European Journal of Operational Research 252 (2016) 334–348

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Online optimization of casualty processing in major incident response: An experimental analysis

Duncan T. Wilson*, Glenn I. Hawe, Graham Coates, Roger S. Crouch

School of Engineering and Computing Sciences, Durham University, South Road, Durham, DH1 3LE, UK

A R T I C L E I N F O

Article history:
Received 18 July 2013
Accepted 13 January 2016
Available online 22 January 2016

Keywords:
Scheduling
Combinatorial optimization
Emergency response
Disaster management
Dynamic

A B S T R A C T

When designing an optimization model for use in mass casualty incident (MCI) response, the dynamic and uncertain nature of the problem environment poses a significant challenge. Many key problem parameters, such as the number of casualties to be processed, will typically change as the response operation progresses. Other parameters, such as the time required to complete key response tasks, must be estimated and are therefore prone to errors. In this work we extend a multi-objective combinatorial optimization model for MCI response to improve performance in dynamic and uncertain environments. The model is developed to allow for use in real time, with continuous communication between the optimization model and problem environment. A simulation of this problem environment is described, allowing for a series of computational experiments evaluating how model utility is influenced by a range of key dynamic or uncertain problem and model characteristics. It is demonstrated that the move to an online system mitigates against poor communication speed, while errors in the estimation of task duration parameters are shown to significantly reduce model utility.

1. Introduction

In the period immediately following a mass casualty incident (MCI), such as the London Bombings of July 7th 2005 (London Assembly, 2006), many decisions need to be made in a fast and effective manner within a high pressure environment (Paton & Flin, 1999). Within emergency response organizations such as the Ambulance Service and the Fire and Rescue Service, decision makers must decide how best to allocate their limited resources amongst the various sources of demand. This problem environment exhibits a large amount of structure, with well defined roles and responsibilities and a clear decision making system as defined through the command and control system (Wallace & de Balogh, 1985). In this respect, the problem represents a strong candidate for the application of mathematical modeling and optimization. However, significant challenges remain, particularly with respect to the volatile nature of the problem environment. That is, the nature of any decision problem is likely to change over time as the problem evolves, and the available information upon which a model can be built will typically be subject to a significant level of uncertainty (Galindo & Batta, 2013).

In the modeling of MCI response, as with the design of any optimization model, it is necessary to make certain assumptions in order to ensure the implementation remains feasible. In this paper we seek to gain a better understanding of several characteristics of the response problem, their associated assumptions, and the extent to which they affect the utility of a scheduling-based optimization model. In order to proceed we first discuss a number of assumptions common to optimization models for MCI response. We cover the modeling of casualty health, their allocation to hospitals for treatment, the transportation of casualties and responders around the response environment, and the representation of tasks which responders must carry out. We go on to focus on how others have considered the dynamic and uncertain nature of the response environment in their models. Based on our findings, we identify gaps that remain uncovered in the literature and we discuss how our research contributes to fill such gaps.

1.1. Common modeling assumptions

Some common assumptions made in the design of operational research models for disaster operations management are identified in Galindo and Batta (2013). Further common assumptions covering the more general area of disaster planning are listed in Auf der Heide (2006).

Depending on the general form of the model, the parameters needed to specify its form can include variables such as

* Corresponding author. Tel.: +44 1904 659005.
E-mail address: d.t.wilson@dur.ac.uk, d.t.wilson@leeds.ac.uk, duncantwilson@gmail.com (D.T. Wilson).

http://dx.doi.org/10.1016/j.ejor.2016.01.021
0377-2217/

commodity supply and demand levels, resource requirements for specific tasks, and the number and nature of casualties. As noted in Galindo and Batta (2013), it is common for models to assume that

1. the information needed to deduce these parameters is available and accurate upon initialization of the model, and
2. the parameters are not required to change over time.

The extent to which these assumptions are justified depends on the specific problem under consideration, but will often be limited by the intrinsic uncertainty and volatility common to all emergency response problems. Some specific examples follow.

1.1.1. Casualty health

Some authors assume there are no meaningful differences between the health levels of casualties (Barbarosoglu & Arda, 2004; Barbarosoglu, Ozdamar, & Cevik, 2002; Mete & Zabinsky, 2010; Rolland, Patterson, Ward, & Dodin, 2010; Wex, Schryen, & Neumann, 2011; 2012). Where differences are acknowledged, it is common to assume all casualties have been partitioned into discrete categories reflecting the urgency of their treatment (Galindo & Batta, 2013), as in the work of Chiu and Zheng (2007); Gong and Batta (2007); Yi and Ozdamar (2007). This is reasonable, as it is normal for an assessment of the health of each casualty (known as triage) to be completed before the remainder of the response is enacted (Group, 2011). It is often assumed that individual casualty health will not change over time, and that assessments of health are always accurate. The attraction of the former assumption is understandable, as the task of accurately forecasting the changing health of casualties in these environments is challenging. Some attempts are described in Cotta (2011); Fiedrich, Gehbauer, and Rickers (2000); Tatomir and Rothkrantz (2006). These models, however, do not provide any way to correct errors in prediction, an occurrence which we can assume to be likely due to the complexity of the underlying process.

