Expert Systems With Applications 59 (2016) 47–59

Contents lists available at ScienceDirect

Expert Systems With Applications

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

A decomposition-based heuristic for stochastic emergency routing problems

Belleh Fontem[a,*], Sharif H. Melouk[b], Burcu B. Keskin[c], Naeem Bajwa[d]

[a] College of Business, University of Mary Washington, 1301 College Avenue, Fredericksburg, VA 22401, United States
[b] Department of Information Systems, Statistics and Management Science, University of Alabama, Tuscaloosa, AL 35487-0226, United States
[c] Department of Information Systems, Statistics and Management Science, University of Alabama, Tuscaloosa, AL 35487-0226, United States
[d] Department of Management, College of Business, University of Arkansas at Little Rock, Little Rock, AR 72204, United States

ARTICLE INFO

Article history:
Received 20 September 2015
Revised 27 December 2015
Accepted 2 April 2016
Available online 21 April 2016

Keywords:
Robust emergency planning
Routing under uncertainty
Decomposition-based algorithms
Crisis management

ABSTRACT

This paper proposes a decomposition-based heuristic for a network delivery problem in which relief workers acquire valuable emergency supplies from relief warehouses, and transport them to meet the urgent needs of distressed population centres. The problem context dictates that the relief items reach these population centres before critical deadlines. However, co-ordination challenges and random disruptions introduce uncertainty in both network travel times and the destination deadlines. Hence, relief workers have to negotiate the tension between ensuring a high probability of punctual delivery and maximising the combined value of the relief supplies delivered. For an arbitrary routing scheme which guarantees punctual delivery in an uncertainty-free state of nature, the heuristic yields an upper bound on the probability that, under uncertainty, the routing scheme described will lead to tardy delivery. We demonstrate our solution approach on a small numerical example and glean insights from experiments on a realistically sized problem. Overall, our central model and proposed solution approach are useful to managers who need to evaluate routing options and devise effective operational delivery plans in humanitarian crisis situations.

1. Introduction

Problems in humanitarian and disaster logistics have been studied for a long time, but have recently gained prominence in light of the increased threat of disasters (floods, hurricanes, tornadoes), which have been made more likely by accelerating climate change. The experience of events like the Japanese tsunami of 2011 indicates that such disasters can spark a chain reaction of problems that can linger on and spawn harmful societal consequences long after the main disaster is over. Moreover, over the next half-century, it is expected that disasters of this nature will increase five-fold (Thomas & Kopczak, 2005).

The use of mathematical modelling to tackle disaster management problems dates back to the 1950s, with work on determining optimal locations for fire-fighting resources (Valinsky, 1955). Knott (1987) and Knott (1988) developed pioneering vehicle routing models for relief management, whereas Fiedrich, Gehbauer, and Rickers (2000) used heuristics to allocate resources so as to minimise the fatality rate during the first few days after an earthquake. Barbarosoglu and Arda (2004) proposed a two-stage stochastic programming framework to plan delivery of supplies to areas affected by a disaster. The trend over time has been for models in this domain to incorporate multiple objectives. Examples of these include the model proposed by Tzeng, Cheng, and Huang (2007), with a triple minimisation objective (cost of service, unsatisfied demand and total response time) or that proposed by Doerner and Focke (2007), with the dual objective of minimising cost of service and unsatisfied demand. Minimising total response time has also gained currency as an objective and is widely incorporated in more recent multi-objective models in Huang, Smilowitz, and Balcik (2012), Nolz et al. (2010), and Van Hentenryck et al. (2010). Generally, in the literature, good strategic prepositioning of relief resources is recognised as the main ingredient in achieving these multiple objectives and coping well in the wake of a major catastrophe. For instance, the shortage of food and medical supplies observed in the aftermath of Hurricane Katrina

* Corresponding author. Tel.: +15406541729.
E-mail addresses: bfontem@umw.edu (B. Fontem), smelouk@cba.ua.edu (S.H. Melouk), bkeskin@cba.ua.edu (B.B. Keskin), anbajwa@ualr.edu (N. Bajwa).

http://dx.doi.org/10.1016/j.eswa.2016.04.002
0957-4174/

in 2005 was widely attributed to poor location planning (Holguin-Veras, Perez, Ukkusuri, Wachtendorf, & Brown, 2012). Location decisions of this nature strongly influence the performance of an emergency response network in terms of the response time and the distribution costs incurred.

However, a frequent assumption in many emergency response models is that for any specific relief network configuration, the routing times between nodes or the deadlines before which the relief workers must reach target nodes are either known beforehand, or at least have distributions that are known with absolute certainty. This is likely to be true only for disasters that do not induce continuing aftershocks in the midst of recovery efforts, that is, disasters which spark very few or very weak secondary disasters. This assumption is central because in the presence of secondary disasters, the parameters of any emergency logistics model are subject to random fluctuations. Hence, current emergency response models usually fail to incorporate the inherent risk associated with secondary disasters and the potential fluidity of disaster relief operations that may result. Secondary disasters include aftershocks (or, in worse circumstances, tsunamis) triggered by earthquakes (for instance, the 2004 Indian Ocean tsunami), landslides triggered by floods, avalanches induced by winds, or even rapidly spreading diseases caused by biological agents. Zhang, Li, and Liu (2012) proposed the only known model incorporating the impact of secondary disasters. They introduced a single-period problem with primary and secondary disaster locations with the overall objective of minimising the cost of allocating emergency resources to the primary disaster areas as well as to the secondary disaster locations. However, they still did not consider the navigational uncertainty that secondary disasters could engender.

