Kent Academic Repository
Full text document (pdf)

DOI

Link to record in KAR

https://kar.kent.ac.uk/25484/

Document Version

UNSPECIFIED

Protecting supply systems to mitigate potential disaster: a model to fortify capacitated facilities

Maria Paola Scaparra[1]

Kent Business School
University of Kent
Canterbury CT2 7PE, UK

Richard L. Church

Department of Geography
University of California, Santa Barbara
Santa Barbara, CA 93106-4060

Abstract

Planning to mitigate the impacts of a disaster can be an important activity for both private companies and public agencies. In this paper we consider a supply system that provides needed goods or services to a region that may be the subject of some type of disaster, such as an attack by a terrorist or the result of a natural event or accident. The supply system is represented by a set of existing capacitated facilities. We assume that the loss of one or more facilities to a disaster will tighten available supply and increase the distances over which the service or good must be delivered, thereby increasing operation costs and reducing service. Such a disaster may even reduce the capacity of the supply/storage to the extent that the goods must be rationed as remaining supply may be outstripped by demand. We consider the case where resources may be available to mitigate some of the impacts of a possible disaster by the advanced protection of one or more facilities. We show how this problem can be formulated as a "tri-level" optimization model and propose a solution approach based on a tree search strategy. We demonstrate the policy implications of this model using a hypothetical planning problem. Through this example, we show how the results of our model can be used to inform planners and policy makers in disaster mitigation planning.

---

[1] Corresponding author: Phone: +44-1227-824556; Fax: +44-1227-761187; mail: m.p.scaparra@kent.ac.uk

1. Introduction

Advanced planning is an important activity within the context of disaster management. The major objective of such planning should be to identify the possible impacts of a given type of disaster and then to determine the best plan of action for mitigating the effects of such an event. For example, when Hurricane Katrina hit the gulf coast of Louisiana and Alabama, disaster was not completely avoided, but it was mitigated by advanced planning. Even though the U.S. Army Corp of Engineers had requested funding over a number of years to upgrade the levee and pump system of New Orleans, the U.S. Congress had not allocated funds for the project and New Orleans was left with a protection system that failed. To look at the immense damage caused by the storm as well as the number of stranded people overlooks the fact that the majority of the people safely departed the city based upon an advanced plan involving reversing the direction of city bound lanes of a major interstate highway. This "contraflow" evacuation plan helped evacuate the majority of the inhabitants in a timely and efficient way (Wolshon 2006). Without that plan, this disaster would have been much worse.

It is important to recognize that advanced disaster planning is an important task for government agencies as well as private companies, although the objectives may be quite different. Whereas public agencies are principally concerned with saving lives and preventing damages like flooding, private companies are concerned with protecting their facilities and keeping goods and services flowing. Large retailers, as an example, often plan for hurricane season in the southeast U.S. by stockpiling goods like water and plywood (Albright 2009). Without this advanced planning, it is quite possible that local demand would cause shortages of critical supplies. The important issue is that, although the objectives may be quite different, advanced planning may be strategically important across the spectrum of government agencies and private corporations in order to deal with a potential disaster.

Over the last few years there has been an increased interest in modeling the fragility of supply and service systems due to some disruptive event (see for example Brown et al. 2006; Snyder et al. 2006; Murray et al. 2008; Church and Scaparra 2006). This focus has been split between network fragility (Wood 1993; Wollmer 1964; O'Kelly and Kim 2007; Murray et al. 2003; Zhuang and Bier 2007; Peterson and Church 2008) and facility system fragility (Church et al. 2004; Rawls and Turnquist 2006) over a wide range of events from natural disasters to intentional strikes. Although the literature of system

disruption is historically rooted in the military problem of interdiction (Wollmer, 1964; Israeli and Wood 2002), recent disasters such as 9-11 have forced planners and policy makers to ask what might happen (Haimes and Longstaff 2002; Haimes 2006; Grubesic et al. 2003) as well as how to plan within the uncertainty of what might befall a system (Altay and Green 2006; Church and Scaparra 2006; Rawls and Turnquist 2006; Ukkusuri and Yushimito 2008; Alcada-Almeida et al. 2009). A few recent papers have focused on how to optimize system protection within a limited budget (Azaiez and Bier 2007, Brown et al. 2006, Church and Scaparra 2006). Notably lacking in past model development is the optimal protection of a system of capacitated facilities. Since most facilities operate within defined levels of operation or capacity, most of the past work cannot be directly applied in real world settings. Our objective of this paper is to propose a model construct which optimizes a limited amount of protection resources among a set of capacitated facilities in order to mitigate a worst case disaster event. We show how this model can be used to help inform planners and policy makers in disaster planning.

In the next section, we present a short review of past work that is related to the problem that we address in this paper. We then formulate a planning model for a capacitated logistics system comprised of a set of demands and facilities, where one or more facilities may be rendered inoperable due to a disaster. We assume that resources are available to fortify or protect some of the facilities and the problem is to allocate the protection resources so that the system is as resilient as possible in the event of a disastrous intentional strike or natural event. Protection resources and strategies can vary considerably depending on the type of system and possible disruption. For example, protecting against an earthquake may involve seismic upgrading, building a new approach bridge to an industrial site, or providing a backup power generator system so that the facility can operate during the loss of electrical power. Protecting against flooding may involve relocating a facility to higher ground, providing pumps and generators to keep a facility dry, or even developing a storm wall. Protecting against an interdictor may involve added security/guards, perimeter fencing, surveillance cameras, hardened internet and communication systems, etc. Whatever the type of disaster, we assume that there is some set of possible actions which can be taken to enhance protection and keep a facility in operation. Although the proposed model is designed to represent a generic system, we expect that the proposed construct here will be modified for specific applications. The important element is that the model is extensible and is representative of the general planning problem for protection in light of a potential disaster. After giving a model formulation and discussing how it can be solved, we present an

application to a hypothetical data set to underscore the insights that may be generated by using the model in practice to support disaster planning. Finally, we conclude with a short discussion on possible refinements.

