A Multi-objective Integrated Facility Location-Hardening Model: Analyzing the Pre- and Post-Disruption Tradeoff

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

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

Abstract

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

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

1 Introduction

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

*Corresponding author. Tel: 662-325-3923.
Email addresses: hugh.medal@msstate.edu (Medal), epohl@uark.edu (Pohl), and rossetti@uark.edu (Rossetti).

Distribution networks, such as power networks and supply chains, are ubiquitous throughout the world. These networks consist of a set of facilities (power sub-stations, ports, distribution centers, etc.) and set of customers that rely on the facilities. Because these facilities form the backbone of distribution networks, facility disruptions often result in severe consequences. One recent example is the 2011 Tohoku earthquake in Japan, which disrupted manufacturing facilities and caused several Japanese automakers to halt car production for up to six months [16]. These severe consequences have forced decision-makers to consider the possibility of facility disruptions when they design their network of facilities. In addition, decision-makers also may choose to harden facilities to protect them from disruptions. Facility hardening, a special case of facility protection, involves allocating resources to a facility (e.g., additional security, retrofitting, etc.) to make it immune to failure [30, 31, 33]. This paper serves to help decision-makers make better facility location and hardening decisions by providing a mathematical model of these decisions and using this model to generate insights about these decisions.

This research focuses on the maximum distance objective, which is also used in the classic p-center problem [12]. Since this objective is concerned with minimizing the worst service experienced by a demand point, it is appropriate for the public sector [9]. Researchers have cited numerous potential applications for the maximum distance objective such as locating emergency vehicles and facilities [13, 23, 32] and locating warning sirens [35]. This article also seeks to minimize the worst-case disruption, i.e., the worst-case risk measure. The worst-case risk measure is well-suited for risk-averse decision-makers, especially in critical infrastructure protection [5, 31, 29].

A growing amount of research exists on locating facilities subject to disruptions and hardening facilities subject to disruptions. Several authors [7, 10, 34, 27] have developed models for locating facilities subject to random disruptions. Other works have considered that located facilities are subject to interdiction, i.e., intentional and calculated attacks. O'Hanley and Church [25] developed a bi-level model for the problem of locating facilities to minimize the post-interdiction total weighted covered demand. Drezner [10] developed a method to minimize the post-interdiction maximum distance from a demand point to its closest located and operating facility. O'Hanley and Church [25] optimized a weighted combination of the system performance before and after interdiction for a facility location problem.

Rather than locating facilities, others have examined the problem of hardening a set of existing facilities. O'Hanley et al. [26] and Li et al. [17] have presented models for hardening facilities subject to random failures. Church and Scaparra [5, 30, 31] have studied the problem of how to harden facilities in order to minimize the post-interdiction total weighted distance. O'Hanley et al. [24] presented a bi-level model to optimally harden facilities in order to minimize the post-interdiction total weighted covered demand.

Some researchers have developed models that include both facility location and facility hardening. Snyder and Daskin [34] present extensions of the p-median and warehouse location models and include perfectly reliable, i.e., hardened, and unreliable facility locations in their model. Specifically, a facility is perfectly reliable if and only if it is located at a perfectly reliable location. Although their study focuses on location, it would be possible to integrate location and hardening decisions in their model if every geographical site had both a reliable and an unreliable location. However, it is unclear whether using their model in this way, which would double its size, would be computationally tractable. Lim et al. [18] were the first to explicitly include both location and hardening decisions in a single model. They present an extension of the warehouse location problem in which the decision maker chooses between locating unreliable facilities and locating perfectly reliable, i.e., hardened, backup facilities at a higher cost. The authors assume one layer of supplier backup. Thus, if a demand point's primary facility fails, the demand point is then immediately assigned to its hardened backup without checking if there is a closer operating facility. This assumption simplifies the model and allows the authors to provide several useful analytical results. Li et al. [17] extend the work of Lim et al. [18] but still assume one layer of supplier backup. The research presented in this paper considers multiple layers of backup, allowing a demand point to be assigned to its closest operating facility after a disruption. Aksen et al. [1] study an extension of the p-median problem in which facilities are susceptible to interdiction. They present a tri-level version of the budget-constrained median location model in which a defender locates and hardens facilities and then an attacker destroys a number of unhardened facilities. Their works extends the work of Lim et al. [18] by modeling multiple layers of backup. Aksen et al. [1] study several methods for solving their problem including a tabu search algorithm and a two-phase heuristic. The research in this paper builds on the work of Aksen et al. [1] by providing an exact procedure for solving the integrated location-hardening problem, rather than a heuristic procedure.

