Scheduling Response Operations under Transport Network Disruptions

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

Modeling the complex decision problems faced in the coordination of disaster response as a scheduling problem to be solved using an optimization algorithm has the potential to deliver efficient and effective support to decision makers. However, much of the utility of such a model lies in its ability to accurately predict the outcome of any proposed solution. The stochastic nature of the disaster response environment can make such prediction difficult. In this paper we examine the effect of unknown disruptions to the road transport network on the utility of a disaster response scheduling model. The effects of several levels of disruption are measured empirically and the potential of using real-time information to revise model parameters, and thereby improve predictive performance, is evaluated.

Keywords

Routing, scheduling, optimization, disaster response

INTRODUCTION

Many decision problems encountered in disaster response require consideration of how emergency responders move around the affected geographic area. In making such decisions we may be concerned with two related aspects. Firstly, the routing problem - which route should a responder take in order to move from location *A* to location *B* in as short a time as possible? Secondly, the prediction problem - how can we predict how long it will take a responder to move from *A* to location *B*? The degree to which these sub-problems can be solved effectively may have a significant impact on the overall decision problem, and so deserve our attention.

Under certain assumptions regarding the nature of the disaster response environment both the routing and prediction sub-problems can be solved quickly and effectively using well-established algorithms from the fields of computer science and operational research. In particular, if one can represent the affected geographical area as a set of nodes, each corresponding to a location of interest (for example, a hospital) or a road junction, together with a set of arcs linking nodes together, each with a parameter describing its distance, shortest path algorithms such as Djikstra's algorithm can be employed to find the optimal route between any two locations. If one assumes a constant speed of travel for responders along such routes, the associated travel time can then be derived. However, this assumption may be hard to justify in a disaster environment. Given a specified route and its distance it may not be possible to accurately predict the actual time needed to travel along the route as parts of it may be subject to unknown disruptions of some form, such as damage caused directly by the disaster or severe traffic congestion caused by the unpredictable behavior of civilians. Under such conditions both the routing and prediction sub-problems become significantly more difficult to solve.

Recent research in the development of optimization models for disaster management has taken a range of approaches in representing the transport network. Considering the problem of prepositioning relief supplies prior to the occurrence of a disaster, Tzeng, Cheng and Huang (2007) assume the "availability and accessibility of information", including the state of the transport network. In contrast, Rawls and Turnquist (2010) propose a two-stage stochastic programming model to assist in its solution. Such an approach allows for the consideration of uncertainty in the transport network parameters following a disaster. A similar approach is employed by Mete and Zabinsky (2010) in a model which aims to minimize the time needed to transport relief commodities to sources of demand. The routes used in this transportation are pre-computed before the disaster and the travel times on each route are assumed known.

The problem of assigning ambulances to clusters of casualties is addressed in (Gong and Batta, 2007), where travel times of ambulances along routes are assumed to be constant. The model is extended in (Jotshi, Gong and Batta, 2009) to acknowledge the uncertain and dynamic nature of the transport network. A data fusion approach

is used in estimating the level of damage and disruption on each road link, as categorized into one of five discrete levels. The transportation of casualties to hospital is also considered by Yi and Ozdemar (2007), where the authors employ a vehicle routing formulation assuming all transport network parameters are known.

The transport network considered in the scheduling model proposed by Fiedrich, Gehbauer and Rickers (2000) is of a dynamic nature in that nodes and links can be added and removed to reflect further developments of the disaster. The model does assume, however, that all such information is known with certainty in the scheduling model. Stochastic travel times are modeled in the work of Wex, Schryen and Neumann (2012), where travel times between all incidents are modeled as normally distributed random variables. Fuzzy logic is employed in the optimization model to account for such uncertainty. A similar scheduling approach is employed by Wilson, Hawe, Coates and Crouch (2012), where travel times are assumed known and constant while the durations of response tasks are of a stochastic nature. The empirical analysis presented demonstrates the significant impact of such stochastic durations upon the utility of the optimization model, which suggests other temporal variability (such as that arising through transport network disruption) would also have a significant effect.