2. Background

Optimization has been used in the design and operation of many logistics systems, ranging from public systems (e.g. optimal fire station location) to private sector systems (e.g. warehouse location and capacity allocation). Although the common emphasis is directed towards modeling for an optimal design or an efficient operation, some efforts have been directed at determining weakness in system design and operation. An example of this is the identification of the critical links of a network, associated with a network's ability to handle traffic flow. Historically, modelers have approached the issue of detecting system weakness with the perspective of an "interdictor" who plans to strike a system. Assume that the "interdictor" has enough resources to strike and render useless K arcs of a network (Wollmer 1964). The question is: "what K arcs when removed have the greatest impact on the remaining system?" Suppose the underlying operation is to ship items along the shortest path between an origin and a destination. The objective of the interdictor would be to take out the arcs which would maximize the length of the shortest route from the origin to the destination (Israeli and Wood 2002). Since the shortest route may well change when an arc is removed, the problem is somewhat complicated to solve (Israeli and Woods 2002). The solution to this type of interdiction model gives the analyst and planner an understanding as to which links are critical as well as to what extent system operation can be compromised by losing a specific number of network links. Thus, the solution represents a worst case, as the model seeks to find the most disruptive solution associated with an event level of some specified magnitude, whether natural or intended. In the next few paragraphs we briefly characterize past research on interdiction associated with networks and facilities. The interested reader should consult Snyder et al. (2006) for a review on system reliability and disruption and Brachman and Church (2009) and Altay and Green (2006) for a review of modeling for disaster management.

Network interdiction models were first developed in the 1960's (see Church et al. 2004 for a review). Wollmer (1964) proposed a network interdiction model to optimally reduce the maximum flow possible through a graph. This particular problem has remained a key research problem (Lim and Smith 2007). Besides modeling the interruption of flow and paths, there has also been an interest in modeling overall

connectivity and possible cases of failure. A good example of this type of analysis is that of Grubesic et al. (2003) where the problem was to analyze topologically the loss of internet services and communication connectivity when certain internet backbone links are compromised. It should also be mentioned that interdiction may be based upon a cost. For example, some arcs may be easier to knock out than others. When interdiction costs vary, then interdiction is assumed to be resource limited and is subject to a budget constraint.

Interdiction models have also been developed to involve a system of facilities. The first facility interdiction model involved the well known p-median location problem. The p-median model involves the location of a set of p facilities (uncapacitated), in order to minimize the weighted distance of serving a set of weighted demand points. It is assumed that each demand will be served by their closest facility. The related facility interdiction model involves finding which r facilities, when removed, increases the weighted distance of service the most (Church et al. 2004). Although most of the past research has been posed where the act of interdiction is always successful when an element is selected (e.g. facility or network link), it has also been modeled probabilistically. Suppose that when a facility is struck by an interdictor, the interdictor fails a certain percentage of the time. We can then define the probabilistic facility interdiction problem as: what is the worst-case possible loss of a facility system when r facility strikes are allowed, where the probability of any successful interdiction is set at some value less than or equal to 1? Using this context, Church and Scaparra (2007) were able to show how to generate an operational "reliability envelop" for a system of facilities subject to natural or intended harm. Note that, if the probability of success is somewhat low, an interdictor may strike one facility many times in order to increase the overall probability of successfully eliminating a key facility (Losada et al. 2009). In a related work Rawls and Turnquist (2006) have developed an emergency supply system location model based upon a set of predefined scenarios. Each scenario is based upon a set of failures, where specific nodes and links are not available for that scenario. For their model they defined each scenario on a specific hurricane storm track and assumed that specific locations or routes in that scenario would be knocked out by the storm. The overall objective was to locate a set of facilities, such that the expected cost of resupply over all scenarios is minimized.

Research on supply systems deployment has also included the possibility of system protection. Interdiction can be thought of as an act by nature or man. In either case it may be successful without taking some type of protective action. The first model to

incorporate protection as an action in a facility operations model involved the median problem, where it was assumed that a limited number of facilities, q, could be hardened or protected from a disaster or interdictor (Church and Scaparra 2007). The objective was to optimally allocate the protective resources in order to reduce the losses that could happen by interdiction. This model was called the r-interdiction median problem with fortification. Network link protection has been considered in two types of problems: the shortest path interdiction protection (Cappanera and Scaparra 2008) and the design of a seismically protected road network to ensure the safe distribution of supplies after an earthquake (Viswanath and Peeta 2003). A review of many of the models that have been defined specifically for emergency management can be found in Brachman and Church (2009).

It is important to mention that many of the models that have been developed to solve for the most destructive case or to identify how to best protect a system with limited resources are posed within the context of bi-level optimization. Such models are often described as a two person game (Stackelberg 1952), where the top level is the leader and the bottom level is the follower (Dempe 2002; Colson et al. 2007). For example, the top level in a model that optimizes protective resources, involves selecting which links or facilities to be fortified. The bottom level involves the "interdictor" which will attempt to do the most harm given that some elements are fortified and others are not. What makes this problem hard to solve is that the bottom level solution is a function of the top level resource allocation and the top level resource allocation is a function of the bottom level disruption. Bi-level problems by their very nature can be very difficult to solve as each level contains an objective that is the exact opposite of the next. This is especially true if integer decision variables appear in both levels (Moore and Bard 1990). In fact, in only a few cases has it been possible to cast a mix-integer bi-level problem as a single level problem. Examples of casting a bi-level problem as a single level problem can be found in Church and Scaparra (2006) and Scaparra and Church (2008a). The fact that the combined problems of protection and destruction are hard to model in a combined fashion has served as a major impediment or barrier to progress in this research. Recent work (Brown et al. 2006) has helped to provide insights as to how these complicated models might be solved.

