European Journal of Operational Research 237 (2014) 257–270

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

A multi-objective integrated facility location-hardening model: Analyzing the pre- and post-disruption tradeoff

Hugh R. Medal [a,*], Edward A. Pohl [b], Manuel D. Rossetti [b]

[a] Dept. of Industrial and Systems Engineering, Mississippi State University, Starkville, MS 39762, USA
[b] Dept. of Industrial Engineering, University of Arkansas, Fayetteville, AR 72701, USA

ARTICLE INFO

Article history:
Received 12 February 2013
Accepted 18 January 2014
Available online 6 February 2014

Keywords:
Location
Interdiction
Multiple objective programming
Catastrophe planning and management
Discrete optimization

ABSTRACT

Two methods of reducing the risk of disruptions to distribution systems are (1) strategically locating facilities to mitigate against disruptions and (2) hardening facilities. These two activities have been treated separately in most of the academic literature. This article integrates facility location and facility hardening decisions by studying the minimax facility location and hardening problem (MFLHP), which seeks to minimize the maximum distance from a demand point to its closest located facility after facility disruptions. The formulation assumes that the decision maker is risk averse and thus interested in mitigating against the facility disruption scenario with the largest consequence, an objective that is appropriate for modeling facility interdiction. By taking advantage of the MFLHP's structure, a natural three-stage formulation is reformulated as a single-stage mixed-integer program (MIP). Rather than solving the MIP directly, the MFLHP can be decomposed into sub-problems and solved using a binary search algorithm. This binary search algorithm is the basis for a multi-objective algorithm, which computes the Pareto-efficient set for the pre- and post-disruption maximum distance. The multi-objective algorithm is illustrated in a numerical example, and experimental results are presented that analyze the tradeoff between objectives.

1. Introduction

This article addresses the problem of finding a set of facilities to locate and a set to protect in order to optimally mitigate against facility disruptions. In particular, the objective of the problem is to minimize the worst-case consequence incurred due to the disruption of facilities. Thus, this objective is appropriate for a situation in which facilities are subject to interdiction, i.e., attacks by an intelligent adversary. After a disruption occurs, a set of demand points are each assigned to their closest non-disrupted facility. Thus, the consequence of a disruption is measured as the maximum travel distance, i.e., the maximum distance from any demand point to its closest located and operating facility. We call this problem the minimax facility location-hardening problem (MFLHP). Further, this article also analyzes a bi-objective version of the MFLHP, simultaneously considering the maximum travel distance both with and without disruptions.

Distribution networks, such as power networks and supply chains, are ubiquitous throughout the world. These networks consist of a set of facilities (power sub-stations, ports, distribution centers, etc.) and set of customers that rely on the facilities. Because these facilities form the backbone of distribution networks, facility disruptions often result in severe consequences. One recent example is the 2011 Tohoku earthquake in Japan, which disrupted manufacturing facilities and caused several Japanese automakers to halt car production for up to six months (Kim, 2012). These severe consequences have forced decision-makers to consider the possibility of facility disruptions when they design their network of facilities. In addition, decision-makers also may choose to harden facilities to protect them from disruptions. Facility hardening, a special case of facility protection, involves allocating resources to a facility (e.g., additional security, retrofitting, etc.) to make it immune to failure (Scaparra & Church, 2008a, 2008b; Smith, 2011). This paper serves to help decision-makers make better facility location and hardening decisions by providing a mathematical model of these decisions and using this model to generate insights about these decisions.

This research focuses on the maximum distance objective, which is also used in the classic p-center problem (Hakimi, 1965). Since this objective is concerned with minimizing the worst service experienced by a demand point, it is appropriate for the public sector (Daskin, 2000). Researchers have cited numerous po-

* Corresponding author. Tel.: +1 662 325 3923.
E-mail addresses: hugh.medal@msstate.edu (H.R. Medal), epohl@uark.edu (E.A. Pohl), rossetti@uark.edu (M.D. Rossetti).

http://dx.doi.org/10.1016/j.ejor.2014.01.040
0377-2217/

tential applications for the maximum distance objective such as locating emergency vehicles and facilities (Hochbaum & Pathria, 1997; Mladenovic, Labbe, & Hansen, 2003; Scaparra, Pallottino, & Scutella, 2004) and locating warning sirens (Suzuki & Drezner, 1996). This article also seeks to minimize the worst-case disruption, i.e., the worst-case risk measure. The worst-case risk measure is well-suited for risk-averse decision-makers, especially in critical infrastructure protection (Church & Scaparra, 2007; Salmeron, Wood, & Baldick, 2009; Scaparra & Church, 2008b).

A growing amount of research exists on locating facilities subject to disruptions and hardening facilities subject to disruptions. Several authors (Cui, Ouyang, & Shen, 2011; Drezner, 1987; O'Hanley, Paola Scaparra, & Garcia, 2013; Snyder & Daskin, 2005) have developed models for locating facilities subject to random disruptions. Other works have considered that located facilities are subject to interdiction, i.e., intentional and calculated attacks. O'Hanley and Church (2011) developed a bi-level model for the problem of locating facilities to minimize the post-interdiction total weighted covered demand. Drezner (1987) developed a method to minimize the post-interdiction maximum distance from a demand point to its closest located and operating facility. O'Hanley and Church (2011) optimized a weighted combination of the system performance before and after interdiction for a facility location problem.

