Kent

Business

School

Working Paper Series

An optimal approach for the interdiction median problem with fortification

Maria P. Scaparra, Kent Business School

Richard L. Church

Working Paper No. 78

February 2005

An optimal modeling approach for the interdiction median problem with fortification

Maria P. Scaparra and Richard L. Church

Department of Geography

University of California, Santa Barbara

Santa Barbara, California 93106-4060

Abstract: Systematic approaches to security investment decisions, from intelligence and surveillance to fortification, are crucial for improved homeland security. We present an optimization modeling approach for allocating protection resources among a system of facilities so that the disruptive effects of possible intentional attacks to the system are minimized. This paper is based upon the *p*-median service protocol for an operating set of *p*-facilities. The primary objective is to identify the subset of *q* facilities which, when fortified, provides the best protection against the worst-case loss of *r* non-fortified facilities. This problem, known as the *r*-interdiction median problem with fortification (IMF), was first formulated as a mixed integer program by Church and Scaparra [6]. In this paper, we reformulate the IMF as a maximal covering problem with precedence constraints, which is amenable to a new solution approach. This new approach produces good approximations to the best fortification strategies. Furthermore, it provides upper and lower bounds that can be used to reduce the size of the original model. The reduced model can readily be solved to optimality by CPLEX. Computational results on two geographical data sets with different structural characteristics show the effectiveness of the proposed methodology for solving IMF instances of considerable size.

Keywords: *p*-median, interdiction, fortification, maximal covering

1. INTRODUCTION

Recent world events have clearly demonstrated that facilities are vulnerable to terrorism. In response to this, a number of governmental agencies in the U.S. have worked to identify the elements of infrastructure that are critical to safeguarding life and promoting safety. Some assets may not be critical in terms of economics, lifelines, or defense, but may be important as national symbols (e.g. the Statue of Liberty in New York Harbor). During times of heightened terrorist alerts, some of these assets have been the subject of increased surveillance and policing. Protection of other assets has been increased by hardening facility perimeters with vehicle barriers, limiting access, moving critical functions to interior areas, developing backup power systems, etc. It is important to identify critical infrastructure and to analyze alternatives of hardening/fortifying critical elements. In short, systematic approaches for making security investment decisions, including intelligence and surveillance, prevention, protection, emergency response and recovery among others, are crucial for guaranteeing the maximum effectiveness of protective efforts.

Identifying critical elements of a supply system has been the subject of military planners for quite some time. The focus of such planning is often directed at identifying the best places to disrupt or interdict an enemy's supply system by offensive actions such as a bombing mission. Interdiction modeling can be used to identify the weakest elements of a system or the worst case intentional sabotage of a system. Wollmer [21] was one of the first to model interdiction of supply lines as an optimization model. Since then, numerous papers have appeared, mainly dealing with the interdiction of transportation networks. The majority of past work is based on network optimization theory and generally aims at interdicting arcs in order to minimize network flow capacity (McMasters and Mustin [17], Ghare, Montgomery and Turner [10], Wood [22]) or maximize the shortest path between a specified origin and destination (Fulkerson and Harding [8], Golden [11], Israeli and Wood [16]). Interdiction of supply and emergency response facilities has recently been modeled (Church, Scaparra and Middleton [5]).

Interdiction models are most useful in a setting involving intentional disruption. As such they can be used to identify the critical links or assets in a system. Critical links of a system may be the best places in which to harden and fortify, so that they are less likely to be destroyed or lost in an intentional attack. Once vulnerabilities of a system have been identified, then it is logical to identify a protection plan (see for example Salmeron, Wood, and Baldick [19]). Such a plan would involve hardening certain assets so that a system operates as best as possible under cases of interdiction. The objective of this paper is to model the optimal allocation of asset fortification given the possibility of interdiction.

There are a number of different types of supply system designs. The one that we address in this paper draws upon the well-known p-median location model. The p-median problem has been applied in a number of public and private settings and was originally defined within the context of communication switching center location by Hakimi [13, 14]. The p-median model involves the optimal location of p-facilities that provide service or supply to a set of demands. Assuming that no capacity restrictions exist, the objective is to minimize the cost or the weighted distance of supplying all demand, where each demand is assigned to its closest facility (in terms of cost or distance). In this context, we assume that a supply system, where user demands are entirely supplied by their closest facilities, already exists and that no decision has to be made concerning the location of the facilities. However, the p facilities currently in the system are susceptible to deliberate sabotage by external *attackers,* unless protective measures are taken to prevent their interdiction. When an unprotected (i.e. non fortified) facility is interdicted, it is considered inoperable. This means that users must be reassigned to more distant facilities, with a commensurate reduction of system efficiency. It is possible that damages caused by interdiction can be repaired, but we assume here that the time that it would take to reestablish a facility or replace the supply is significant enough that the system would operate in an inefficient state for some period of time. We consider the case where fortification resources are limited. We assume that at most q of the existing p facilities can be fortified and that interdiction involves the loss of r of the p-q unfortified facilities. Our objective is to identify the optimal set of q facilities to fortify or harden in order to hedge against the most disruptive interdiction of r facilities.

The fortification/interdiction problem just described was first introduced by Church and Scaparra [6] and referred to as the r-interdiction median problem with fortification (IMF). In [6] the authors demonstrated that protecting the most vulnerable facilities or predictable targets is not necessarily the most cost-effective way of confronting threats. Fortification patterns which take into account the interdependency among the system components and the effect of multiple, simultaneous losses can produce better and more resilient protection plans. Mathematical models are hence needed which are able to capture these independencies and identify the most efficient protection investment plans.

A natural way of looking at the fortification/interdiction problem is within the context of a leader-follower or Stackelberg game, where the entity responsible for coordinating the fortification activity is the leader and the interdictor is the follower. Such a game can be expressed mathematically as a bilevel programming problem (Dempe [7]): the upper level problem involves decisions to determine which facilities to harden, whereas the lower level problem entails the interdictor responses on which unprotected facilities to attack. Even if in practice we cannot assume that the attacker is aware of the damage he can inflict on a system and, consequently, that he is able to identify the best attacking strategy, the assumption that the interdictor attacks in an optimal way is used as a tool to model worst case scenarios and estimate the worst case efficiency loss in response to a given fortification strategy.

The bilevel formulation of the fortification/interdiction problem is provided in Scaparra and Church [20]. In general, solving bilevel problems is a difficult task even in their simplest version, i.e. when both the lower level problem and the upper level problem only contain continuous variables. Such a case has been proved to be strongly NP-hard by Hansen, Jaumard and Savard [15]. Additional difficulties arise when the decisions at both levels require integrality constraints, as it is in the fortification/interdiction problem treated in this context (the reader is referred to Bard [2] and Moore and Bard [18] for a detailed treatment of integer bilevel programs). The development of efficient exact techniques to solve discrete bilevel programs is currently a fertile area of research and no universal algorithm exists for their solutions.

In order to solve the fortification/interdiction problem to optimality, Church and Scaparra [6] showed that, under specific hypothesis on the size of the problem parameters, the IMF problem can be formulated as a single-level mixed integer model and solved directly through commercial optimization software. However, the applicability of such an integer model is confined to problem instances with a few facilities and with modest interdiction and fortification resources. Many real distribution, supply and emergency response systems, such as electric utility company or fire station networks, may contain a much larger number of vulnerable facilities. Our objective here is to develop an alternative formulation for IMF and devise an efficient solution technique tailored to the new mathematical structure of the problem for solving instances of more realistic size.

The remainder of the paper is organized as follows. In the next section, we introduce some notation and present the original formulation of the r-interdiction median problem with fortification (IMF). In Section 3, we reformulate IMF as a special type of maximal covering problem, which involves precedence constraints. A heuristic approach for solving the new model is provided in Section 4. Section 5 describes an interval search improvement procedure. The complete modeling approach, which involves a reduced model and exploits the findings in the previous sections, is illustrated in Section 6. Numerical results obtained by solving the reduced model with CPLEX are presented in Section 7, followed by conclusions and recommendations for future work.