While uncertain disruption in the transport network of the environment following a disaster has been modeled to varying degrees in recent research, there has yet to be an empirical analysis into the effect of such disruption on the utility of an optimization model. In this short paper we propose to extend the work presented in (Wilson et al., 2012) and perform such an analysis, in the context of a scheduling model designed for use during a Mass Casualty Incident (MCI) response, with a particular focus on the prediction sub-problem.

MODEL AND ANALYSIS

The scheduling model & an example problem

In this paper we employ the same scheduling model of an MCI response described by Wilson et al. (2012). Briefly, the combinatorial optimization model defines a solution to the response problem through three components: an allocation of tasks to available responders; an ordering of all tasks assigned to each responder; and an allocation of each casualty to one of the available hospitals. Solutions are evaluated through a multi-objective function which predicts the number of fatalities resulting from the proposed solution together with a dimensionless measure of the suffering endured by casualties, which involves measuring how long each casualty must wait before being admitted to a hospital together with the level of care available at that hospital. We denote these measures by f_1 (fatalities) and f_2 (suffering), omitting details of their form and instead referring the reader to (Wilson et al., 2012) for further details.

Throughout this paper we shall refer to a single example scenario used when conducting all computational experiments. The scenario involves three separate incident sites across central London, with each incident resulting in 70 casualties. The response resources available consist of 53 ambulances (with crew) and 27 fire appliances (with crew). The environment includes a graph representing the central London road network at a fine level of detail. In solving the problem using the scheduling model described earlier, a number of tasks relating to each casualty (namely their extrication, treatment and transportation to hospital) must be assigned to appropriate responders and ordered in such a way as to minimize the objectives f_1 and f_2.
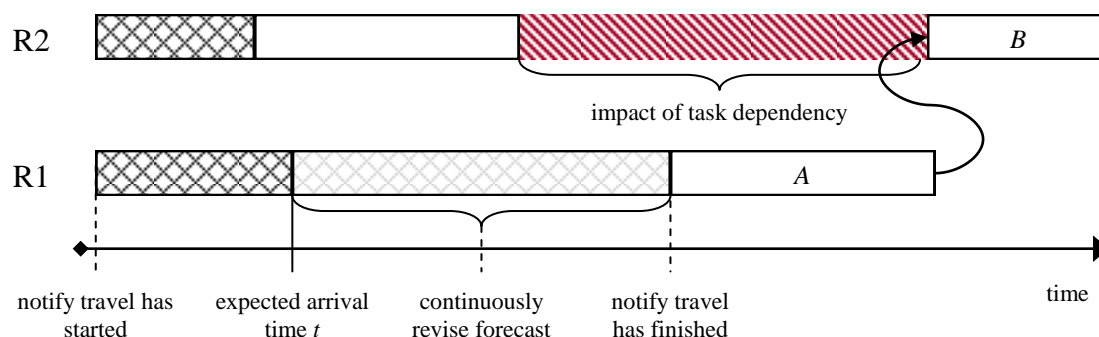

Figure 1. A possible consequence of inaccurate travel time forecasts due to transport network disruption, where tasks assigned to responders R1 and R2 are represented by clear boxes and periods of travel are represented by crisscrossed boxes.

Proceedings of the 10ᵗʰ International ISCRAM Conference – Baden-Baden, Germany, May 2013
T. Comes, F. Fiedrich, S. Fortier, J. Geldermann and T.Müller, eds.

684

Predicting and simulating travel times

In order to calculate the objective values of any given solution we must first forecast a schedule of the response operation by predicting the start and end time of each task. One consideration in making such predictions is the effect of routing and travel times. As illustrated in Figure 1, an inaccurate initial prediction of a travel time can impact the finishing time of related tasks assigned to other responders as some tasks are constrained through dependency relations, as in the case of tasks *A* and *B*. That is, a single error in the prediction of a travel time may propagate through a large section of the remaining schedule, leading to a significant negative impact.