Rather than locating facilities, others have examined the problem of hardening a set of existing facilities. O'Hanley, Church, and Gilless (2007b) and Li, Zeng, and Savachkin (2013) have presented models for hardening facilities subject to random failures. Church and Scaparra (2007) and Scaparra and Church (2008a, 2008b) have studied the problem of how to harden facilities in order to minimize the post-interdiction total weighted distance. O'Hanley, Church, and Gilless (2007a) presented a bi-level model to optimally harden facilities in order to minimize the post-interdiction total weighted covered demand.

Some researchers have developed models that include both facility location and facility hardening. Snyder and Daskin (2005) present extensions of the *p*-median and warehouse location models and include perfectly reliable, i.e., hardened, and unreliable facility locations in their model. Specifically, a facility is perfectly reliable if and only if it is located at a perfectly reliable location. Although their study focuses on location, it would be possible to integrate location and hardening decisions in their model if every geographical site had both a reliable and an unreliable location. However, it is unclear whether using their model in this way, which would double its size, would be computationally tractable. Lim, Daskin, Chopra, and Bassamboo (2010) were the first to explicitly include both location and hardening decisions in a single model. They present an extension of the warehouse location problem in which the decision maker chooses between locating unreliable facilities and locating perfectly reliable, i.e., hardened, backup facilities at a higher cost. The authors assume one layer of supplier backup. Thus, if a demand point's primary facility fails, the demand point is then immediately assigned to its hardened backup without checking if there is a closer operating facility. This assumption simplifies the model and allows the authors to provide several useful analytical results. Li et al. (2013) extend the work of Lim et al. (2010) but still assume one layer of supplier backup. The research presented in this paper considers multiple layers of backup, allowing a demand point to be assigned to its closest operating facility after a disruption. Aksen, Aras, and Piyade (2013) study an extension of the *p*-median problem in which facilities are susceptible to interdiction. They present a tri-level version of the budget-constrained median location model in which a defender locates and hardens facilities and then an attacker destroys a number of unhardened facilities. Their works extends the work of Lim et al. (2010) by modeling multiple layers of backup. Aksen et al. (2013) study several methods for solving their problem including a tabu search algorithm and a two-phase heuristic. The research in this paper builds on the work of Aksen et al. (2013) by providing an exact procedure for solving the integrated location-hardening problem, rather than a heuristic procedure.

This article builds upon the facility location and facility hardening literatures by making the following main contributions. (1) A new model for integrating facility location and hardening decisions; in particular, a natural three-level formulation is converted to a single-level mixed-integer program (MIP) by taking advantage of the structure of the MFLHP. This model is accompanied by a binary search solution procedure along with a method for obtaining a lower bound. (2) This integrated model and solution method forms the basis of an algorithm that computes the complete Pareto-efficient set for the pre- and post-disruption maximum distances. This algorithm is based on a method from Medal, Rainwater, Pohl, and Rossetti (2013) that optimizes facility location decisions but does not model facility hardening. (3) A set of computational experiments provide results that should help decision-makers better understand the tradeoff between the pre- and post-disruption maximum distances when making the decision to locate and harden facilities subject to disruptions. Toward this end we present the following analyses: (i) an analysis of the Pareto-efficient set between the pre- and post-disruption maximum distances; (ii) an analysis of the penalty incurred for optimizing either the pre- or post-disruption radius in isolation; and (iii) an analysis of the benefit of considering facility hardening when locating facilities subject to disruptions.

The remainder of this article is as follows. In Section 2 the MFLHP is described and a three-level model of the problem is converted to a single-level MIP. In Section 3, two algorithms are presented for single- and bi-objective versions of the MFLHP. An example that demonstrates the bi-objective MFLHP is given in Section 4. Section 5 reports the results of computational experiments on the single- and bi-objective MFLHP. Section 5 concludes the article with a summary and a discussion of future work.

2. Problem description and models

The purpose of the MFLHP model is to

locate a set of facilities and harden a subset of the located facilities in order to minimize the worst system performance over all possible disruption scenarios consisting of the disruption of r facilities. The system performance for a disruption scenario is the maximum distance from a demand point to its closest located and operating facility.

The MFLHP model is appropriate for two situations: (1) facilities are vulnerable to naturally-caused disruptions and the decision-maker wishes to mitigate against the worst-case consequence due to the loss of r facilities and (2) facilities are subject to a strategic attacker who seeks to attack up to r facilities in order to generate the largest consequence possible and the decision-maker wishes to mitigate against these attacks.

To understand the model, it may help to divide it into three stages: (1) the mitigation, (2) the disruption, and (3) the response. We use the generic term *facility* to refer to a physical entity that we are locating and hardening. The mitigation stage, which happens before the disruption occurs, involves actions taken to mitigate against the disruption. The mitigation decisions in our model concern where to locate facilities and which facilities to harden, and these decisions are made simultaneously. If a facility is hardened in our model, it is always available to serve demand points. In