2. BACKGROUND AND THE INTERDICTION MEDIAN PROBLEM WITH FORTIFICATION

The scientific literature on network interdiction was recently surveyed in Church, Scaparra and Middleton [5]. They broaden the focus from the interdiction of arcs or links to the interdiction of facilities, such as power plants, hospitals, and emergency response facilities. They also proposed two new interdiction models that are based upon two different service protocols. Whereas the study of interdiction problems dates back several decades and is documented in a number of papers, little attention has been paid to the problem of asset protection for mitigating system degradation caused by interdiction. In a recent paper, Salmeron et al. [19] discussed the issue of hardening components of a electrical power system in order to improve the security of electrical supply against disruption caused by terrorist attacks. However, they did not formalize an approach for identifying the optimal set of components to harden. They suggested using the outcome of an interdiction model as an indication of the critical components to be fortified, but left the modeling of protective measures as a topic for future research. Church and Scaparra [6] have recently incorporated the option of fortification or asset hardening in a location model. The new model considers a p-median configuration of facilities, where q of p facilities can be fortified and where r of the remaining p-q unprotected facilities are subject to interdiction. They designated this as the interdiction median problem with fortification (IMF). They also demonstrated that protection practices, which do not account explicitly for fortification efforts, rarely deploy security resources to the greatest advantage. The IMF model and its equivalent formulation proposed in this paper, aim at capturing this additional degree of complexity in a single mixed integer program.

Let N, indexed by i, represent the set of demand nodes and denote by a_i the demand for service at node i. Assume that in the current system there are p operating facilities and that each demand is fully supplied by its closest facility. Let F, indexed by j, represent the facility set and let d_{ij} be the shortest distance between the facility at j and demand node i. In accordance with the median paradigm, system efficiency is measured as the sum of the shortest demand-weighted distances between each demand point and its closest facility. We assume that interdiction resources are limited so that the value of r is relatively small in value. We further assume that it is possible to enumerate all possible ways in which an existing facility pattern of p facilities can be interdicted r times. We define H, indexed by h, as the set of all possible interdiction patterns, and I_h as the set of interdicted facilities in pattern h. We can associate with each interdiction pattern h the value WD_h, which denotes the optimal p-median objective function value after the interdiction of the facilities comprising the interdiction set I_h. The value WD_h can be easily computed by simply reassigning the demands currently served by the facilities in I_h to their closest facility in $F \setminus I_h$. More specifically, let d_i^h be the shortest distance between demand i and its closest non-interdicted facility given the interdiction pattern h. Then, with no fortification,

$$WD_h = \sum_{i=1}^{n} a_i d_i^h \tag{1}$$

Finally, for each demand i and for each interdiction pattern h, we define the set $B_i^h = \{ j \in I_h \mid d_{ij} < d_i^h \}$. B_i^h represents the set of closest sites to i that have been interdicted in pattern h. An interdiction pattern h is partially thwarted if any of the facilities in I_h is fortified.

In order to formulate the r-interdiction median problem with fortification (IMF) as proposed in Church and Scaparra [6], we consider the following decision variables:

$$z_j = \begin{cases} 1, & \text{if a facility located at } j \text{ is fortified} \\ 0, & \text{otherwise} \end{cases}$$

$$x_{ij}^h = \begin{cases} 1, & \text{if demand } i \text{ assigns to fortified facility } j \text{ in interdiction pattern } h \\ 0, & \text{otherwise} \end{cases}$$

The IMF model is then:

$$\min\ W \tag{2}$$

s.t. $$x_{ij}^h \le z_j \quad \text{for all } i \in N, \text{for all } h \in H, \text{and for all } j \in B_i^h \tag{3}$$

$$\sum_{j \in B_i^h} x_{ij}^h \le 1 \quad \text{for all } i \in N, \text{and for all } h \in H \tag{4}$$

$$W \ge WD_h - \sum_{i \in N} \sum_{j \in B_i^h} a_i (d_i^h - d_{ij}) x_{ij}^h \quad \text{for all } h \in H \tag{5}$$

$$\sum_{j \in F} z_j = q \tag{6}$$

$$z_j \in \{0,1\} \quad \text{for all } j \in F \tag{7}$$

$$x_{ij}^h \in \{0,1\} \quad \text{for all } i \in N, \text{for all } h \in H, \text{and for all } j \in B_i^h \tag{8}$$

where:

W = the weighted distance resulting from the fortification of q facilities, assuming that the worst-case interdiction of r non-fortified facilities occurs.

The above model involves minimizing the weighted distance of serving all demand assuming that each demand is served by its closest non-interdicted site. It is based upon the assumption that given a pattern of fortification, the worst case of r-interdiction occurs. That is, unprotected facilities will be interdicted in such a manner that the weighted distance will be maximized given a fortification plan. The overall objective is to find the fortification plan that minimizes the impact of interdiction. To accomplish this, the model keeps track of the impact of each possible interdiction pattern. A given interdiction pattern can be thwarted in its effectiveness by fortifying one or more sites in that pattern. The impact of each possible interdiction pattern is based upon which facilities are fortified. Constraints (3) precludes the assignment of a demand i to a facility j which is interdicted in pattern h, unless j is fortified. Constraints (4) ensure that each demand i is assigned to at most one facility, given that interdiction pattern h occurs. Such a facility is either the closest non-interdicted site (if $B_i^h = \emptyset$ or none of the facilities in B_i^h is fortified) or the closest fortified site. Assignment to the closest site is guaranteed by the minimization of the problem. Note that no assignment variables x_{ij}^h exist if the closest facility to i is not interdicted in interdiction pattern h. Thus, there are no constraints (3) and (4) corresponding to such a case. Constraints (5) define the weighted distance after interdiction for each pattern h in terms of the new assignment variables. If none of the facilities in interdiction pattern h is fortified, or if the fortification does not affect the customer assignments, then the weighted distance WD_h associated with interdiction pattern h remains unchanged. Otherwise, the second term in the right-hand side of inequalities (5) represents the improvement in weighted distance derived by reassigning demands to the fortified facilities. Constraints (5) simply force the variable W to be the worst-case weighted distance after interdiction of non-fortified facilities. Hence, the model optimizes the weighted distance after the most disruptive interdiction in response to the selected fortifications. Constraint (6) simply states that q facilities are fortified. Finally, constraints (7) and (8) force integrality of the decision variables.

Church and Scaparra [6] discuss ways of reducing the integer program (2)-(8). Such reductions, including variable replacements and constraint eliminations, are based upon considerations on the size of the sets B_i^h, and upon elements of the COBRA formulation of the p-median problem given in Church [4] that can be utilized in IMF. We will show in Section 7 that the solution procedure that we introduce in this paper is considerably faster than solving the reduced form of the IMF formulation given in Church and Scaparra [6].

3. A MAXIMUM COVERING TYPE FORMULATION

The mathematical formulation of IMF given in section 2 is based upon the assumption that the number of interdictions, r, is relatively small and that, consequently, it is possible to identify in advance all the p choose r interdiction patterns. In this context, we make the same assumption. We also assume for the remainder of this paper that $q + r < p$. Additionally, we assume that the interdiction patterns are sorted in non-increasing order of the median objectives after interdiction,

WD_h, so that pattern 1 is the most disruptive interdiction pattern and pattern $|H|$ is the least disruptive interdiction pattern. We now make the following observation.

Observation 1. The worst-case interdiction will occur for an interdiction pattern attacking solely unprotected sites. If an interdiction pattern I_h is partially thwarted by the fortification of one or more facilities in that pattern, an interdiction pattern I_k exists that consists of the unprotected sites of I_h and other unprotected sites such that the weighted distance of I_k exceeds that of I_h. This will always be true given that $q + r < p$. Thus, fortification of one or more sites of an interdiction pattern means that that interdiction pattern will not represent the worst case.

If the interdiction patterns can be enumerated and sorted, then the worst-case interdiction of r facilities can be easily identified in response to any given set of q fortifications, as stated below.

Proposition 1 Let S be a set of q candidate facilities for fortification and assume that the interdiction patterns have been reindexed so that $WD_1 \geq \ldots \geq WD_h \geq \ldots \geq WD_{|H|}$. Then, the worst-case loss of r facilities after the fortification of S coincides with the set I_{h^*}, where $h^* = \min\{h \in H \mid I_h \cap S = \emptyset\}$.