Although the above review has only listed a small number of the many contributions in the systems literature involving disasters and protection, it does characterize the extent to which specific actions, both disruptive and protective, have been optimized within the

context of service and supply systems. It should be mentioned here that past work has concentrated principally on capacitated network flow and shortest path interdiction or uncapacitated facility interdiction and protection. Unfortunately, most systems operate with capacity limits, and when a facility has been struck by disaster, alternate facilities may not have enough combined capacity to handle the supply needs of those originally served by the damaged facility. That is, most of the past facility optimization work does not apply to the type of problem that one might find in disaster supply management. Our objective, here, is to propose a model construct which will help to fill that gap. In the next section we propose a model which optimizes a limited amount of protective resources among a set of capacitated facilities serving a set of demand points in order to mitigate a worst case disaster event involving the loss of one or more facilities. Conceptually, this model contains three levels: the top level- allocates protective resources; the intermediate level - interdicts one or more facilities; and the bottom level- operates as efficiently as possible the remaining facilities and serves demand as best as possible. Essentially, there are three types of decisions being made by two different actors: 1) the system operator decides what to protect and operates what is left as efficiently as possible, and 2) the interdictor decides what facilities to disrupt as a worst case event. Although we use the term "interdictor" to represent the role of disruption, it represents in general some event doing harm, e.g. a terrorist or a natural disaster, which disrupts the system maximally within limits (even hurricanes and major natural disasters do not destroy everything). We will use the term "fortification" to represent the type of protective action that is taken to fend off a destructive strike, whether natural or man-made. Thus, the model that we propose is somewhat general in its scope, and we expect that it will be tailored for specific applications. Finally, we will show how results from this general model can be used to help inform planners and policy makers in disaster planning based upon a hypothetical case study.

3. Protecting facilities in order to deal with a disaster: Model development

In this section, we present the mathematical formulation of the three-level Capacitated r-Interdiction Median problem with Fortification, referred to in the following as CRIMF. The formulation uses the following set of parameters and decision variables.

Parameters

N set of n demand nodes (indexed by i)

F set of p existing facilities (indexed by j)

a_i demand at node i

7

c_j capacity at facility j

d_{ij} cost for serving customer i from facility j

θ_i penalty for not serving customer i (per unit of demand)

r   number of possible losses

q   number of facilities that can be hardened

Decision Variables

z_j 1 if facility j is protected; 0 otherwise

s_j 1 if facility j is interdicted; 0 otherwise

x_{ij} fraction of demand i served by facility j after interdiction

u_i fraction of demand i which is not served after interdiction

The variables z_j and s_j represent the upper level protection variables and the middle level interdiction variables, respectively. The variables x_{ij} and u_i are used in the lower operational level to evaluate the system efficiency after protection and interdiction.

The trilevel formulation of CRIMF is as follows:

$$\max \quad K(z) \tag{1}$$

$$\sum_{j \in F} z_j = q \tag{2}$$

$$z_j \in \{0,1\}, \quad \forall j \in F \tag{3}$$

$$K(z) = \max \quad H(s) \tag{4}$$

$$\sum_{j \in F} s_j = r \tag{5}$$

$$s_j \leq 1 - z_j, \quad \forall j \in F \tag{6}$$

$$s_j \in \{0,1\}, \quad \forall j \in F \tag{7}$$

$$H(S) = \min \sum_{i \in N} \sum_{j \in F} d_{ij} x_{ij} + \sum_{i \in N} \theta_i u_i \tag{8}$$

$$\sum_{j \in F} x_{ij} + u_i \geq a_i, \quad \forall i \in N \tag{9}$$

$$\sum_{i \in N} x_{ij} \leq (1 - s_j) c_j, \quad \forall j \in F \tag{10}$$

$$x_{ij} \geq 0, \quad \forall j \in F, i \in N \tag{11}$$

$$u_i \geq 0, \quad \forall i \in N \tag{12}$$

8

The CRIMF model identifies the optimal allocation of q protection resources (2) which minimizes the cost function, K (1). This function represents the maximum damage, in terms of service and lost sale costs, that can be inflicted to the system after the loss of r facilities. The correct value of K is computed in the middle level problem, where the interdictor allocates r offensive resources (5) among the unprotected facilities (6) so as to maximize the system's cost (4). In turn, the computation of the system costs after interdiction, H, requires solving the lower level operational problem. At the operational level, the objective is to minimize the total cost (8), which includes the cost for serving the customers and the lost sale cost incurred when some customer demands cannot be met due to the insufficient capacity of the system after interdiction. Constraints (9) state that all customer demand needs to be accounted for, either as demand allocated to some operational facility or as unmet demand. The facility capacity restrictions are modeled in constraints (10), which also prevent the allocation of customers to interdicted facilities. The integrality of the protection and interdiction variables is enforced by constraints (3) and (7) respectively, whereas constraints (11) and (12) state the non-negativity of the operational variables. Note that the cardinality constraints on the offensive and defensive resources, (2) and (5) respectively, can be replaced by budget constraints if the system facilities have different protection and interdiction costs. The solution methodology described in the next sections can be easily adjusted to handle this case.

