Estimating the Value of Casualty Health Information to Optimization-Based Decision Support in Response to Major Incidents

Duncan T. Wilson, Glenn I. Hawe, Graham Coates and Roger S. Crouch
School of Engineering and Computing Sciences, Durham University, Durham, UK DH1 3LE
{d.t.wilson, g.i.hawe, graham.coates, r.s.crouch}@durham.ac.uk

ABSTRACT

In this paper we describe a work-in-progress decision support program designed for use in the response to major incidents in the UK. The proposed program is designed for use in a continuous fashion, where the updating of its model, the search for solutions to the model through an optimization algorithm, and the issuing of these solutions are carried out concurrently. The model facilitates the inclusion of dynamic and uncertain features of emergency response. The potential of such an approach to deliver high-quality response plans through enabling more accurate modeling is evaluated through focusing on the case of casualty health information. Computational experiments show there is significant value in monitoring the dynamic and uncertain health progression of casualties and updating the model accordingly.

Keywords

Decision Support, combinatorial optimization, emergency response

INTRODUCTION

When designing a decision support program for use in disaster response, we aim to produce a tool which will supply the decision makers with advice which will assist in the formation of a high quality response operation. In an optimization based program, two components are of fundamental importance – the model and the optimization algorithm used to solve the model. When considering our aim of delivering high quality advice, development on both components can contribute. The contribution of the optimization algorithm is particularly clear, where increasingly sophisticated algorithms can find higher quality solutions in a shorter time. However, the potential for focused model development to increase performance should not be overlooked. Poorly designed models which have neglected to include pertinent details or rely on invalid assumptions will, regardless of the optimization algorithm employed, lead to unrealistic and/or irrelevant advice being passed to the decision maker which, if followed, will result in poor performance. The potential benefit arising from the inclusion of a particular detail or feature into the model can be quantified through computational experiments, and therefore are directly comparable with any benefits afforded through increasingly sophisticated algorithms.

A subset of features in disaster response of particular interest, due to the modeling challenge they pose, are those of a dynamic and uncertain nature where some information important to the decision making process is known to vary over the course of the response operation but the variation is not known with certainty. One example of such information is the number of casualties under consideration. This figure is not always constant, and variations in it may be completely unpredictable. For example, the terrorist attack of London in July 2005 consisted initially of three coordinated, simultaneous explosive attacks followed around one hour later by a fourth, leading to a new incident and set of casualties to be considered by decision makers defining the emergency response operation (London Assembly, 2006). Similarly, the reliability of sections of the transportation network and the resources available at nearby hospitals may also vary in uncertain or unpredictable ways.

The importance of incorporating such dynamic and uncertain features into decision support models has been recognized in previous work, with various techniques employed in an effort to do so. Examples of two-stage stochastic programming being used include (Barbarosoglu & Arda, 2004), where the transportation of resources to points of demand is considered, and (Mete & Zabinsky 2010) where the problem is that of distributing medical supplies. In both cases, the authors identify characteristics for which the dynamic variation consists of a single point in time where previously unknown information is revealed. Prior to the revealing of this information, possible realizations are considered and assigned probabilities, thus acknowledging the uncertain nature of the information. In (Fiedrich et. al, 2000) the uncertain nature of the health of casualties is recognized and an objective function which estimates the expected number of fatalities incurred over the course of the response operation is proposed to address this.

Work such as (Yi & Ozdamar, 2007) have focused on the development and implementation of fast solution algorithms to be used in solving their models, enabling it to be updated and re-solved in a timely fashion and thus facilitate re-planning. Similarly (Rolland et. al, 2010) looks to develop an optimization algorithm capable of solving the proposed model in near real-time, citing that "*This ability… allows managers to re-solve a particular response or recovery problem as conditions at the scene change*". This approach of building a model based on the information currently available, solving this model, and then re-building at a later stage to incorporate new information requires consideration to be given to how long the intervals between each iteration should be. In (Gong & Batta, 2007), where such an approach is adopted within a resource allocation framework, the authors note that too long an interval will lead to out-of-date models, whereas too short a period could lead to frequent reallocation of resources and therefore wasted time.