1.1.2. Hospitals

Many models assume that the allocation of casualties to hospitals will be done automatically and appropriately. Limited examples of including hospital allocation into a wider decision problem can be found in Jotshi, Gong, and Batta (2009); Mysore et al. (2005); Wilson, Hawe, Coates, and Crouch (2013a). In Wilson et al. (2013a) an often ignored aspect of casualty management, *self presentation*, is discussed. It is often assumed that all casualties are transported to hospital by the Ambulance Service only (Auf der Heide, 2006), with the casualty undergoing triage and treatment operations prior to this. In reality, it is common for some casualties to remove themselves from the incident site and transport themselves to a hospital of their choosing. In Wilson et al. (2013a) it is assumed that this process could be predicted accurately. In scenarios where this is not possible, a dynamic approach, updating the model regarding the number of casualties who have left the incident scene and who have arrived at each hospital, may be effective.

1.1.3. Transportation

The transport network within the problem environment is often assumed to be known, both in terms of topology and the travel times between locations (Yi & Kumar, 2007; Zhang, Li, & Liu, 2012). As noted in Galindo and Batta (2013), the former assumption is more justified than the latter. Examples of removing the latter assumption include (Wilson, Hawe, Coates, & Crouch, 2013b). In this work it is demonstrated that disruption to the network resulting in uncertainty in travel times can have a significant effect on the performance of an optimization model. As such, this problem characteristic should not be ignored.

Uncertainty in the disruption of the transport network has been incorporated to a limited extent using stochastic programming formulations. Examples include (Barbarosoglu & Arda, 2004; Mete & Zabinsky, 2010; Rawls & Turnquist, 2010), which consider a finite number of scenarios, each with assigned probability and associated network parametrization. Uncertainty is also acknowledged in the work of Jotshi et al. (2009), which extends the ambulance allocation model presented in Gong and Batta (2007) by including a data fusion step to estimate the level of damage and disruption on each road link. A solution methodology for finding optimal paths in a disrupted network following a disaster is presented in (Zhang, Zhang, Zhang, Wei, & Deng, 2013). The authors employ the network representation described by Yuan and Wang (2009), where the travel time associated with each edge of the transport network is assumed to increase over time in a manner which reflects its proximity to the disaster. A dynamic transport network structure is also modeled in the work of Fiedrich et al. (2000), with nodes and edges being added or removed to reflect the impact of both the disaster and the response operation.

1.1.4. Task durations

Where the modeling methodology involves the allocation of discrete tasks to available responder units, the times needed to complete these tasks are necessary problem parameters. Examples include the scheduling models presented in Rolland et al. (2010) and Wex et al. (2011). In the former, the authors propose a specific solution algorithm which, through its fast execution, is designed to facilitate the solving of their proposed model in near-real time. The authors argue this will allow decision makers to re-solve any particular response problem when conditions change, although this capability is not explicitly tested and evaluated. In Wex et al. (2011) a similar modeling methodology is proposed, where all necessary parameters are assumed to be fixed and known upon model initialization. This model is extended in Wex, Schryen, and Neumann (2012), allowing for task durations to be represented by fuzzy values in an effort to acknowledge the uncertainty inherent in available information. The authors suggest the model should be regularly rebuilt and solved when the problem environment has evolved by some significant degree.

1.2. Modeling uncertainty and dynamicity

All the assumptions mentioned relate to model parameters which change over time, either because they are estimates of unknown real values and can therefore be revised as new information comes to light, or because the real values themselves are of a dynamic nature, or both. In the worst cases these assumptions will render a model unusable in many realistic scenarios. General strategies to their removal tend to take either a stochastic yet static approach, applying stochastic (Barbarosoglu & Arda, 2004; Chang, Tseng, & Chen, 2007; Mete & Zabinsky, 2010) or robust (Bozorgi-Amiri, Jabalameli, Alinaghian, & Heydari, 2012) programming to find solutions which will remain valid as the problem evolves over time, or a dynamic approach, allowing for the model to be updated at a number of set length intervals to help ensure it remains applicable (see, for example, Lee, Ghosh, & Ettl, 2009; Ozdamar, Ekinci, & Kucukyazici, 2004; Yi & Kumar, 2007). Only limited steps have been taken with the latter approach. In the context of manufacturer or retailer response to hurricanes, the supply chain models proposed in Lodree and Taskin (2009); Taskin and Lodree (2011) employ a Bayesian approach to allow for dynamic information to be incorporated into future decisions. In Gong and Batta (2007) the authors note that determining the appropriate length of update interval is crucial to performance, proposing that future work should