4. Integer bilinear reformulation of the interdictor-user problem

In this section, we show how the two bottom level problems (interdictor and user) can be collapsed into a single level problem. This level reduction is based upon the observation that, given any solution s to the interdiction problem, the lower level operational problem is a transportation problem in the continuous variables x and u only. It is therefore possible to take the dual of this problem to obtain a single level interdiction problem which inherently incorporates the optimal system operations. Duality techniques have been extensively used to solve bilevel programs with linear inner problems (see for example Wood 1993, Israeli and Wood 2002, and more recently Lim and Smith 2007, Bayrak and Bailey 2008 and Losada et al. 2009). In our case, duality is used to reduce the initial tri-level defender-interdictor-user model to a bilevel program.

Let ρ_i and π_j be the dual variables associated with the demand constraints (9) and the capacity constraints (10) respectively. Then, the interdiction-user problem can be reformulated as a mixed-integer bilinear problem as follows:

$$K(Z) = \max \sum_{i \in N} a_i \rho_i - \sum_{j \in F} c_j (1 - s_j) \pi_j \tag{13}$$

$$\rho_i - \pi_j \le d_{ij}, \quad \forall i \in N, j \in F \tag{14}$$

$$\rho_i \le \vartheta_i, \quad \forall i \in N \tag{15}$$

$$\sum_{j \in F} s_j = r \tag{16}$$

$$s_j \le 1 - z_j, \quad \forall j \in F \tag{17}$$

$$s_j \in \{0,1\}, \quad \forall j \in F \tag{18}$$

$$\rho_i \ge 0, \quad \forall i \in N \tag{19}$$

$$\pi_j \ge 0, \quad \forall j \in F \tag{20}$$

To resolve the non-linearity in the objective function introduced by the dualization, we use a simple linearization technique which consists of replacing each bilinear term $(1 - s_j)\pi_j$ with a new variable v_j (with $v_j \ge 0$) and adding the following linearization constraints to the model:

$$v_j \le (1 - s_j) M \tag{21}$$

$$v_j \ge \pi_j - s_j M \tag{22}$$

The use of constraints (21) and (22) ensures that the new model using the replacement variables v is equivalent to the bilinear model (13)-(20). In fact, if $s_j = 1$ in the bilinear model (and hence the bilinear term is zero), from constraints (21) and (22) we have that $-M \le v_j \le 0$. As the variables v_j are non-negative, it must be that $v_j = 0$. Conversely, if $s_j = 0$ and the bilinear term is equal to π_j, from constraints (21) and (22) we have that $\pi_j \le v_j \le M$. As the replacement variable v_j appears in the objective of a maximization problem with a negative coefficient, at optimality it will take the smallest possible value, i.e. π_j. A possible value for the constant M to tighten the formulation is $M = max(max_i \theta_i, max_{ij} d_{ij})$.

The resulting single level formulation of the interdictor-user problem is:

$$K(Z) = \max \sum_{i \in N} a_i \rho_i - \sum_{j \in F} c_j v_j \qquad (23)$$

$$\rho_i - \pi_j \le d_{ij}, \quad \forall i \in N, j \in F \qquad (24)$$

$$v_j \ge \pi_j - s_j M, \quad \forall j \in F \qquad (25)$$

$$v_j \le (1 - s_j)M, \quad \forall j \in F \qquad (26)$$

$$\sum_{j \in F} s_j = r \qquad (27)$$

$$s_j \le 1 - z_j, \quad \forall j \in F \qquad (28)$$

$$s_j \in \{0,1\}, \quad \forall j \in F \qquad (29)$$

$$0 \le \rho_i \le \vartheta_i, \quad \forall i \in N \qquad (30)$$

$$\pi_j, v_j \ge 0, \quad \forall j \in F \qquad (31)$$

This formulation is a linear MIP and can hence be solved by standard, off-the-shelf optimization software.

5. Solving the bilevel CRIMF to optimality

As previously mentioned, by collapsing the interdictor-user problem into a single level program, CRIMF becomes a bi-level problem. We can therefore solve it by using the implicit enumeration algorithm originally proposed to solve the uncapacitated version of the r-interdiction median problem with fortification (Scaparra and Church 2008b). Different variants of this implicit enumeration method have been successfully applied for solving other protection problems, which could not be easily solved by standard decomposition methods commonly used for bilevel programs. Examples of successful applications can be found both within the context of uncapacitated facility protection (Liberatore et al. 2009) and network protection (Cappanera and Scaparra 2009).

The core of the approach is based upon the observation that an optimal protection strategy must include at least one of the facilities which are interdicted in the optimal solution to the interdiction problem without fortification (Church and Scaparra 2006). This observation can be used recursively within an implicit enumeration scheme to reduce the number of protection strategies to be evaluated.

The implicit enumeration algorithm is implemented as a search on a binary tree. At the root node, the interdiction problem (23)-(31) is solved by a MIP solver to identify the optimal set of interdictions when no protection takes place. Then a branching procedure is

11

started on the protection variables associated with the facilities in the optimal interdiction set, called the candidate set of branching variables (CS). When a protection variable z_j is selected from the CS of the current tree node and fixed to one, a new tree node is generated and an interdiction problem is solved with an additional constraint which prevents the interdiction of facility j. The optimal solution to this problem generates the candidate set of variables to branch on associated with the new node. When a protection variable z_j is selected from the CS and fixed to zero, a new tree node is generated which inherits the CS from its parent node but without the variable z_j. If the CS of the new node is empty, the node is fathomed; otherwise a new variable is selected from it for branching. A leaf node is reached when exactly q protection variables are fixed to one along a path from the root node to the current node. The algorithm visits the tree nodes according to a depth-first strategy and at each node the variable to branch on is selected at random from the CS.