In this paper we focus on a specific dynamic and uncertain feature of importance to major incident response planning: the health of casualties. The importance of this information is clear from the standard practice within ambulance services (London Ambulance Service, 2007) of carrying out full a full triage before any other operations may commence. The fact that casualties may have different health levels is recognized in work such as (Yi & Ozdamar, 2007), but the fact that these health levels may change frequently throughout the response operation is not. Using a work-in-progress decision support program designed to update its model as and when new information becomes available whilst it is being search for solutions, we estimate the value of incorporating dynamic and uncertain casualty health information.

MODEL AND SOLUTION

In this research, we consider the problem of allocating tasks to emergency responders during the immediate aftermath of a mass-casualty incident. For the purpose of this paper we restrict ourselves to considering response following terrorist attacks, identifying the relevant tasks, responders and environmental features which need to be incorporated into our model. This task-allocation model is linked to a simulation and updated continuously whilst being simultaneously searched by an optimization algorithm, with plans being deployed in a piecewise fashion over a rolling planning horizon.

Environment, tasks and resources

The modeled environment consists primarily of a road transport network, represented as a graph, located on which are sites of interest i.e. hospitals, fire stations and attack sites. Hospitals are given an initial capacity in terms of the number of casualties who may be admitted without the quality of care administered deteriorating. Within the environment, casualties are located at nodes on the transport network and are assigned both an initial health state vector and an indication of whether or not they are trapped under debris and therefore requiring assistance from the Fire & Rescue service through a "rescue" task. Casualty health is represented by a probability vector for each individual $h_c = (p_1, p_2, p_3, p_4)$, where p_i = the estimated probability that the casualty is currently in health state i. Health states correspond to the classifications used in triage sieve operations as implemented by ambulance services in the UK, consisting of Delayed (i.e. walking wounded) (i=1), Urgent (i=2), Immediate (i=3) and Deceased (i=4) (London Ambulance Service, 2007).

The resources we consider are paramedic teams, each of which has an assigned ambulance, and fire-fighter teams, each of which has an assigned fire-appliance. Paramedic teams are initially located at one of the modeled hospitals, while fire-fighter teams are initially placed at fire-stations. The number of resources is set to be constant throughout the response operation. The tasks to be allocated to fire-fighters are rescue tasks, as mentioned above. Paramedic teams may be assigned transportation tasks, detailing which casualty should be taken to which hospital. Paramedics may also be assigned treatment tasks which, unlike rescue and transportation tasks, are not essential to a valid solution and may be omitted. Some paramedics may be trained in HART (Department of Health, 2008) allowing them to perform advanced treatment tasks.

Solution representation, objective function and search algorithm

A solution to this problem is a set of ordered lists, each associated with a specific responder and detailing which tasks they are to carry out and in what order. In addition, for each transportation task the destination hospital must be identified. The resulting set of solutions forms a combinatorial optimization problem. Given a solution, a schedule of the response operation detailing the expected start and end times of each task can be constructed by considering any dependency relations which exist between tasks together with estimates of task duration and travel time.

In order to compare any two solutions, an objective function is required. Having created a schedule for the solution in question, a set of discrete time Markov chains are used to predict changes in the probability health vector of each casualty. Three separate chains are used in order to distinguish between each of the three environments a casualty can find themselves in: trapped, waiting to be rescued; at a Casualty Clearing Station (CCS), waiting to be transported to a hospital; and in an ambulance. The structure and transition probabilities of the three Markov chains, which for the purpose of this paper are assumed, are illustrated in Figure 1.

