Expert Systems With Applications 54 (2016) 251–264

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Metaheuristics for protecting critical components in a service system: A computational study

Chun-Hung Cheng[a], Tze-Wai Lai[b], David Yanyi Yang[a,*], Yuntao Zhu[c]

[a] Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, Hong Kong
[b] Information Technology Services Centre, The Chinese University of Hong Kong, Shatin, Hong Kong
[c] American Express, 18850 N. 56th Street, Phoenix, AZ 85054, USA

A R T I C L E I N F O

Keywords:
Metaheuristics
Location
Interdiction median problem
Tabu search
Simulated annealing
Genetic algorithms

A B S T R A C T

Deliberate sabotages and terrorist attacks are major threats to the safety of modern societies. These attacks often target at important infrastructures such as energy production and transmission systems, food and water supply networks, telecommunications networks, transportation networks, etc. In such systems, some components are critical as their malfunction may adversely affect the operations of the whole systems. This research examines several models based on the median problem for identifying these components in a service system. In addition to the existing exact solution methods, we propose meta-heuristics to tackle this computationally hard problem. Our hybrid approach combines the strengths of both meta-heuristics and exact solution methods. The experiment shows that the combination of solution methods significantly cut down the computational requirement for finding the critical components in a service network for protection.

1. Introduction

The objective of a typical planning model is to design a service system as efficient as possible that is to either maximize coverage or minimize cost. However, terrorist attacks such as the World Trade Center attack in 2001, the Madrid train bombing in 2004, the London public-transport bombing in 2005, and the Mumbai train bombing in 2006, force us to consider disruptions caused by these incidents in the study of location models. The identification and protection of critical components are of utmost importance as they enhance the system resilience and mitigate the detrimental effect on the system. In fact, identifying vital components has been recognized as one of several national research priorities after the 9/11 attacks according to Cutter, Richardson, and Wilbanks (2003). In the context of location modeling, this motivates the investigation of the interdiction and fortification aspects of different location planning models.

The study of interdiction models actually has long been the interest of military planners since McMasters and Mustin (1970). In these models, it is assumed that an intelligent attacker exists. The attacker wants to cause maximum harm to an existing service system utilizing his/her resources. Consequently, in some sense, the critical or key elements of a system can be identified by analyzing the behavior of the rational attacker. However, the majority of such interdiction models focus on arc interdiction instead of facility interdiction. To the best of our knowledge, Church, Scaparra, and Middleton (2004) are the very first group of researchers who examined the interdiction aspect of two classic location models, namely the p-median and the maximum coverage problems.

Church and Scaparra (2007) extended the interdiction version of the *p*-median location model by incorporating a fortification layer into the interdiction model. By studying the attacker's strategy from the standpoint of the defender, the fortification model considers explicitly which sites should be fortified or strengthened to best use the available resources for protection against possible disruptions. Several exact solution methods for solving the fortification version of the *p*-median location problem are developed and several extensions are proposed.

In this work, we attempt to implement and experiment three different meta-heuristics to reduce the solution time of these types of problems. Our goal is to make large instances solvable as opposed to existing solution methods. Problem instances are generated using the London data set, which has been used extensively in the location literature. Our experimental results suggest that the performance of the exact solution methods is unsatisfactory even for the problems of modest size. The hybrid approach combining meta-heuristics and optimization methods are demonstrated to use notably less time. They are also astonishingly effective and can often find optimal solutions in many test instances.

* Corresponding author. Tel.: +852 64282743.
E-mail addresses: chcheng@se.cuhk.edu.hk (C.-H. Cheng), twlai@cuhk.edu.hk, tszwailai@alumni.cuhk.net (T.-W. Lai), yyyang@se.cuhk.edu.hk, yanyi.david@gmail.com (D.Y. Yang), yuntao.zhu@aexp.com (Y. Zhu).

http://dx.doi.org/10.1016/j.eswa.2015.10.029
0957-4174/

As we will discuss in the next section, many solution methods are available for the interdiction problem. However, many of them cannot deal with large problem instances. In reality, many service systems such as energy production and transmission systems, food and water supply networks, telecommunications networks, transportation networks, etc. are large. There is a need for computationally viable methods. Our research will provide a hybrid approach that is capable of solving larger problem instances.

The remainder of the paper is organized as follows. The related and previous work is summarized in Section 2. The mathematical formulations are presented in Section 3. The problem-solving methodologies are discussed in Section 4. Results from computational examples are illustrated in Section 5. Finally, Section 6 concludes the paper and gives some guidelines for future research.

2. Background

There is a variety of quantitative approaches developed to assist the design of a service system with possible disruptions in mind. In this research, we assume there is an existing system and the goal is to select a subset of facility sites to be protected with the available resources against disruptions. The underlying model is a *p*-median problem. In other words, the metric that is used to evaluate the efficiency of the system is the total demand-weighted distance. Furthermore, the disruptions being considered are deliberate attacks. We assume that there is a rational attacker, trying to inflict maximum harm to the system within a given amount of the resources.