A nice feature of this approach is that the size of the tree and, hence, the number of interdiction problems (23)-(31) which are solved during the execution, is not affected by the number of facilities in the initial configuration, p, but only by the parameters q and r (the reader is referred to Scaparra and Church, 2008 for a formal proof). Additionally, at termination, the algorithm identifies all optimal protection strategies, if more than one exists, thus providing a system planner with the flexibility of choosing the preferred alternative on the base of other managerial and planning criteria.

6. Experimental Results

We have conducted a set of computational experiments to assess the overall effectiveness of the CRIMF modelling approach and to provide some managerial insights on the potential impact of protection strategies on system efficiency improvements. Specifically, we have focused the analysis on 1) the tractability of the CRIMF model; 2) the impact of protective resources on overall cost reduction in case of disruption; and 3) the sensitivity of the protection strategies to the number of potential losses.

6.1 Data Sets and Experimental Setup

The computational experiments were conducted on a benchmark data set, frequently used in location analysis: the London, Ontario data set (Goodchild and Noronha 1983). This data set contains 150 demand nodes, and distances are based upon a road network. Using this data set, we have solved problem instances for different combinations of the parameters p, q and r. Namely, we considered values of p in {20, 25, 30}, q in {1, 2, 3, 4,

5} and r in {1, 2, 3, 4, 5}. In our testing, we assumed that the p facilities in the initial configuration are located at the optimal sites identified by solving a capacitated p-median problem. As the London data set does not contain information about facility capacities, for each value of p, we set the capacity of the facilities equal to $\frac{\Sigma_i a_i}{0.9p}$. This choice is equivalent to assuming that the capacity utilization of the system under normal condition is about 90% or, in other words, that the system has a 10% built-in idle capacity. Finally, we assumed that the lost sale cost incurred if a customer cannot receive service after a disruption is larger than the cost of serving any customer from any facility in the system. More specifically, for each customer i, we set $\vartheta_i = 1.5max_{ij}d_{ij}$.

The experiments were run on a PC equipped with an Intel Core 2 CPU @ 2.4 GHz, 3GB of RAM and Windows XP Professional operating system. The implicit enumeration algorithm described in Section 5 was implemented in C++ and compiled using Microsoft Visual C++ .NET 2003. The interdiction problems (23)-(31) at each node of the enumeration tree were solved using the MIP solver Cplex 11.

6.2 Computational Analysis

The results of our computational investigation are summarized in Table 1 and Table 2. Table 1 displays the objective function value of the optimal protection strategies for various combinations of the parameters. Reading the table by row, we can observe for each system size, p, and protection resource level, q, the impact of additional losses on the system operational cost. Conversely, each column highlights the impact on the objective of increasing the protection resources. A more in-depth analysis of the latter issue is provided in the following section.

For the same combinations of parameters, Table 2 shows the computing time in seconds. The solution times suggest that we can solve to optimality CRIMF instances of realistic size with relative ease, if the number of possible losses is small. As expected, the computing effort increases appreciably for larger values of r, larger systems and higher levels of protection resources. Whereas all the instances with r < 2 were solved in a matter of seconds, larger values of this parameter led to much higher running times. The most difficult instance (p = 30, q = 5, r = 5) required more than a day to be solved to optimality. However, given that this kind of protection models is strategic in nature, computing times in this order of magnitude are not overly critical and do no represent an impediment to the actual implementation of the optimal protection strategies. Overall, we can conclude that the CRIMF model is tractable for all the selected combinations of

parameters. Also, we believe that these ranges of parameter values are the most interesting from a practical point of view. A coordinated attack on or the simultaneous loss of a much larger number of facilities than the ones considered here seems unrealistic or, at least, very unlikely to occur. Also, the effects of interdiction and fortification are much less critical in systems with a very large number of facilities, due to greater levels of redundancies present in such systems. Thus the case handled here with systems of around 30 facilities and possible losses of 4-5 facilities seems quite realistic. On the other side, fortification resources could vary considerably. We have reported results for protection strategies involving the fortification of up to 5 facilities, but our modelling approach can handle larger values as well, especially if the number of interdictions, which is the real critical element of the approach, is small.

Table 1. Objective values for the London, Ont. data set and different values of p, q and r

p		r = 1	r = 2	r = 3	r = 4	r = 5
	q = 1	146,091.49	181,437.34	374,022.61	587,408.95	803,568.21
	q = 2	143,269.58	178,895.82	373,906.37	584,615.28	798,523.83
20	q = 3	140,640.15	171,150.05	370,833.41	582,509.06	795,972.70
	q = 4	140,409.57	170,172.52	369,481.99	580,218.68	792,425.57
	q = 5	140,406.91	170,107.07	368,105.01	579,142.22	789,787.80
	q = 1	127,892.53	153,265.74	228,495.65	393,654.45	564,483.89
	q = 2	124,960.03	150,999.82	225,191.95	390,720.46	562,591.84
25	q = 3	124,169.93	148,909.87	222,876.47	389,744.09	561,136.40
	q = 4	123,362.77	144,310.84	220,408.47	388,750.33	558,391.70
	q = 5	121,180.59	142,107.89	218,074.57	386,808.49	557,579.34
	q = 1	110,621.32	125,830.69	147,235.51	251,914.07	395,240.27
	q = 2	106,093.07	119,684.24	145,398.43	250,496.84	393,437.48
30	q = 3	104,969.78	118,077.38	140,471.42	248,534.05	390,634.70
	q = 4	104,853.24	117,964.80	138,018.79	246,463.88	387,888.26
	q = 5	104,637.45	117,442.56	137,346.00	246,410.57	387,798.18