Proof. The proof is straightforward. Namely, every pattern preceding h^* in the ordering is thwarted by the fortification of S since by definition h^* is the index of the first interdiction pattern in the ordering which has no facility in the fortification set. Every pattern following pattern h^* in the ordering is less disruptive than h^*. Hence I_{h^*} is the worst-case loss in response to S.

Observation 2. The IMF problem can then be restated as the problem of finding the set of fortifications S such that the index h^* of the worst interdiction pattern in response to S is as large as possible. This implies that the corresponding median objective after the worst-case interdiction, WD_{h^*}, is as small as possible.

Proposition 1, together with the two observations imply that the IMF model can be formulated in a different manner. The effectiveness of a fortification pattern is based upon whether one or more sites in the most disruptive interdiction patterns have been fortified. The basic premise of this new model is to *cover* the greatest number of most disruptive interdiction patterns with q fortifications. An interdiction pattern h is *covered* by the fortification of a facility j if j belongs to

I_h. To define the problem formally, we use the same notation given in Section 2. Additionally, we introduce the following decision variables:

$$y_h = \begin{cases} 1, \text{if interdiction pattern } h \text{ is thwarted (covered)} \\ 0, \text{otherwise} \end{cases}$$

Under the assumption that the interdiction patterns are ordered, IMF can be formulated as the following maximal covering problem (MCP):

$$\max \quad \sum_{h \in H} w_h y_h \tag{9}$$

$$\text{s.t.} \quad \sum_{j \in I_h} z_j \geq y_h \quad \text{for all } h \in H \tag{10}$$

$$\sum_{j \in F} z_j = q \tag{11}$$

$$z_j \in \{0,1\} \quad \text{for all } j \in F \tag{12}$$

$$0 \leq y_h \leq 1 \quad \text{for all } h \in H \tag{13}$$

where the objective weights w_h are recursively defined as follows:

$$w_h = \begin{cases} g & , \text{ if } h = H \\ \sum_{k>h} w_k + \varepsilon & , \text{ if } h = H-1,\ldots,1 \end{cases}$$

In the above definition, g and e are small, positive real numbers. This specific choice of the objective weights ensures that, in an optimal solution, the coverage of a pattern h is preferred to the coverage of all patterns following h in the ordering, i.e., each pattern will have higher coverage priority than the sum of all the patterns that are less disruptive. Consequently, a solution to problem (9)-(13) identifies the set of q fortifications that covers the greatest number of most disruptive interdiction patterns. Note that the integrality restriction on the y_h variables can be relaxed as in the standard formulation of the maximal covering problem. The next proposition establishes the equivalence between IMF and MCP.

Proposition 2. Let (z^*, y^*) be an optimal solution to MCP and $h^* = \min\{h \in H \mid y_h^* = 0\}$. Also, let W^* be the objective function value corresponding to the optimal solution to IMF. Then, $WD_{h^*} = W^*$.

Proof. The proof follows directly from Proposition 1.

Unfortunately, even when the number of interdiction patterns is relatively modest, the use of weights to force an ordered coverage of the interdiction patterns becomes impractical. This is due to the exponential growth in the magnitude of the weights, and is true also for very small values of the base weight g and of the weight increment e. As an example, consider a small problem instance with 10 operating facilities and resources to interdict 4 facilities, and assume that both g and e are fixed to 0.001. The number of possible interdiction patterns is only 210, but the weight associated with the most disruptive interdiction pattern is $w_1 = 8.2\times10^{59}$. It is clear that the use of weights represents a major limitation to the applicability of the maximal covering model for solving IMF instances of relevant size. However, as an alternative to the use of weights, priorities in covering the interdiction patterns can be enforced through the use of the following precedence constraints:

$$y_h \geq y_{h+1} \quad \text{for all } h = 1..|H|-1, \tag{14}$$

where the variables y_h are redefined as:

$$y_h = \begin{cases} 1, \text{ if all the interdiction patterns from 1 to } h \text{ are thwarted (covered)} \\ 0, \text{ otherwise} \end{cases}$$

We will refer to the maximal covering model with precedence constraints (14) as MCPC. In MCPC, the objective coefficients can all be set to 1. The resulting model has p integer variables, $|H|$ continuous variables and $2\times|H|$ constraints. Clearly, as the number of patterns increases (as a consequence of increased values of the parameters p and r), solving this model using general-purpose optimization software may become computationally prohibitive. In the next section, we describe a greedy algorithm, which finds approximate solutions to the mixed integer program (9)-(14). We then show how the approximate solutions can be improved by searching for the optimal value h^* in a subinterval of $[1, |H|]$, and show how the information obtained at the end of this process can then be employed to reduce the size of the MCPC.

4. AN APPROXIMATE SOLUTION ALGORITHM

We first developed a simple heuristic technique for solving the MCPC that could be used to calculate a valid bound on fortification. This approach consists of a greedy strategy that starts with an empty set of fortifications and iteratively adds facilities to it according to a greedy rule. At each iteration, the newly selected facility is the one which, when fortified, thwarts the greatest number of uncovered, most disruptive interdiction patterns. The steps of the greedy process are defined below. Note that the interdiction patterns are still assumed to be in non-increasing order of WD_h.

Greedy Procedure

1. Set $k = 1$ and $S = \emptyset$.
2. For each facility j in interdiction pattern k, compute the number n_j of consecutive patterns after pattern k which are either already covered by some facility in S or that would be covered by the fortification of j. Let j^* be the facility for which this number is the largest, i.e., $n_{j^*} = max\ \{n_j \mid j \in I_k\}$.
3. Set $S = S \cup \{j^*\}$ and $k = k + n_{j^*} + 1$.
4. If $|S| < q$, repeat steps 2-3. Otherwise, set $L = k - 1$ and stop.

At termination, L represents the index of the last pattern in the ordering that is thwarted by hardening the facilities in the fortification set S. Thus, L is a lower bound estimate of the optimal number of worst-case interdiction patterns that can be prevented with q fortifications. By proposition 1, I_{L+1} is the worst-case interdiction set in response to S, and WD_{L+1} is an upper bound to the optimal objective W^*.

The greedy heuristic can be implemented to run in $O(r|H|)$ time, provided that the interdiction patterns have been previously sorted. In fact, the main step of the algorithm (step 2) is executed only q times. Each repetition of step 2 involves the identification of an additional facility j^* to be inserted in S and requires checking the coverage of at most $n_{j^*} + 1$ patterns for each of the r

interdicted facilities of the pattern under exam. However, when considering the computational work involved for a completed greedy solution, each interdiction pattern is checked for coverage at most r times. This checking operation can be performed in constant time by using appropriate data structures. Hence, the overall time required by the greedy algorithm is $O(r|H|)$. The sorting of the patterns can be implemented in $O(|H|\log|H|)$. In the next section we show how the greedy result can be improved.

5. IMPROVING THE BOUNDS BY INTERVAL SEARCH

The procedure for enhancing the greedy solution and for reducing the size of MCPC consists of an interval search performed over a subset of the interdiction patterns. The interval search aims at: 1) improving the current lower estimate, L, on the optimal number of most disruptive interdiction patterns that can be thwarted with q fortifications, denoted by k^*, and 2) providing an upper estimate, U, on this number. The search starts by considering the interval $[L, |H|]$. We know that the optimal number of thwarted patterns k^* falls within this interval At each subsequent step, a new trial point k is selected in the interval, and is evaluated by determining the minimum number of fortifications required to thwart all of the first k ordered patterns (by solving a set covering model in step 2 of the search procedure given below). This information is then used to obtain a new, smaller bracketing interval. The selection of each trial point uses linear interpolation for estimating the optimal k^* from known lower and upper values. In the following, we denote by $[L, U]$ the current bracketing interval, and by q_L and q_U the number of fortifications needed to disrupt the first L and U interdiction patterns, respectively.