Figure 1 illustrates the fact that the problem considered is of an inherently dynamic nature, with information being revealed in a gradual manner in real time. As such, the optimization model described has been designed for use in a real time, online manner, where the model continuously updates relevant parameters as and when information becomes available. For example, the scheduling model predicts task A to begin at time *t*. Upon reaching time *t*, the scheduling model continuously updates this expected starting time until confirmation is received that responder R2 has reached its target destination and can begin work on task A.

Predicting travel times

Travel times are estimated using the model described by Kolesar, Walker and Hausner, (1975), as recently validated by Budge, Ingolfsson and Zerom, (2010). The formula gives the median travel time of an ambulance travelling from A to B given the distance travelled. In order to find this distance, *d*, Djikstra's algorithm is used in conjunction with the transport network represented as a graph, with the weight of each edge corresponding to its distance. This distance is then used to give the median travel time,

$$m(d) = \begin{cases} 2\sqrt{d/a} & d \le 2d_c \\ v_c/a + d/v_c & d > 2d_c \end{cases} \qquad (1)$$

where $d_c = v_c^2/(2a)$ denotes the distance required to travel in order to reach "cruise speed" v_c and *a* is the average acceleration of the vehicle as it increases speed to v_c . The values of these parameters are taken from the analysis of ambulance travel times in Calgary, Canada, presented in (Budge et al., 2010).

Simulating travel times

Whereas the scheduling model uses the median travel time of a given route in its predictions, in a simulation of a response operation the actual travel times will be subject to random variation around this median. As discussed by Westgate, Woodward, Matteson and Henderson (2011), travel times may be modeled as random variables X following a lognormal distribution $X \sim logN(\mu, \tau)$ with an assumed $\tau = 0.00227$. Noting that the median of such a distribution is given by $m = e^{\mu}$ (Johnson, Kotz and Balakrishnan, 1995), the parameter μ can be calculated using the median value obtained in (1).

In order to model the effect of transport network disruption of a given route we modify the distance parameter of each of the links involved. Specifically, a random variable *Y~exp(ρ)* is sampled for each link, the distance of which is then multiplied by the factor (1+*Y*). As such, we can interpret *ρ* as a parameter representing the level of disruption to the transport network. For example, setting *ρ* = 0.5 will lead to the distance parameter of each link on the road network being increased on average by a factor of (1+E[*Y*]) = (1 + 2) = 3. We note that this approach is a pragmatic one with the purpose of generating random, parameterized disruption to travel times and optimal routes, and is not intended to realistically represent the details of a transport network disruption.

Improving prediction through online learning

In order to improve prediction we note that, while a perfect knowledge of the effects of disruption from time t = 0 may not be feasible, we can gather information as the response operation progresses and use this information to assist in better predictions for the remainder of the operation.

As stated previously, travel times along any route with a given level of disruption are assumed to follow a lognormal distribution with an assumed, constant precision$\tau, X \sim logN(\mu, \tau)$. In order to incorporate knowledge of observed travel times we employ a Bayesian approach in estimating the unknown parameter μ. Specifically, using the conjugate prior distribution for $\mu, \mu \sim N(\mu_0, \tau_0)$, we can calculate the posterior distribution following the observation of *n* data $x_i, \mu \sim N(\mu_n, \tau_n)$ where

$$\mu_n = \frac{\tau_0\mu_0 + \tau \sum_{i=1}^{n} \ln(x_i)}{\tau_0 + n\tau}, \quad \tau_n = \tau_0 + n\tau.$$

Proceedings of the 10th International ISCRAM Conference – Baden-Baden, Germany, May 2013
T. Comes, F. Fiedrich, S. Fortier, J. Geldermann and T.Müller, eds.

685

The expectation of this posterior distribution, $\hat{\mu} = E(\mu)$, is then used as an estimate of μ, giving $X \sim logN(\hat{\mu}, \tau)$. As noted previously, the median travel time for the route in question can then be estimated as $m = e^{\hat{\mu}}$. This routine is carried out for each single observation x immediately upon its observation.