Table 2. Running times (sec.) for the London, Ont. data set and different values of p, q and r

p		r = 1	r = 2	r = 3	r = 4	r = 5
	q = 1	1.23	7.89	29.02	32.78	27.28
	q = 2	1.97	16.56	83.10	78.58	77.56
20	q = 3	2.80	29.36	255.84	173.64	251.94
	q = 4	3.44	46.30	536.99	403.03	555.38
	q = 5	4.06	65.64	935.79	732.54	1,131.21
	q = 1	1.89	23.38	179.13	283.32	283.95
	q = 2	2.78	41.97	424.39	810.51	1,173.43
25	q = 3	3.67	73.41	796.16	1,626.42	3,645.48
	q = 4	6.98	112.50	1,203.80	2,961.13	8,398.46
	q = 5	7.86	154.42	1,993.71	5,523.97	17,155.18
	q = 1	2.47	49.48	362.14	1,311.12	1,822.96
	q = 2	4.13	123.63	900.10	3,772.30	7,273.92
30	q = 3	5.88	242.49	1,802.05	9,083.70	23,750.22
	q = 4	6.91	394.93	3,270.84	21,240.75	58,955.12
	q = 5	11.70	551.72	6,288.59	47,558.06	119,013.35

6.3 Marginal Efficiency Analysis

The purpose of this section is to present a practical analysis framework to help system analysts and policy makers identify efficient employments of protective resources, using the results of the CRIMF model. Specifically, we analyze the impact that additional protection investments may have on the overall system efficiency (or cost reduction) in case of disruptions of different magnitude. We illustrate the analysis using the London, Ontario data set with 20, 25 and 30 operating facilities. The results of the analysis are displayed in Fig. 1. The column graph shows the percentage marginal efficiency improvement, in terms of reduced service and lost sale costs, derived from any individual protection. It can be noted that the marginal efficiency gains are not monotonically decreasing in the number of fortifications. Each column presents the marginal improvement in system efficiency associated with each added level of fortification. As an example, in a system with 20 facilities (p = 20) and 2 possible losses (r = 2), the optimal protection of a single facility (q = 1) results in an efficiency increase of slightly more than 2.5% (the lowest segment of the second column in the graph). The protection of the second facility contributes an additional 1.5% improvement (resulting in a total improvement of 4% for the first two facilities protected). However, the greatest

percentage marginal benefit is obtained when the third facility is protected. This yields a further 4% improvement with an overall efficiency gain of more than 8%.

The graph in Fig. 1 also highlights possible resource wastage. For instance, when p = 20 and r = 3 (third column), there is almost no added value in protecting two facilities rather than one, whereas the protection of a third facility more than triples the return on investment compared to the second fortification. In such a situation, security investments are warranted only to protect one or three facilities, depending upon the resource availability, risk tolerance and minimum efficiency requirement. However, security investment for protecting two facilities should be avoided for this case. Finally, it is worth noticing that the percentage efficiency improvements are more noticeable when the system is still able to provide service to all the customers in spite of the disruption, whereas they are less pronounced when lost sale costs are incurred. As an example, the system with 20 facilities is still able to satisfy all customer demands if one or two facilities are lost due to disruption (columns 1 and 2 in the graph). However, if three or more facilities are disrupted, the remaining facilities cannot supply all the demand. The percentage efficiency improvement in these cases is somewhat less noticeable (columns 3, 4 and 5). A similar behaviour can be noticed for the systems with 25 and 30 facilities.
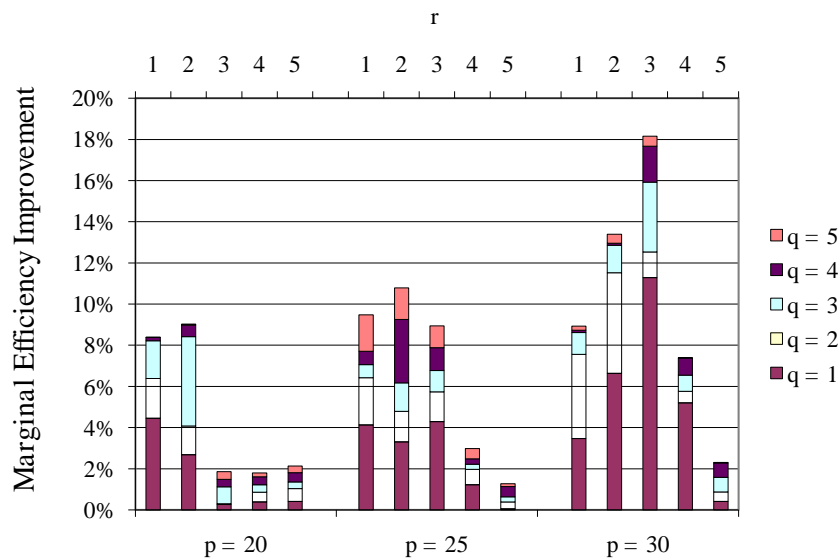

Fig. 1. Marginal percentage improvement in efficiency due to any additional protection

