed to find out how the demand varies with time. Common components in these models are trends, cyclical variations and seasonal variations. In a demand function for urgent ambulance health care, the seasonal variations will be prominent, since the demand in a zone might vary substantially depending on the time of the day, which day of the week it is, or which month it is. Judgemental methods are mainly built on knowledge and qualified guesses of the present and of future events; e.g. the knowledge that a big sports event will take place in a zone will of course affect the expected demand for ambulances. In order to obtain a fair estimation of the expected operational demand for ambulances in a zone, a combination of the three types of methods may be needed.

When all the necessary data has been collected, proper values for the parameters have to be found. In the case of the preparedness measure (1), there are two different kinds of parameters; the weights c_j and the contribution factors

γ^l. A way of calibrating these, as it was done for the county of Stockholm in Sweden, is described in Andersson *et al* [3].

Verifying that the preparedness measure (1) actually mirrors what the ambulance operators refer to as preparedness, can only be done by the operators themselves. This can be facilitated by visualising the calculations in the GIS and let the operators determine whether the calculations are correct. In this manner, it will also be made clear if the operators have any use of the measure, i.e. they evaluate the system enhancement at the same time as they verify that it works as intended. However, if the evaluations of the automatic dispatch and the automatic relocation applications show that they have a positive impact on the call service times, this also verifies that the preparedness measure is useful.

The call service times are the most important quality measure of the ambulance logistics service, and they are calculated as the time span from when the call is received by the SOS operator until the ambulance personnel reach the patient. In Andersson *et al* [4] it is shown that the use of dynamic relocations, which are based on the preparedness measure, decreases the call service times. This verifies that the preparedness measure, and the automatic relocation application, works as intended.

It should also be verified that the decisions made, or proposed, by the two automatic suggestion applications are attractive from a user perspective, i.e. that the ambulance operators agree that the decisions are logical. If they do not look like decisions that could have been made by one of the operators, the users will probably never trust the applications enough to use them.

The verification that a simulation model actually simulates the system it is built for, is usually done by feeding the model with historical data and checking that the simulated actions and results correspond to the historical outcome. However, that would be difficult in this case, since there are numerous situations where the decisions made by the simulator may not be the same as the ones made historically, without the simulated decisions being wrong. For example, for a new Prio 2 call there are two ambulances equally suited to serve it; the simulator picks one while the ambulance operator in reality chose the other one. At the time the decision is made, neither the operator nor the simulator can know the full consequences of the assignment, which may differ substantially depending on where the next set of calls sites are located. They can only pick the ambulance that they think is best suited for serving the call, the operator bearing in mind the need for an adequate preparedness, while the simulator uses the preparedness calculations to reach its decision.

The verification that the simulator works therefore has to be done by verifying that the components work as intended, i.e. that the call generator can produce valid scenarios, that the ambulances travel as they should and that the automatic dispatch and relocation applications make logical decisions.

5 Final remarks

In this paper, a number of enhancements to the GIS in the SOS centrals in Sweden has been described and discussed. These include a way of calculating

Management Information Systems, C. A. Brebbia (Editor)

the preparedness in different zones in the area of responsibility, as well as a simulation module and two decision support tools for operational ambulance control.

Currently SOS Alarm collaborates with Carmenta and Linköping University in a project aiming to install the enhancements as external and internal modules to the GIS in the SOS central in Stockholm. Computational testing and theoretical reasoning has shown us the benefits of the new applications, but it is not until the intended users have evaluated them, that their true worth will be clear.

It may be noted though, that the developed applications are not primarily intended for the ambulance controllers who are top of the class with years of experience. They usually do not need this type of decision support, even if their work performance should not suffer from having it. The enhancements are primarily intended for the less experienced controller who has not yet been able to acquire and memorise the vast amount of information required to make a quick and faultless decision every time there is a call for an ambulance. As described earlier, the new applications are also intended for evaluating strategic decisions, which are generally made by the county councils. The intended users in this case are still SOS Alarm employees, since SOS Alarm provides the city councils with the background information necessary for making the decisions.
Even if it the effects from the new applications proves to be small, i.e. the reduction of the call service times are only reduced marginally, a mean reduction of a few seconds for a set of calls can be the difference between life and death when studying as single call. This undoubtedly makes it worthwhile to continuously improve the quality of the ambulance logistics service.

References

[1] Alsaloumm, O., I. & Rand, G., K., A goal-programming model applied to the EMS system at Riyadh City, Saudi Arabia. Working Paper LUMSWP2003/035, Lancaster University Management School, Lancaster, UK, 2003.

[2] Andersson, T., Petersson, S. & Värbrand, P., OPAL – Optimized Ambulance Logistics, *Proc. of TRISTAN V: The Fifth Triennial Symposium on Transportation Analysis*, 2004.

[3] Andersson, T., Petersson, S. & Värbrand, P., Calculating the preparedness for an efficient ambulance health care, Working Paper, Dep. of Science and Technology, Linköping University, 2004.

[4] Andersson, T., Petersson, S. & Värbrand, P., Dynamic ambulance relocation for a higher preparedness, *Proc. of the 35th Annual Meeting of the Decision Sciences Institute*, 2004.

[5] Carmenta, www.carmenta.se.

[6] Ericsson, www.ericsson.com.

[7] Gendreau M., Laporte G. & Semet F., A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. *Parallel Computing*, **27**, pp. 1641-1653, 2001.

Management Information Systems, C. A. Brebbia (Editor)

[8] Goldberg, J., Operations research models for the deployment of emergency service vehicles. *EMS Management Journal*, **1**, pp. 20-39, 2004.
[9] Kamenetzky, R., Suman, L. & Wolfe, H., Estimating the need and demand for prehospital care. *Operations Research*, **30(6)**, pp. 1148-1167, 1982.
[10] Mobitex Association, www.mobitex.org.
[11] SOS Alarm, www.sosalarm.se.
[12] Weintraub, A., Aboud, J., Fernandez, C., Laporte, G. & Ramirez, E., An emergency vehicle dispatching system for an electric utility in Chile. *Journal of the Operational Research Society*, **50**, pp. 690-696, 1999.

Management Information Systems, C. A. Brebbia (Editor)
© 2004 WIT Press, www.witpress.com, ISBN 1-85312-728-0