The main contribution of this research is a decomposition-based heuristic to solve a delivery problem characterised by *both* uncertain travel times and uncertain deadlines in an emergency situation that due to secondary disasters, is non-stationary (i.e., has yet to settle into a steady-state). Travel time uncertainty directly stems from random network congestion that is typical of such situations. Deadline uncertainty by contrast, is a result of spontaneously occurring crisis demands which are modellable as non-stationary random processes (see for example the demand model in Roy, Sebastian, and Sharma (2013)). Because achieving adequate levels of situational awareness is challenging in dynamic disaster environments (Harrald & Jefferson, 2007), the duration for which new victims in such environments may be capable of coping without emergency assistance is also likely to be uncertain. Hence, the requirement to deliver necessary emergency services will be coupled with uncertainty as to when such services must be delivered in order to avoid outcomes that are deemed unacceptable. Alternatively stated, sources of demand in fluid emergency situations will likely have implicit random deadlines on their waiting times. Pavone and Frazzoli (2010) identify other routing contexts (e.g., surveillance missions by fleets of unmanned aerial vehicles and automatic robotic payload delivery networks) in which stochastic deadline constraints also play a vital role. Crucially, the only conditions we impose on the distributional nature of the random data (i.e., random travel times and random deadlines) corrupting the system are those of boundedness and symmetricity. An important advantage conferred by our solution framework is that it enables the decision-maker to create a menu of routing solutions of varying risk and quality from which she can choose. In addition, computational experience suggests that the solution framework we propose is viable for problems of realistic size.

This paper is organised as follows. Section 2 describes the problem setting. Section 3 reviews the related literature and in Section 4, we outline and explain our heuristic approach. In Section 5, we apply our heuristic on a small example, and also carry out a larger computational study. Section 6, concludes the paper and suggests avenues for further investigation.

2. Problem description

The emergency response problem that we consider is a static open routing distribution problem on a network composed of a set of relief warehouse nodes and another set of target nodes representing populations in distress. Open routing means that feasible paths on the relief network do not form closed loops. The relief warehouses are nodes from which emergency workers leave with valuable emergency supplies destined for other relief warehouses or for target population nodes. With each relief warehouse, a fixed scalar is associated, which represents the value of the relief supplies available at that relief warehouse. With each population location, a random deadline is associated, by which the supplies must be delivered. These stochastic deadlines model the fact that at any given moment, the urgency of the emergency situation at a population node is likely to be randomly affected by secondary disasters and aftershocks. We assume the properties of boundedness and symmetricity for the distribution governing the deadline uncertainty. Similarly, the travel times between the network nodes are also uncertain, but assumed to be describable by probability distributions that are bounded and symmetric. These assumptions render the analysis more tractable, but they sacrifice little in terms of generalisability because most distributions in practice can be easily approximated by distributions possessing these attributes.

Given a beginning relief warehouse and a target population node, the goal in our model is to find an open route which maximises the value of the relief supplies without allowing the probability of missing the target population node's deadline to exceed some pre-specified positive probability. The problem that we study thus belongs to the family of stochastic generalisations of a class of deterministic routing problems called *orienteering problems*. In the deterministic Orienteering Problem, one seeks a route to maximise the combined value of a set of visited nodes while adhering to a collection of constraints involving upper bounds on time (i.e., deadlines), fuel consumption or total distance travelled. In contrast, deterministic *vehicle routing problems* (e.g., Morais, Mateus, and Noronha (2014), Hernandez, Feillet, Giroudeau, and Naud (2006)) and their stochastic siblings (e.g., Wang, Ma, Xu, Wang, and Liu (2015), Cardoso, Schütz, Mazayev, Emanuel, and Corrêa (2015), Zhang, Lam, and Chen (2016)) are typically concerned with cost minimisation under demand satisfaction and other problem-specific requirements. For good resources on stochastic extensions to the deterministic Orienteering Problem, we point the reader to Lau, Yeoh, Varakantham, Nguyen, and Chen (2012), Varakantham and Akshat (2013), Evers, Glorie, van der, Barros, and Monsuur (2014) and Gupta, Krishnaswamy, Nagarajan, and Ravi (2015). Each of these stochastic orienteering problems considers uncertainty in one way or another although none is quite congruent with the problem treated in this study.

2.1. Model structure

Consider a relief network with a set I of relief warehouses and another set J of distressed population centres. Beginning from relief warehouse node $i' \in I$, the relief worker chooses a sequence of links connecting relief warehouses to each other with the goal of ending up at a target distressed population node $j \in J$. It takes $\tilde{t}_{il}$ time units (where $i \in I$, $l \in I \cup J$) to navigate each link on the relief network. The parameter $\tilde{t}_{il}$ is random, but has a probably density function that is bounded and symmetric on the interval $[t_{il} - \hat{t}_{il}, t_{il} + \hat{t}_{il}]$ where t_{il} is some *nominal* (i.e., mean) quantity, and $\hat{t}_{il}$ represents the maximum deviation from the mean travel time. There is a stochastic deadline $\tilde{\theta}_j$ by which it is desirable to bring