6.4 Solutions Sensitivity to the Number of Potential Losses

In this final section, we address the issue of estimating the parameter r and analyze the impacts of making a miscalculation/estimation of this parameter in terms of which optimal protection strategy is selected. The choice of r is a critical one and fixing it to a specific value may seem subjective. The number of terrorist targets or potential losses due to catastrophic events, in fact, cannot be estimated on the base of historical data and is therefore difficult to predict accurately. To overcome this potential shortcoming of our modeling approach, we have solved the CRIMF model for different values of r. We then used the results obtained to infer core sets of key facilities to harden and, eventually, to identify the best protection strategy across the range of r values considered. An example of this analysis is depicted in Tables 3, 4 and 5 for the London, Ontario data set with 30 facilities and assuming that resources are available to protect 5 facilities (q = 5). Table 3 displays the optimal protections and interdiction responses identified by solving CRIMF with values of r ranging between 1 and 5. It can be seen that there are some critical facilities (e.g., facility 47) that are selected in every protection plan, independently on the number of expected losses. On the other side of the spectrum, some facilities only appear in one protection or interdiction set (e.g. 111 and 117).

To identify the overall best protection set, we consider the optimal protection plan computed for a giver r and evaluate the cost of that plan if a different number of losses occur in practice. The results are displayed in Table 4. The minimum cost in each column is achieved along the diagonal, since in this case the assumed number of losses is the one occurring in reality. For all the other cases, it is possible to compute the extent to which the cost may increase if the r is misestimated. The absolute cost increase for each pair of assumed-actual number of losses is displayed in Table 5. The maximum and average cost increase for each assumed value of r are also displayed in the last two columns. This analysis suggests that the best protection strategy for this problem instance is the one obtained with r = 2 as it results in the minimum maximum and average absolute increases in cost. Note that the protection plans identified with high values of r (e.g. r = 4 or r = 5) can perform quite poorly if the extent of the disruption is smaller. In this case, in fact, cost increases of about 7% can be observed as compared to the optimal protection plans (e.g. if the actual r is 2 or 3).

Although stochastic models which explicitly take into account expected numbers of losses have been developed for the uncapacitated RIMF (Liberatore et al. 2008, Liberatore and Scaparra 2009), these models cannot be easily adjusted to deal with

capacity restrictions. In the absence of more advanced stochastic models, the type of analysis illustrated above may represent a viable alternative for the identification of sound protection strategies against disruptions of unknown magnitude.

Table 3. Optimal protections and interdictions for the London, Ont. data set with p=30 and q = 5.

r	Optimal Protection Set					Optimal Interdiction Set				
1	47	56	66	119	149	141				
2	47	56	84	141	149	66	108			
3	8	47	119	141	149	56	66	84		
4	47	66	84	117	141	56	70	108	111	
5	47	56	66	70	84	39	103	141	144	149

Table 4. Cross-comparison of r-optimized CRIMF solutions across different number of facility losses. Objective function values.

Assumed number of losses	Actual number of losses				
	r = 1	r = 2	r = 3	r = 4	r = 5
r = 1	104,637.45	119,684.24	145,398.43	250,705.21	391,321.88
r = 2	104,969.78	117,442.56	140,471.42	247,382.91	391,133.97
r = 3	106,093.07	118,077.38	137,346.00	250,496.84	395,240.27
r = 4	110,621.32	125,830.69	146,662.16	246,410.57	390,634.70
r = 5	110,621.32	125,830.69	147,235.51	251,914.07	387,798.18

Table 5. Cross-comparison of r-optimized CRIMF solutions across different number of facility losses. Absolute increase in costs.

Assumed number of losses	Actual number of losses						
	r = 1	r = 2	r = 3	r = 4	r = 5	Max	Avg
r = 1	0.00	2,241.68	8,052.43	4,294.64	3,523.70	8,052.43	3,622.49
r = 2	332.33	0.00	3,125.42	972.34	3,335.79	3,335.79	1,553.18
r = 3	1,455.62	634.82	0.00	4,086.27	7,442.09	7,442.09	2,723.76
r = 4	5,983.87	8,388.13	9,316.16	0.00	2,836.52	9,316.16	5,304.94
r = 5	5,983.87	8,388.13	9,889.51	5,503.50	0.00	9,889.51	5,953.00

7. Summary and Conclusions

Policies to mitigate possible emergency events such as natural and man-made disasters are facing increasing scrutiny. Mitigation strategies include among others the protection or hardening of physical infrastructure components so as to make service and supply systems more resilient to external disruptions and circumscribe the harmful ripple effects of disruptive events. In this paper, we have proposed a new modeling approach to identify sound protection strategies for capacitated facility systems, so that limited protection resources are utilized in the most trenchant way possible. The introduction of capacity constraints within a protection-interdiction model greatly increases the complexity of the problem and requires the formulation of a three-level program. We showed how this multi-level program can be reduced to a bi-level model by dualization of the inner level user problem. This bi-level model can be solved by the implicit enumeration algorithm proposed in Scaparra and Church 2008b for the uncapacitated problem. We have also presented an application of this model and suggested how model results could be used in advanced planning.

In terms of future work, we envision several possible extensions and variations of the proposed protection model. These may include: stochastic components, multi-period planning models, multi-tier supply systems and inventory management issues. Analogous models which optimize protection strategies against expected losses rather than worst-case losses will also be investigated.

Acknowledgments

References

Albright M (2009) Big retailers plan for hurricane season with military precision. St. Petersburg Times (June 7, 2009), St. Petersburg, Florida.

Alcada-Almeida L, Tralhao L, Santos L, Countinho-Rodrigues J (2009) A Multiobjective Approach to Locate Emergency Shelters and Identify Evacuation Routes in Urban Areas. Geographical Analysis 41(1): 9-27.