Experimental analysis

In order to assess how sensitive the performance of the scheduling model is to uncertain disruptions in the transport network, a number of experiments were conducted according to the Monte Carlo method with various values of ρ. In each experiment the scheduling model was run over the course of an entire response operation. All routes between locations of interest (such as incident sites and hospitals) were found on initialization using the standard (undisrupted) network parameterization. The network disruption was then applied in the manner described earlier, resulting in all travel times being extended by a random amount. As the operation progresses the scheduling model is notified of when each trip is completed and can therefore maintain an accurate picture of all travel times up to the current point in time. All future travel times continue to be predicted using the pre-calculated values. The results of these experiments are contrasted with corresponding cases where the scheduling model is assumed to know of the disrupted travel times, and can therefore make more accurate predictions. The results are shown in Figure 2.
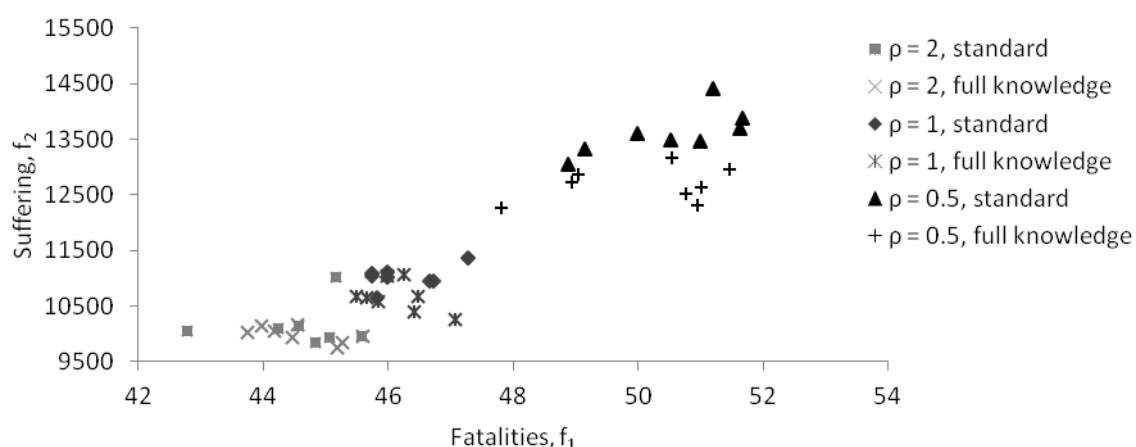

Figure 2. Comparisons of standard and full knowledge regarding network parameters for three levels of disruption.

The results illustrate the benefit to the scheduling model of possessing full knowledge in terms of transport network disruption. For the cases $\rho = 2, 1$ the observed difference in performance in both objective terms was not statistically significant under t-tests for means with assumed unequal variance. Moving to $\rho = 0.5$, however, results in a statistically significant difference in means from 13615 to 12682 in the case of f_2 $(p \approx 0)$. We emphasize here that the difference in performance is due entirely to the ability to solve the prediction sub-problem, as the routes taken by responders to get from A to B are always the same.

In the preceding analysis we have shown that the availability of information regarding the state of disruption of the transport network has a significant effect on the utility of the scheduling model when the disruption parameter is set to 0.5. In order to evaluate the potential of online learning to increase the performance of the scheduling model through facilitating better predictions of travel times, a number of experiments employing the methodology on the same example problem described previously were carried out. The average value of f_1 under this approach was 50.8 with a standard deviation of 1.0. The corresponding values of f_2 were 13348 and 299. In comparison to those value observed for the same problem but where no information of network disruption was gathered, this results no statistically significant difference in f_1. In terms of f_2, an average difference of 267 was observed with a p value $p = 0.079$. This mild improvement suggests that information is most beneficial in the early stages of scheduling, before many trips have been made and travel times recorded.

CONCLUSIONS AND FURTHER WORK

The analysis presented in this paper has demonstrated the importance of acknowledging the possibility of disruptions to the transport network when attempting to deliver decision support during an MCI response through a scheduling model. In particular, we have demonstrated the benefit of obtaining information regarding the extent of network disruption, both in the idealized case of full knowledge and a realistic case where travel

Proceedings of the 10th International ISCRAM Conference – Baden-Baden, Germany, May 2013
T. Comes, F. Fiedrich, S. Fortier, J. Geldermann and T.Müller, eds.