Linear Interpolation Search Procedure

1. *Initialization.* Set $q_L = q$, $U = |H|$, $q_U = p - r + 1$, and $k = L$.
2. *Evaluation.* Solve the following set covering problem SC(k) :

$$\begin{aligned} \min \quad & \sum_{j \in F} z_j \\ \text{s.t.} \quad & \sum_{j \in I_h} z_j \geq 1 \quad \text{for all } h \in 1..k \qquad \text{SC}(k) \\ & z_j \in \{0,1\} \quad \text{for all } j \in F \end{aligned}$$

Let $\hat{q}(k)$ be the optimal solution to SC(k).

3. *Interval Reduction.* If $\hat{q}(k) > q$, then set $U = k$, $q_U = \hat{q}(k)$. Otherwise, $L = k$, $q_L = \hat{q}(k)$.
4. *Point Selection.* Compute:

$$x = \begin{cases} (q - q_L)\dfrac{U-L}{q_U - q_L} + L, & \text{if } q_L \neq q \\ \\ (q+1-q_L)\dfrac{U-L}{q_U - q_L} + L, & \text{if } q_L = q \text{ and } q_U > q+1 \\ \\ L + 0.5(U - L), & \text{if } q_L = q \text{ and } q_U = q+1. \end{cases} \tag{15}$$

   Set $k = \lfloor x \rfloor$, where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x.
5. *Termination.* Repeat steps 2-4 until some stopping criterion is met.

Step 1 initializes the interval bounds with the corresponding values q_L and q_U, and the initial trial number of most disruptive patterns that can be covered with q fortifications, k. Note that the initial values of q_L and q_U are upper bounds on the minimum number of fortifications needed to disrupt the first L and U patterns, respectively. More specifically, q_L is the solution returned by the greedy heuristic, while q_U is assigned an initial value of $p - r + 1$, since fortifying that number of facilities ensures that all the $|H|$ interdiction patterns are disrupted. In step 2, a set-covering problem is solved to determine the minimum number of facilities that must be fortified to thwart the first k most disruptive patterns. If the disruption of the worst k interdiction patterns requires more than the available q fortifications, then k becomes the new upper estimate U. Conversely, if it is possible to cover all of the worst k patterns with q fortifications, k becomes the new lower estimate L. At step 3, the new trial number k is appropriately chosen in the reduced interval [L, U]. The choice of k is based upon the assumption that the number of covered patterns grows linearly as a function of the number of fortifications. That is, the new value k is interpolated so that the point (k, q) lies on the line segment connecting the points (L, q_L) and (U, q_U). The first equation in (15) represents a standard linear interpolation formula. This formula needs to be modified in two cases, which both occur when $q_L = q$. When this is case, the first interpolation formula given in (15) would return the point L as the next estimate for k*. If q_U is greater than q + 1, the problem can be overcome by estimating the new value at the point $q + 1$ instead of q, as stated by the second equation in (15). On the other hand, if $q_U = q + 1$, the interpolation at $q + 1$

would return U as the next estimate. In this case, we use bisection instead of interpolation to compute the next k, according to the last part of equation (15).

Figure 1 provides an example of these two special cases for a simple problem with $p = 15$, $r = 4$, and $q = 5$. The piecewise non-continuous step shaped function in the figure represents the minimum number of fortifications needed to cover any given number of the most disruptive interdiction patterns. The function is depicted to illustrate the search process, even though it is a solution on this function that is the object of our search. The point on the function marked with the asterisk corresponds to the optimal number of patterns that can be thwarted with the 5 available fortifications, i.e. 311. Identifying this number or a small bracketing interval around it is the primary objective of the interval search. The initial bracketing interval is set to $[L, U] = [192, 1365]$, where 192 is the number of patterns covered with 5 fortifications by the greedy solution, and 1,365 is the total number of patterns. The piecewise linear function is approximated by the line joining the points (192, 5) and (1365, 12), where 12 is an upper bound on the number of fortifications needed to disrupt the 1,365 interdiction patterns. The two points are depicted with black solid triangles in the figure. Since $q_{192} = 5 = q$, the new trial value k is chosen so that the point $(k, q + 1)$ lies on this line. The resulting value is $k = 358$. The corresponding point (358, 6) is represented as a non-solid triangle in the figure. Note that had we chosen to interpolate the new value using q, we would have obtained the previous lower estimate 192. At the next iteration, problem SC(358) is solved to compute the actual number of fortifications needed to disrupt the first 358 patterns. Since the optimal number is, in fact, 6, which is greater than 5, the upper bound is set to 358 and the interval is reduced to [192, 358]. At this point, no further interpolation is possible since the evaluations at q and $q + 1$ would only return the interval extremes. Hence, we resort to bisection, with the new trial value fixed to 274, i.e. the midpoint between 192 and 358. The procedure is repeated in this same manner until a given termination criterion is met.

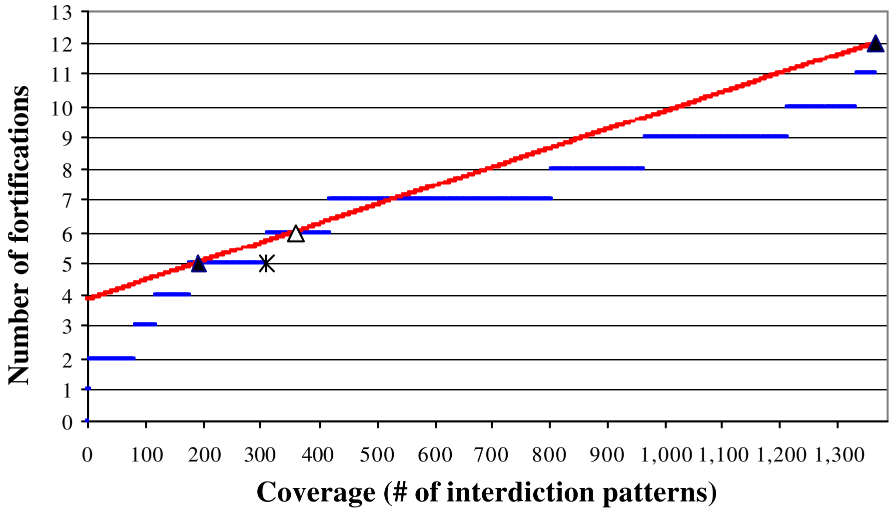

Fig. 1. Interval search with linear interpolation.

Note that in this example, the search interval was already reduced by more than 70% after the first iteration. We found that this type of rapid reduction can be attributed to the use of the interpolation process. In preliminary experiments, we employed a straightforward binary search to narrow the search interval. As a binary search does not take advantage of the fact that function is monotonically increasing, it is not surprising that a binary search performed poorly in comparison to the less myopic linear interpolation routine.

With respect to the stopping criterion, we considered the simultaneous use of three different rules for terminating the interval search. They are based upon the following conditions:

i) *MaxIter* iterations have been performed

ii) $U - L < \beta$.

iii) $(WD_{L+1} - WD_U)/WD_U \leq \gamma$ (i.e. the relative optimality gap falls below a given threshold)

According to rule (i), the search is stopped after a predefined number of iterations, *MaxIter*. This stopping rule limits a priori how many set-covering problems will be solved, and hence provides control on the computational effort spent during the interval search phase. According to rule (ii), steps (2)-(4) are repeated until the interval width is reduced to a desired value β. Since the MCPC reductions described in the next section depend on the final values of *L* and *U,* rule (ii) provides control on the size of the reduced mixed-integer model. Finally, rule (iii) terminates the search if the relative optimality gap falls below a given value, γ, and hence it guarantees a given level of solution accuracy. Basically, rule (ii) is more suitable if the line search is used for reducing the size of the maximal covering model, where the principal objective is to identify an optimal solution to a MCPC. In contrast, rule (iii) is more suitable if the interval search is used as a stand-alone improvement method, where the principal objective is to find a good approximate solution to the MCPC. With both rules, it is difficult to anticipate how many set-covering models need to be solved. As a consequence, it is difficult to fine-tune the parameters β and γ in such a way that there is a clear benefit in continuing the interval search process rather than terminating the search and proceeding to the solution of the resulting reduced version of MCPC. Consequently, Rule (i) is the safest way of monitoring the computing time spent in the search phase. The computational experience reported in Section 7 is based upon the combined use of rules (i) through (iii), where *MaxIter* = 7, $\gamma = 2.5\%$, $\beta = 50$. Basically, the search phase is terminated as soon as one of the three conditions is met, but only after it has been repeated a minimum of 5 times. Preliminary results aimed at fine-tuning the stopping rules demonstrated that this particular choice represents a good trade-off between the computational effort required by the search and the extent of the reduction of the final model.