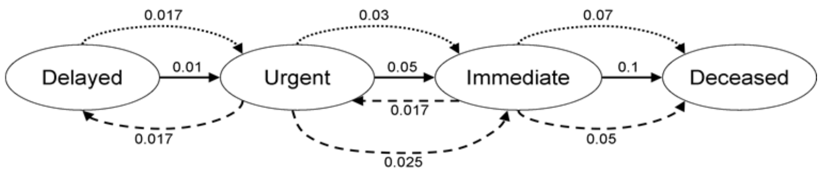

Figure 1. Markov chain structure and transition probabilities, where the solid line corresponds to the "trapped" environment, the dotted line to the CCS environment, and the dashed to the "in ambulance" environment.

Given a health state vector $h_c(t)$ of casualty c at time *t*, if we know they will be in environment *k* for δ minutes then we use the Markov chain MC_k to predict the health state vector of casualty c at time $t + \delta$,

$$h_c(t + \delta) = MC_k(h_c(t), \delta).$$

Given an initial health state vector and denoting by t_c the scheduled hospital arrival time, we can estimate $h_c(t_c)$ for each casualty. We then define our objective function $f(s)$ as the expected number of fatalities after all casualties have been transported to hospital:

$$f(s) = (0\ 0\ 0\ 1) * \sum_c h_c(t_c).$$

We note that this calculation assumes that if a casualty is alive at the point of entering a hospital, they do not go on to die. In order to search the model's solution space, a simple "best improvement" local search metaheuristic has been implemented, where at each iteration all neighbors of the current best known solution are evaluated and the best improvement is selected.

Online framework and simulation

Our model is designed for use in an online capacity, where the computation timeline associated with the optimization algorithm is tied to the timeline of the response operation. Our framework continually updates the model as and when new information is received while it searches the solution space. As opposed to starting the search process, waiting for some termination criteria to be met and then issuing the full resulting schedule, in the online framework a responder queries the program as it nears the completion of its current task. At this point the program consults the best schedule found so far to determine which task should next be issued to the responder in question. This process is illustrated in Figure 2, where we see tasks being issued at time points t*, after which point they are fixed in the schedule while the remaining tasks can still be adjusted in the optimization process.

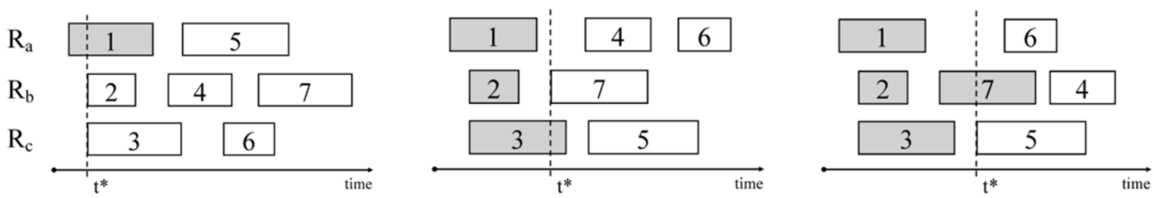

Figure 2. An illustration of a rolling schedule being issued to three resources in real time.

In order to test the online framework, a source of real time information is required to replicate that which might be available during an actual response operation. For the purpose of this paper, a relatively simple simulation has been implemented to achieve this. The simulation works alongside the decision support program, receiving notification of tasks being issued to responders. At these points, initial estimates of the duration of the task are revised and improved to reflect the fact that the responder at the scene will be better equipped to make such estimates. Task duration estimates are updated once again upon the completion of each task. Finally, the health of each casualty is also simulated in a discrete sequence, using the same Markov chains used by the objective function and shown in Figure 1.

RESULTS

One hypothetical problem was used for all experiments, and was set-up using the *Scenario Designer* of the STORMI package described in (Hawe et. al, 2012). The scenario involves five separate simultaneous terrorist attacks across London, each resulting in 18 casualties. The first thirty minutes of the response operation are considered, involving eleven paramedic teams together with five fire-fighter teams.