686

times are recorded and the model parameters are updated in real time.

The presented analysis has assumed throughout that the route used in moving between any two locations can always be calculated prior to the incident and is based on finding the route of minimum distance. Following a transport network disruption, however, alternative routes with lower expected travel time may exist. Future work will focus on reviewing and developing appropriate methodology to find these improved routes, in particular considering a pragmatic approach where, rather than instructing each responder on which route to take, we let responders decide for themselves. We assume that responders would combine local knowledge with exploration of the environment to gradually find routes of decreasing median travel time – indeed, this is exactly the situation recorded from a responder to the Haiti earthquake by de la Torre, Dolinskaya and Smilowitz (2012), who note that drivers had "no maps with updated information and had to discover the best routes by driving and exploring". This approach will likely lead to shorter travel times. However, by releasing control over routing decisions it may also lead to increased variability and therefore make prediction more difficult.

ACKNOWLEDGEMENTS

REFERENCES

1. Budge, S., Ingolfsson, A. and Zerom, D. (2010) Empirical Analysis of Ambulance Travel Times: The Case of Calgary Emergency Medical Services. *Management Science*, 56, 716-723.

2. de la Torre, L. E., Dolinskaya, I. S. and Smilowitz, K. R. (2012) Disaster Relief Routing: Integrating Research and Practice. *Socio-Economic Planning Sciences*, 46, 88-97.

3. Fiedrich, F., Gehbauer, F. and Rickers, U. (2000) Optimized Resource Allocation for Emergency Response After Earthquake Disasters. *Safety Science*, 35, 41-57.

4. Gong, Q. and Batta, R. (2007) Allocation and Reallocation of Ambulances to Casualty Clusters in a Disaster Relief Operation. *IIE Transactions*, 39, 27-39.

5. Johnson, N. L., Kotz, S. and Balakrishnan, N. (1995) Continuous Univariate Distributions. Wiley, New York.

6. Jotshi, A., Gong, Q. and Batta, R. (2009) Dispatching and Routing of Emergency Vehicles in Disaster Mitigation using Data Fusion. *Socio-Economic Planning Sciences,* 43, 1-24.

7. Kolesar, P., Walker, W. and Hausner, J. (1975) Determining the Relation between Fire Engine Travel Times and Travel Distances in New York City. *Operations Research*, 23, 614-627.

8. 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.

9. Rawls, C. G. and Turnquist, M. A. (2010) Pre-Positioning of Emergency Supplies for Disaster Response. *Transportation Research Part B: Methodological,* 44, 521-534.

10. Tzeng, G-H., Cheng, H-J. and Huang, T. D. (2007) Multi-Objective Optimal Planning for Designing Relief Delivery Systems. *Transportation Research Part E: Logistics and Transportation Review*, 43, 673-686.

11. Westgate, B. S., Woodward, D. B., Matteson, D. S. and Henderson, S. G. (2011) Travel Time Estimation for Ambulances using Bayesian Data Augmentation. Technical report, Cornell University.

12. Wex, F., Schryen, G. and Neumann, D. (2012) Operational Emergency Response under Informational Uncertainty: A Fuzzy Optimization Model for Scheduling and Allocating Rescue Units, *Proc. of ISCRAM 2012,* Vancouver.

13. Wilson, D. T., Hawe, G. I., Coates, G., Crouch, R. S. (2012) Stochastic Task Durations in the Scheduling of Mass Casualty Incident Response Operations, *Proc. of PlanSIG 2012,* Middlesborough.

14. 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.

15. Zhang, J-H., Li, J. and Liu, Z-P. (2012) Multiple-Resource and Multiple-Depot Emergency Response Problem Considering Secondary Disasters. *Expert Systems with Applications,* 39, 11066-11071.

Proceedings of the 10th International ISCRAM Conference – Baden-Baden, Germany, May 2013
T. Comes, F. Fiedrich, S. Fortier, J. Geldermann and T.Müller, eds.

687