The set covering problem defined at step 2 of the procedure is itself a difficult problem to solve (Garey and Johnson [9]). However, the size of each SC(k) is usually small, having p variables and k constraints, as compared to the size of the maximal covering problem (9)-(14). Furthermore, the set covering problems are not solved from scratch at each iteration. Rather, each problem SC(k) is generated from the problem solved at the previous iteration by either adding or deleting constraints, depending on whether k is increased or decreased. The optimal solution to the previous problem can then be used as a starting solution for the new problem to save computing time. As noted in Caprara, Fischetti and Toth [3], general-purpose linear programming solvers based on branch and bound are very competitive approaches for solving the set covering problem to optimality, and usually outperform the best exact algorithms presented in the literature. In our empirical investigations, we used the branch and bound based solver CPLEX to solve the set covering models.

6. REDUCED MAXIMAL COVERING MODEL WITH PRECEDENCE CONSTRAINTS

The bounds U and L returned by the linear interpolation search procedure can be used to reduce the number of variables and constraints in the original model (9)-(14). More specifically, the first L continuous variables can be fixed to 1, as the optimal solution must cover at least the first L ordered interdiction patterns. The interval search process has also identified that it is not possible to thwart the last $|H|$-U ordered interdiction patterns. Thus, the variables indexed from $U + 1$ to $|H|$ associated with the last $|H|$-U ordered interdiction patterns can be eliminated together with the coverage constraints (10) and the precedence constraints (14) associated with these variables. The resulting reduced model RMCPC is:

$$\max \sum_{h=L+1}^{U} y_h \tag{16}$$

$$\text{s.t.} \quad \sum_{j \in I_h} z_j \geq 1 \quad \text{for all } h = 1...L \tag{17}$$

$$\sum_{j \in I_h} z_j \geq y_h \quad \text{for all } h = L+1,...,U \tag{18}$$

$$y_h \geq y_{h+1} \quad \text{for all } h = L+1,...,U-1 \tag{19}$$

$$\sum_{j \in F} z_j = q \tag{20}$$

$$z_j \in \{0,1\} \quad \text{for all } j \in F \tag{21}$$

$$0 \leq y_h \leq 1 \quad \text{for all } h = L+1..U \tag{22}$$

The resulting model has $p + U - L$ variables and $2U - L - 1$ constraints. Typically, the problems were reduced to 1-15% of their original size, after solving between 5 and 7 set covering problems. The extent of the reduction is further discussed in the next section.

Note that the proposed solution approach can be easily extended to the more general problem where each facility has a different protection cost and the constraint fixing the number of possible fortifications is replaced by a budget constraint. In such a case, the greedy procedure would need to be modified so that each newly selected facility at step 2 is the one which covers the greatest number of uncovered patterns per unit cost. The interval search procedure requires two modifications: 1) the minimum-cardinality set-covering problems solved in the evaluation step are replaced by minimum-cost set-covering problems; 2) the linear interpolation is performed by using fortification costs and a budget instead of number of fortifications and q in formula (15). Finally, constraint (20) in the reduced formulation is simply replaced by a budget constraint.

7. COMPUTATIONAL STUDY

In this section we report on the computational results obtained by applying the modeling approach described in the previous sections to two different geographical data sets: the 150 node London, Ontario data set (Goodchild and Noronha [12]) and the 316 node Alberta data set [1]. Each problem was solved for different combinations of the parameters p, q, and r. The code was programmed in C++ and the tests were run on a PC with a Pentium 4, 1.8Ghtz processor with 512 MB of RAM. The set covering problems in the search phase and the reduced MCPC were solved using CPLEX 7.0 with user-specified parameter settings fine-tuned to improve performance. For instance, we forced CPLEX to use the primal optimizer at the initial root node and the dual simplex method after branching. Also, we set the tree search strategy of CPLEX to emphasize feasibility rather than optimality. Several experiments with different parameter settings proved that this combination yielded overall the best results.

The first set of experiments was conducted on the London data set to test the computational performance of the new MCPC model as compared to the original formulation IMF (Church and Scaparra [6]). The problem was solved for $p = 20$ operating facilities, for q ranging between 6 and 10, and r ranging between 2 and 4. Thus, we ran the experiments on the 15 largest problem instances solved by Church and Scaparra [6]. In the initial configuration, the 20 existing facilities were located by solving the unrestricted p-median problem, so that the base solution was the optimal median solution. Table 1 presents the results for the original IMF model, the results for the new MCPC, and for the reduced version of the MCPC (RMCPC). Each model was solved by using CPLEX. The RMCPC model was constructed after performing both the greedy heuristic and interval search procedures. The first three columns in Table 1 list the parameters values for each problem (p, q, and r). The fourth column gives the total number of possible interdiction patterns (i.e. |H|), followed by the optimal objective value of the interdiction median problem with fortification. Finally, the last three columns report the computational times required for solving the three models. Note that the times reported for the MCPC include the times to generate and sort all possible interdictions, in addition to the CPLEX execution time. The time in solving the RMCPC includes the time to generate and sort all possible interdiction patterns, the computation time for the greedy heuristic and the interval search, and the time for solving the reduced model with CPLEX. When we compare the computational performance of the three model options, we observe that the solution of MCPC does not yield significant gains with respect to IMF, although some improvements can be noticed for larger values of r and q. On the other hand, the reduced model yields significant computational savings for every problem solved. On average, solving RMCPC was about 190 times faster than solving MCPC and 240 times faster than solving IMF. Even more important than the average gain in solution efficiency is that the most significant time gains were obtained for the largest value of the parameter r ($r = 4$), which was identified in Church and Scaparra [6] as the main limitation of the applicability of IMF.

Given the effectiveness of the proposed methodology in solving this first set of problems, we extended our empirical investigation for the RMCPC modeling approach to problem instances with larger parameter values. In particular, we considered problem instances with 25 and 30 facilities in the existing configuration, and allowed up to 7 interdictions and 7 fortifications. Globally, we solved 24 instances for the two different spatial data sets. The detailed results for the London data set are presented in Table 2, while Table 3 provides the results for the Alberta data set. Tables 2 and 3 have the same structure, with the first 4 columns having the same meaning as in Table 1. The next four columns indicate respectively: the number of most disruptive patterns covered by the greedy heuristic; the lower bound (L) and upper bound (U) to the optimal coverage produced by the improvement procedure; and the number of most disruptive patterns covered in the optimal solution to the RMCPC. The next column gives the optimal objective values. The following four columns provide some statistical information about the performance of the greedy procedure and the interval search procedure. Namely, they report the percentage of the optimal number of patterns that were covered at the end of each of the two phases (columns %Cov). They also give the percentage error of the greedy heuristic objective values compared to the optimal values (column %Err), and the percent relative optimality gap obtained at the end of the search phase (column %GAP). Finally, the last 5 columns provide the running times spent in each phase of this new solution approach. The last column presents the total computation time required for all steps of the process.

The results presented in Table 2 present a clear picture concerning the performance of the individual phases of the algorithm. First, we notice that 4 of the 24 problems were already solved to optimality by the greedy heuristic. On average, the greedy error was 2.98% for all test problems. The worst performance was obtained for the last problem, where the greedy covered only around 3% of the patterns covered in the optimal solution, resulting in a 10.26% error. Note that, in general, the percentage error for the objective function is quite small, even when the percentage of covered patterns is relatively low. For example, the greedy heuristic covered less than 40% of the optimal number of patterns for the problem with p = 30, q = 3, and r = 4. Nevertheless, the error in terms of the WD values associated with these interdiction patterns was less than 2%. This is mainly due to the fact that even if the greedy heuristic fails in covering a large number of interdiction patterns, it does cover the most disruptive ones. Since the marginal gain in the median objective due to the coverage of additional patterns is likely to diminish as the number of covered patterns increases, the failure in covering the “tail” patterns has a limited impact on the overall objective. Also, note that the greedy percentage error is achieved with very little computational effort, less than a second for every problem. By comparison, much more computational time is needed for sorting the interdiction patterns (as was expected due to the complexity of the sorting process).