In order to estimate the value of incorporating dynamic casualty health modeling, two policies were compared. In the first, the health of each casualty is revealed only once, at the outset of the response operation. In the second case, updated information revealing the current health state of each casualty is received every minute, resulting in thirty updates per casualty over the course of the thirty minutes considered. The progression of the search process over the first five minutes of a single run, for both cases, is shown in Figure 3 for illustrative purposes, demonstrating the effect on the objective of regular information updates.

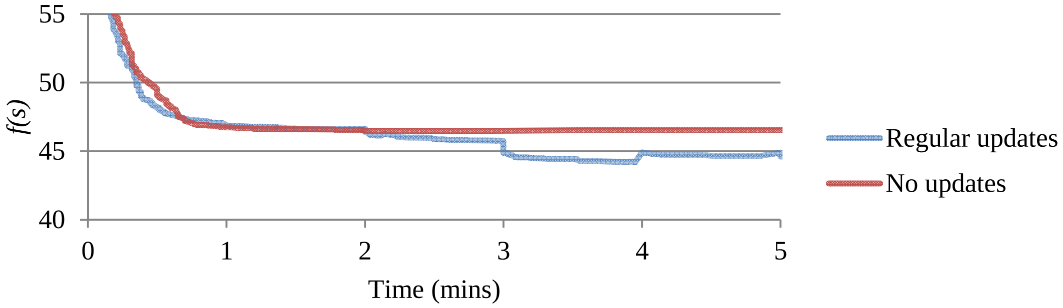

Figure 3. Performance over a single run when updates are received every minute and when no updates are received.

In addition to these extreme cases, several intermediate policies were examined, corresponding to a varying number of evenly spaced updates of casualty health being received over the thirty minutes under consideration. Specifically, we considered policies of two, four and ten updates. In all cases, at the end of the thirty minutes the resulting schedule is evaluated using the actual health states progression of casualties as recorded by the simulation, so that comparisons may be made. Fifty runs for each level of update frequency under consideration were completed, with the resulting mean and standard deviation shown in Figure 4(a).

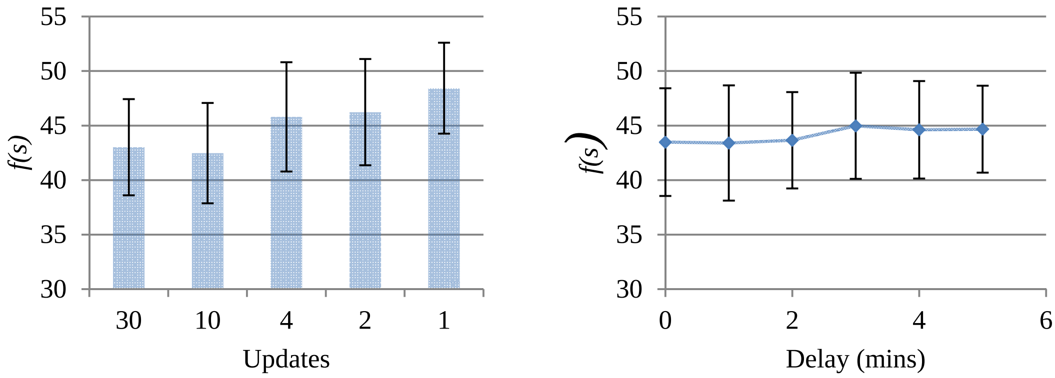

Figure 4. (a) Final objective values against the number of updates over the thirty minute period, (b) final objective value against the delay in updates reaching the model. Error bars in both charts show one standard deviation.

Further experiments were carried out to examine the effect a delay in the updates reaching the model may have on performance. For each instance, a policy of constant updating was employed, with a delay of one to five minutes imposed. The results are shown in Figure 4(b).

We observe that updating every minute reduces the objective value from an average of 48.43 to 43.02 expected fatalities when compared with updating only once, thus providing an estimate of the value of casualty health information. Improvements are also observed across the intermediate policies. The relationship between objective value and the delay in information updates reaching the model is less pronounced, with a five minute delay resulting in 1.20 more expected fatalities, on average.