Altay N, Green WG (2006) OR/MS research in disaster operations management. European Journal of Operational Research 175(1): 475-493.

Azaiez MN, Bier VM (2007) Optimal resource allocation for security in reliability systems. European Journal of Operational Research, 181:773–786, 2007.

Bayrak H, Bailey MD (2008) Shortest path network interdiction with asymmetric information. Networks 52(3):133-140.

Brachman ML, Church RL (2009) Planning for a disaster: A review of the literature with a focus on transportation related issues. FiRST Report, Geotrans Laboratory, UCSB, Santa Barbara CA.

Brown G, Carlyle M, Salmeron J, Wood K (2006) Defending critical infrastructure. Interfaces 36(6):530–544.

Cappanera P, Scaparra MP (2008) Optimal allocation of protective resources in shortest path networks. KBS Working Paper No.177, University of Kent.

Church RL, Scaparra MP (2006) Analysis of facility systems' reliability when subject to attack or a natural disaster. In: Critical Infrastructure: Reliability and Vulnerability. A. Murray and T. Grubesic Eds. Publisher: Springer.

Church RL, Scaparra MP (2007). Protecting critical assets: The r-interdiction median problem with fortification. Geographical Analysis 39:129-146.

Church RL, Scaparra MP, Middleton RS (2004) Identifying critical infrastructure: The median and covering facility interdiction problems. Annals of the Association of American Geographers 94(3) 491-502.

Colson B, Marcotte P, Savard G (2007) An overview of bilevel optimization. Annals of Operations Research 153(1): 235-256.

Dempe S (2002) Foundations of Bilevel Programming. Kluwer Academic Publishers, The Netherlands.

Goodchild MF, Noronha VT (1983) Location-allocation for small computers. Monograph No. 8. Iowa City IA, Department of Geography, The University of Iowa.

Grubesic TH, O'Kelly ME, Murray AT (2003) A geographic perspective on internet survivability. Telematics and Informatics 20: 51-69.

Haimes YY (2006) On the Definition of Vulnerabilities in Measuring Risks to Infrastructures. Risk Analysis 26(2): 293-296.

Haimes YY, Longstaff T (2002) The Role of Risk Analysis in the Protection of Critical Infrastructures Against Terrorism. Risk Analysis 22(3): 439-444.

Israeli E, Wood RK (2002) Shortest-path network interdiction. Networks 40:97-111.

Liberatore F, Scaparra MP, Daskin M (2008) Analysis of facility protection strategies against an uncertain number of attacks: The stochastic R-interdiction median problem with fortification. KBS Working Paper 176, University of Kent. Under review.

Liberatore F, Scaparra MP (2009) Optimizing protection strategies for supply chains: Comparing classic decision making criteria in an uncertain environment. KBS Working Paper 186, University of Kent. Under Review.

Lim C, Smith JC (2007) Algorithms for Discrete and Continuous Multicommodity Flow Network Interdiction Problems. IIE Transactions 39: 15-26.

Losada C, Scaparra MP, Church R, Daskin M (2009) Modeling approaches for the multi-source interdiction median problem. KBS Working Paper No.187, University of Kent. Under review.

Moore JT, Bard JF (1990). The mixed integer linear bilevel programming problem. Operations Research 38 911-921.

Murray AT, Matisziw TC, Grubesic TH (2003) Critical network infrastructure analysis: interdiction and system flow. Journal of Geographical Systems 9(2): 103-117.

Murray AT, Matisziw TC, Grubesic TH (2008) A methodological overview of network vulnerability analysis. Growth and Change 39(4): 573-592.

O'Kelly ME, Kim H (2007) Survivability of Commercial Backbones with Peering: A Case Study of Korean Networks. In: Critical Infrastructure: Reliability and Vulnerability. A. Murray and T. Grubesic Eds. Publisher: Springer.

Peterson SK, Church RL (2008) A framework for modeling rail transport vulnerability. Growth and Change 39(4): 617-641.

Rawls CG, Turnquist MA (2006) Pre-positioning of Emergency Supplies for Disaster Response. Paper presented at the IEEE International Symposium on Technology and Society, Queens, NY.

Scaparra MP, Church RL (2008a) An exact solution approach for the interdiction median problem with fortification. European Journal of Operational Research, 189:76–92.

Scaparra MP, Church RL (2008b) A bilevel mixed integer program for critical infrastructure protection planning. Computers & Operations Research, 35:1905–1923.

Snyder L., Scaparra MP, Daskin M, Church RL (2006) Planning for disruption in supply chain networks. In: TutORials in Operations Research. Models, Methods, and Applications for Innovative Decision Making. M. Johnson, B. Norman, and N. Secomandi Eds. Publisher: INFORMS.

Stackelberg H (1952) The Theory of Market Economy. Oxford University Press.

Ukkusuri SV, Yushimito WF (2008) A Location-Routing Approach for the Humanitarian Pre-Positioning Problem. Transportation Research Record 2089: 18-25.

Viswanath K, Peeta S (2003) Multi-commodity maximal covering network design problem for planning critical routes for earthquake response. Transportation Research Record 1857: 1-10.

Wollmer R (1964) Removing arcs from a network. Operations Research 12 934-940.

Wolshon B (2006) Evacuation planning and engineering for hurricane Katrina. The Bridge 36(1) 27-34.

Wood RK (1993) Deterministic network interdiction. Mathematical and Computer Modeling 17 1-18.

Zhuang J Bier VM (2007) Balancing terrorism and natural disasters - defensive strategy with endogenous attacker effort. Operations Research 55(5):976–991.