After execution of the search procedure, the coverage of ordered interdiction patterns associated with the resulting lower bound L averaged 91% of what can be optimally covered. Even though no optimal solutions were detected during this search phase, the optimality gap measured at the end of the procedure was always below 2% (0.67% on average) and this was always achieved within the preset *MaxIter=7* limit. The time spent in the search phase was, on average, around one third of the total time, while the remaining two thirds of the time were spent on solving the reduced model to optimality. The sorting and greedy times were negligible with respect to the total computing time.

The computation experience generated on the Alberta data set provides somewhat different insights. This data set contains 316 demand points whereas the London data set contained 150 nodes. This fact alone adds to the complexity of the overall problem. The results for the Alberta data set are summarized in Table 3. Notice that 3 of the 24 instances could not be solved to optimality; the objective function values for these problems are preceded by a question mark in the table. However, for all other cases but one, the optimal solution was found by the greedy heuristic. This behavior can be explained by understanding that some of the facilities in the existing p-median configurations are noticeably more important than others in providing efficient service. In other words, the distribution of the demands is very heterogeneous, with a few key facilities supplying significantly more demand than others. Clearly, these key facilities must be fortified to guarantee a high post-interdiction service level, which makes them generally easy to detect using a greedy scheme. Given that the ordered set of interdiction patterns would first involve all interdiction patterns containing these key sites, the number of most disruptive interdiction patterns that can be thwarted is generally very high, when these same key sites are fortified. Consider for instance the last problem in Table 3. At the end of the greedy phase, almost three fourths of the total number of interdiction patterns (almost 1,500,000 patterns) could be covered with the 7 available fortifications, whereas in the London data set the same number of fortifications could thwart only 6,749 patterns at the end of the greedy phase and only 219,411 patterns at optimality. Clearly, as the number of preventable worst-case patterns increases, so does the complexity of the minimum-cardinality set covering problems solved during the improvement phase, since the number of constraints of each SC problem depends upon the upper and lower bounds of the optimal number of patterns which can be thwarted. As a consequence of number of constraints, the improvement phase could not be executed for the three problems with $p = 30$ and $r = 7$, i.e. for the problems with more that two million interdiction patterns.

Table 4 and Table 5 report some information related to the extent of the model reduction obtained at the end of the improvement phase for the London data set and the Alberta data set respectively. For each parameter combination, Table s 4 and 5 provides: the number of variables and constraints of the initial maximal covering problem with precedence constraints (columns 4 and 5); the number of variables and constraints of the reduced RMCPC model (columns 6 and 7); the percentage reduction in model size (columns 8 and 9); and the size of the largest set SC problem solved (column 10). For the London data set, the number of variables was reduced on average by 99.25% and the number of constraints by 95.85%. The computing time to obtain these significant reductions was quite small considering the size of the problems being solved. In the worst case, it took slightly more than half an hour to reduce the problems with more than two million variables and four million constraints to a manageable size. The majority of the instances could not have been solved by CPLEX given their initial size. For the Alberta data set, the variable reduction amounted to 99.63% of the original size and the constraint reduction averaged 78.96%.

8. CONCLUSIONS AND FUTURE RESEARCH

This paper has presented a new modeling approach for determining optimal protection/fortification plans involving a supply system based upon the p-median protocol. The overall objective of this approach is to minimize the loss of system service efficiency caused by worst-case interdiction strikes. This problem was recently introduced by Church and Scaparra [6]. They developed a model of this problem called the interdiction p-median problem with fortification (IMF) and presented results associated with solving that model using general purpose integer-linear programming software. In this paper we have shown that the original IMF model can be reformulated as a maximal covering model with precedence constraints (MCPC) based upon an order set of all possible interdiction patterns. A methodology for solving this model has also been presented that involves a process that can be used to reduce the resulting MCPC model called RMCPC. We have shown that the maximal covering with precedence constraints formulation for the IMF problem can be efficiently solved by CPLEX after implementing a specialized model reduction approach. The proposed modeling approach represents a significant enhancement over direct solutions of the original IMF formulation presented in Church and Scaparra [6] as this new approach significantly reduces the time needed to solve problems as compared to using the original IMF formulation. Additionally the new modeling approach makes it possible to solve larger problem instances than what was possible using the IMF model. We also demonstrated that when demand is concentrated and very heterogeneous, a greedy approach to the MCPC tends to perform well, in terms of final objective value. Overall, fortification strategies could be determined to thwart optimally up to two million possible interdiction responses.

In view of recent world events there is a heightened awareness and concern for protecting infrastructure, supply systems and the public from terrorist attack. Protecting infrastructure is an enormous challenge, since resources are limited and potential risks are high. The model presented in this paper is intended to provide an example of how fortification resources can be optimized to provide effective protection of vital industry assets. Obviously, the scale and complexity of real distribution, supply and emergency response systems presents a significant challenge in planning against possible acts of terrorism. Since it is impossible to protect all assets, it is important to devise approaches for identifying critical elements, optimize the protection of key system features, plan for emergency response, and schedule and plan repair efforts. The models addressed here represent one of the simpler system protocols. Model extensions should be developed that incorporate probabilistic elements, different service/supply protocols, and capacity restrictions on facility supplies. Also, the modeling effort should be accompanied by the development of new specialized solution approaches able to handle large-scale problems. We hope that the framework presented in this paper can serve as a useful preliminary step in that direction and can inspire future work in modeling and solving increasingly complex interdiction and fortification problems.

ACKNOWLEDGMENTS

We wish to acknowledge the support of the Research and Special Programs Administration of the U.S.D.O.T. Also, we wish to thank our colleagues within the National Consortium on Remote Sensing in Transportation, Infrastructure for their helpful comments.

REFERENCES

[1] O. Alp, Z. Drezner, and E. Erkut, An Efficient Genetic Algorithm for the *p*-Median Problem, Annals of Operations Research 122 (2003), 21-42.

[2] Bard, J. F., *Practical Bilevel Optimization: Algorithms and Applications*. Kluwer Academic Publishers, The Netherlands, 1998.

[3] A. Caprara, M. Fischetti, and P. Toth, Algorithms for the set covering problem, Annals of Operations Research 98 (2000), 353-371.

[4] R.L. Church, COBRA: A new formulation of the classic *p*-median location problem, Annals of Operations Research 122 (2003), 103-120.

[5] R.L. Church, M.P. Scaparra, and R.S., Middleton, Identifying critical infrastructure: The median and covering facility interdiction problems, Annals of the Association of American Geographers 94 (2004), 491-502.

[6] R.L. Church, and M.P., Scaparra, Protecting Critical Assets: The *r*-interdiction median problem with fortification, Working paper (2004).

[7] S. Dempe, *Foundations of Bilevel Programming*, Kluwer Academie Publishers, Dordrecht, 2002.

[8] D.R. Fulkerson, and G. C. Harding, Maximizing the minimum source-sink path subject to a budget constraint, Mathematical Programming 13 (1977), 116-118.

[9] M.R. Garey, and D.S. Johnson, Computers and intractability: A guide to the theory of NP-completeness, W.H. Freeman, San Francisco, 1979.

[10] P. M. Ghare, D. C. Montgomery, and W. C. Turner, Optimal interdiction policy for a flow network, Naval Research Logistics Quarterly 18 (1971), 37-45.

[11] B. Golden, A problem of network interdiction, Naval Research Logistics Quarterly 25 (1978), 711-713.

[12] M. Goodchild, and V. Noronha, Location-allocation for small computers, Monograph No. 8., Department of Geography, The University of Iowa, Iowa City, 1983.

[13] Hakimi, S.L. Optimum location of switching centers and the absolute centers and medians of a graph, Operations Research 12 (1964), 450-459.