Church et al. (2004) proposed two interdiction models based on the classic facility location problem and formed a foundation for many facility fortification models. The two models are designated as the *r*-interdiction median problem (RIM) and the *r*-interdiction covering problem (RIC). The former is based on the well-known *p*-median problem (PMP), introduced in Hakimi (1964 and 1965), whereas the latter is the antithesis of the maximal covering location problem (MCLP), formulated by Church and ReVelle (1974). For the model by Keçici, Aras, and Verter (2012), which allows making the facility location, relocation, and protection decisions simultaneously, they developed an exhaustive enumeration method to solve the problem to optimality. Stochastic interdiction problem for median systems were also studied. Losada, Scaparra, Church, and Daskin (2012) assume the disruptions to be uncertain and show that such model with stochastic settings could be reformulated as a deterministic counterpart problem.

Church and Scaparra (2007) made the *r*-interdiction median problem with fortification (RIMF) as an extension of RIM. The extended model addresses the problem of fortifying a subset of sites against possible attacks to ensure that the system performs its intended function as much as possible even after a man-made disaster. It is formulated as a single-level program by explicitly enumerating all possible interdiction scenarios and thus can be readily solved by virtually any commercial optimization software packages. Scaparra and Church (2008b) later reformulated the problem as a maximal covering problem with precedence constraints. They demonstrate that this allows the RIMF problem to be solved more efficiently, though it sacrifices the convenience of using general-purpose MIP solvers. Unfortunately, these solution techniques require complete enumeration of interdiction patterns and hence they can only tackle very small problem instances. Since the explicit enumeration is inefficient and at the same time, fortification problems are well-fitted into the framework of von Stackelberg (1952) game, a natural way to look at fortification problems is a bi-level approach examined by Scaparra and Church (2008a). It relies on the implicit enumeration (IE) algorithm which recursively exploits information from the conditional RIM problems to reduce the size of the search tree.

In addition, many researchers investigated other variants of RIM. Church and Scaparra (2006) demonstrated the probabilistic *r*-interdiction median model in which an attempted attack may not be successful. In this stochastic version of RIM, there is a certain probability that an attack on a facility may fail. Losada, Scaparra, and Church (2010a) set up three variants of the probabilistic *r*-interdiction median models. They identify the parameters that affect the time performance and test the robustness of the models under a number of probability distributions. Aksen and Aras (2012) introduced the concept of partial facility interdiction in which the attacker decides not only what facilities to compromise but also at what fraction with the constraint of a fixed interdiction budget. Later, Aksen, Akca, and Aras (2014) considered partial interdiction decisions with capacitated facilities and outsourcing options. Forghani and Dehghanian (2014) also looked into the partial interdiction model for a capacitated hierarchical system. An exhaustive enumeration solution was used to solve a bi-level problem formulation. Zhang, Zheng, Zhang, and Du (2015) examined an interdiction median model for multi-source supply systems. They made three assumptions: capacitated facilities, partial interdiction, and multi-sourcing deliveries.

Another research direction aims at RIMF. Scaparra and Church (2010) worked further on the capacitated *r*-interdiction median problem with fortification (CRIMF) in which the facility sites have limited capacity. After an attack, some of the demand may no longer be served and hence a penalty is imposed for each unit of unsatisfied demand. Aksen, Piyade, and Aras (2010) augmented the bi-level RIMF and developed a model which is regarded as the budget constrained *r*-interdiction median problem with capacity expansion (BCRIMF-CE). In this model, strengthening different sites require different levels of resources. Furthermore, as the demand points which were originally assigned to an interdicted facility have to be reassigned to a working facility after an attack, BCRIMF-CE also accounts for the capacity expansion required for these surviving facilities. Aksen and Aras (2012) extended the BCRIMF-CE and combined it with a fixed charge facility location problem. Aksen, Aras, and Piyade (2013) further extended the BCRIMF-CE by incorporating also the planning decision into the model. Losada, Scaparra, Church, and Daskin (2010b) formulated a version of RIMF that takes facility recovery into account, which results a mixed integer bi-level linear problem (MIBLP). In their model, a facility can recover from an attack after certain time and is subject to frequent disruptions. Liberatore, Scaparra, and Daskin (2011) developed the stochastic *R*-interdiction median problem with fortification (S-RIMF) which deals with a random number of possible losses under the intentional disruptions. This extension addresses the situation that the exact capabilities of the attacker are unknown to the defender. That is, the number of interdictions is expressed as a probability distribution.

Liberatore, Scaparra, and Daskin (2012) suggested an extension which considers the correlation and ripple effects of disruptions and proposed a model called the wave propagation *r*-interdiction median problem with fortification (WaveRIMF). It addresses a situation where the attack on one facility not only destroys the facility being attacked but also degrades the service level of the adjacent facilities. The above literatures assume that protected facilities are immune to interdiction. However, Zhu, Zheng, Zhang, and Cai (2013) relax this assumption. It uses greedy search to solve the RIM with probabilistic protection model and gets good approximations to optimal fortification strategies with a fast speed.

To the best of our knowledge, few researchers have examined the application of meta-heuristics to the RIMF problem. In some cases, the large problem instances require prohibitively long computational time to solve. There is a need for alternative solution approaches. Our hybrid approach combines the idea of metaheuristics and exact solution methods providing very good solutions