CONCLUSIONS AND FURTHER WORK

An online decision support program for use in response to major incidents has been described. The model employed within the program allows for the inclusion of features pertaining to the dynamic and uncertain characteristics observed in the response environment, such as the variability in casualty health and the duration of tasks. Examining the case of casualty health, computational experiments were carried out to ascertain if inclusion of this detail in our decision support model would result in higher performance, which was confirmed to be the case. Furthermore, the dependency of this increased performance on the speed with which these updates are passed to our model was examined, with relatively small penalties incurred for delays of up to five minutes.

Future research in this project will focus on further experimental analysis into the benefits of including other pertinent details into the proposed model, such as updates on the state of the transport network. We also note that the online framework described in this paper presents challenges in the design of optimization algorithms, due to the solution space constantly changing in both size and shape. Future work will focus on developing adaptive algorithms capable of adjusting to suit these changes.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the funding provided by the UK's EPSRC (EP/G057516/1). The authors also thank practitioners from the Emergency Planning Units of Cleveland and Tyne & Wear, Co. Durham & Darlington Civil Contingencies Unit, Government office of the North East, Fire and Rescue services of Co. Durham & Darlington and Tyne & Wear, North Eat Ambulance Service, and Northumbria Police.

REFERENCES

1. Altaya, N. and Green, W. G. (2006) OR/MS Research in Disaster Operations Management. *European Journal of Operational Research*, 175, 475-493.
2. Barbarosoglu, G. and Arda, Y. (2004) A Two-Stage Stochastic Programming Framework for Transportation Planning in Disaster Response. *Journal of the Operational Research Society*, 55, 43-53.
3. Chiu, Y-C. and Zheng, H. (2007) Real-Time Mobilization Decisions for Multi-Priority Emergency Response Resources and Evacuation Groups: Model Formulation and Solution. *Transportation Research Part E: Logistics and Transportation Review*,43, 710-736.
4. Hawe, G. I., Wilson, D. T., Coates, G. and Crouch, R. S. (2012) STORMI: An Agent-Based Simulation Environment for Evaluating Responses to Major Incidents in the UK, *Proc. of ISCRAM 2012*, Vancouver
5. Department of Health (2008) Ambulance Staff In The Inner cordon: Hazardous Area Response Teams (HART). Available from: www.ambulancehart.or.uk
6. Fiedrich, F., Gehbauer, F. and Rickers, U. (2000) Optimized Resource Allocation for Emergency Response After Earthquake Disasters. *Safety Science*, 35, 41-57
7. Gong, Q. and Batta, R. (2007) Allocation and Reallocation of Ambulances to Casualty Clusters in a Disaster Relief Operation. *IIE Transactions*, 39, 27-39.
8. London Ambulance Service (2007) *Major Incident Plan.* Available from: www.londonambulance.nhs.uk/about-us/what-we-do/dealing -with-major-incidents.aspx
9. London Assembly (2006) *Report of the 7 July Review Committee.* Available from legacy.london.gov.uk/assembly/reports/7july/report.pdf
10. Mete, H. O., Zabinsky, Z.B. (2010) Stochastic Optimization of Medical Supply Location and Distribution in Disaster Management. *International Journal of Production Economics*, 126, 76-84
11. Rolland, E., Patterson, R. A., Ward, K., and Dodin, B. (2010) Decision Support for Disaster Management. *Operations Management Research*, 3, 68-79
12. Simpson, N. C., Hancock, P. G., (2009) Fifty Years of Operational Research and Emergency Response. *Journal of the Operational Research Society*, 60, SI26-SI39
13. Yi, W., Ozdamar, L. (2007) A Dynamic Logistics Coordination Model for Evacuation and Support in Disaster Response Activities. *European Journal of Operational Research*, 179, 1177-1193.