[14] Hakimi, S.L. Optimum distribution of switching centers and some graph related theoretic problems, Operations Research 13 (1965), 462-475.

[15] P. Hansen, B. Jaumard, and G. Savard. New branch and bound rules for linear bilevel programming, SIAM Journal on Scientific and Statistical Computing 13 (1992) 1194-1217.

[16] E. Israeli, and R. K. Wood, Shortest-path network interdiction, Networks 40 (2002), 97-111.

[17] A. W. McMasters, and T. M. Mustin, Optimal interdiction of a supply network, Naval Research Logistics Quarterly 17 (1970), 261-268.

[18] Moore, J.T., J.F. Bard. The mixed integer linear bilevel programming problem. *Operations Research* 38 (1990), 911-921.

[19] J. Salmeron, K. Wood, and R. Baldick, Analysis of electric grid security under terrorist threat, IEEE Transactions on Power Systems 19 (2004), 905-912.

[20] M.P. Scaparra, and R.L. Church, A bilevel mixed integer program for critical infrastructure protection planning, Working paper (2004).

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

[22] R. K. Wood, Deterministic network interdiction, Mathematical and. Computer Modelling 17 (1993) 1-18.

p	*q*	*r*	*|H|*	**Opt. Val.**	**Time (sec.)** IMF	MCPP	RMCPC
20	6	2	190	143,120.20	0.42	0.36	0.02
20	6	3	1140	153,263.89	8.08	9.47	0.08
20	6	4	4845	164,855.27	246.06	396.50	0.53
20	7	2	190	141,722.36	0.41	0.39	0.01
20	7	3	1140	151,219.26	9.94	9.11	0.14
20	7	4	4845	163,083.32	380.94	408.06	1.16
20	8	2	190	140,661.26	0.42	0.39	0.02
20	8	3	1140	150,890.24	10.09	12.53	0.20
20	8	4	4845	161,448.25	611.08	405.70	1.48
20	9	2	190	139,808.49	0.44	0.33	0.02
20	9	3	1140	150,396.96	14.95	19.25	0.31
20	9	4	4845	160,219.78	898.52	496.72	2.55
20	10	2	190	139,075.35	0.47	0.31	0.02
20	10	3	1140	149,219.04	15.86	10.45	0.38
20	10	4	4845	156,646.17	714.61	514.11	5.09
Avg.					194.15	152.25	0.80

Table. 1. Comparison between the IMF, MCPP and RMCPC models.

				Coverage				Obj	Greedy		Search		Time (sec.)				
p	**q**	**r**	**|H|**	**Greedy**	**L**	**U**	**OPT**	**Val**	**% Cov**	**% Err**	**% Cov**	**% GAP**	**Sort**	**Greedy**	**Search**	**Cplex**	**Total**
25	3	4	12,650	64	64	146	64	153,638.54	100.00	0.00	100.00	1.68	0.09	0.00	0.06	0.01	0.16
25	3	5	53,130	49	233	417	243	164,458.35	20.16	3.29	95.88	1.17	0.94	0.00	0.41	0.03	1.39
25	3	6	177,100	70	720	1,371	915	174,942.60	7.65	5.57	78.69	1.74	1.75	0.03	3.89	0.45	6.92
25	3	7	480,700	285	1,223	2,161	1,798	188,282.97	15.85	3.92	68.02	1.37	9.33	0.09	29.75	0.53	44.72
25	5	4	12,650	113	983	1,157	991	143,058.40	11.40	6.20	99.19	0.50	0.09	0.00	0.14	0.08	0.31
25	5	5	53,130	1,928	4,562	4,939	4,810	151,559.19	40.08	3.04	94.84	0.27	0.94	0.02	2.06	1.16	4.18
25	5	6	177,100	12,133	12,133	13,507	12,133	162,485.24	100.00	0.00	100.00	0.38	1.75	0.03	36.42	5.34	43.54
25	5	7	480,700	35,812	35,812	39,784	35,812	171,987.27	100.00	0.00	100.00	0.41	9.33	0.13	265.53	131.44	406.43
25	7	4	12,650	2,049	3,042	3,207	3,164	137,307.81	64.76	1.74	96.14	0.24	0.09	0.00	1.98	0.44	2.51
25	7	5	53,130	2,766	9,061	9,847	9,358	147,589.13	29.56	4.58	96.83	0.35	0.94	0.00	16.19	4.36	21.49
25	7	6	177,100	30,133	30,133	31,546	30,633	156,685.61	98.37	0.07	98.37	0.21	1.75	0.09	186.94	43.11	231.89
25	7	7	480,700	78,208	99,170	103,363	99,823	164,595.84	78.35	1.16	99.35	0.21	9.33	0.28	806.06	274.50	1,090.17
30	3	4	27,405	183	466	608	478	121,378.81	38.28	1.99	97.49	0.54	0.22	0.00	0.06	0.03	0.31
30	3	5	142,506	787	787	1,172	787	132,032.99	100.00	0.00	100.00	0.80	2.80	0.01	1.74	0.08	4.63
30	3	6	593,775	967	1,809	2,651	2,556	140,618.50	37.83	2.52	70.77	0.91	7.11	0.11	16.19	0.27	23.68
30	3	7	2,035,800	285	3,369	6,453	3,933	152,969.85	7.25	5.57	85.66	1.73	47.48	0.91	131.53	1.42	181.34
30	5	4	27,405	950	1,525	1,668	1,580	118,060.47	60.13	1.44	96.52	0.30	0.22	0.00	0.16	0.08	0.46
30	5	5	142,506	1,118	2,189	2,724	2,562	128,667.35	43.64	1.90	85.44	0.55	2.80	0.01	2.25	0.84	5.90
30	5	6	593,775	2,560	7,252	9,598	8,217	137,061.54	31.15	2.59	88.26	0.61	7.11	0.11	29.61	9.63	46.46
30	5	7	2,035,800	3,933	14,683	20,058	19,627	146,299.89	20.04	4.56	74.81	0.84	47.48	0.88	2,129.28	83.66	2,261.30
30	7	4	27,405	1,702	3,538	3,844	3,824	114,789.52	44.51	2.60	92.52	0.25	0.22	0.00	1.20	0.38	1.80
30	7	5	142,506	6,599	16,791	18,490	18,160	121,953.59	36.34	3.06	92.46	0.36	2.80	0.03	32.22	27.53	62.58
30	7	6	593,775	6,386	45,545	50,440	48,873	130,678.87	13.07	5.45	93.19	0.32	7.11	0.11	261.05	132.64	400.91
30	7	7	2,035,800	6,749	206,482	235,016	219,411	136,730.79	3.08	10.26	94.11	0.44	47.48	0.89	2,011.47	11,837.59	13,897.43
Avg.									45.90	2.98	91.61	0.67	8.72	0.16	248.59	523.15	780.85

Table. 2. Results from the application of the RMCPC formulation to the London data set.

				Coverage				Obj	Greedy		Search		Time (sec.)				
p	**q**	**r**	**|H|**	**Greedy**	**L**	**U**	**OPT**	**Val**	**% Cov**	**% Err**	**% Cov**	**% GAP**	**Sort**	**Greedy**	**Search**	**Cplex**	**Total**
25	3	4	12,650	2,128	2,128	2,162	2,128	57,178,804	100.00	0.00	100.00	2.42	0.17	0.00	0.39	0.08	0.64
25	3	5	53,130	11,118	11,118	11,263	11,118	62,785,799	100.00	0.00	100.00	1.69	1.67	0.01	3.91	0.52	1.39
25	3	6	177,100	45,152	45,152	45,394	45,152	67,647,091	100.00	0.00	100.00	1.40	3.41	0.05	41.25	3.48	6.92
25	3	7	480,700	144,331	144,331	145,644	144,331	73,255,693	100.00	0.00	100.00	1.83	15.91	0.25	482.92	15.38	44.72
25	5	4	12,650	3,668	3,668	3,734	3,668	43,265,666	100.00	0.00	100.00	0.71	0.17	0.01	0.95	0.30	1.43
25	5	5	53,130	19,905	19,905	20,164	19,905	46,077,394	100.00	0.00	100.00	0.63	1.67	0.01	22.16	3.91	27.75
25	5	6	177,100	78,382	78,382	79,204	78,382	49,452,379	100.00	0.00	100.00	0.41	3.41	0.09	222.89	31.09	257.48
25	5	7	480,700	246,411	246,411	248,502	246,411	52,315,922	100.00	0.00	100.00	0.29	15.91	0.36	1,188.22	421.23	1,625.72
25	7	4	12,650	8,538	8,538	8,572	8,538	31,746,528	100.00	0.00	100.00	0.39	0.17	0.00	5.88	2.03	8.08
25	7	5	53,130	38,969	38,969	39,095	38,969	34,993,344	100.00	0.00	100.00	0.24	1.67	0.03	89.77	16.70	108.17
25	7	6	177,100	138,437	138,437	138,808	138,437	37,921,475	100.00	0.00	100.00	0.21	3.41	0.14	634.30	306.86	944.71
25	7	7	480,700	397,386	397,386	398,253	397,386	40,381,720	100.00	0.00	100.00	0.15	15.91	0.58	5,643.98	2,080.50	7,740.97
30	3	4	27,405	3,752	3,875	3,936	3,906	46,518,667	96.06	7.43	99.21	2.08	0.38	0.00	0.73	0.47	1.58
30	3	5	142,506	24,977	24,977	25,296	24,977	52,028,642	100.00	0.00	100.00	2.50	4.70	0.05	5.69	1.28	11.72
30	3	6	593,775	123,055	123,055	124,392	123,055	56,889,934	100.00	0.00	100.00	1.59	9.53	0.23	41.91	10.50	62.17
30	3	7	2,035,800	491,537	--	--	--	?60,993,485	--	--	--	--	39.44	1.66	--	--	--
30	5	4	27,405	6,174	6,174	6,294	6,174	36,974,712	100.00	0.00	100.00	0.58	0.38	0.02	1.55	0.55	2.50
30	5	5	142,506	39,122	39,122	39,737	39,122	40,150,939	100.00	0.00	100.00	0.71	4.70	0.06	28.72	8.83	42.31
30	5	6	593,775	201,350	201,350	203,802	201,350	42,611,184	100.00	0.00	100.00	0.63	9.53	0.31	908.19	452.75	1,370.78
30	5	7	2,035,800	784,973	--	--	--	?45,728,325	--	--	--	--	39.44	2.69	--	--	--
30	7	4	27,405	14,186	14,186	14,268	14,186	27,604,698	100.00	0.00	100.00	0.56	0.38	0.00	11.78	1.38	13.54
30	7	5	142,506	85,428	85,428	85,803	85,428	30,064,943	100.00	0.00	100.00	0.45	4.70	0.11	267.66	134.70	407.17
30	7	6	593,775	390,126	390,126	391,540	390,126	32,569,976	100.00	0.00	100.00	0.24	9.53	0.55	7,073.27	1,092.89	8,176.24
30	7	7	2,035,800	1,433,807	--	--	--	?34,767,116	--	--	--	--	39.44	4.44	--	--	--
Avg.									99.81	0.35	99.96	0.94	9.40	0.49	794.10	218.35	993.14

Table. 3. Results from the application of the RMCPP formulation to the Alberta data set.

			Before Reduction		After Reduction		% Reduction		Largest
p	**q**	**r**	**Vars**	**Constrs**	**Vars**	**Constrs**	**Vars**	**Constrs**	**SC solved**
25	3	4	12,675	25,300	107	227	99.16	99.10	25 × 725
25	3	5	53,155	106,260	209	600	99.61	99.44	25 × 2,996
25	3	6	177,125	354,200	676	2,021	99.62	99.43	25 × 10,482
25	3	7	480,725	961,400	963	3,098	99.80	99.68	25 × 30,309
25	5	4	12,675	25,300	199	1,330	98.43	94.74	25 × 1,504
25	5	5	53,155	106,260	402	5,315	99.24	95.00	25 × 4,938
25	5	6	177,125	354,200	1,399	14,880	99.21	95.80	25 × 23,129
25	5	7	480,725	961,400	3,997	43,755	99.17	95.45	25 × 67,588
25	7	4	12,675	25,300	190	3,371	98.50	86.68	25 × 3,372
25	7	5	53,155	106,260	811	10,632	98.47	89.99	25 × 12,207
25	7	6	177,125	354,200	1,438	32,958	99.19	90.70	25 × 41,437
25	7	7	480,725	961,400	4,218	107,555	99.12	88.81	25 × 111,748
30	3	4	27,435	54,810	172	749	99.37	98.63	30 × 1,316
30	3	5	142,536	285,012	415	1,556	99.71	99.45	30 × 6,947
30	3	6	593,805	1,187,550	872	3,492	99.85	99.71	30 × 27,911
30	3	7	2,035,830	4,071,600	3,114	9,536	99.85	99.77	30 × 92,807
30	5	4	27,435	54,810	173	1,810	99.37	96.70	30 × 2,099
30	5	5	142,536	285,012	565	3,258	99.60	98.86	30 × 7,543
30	5	6	593,805	1,187,550	2,376	11,943	99.60	98.99	30 × 30,712
30	5	7	2,035,830	4,071,600	5,405	25,432	99.73	99.38	30 × 197,443
30	7	4	27,435	54,810	336	4,149	98.78	92.43	30 × 4,148
30	7	5	142,536	285,012	1,729	20,188	98.79	92.92	30 × 20,188
30	7	6	593,805	1,187,550	4,925	55,334	99.17	95.34	30 × 65,123
30	7	7	2,035,830	4,071,600	28,564	263,549	98.60	93.53	30 × 311,105
Avg.							99.25	95.85	

Table. 4. Extent of the model reduction for RMCPC model applied to the London data set.

			Before Reduction		After Reduction		% Reduction		Largest
p	**q**	**r**	**Vars**	**Constrs**	**Vars**	**Constrs**	**Vars**	**Constrs**	**SC solved**
25	3	4	12,675	25,300	59	2,195	99.53	91.32	25 × 2,680
25	3	5	53,155	106,260	170	11,407	99.68	89.27	25 × 13,451
25	3	6	177,125	354,200	267	45,635	99.85	87.12	25 × 52,912
25	3	7	480,725	961,400	1,338	146,956	99.72	84.71	25 × 165,353
25	5	4	12,675	25,300	91	3,799	99.28	84.98	25 × 4,195
25	5	5	53,155	106,260	284	20,422	99.47	80.78	25 × 21,980
25	5	6	177,125	354,200	847	80,025	99.52	77.41	25 × 84,962
25	5	7	480,725	961,400	2,116	250,592	99.56	73.93	25 × 263,144
25	7	4	12,675	25,300	59	8,605	99.53	65.99	25 × 8,811
25	7	5	53,155	106,260	151	39,220	99.72	63.09	25 × 39,979
25	7	6	177,125	354,200	396	139,178	99.78	60.71	25 × 141,410
25	7	7	480,725	961,400	892	399,119	99.81	58.49	25 × 404,327
30	3	4	27,435	54,810	91	3,996	99.67	92.71	30 × 4,736
30	3	5	142,536	285,012	349	25,614	99.76	91.01	30 × 30,085
30	3	6	593,805	1,187,550	1,367	125,728	99.77	89.41	30 × 144,450
30	3	7	--	--	--	--	--	--	--
30	5	4	27,435	54,810	150	6,413	99.45	88.30	30 × 7,138
30	5	5	142,536	285,012	645	40,351	99.55	85.84	30 × 44,044
30	5	6	593,805	1,187,550	2,482	206,253	99.58	82.63	30 × 220,970
30	5	7	--	--	--	--	--	--	--
30	7	4	27,435	54,810	112	14,349	99.59	73.82	30 × 14,845
30	7	5	142,536	285,012	405	86,177	99.72	69.76	30 × 88,431
30	7	6	593,805	1,187,550	1,444	392,953	99.76	66.91	30 × 401,438
30	7	7	--	--	--	--	--	--	--
Avg.							99.63	78.96	

Table. 5. Extent of the model reduction for RMCPC model applied to the Alberta data set.

UNIVERSITY OF KENT