This article builds upon the facility location and facility hardening literatures by making the following main contributions. (1) A new model for integrating facility location and hardening decisions; in particular, a natural three-level formulation is converted to a single-level mixed-integer program (MIP) by taking advantage of the structure of the MFLHP. This model is accompanied by a binary search solution procedure along with a method for obtaining a lower bound. (2) This integrated model and solution method forms the basis of an algorithm that computes the complete Pareto-efficient set for the pre- and post-disruption maximum distances. This algorithm is based on a method from Medal et al. [22] that optimizes facility location decisions but does not model facility hardening. (3) A set of computational experiments provide results that should help decision-makers better understand the tradeoff between the pre- and post-disruption maximum distances when making the decision to locate and harden facilities subject to disruptions. Toward this end we present the following analyses: (i) an analysis of the Pareto-efficient set between the pre- and post-disruption maximum distances; (ii) an analysis of the penalty incurred for optimizing either the pre- or post-disruption radius in isolation; and (iii) an analysis of the benefit of considering facility hardening when locating facilities subject to disruptions.

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

2 Problem Description and Models

The purpose of the MFLHP model is to

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

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

To understand the model, it may help to divide it into three stages: 1) the mitigation, 2) the disruption, and 3) the response. We use the generic term *facility* to refer to a physical entity that we are locating and hardening. The mitigation stage, which happens before the disruption occurs, involves actions taken to mitigate against the disruption. The mitigation decisions in our model concern where to locate facilities and which facilities to harden, and these decisions are made simultaneously. If a facility is hardened in our model, it is always available to serve demand points. In the disruption stage, the disruption causes exactly r facilities to fail. Thus, if p facilities are located, there are $\binom{p}{r}$ combinations of facility disruptions. (Later, we show that we do not have to consider all combinations.) In the response stage, demand points are served by their closest located facility. Since the decision in this stage is so simple–find the closest facility to each demand point–this stage will be implicit in the model.

To understand the three-stage model, it is helpful to think of it as consisting of three players acting in sequence: a defender, an attacker, and an operator. In the first stage, the defender mitigates against the actions of the attacker by strategically locating and hardening facilities. The defender's objective is to minimize the attacker's objective. The attacker, knowing the location and hardening actions taken by the defender, then destroys r facilities. The objective of the attacker is to maximize the operator's objective, i.e., maximize the post-disruption radius. The operator, observing the actions of the attacker, pairs each demand point with its closest available facility in order to minimize the post-disruption radius.

The following notation will be used to describe the problem. Let $\mathcal{I}$ be a set of potential facility locations, $\mathcal{L}$ be a set of located facilities, and $\mathcal{H}$ be a set of hardened facilities. Since only located facilities can be hardened, $\mathcal{H} \subseteq \mathcal{L}$. Denote by $\chi(\mathcal{L}, \mathcal{H})$ the cost of all location and hardening activities, which is subject to a budget b. (Note that b must be large enough to ensure that there exists a feasible solution $(\mathcal{L}, \mathcal{H})$ such that$|\mathcal{L}| \geq r + 1$ or $|\mathcal{H}| \geq 1$.) The set $\mathcal{O}$ is the set of located facilities that fail in a disruption; at most r facilities can fail in a disruption. A set of demand points is represented by the set $\mathcal{J}$. We measure how effectively a facility located at i can serve the demand point located at j using a measure ϕ_{ij}. The measure ϕ_{ij} could be the distance between i and j or a function of the distance between i and j. It could also represent the distance multiplied by the demand weight w_j. Let Φ be a map that holds the assignments of demand points to facilities; thus, $\Phi(j)$ is the facility assigned to demand point j.

Thus, the defender's optimal objective value, $U^{(r)*}$, is equal to the optimal value of the following three-level optimization problem:

$$\min_{\mathcal{H} \subseteq \mathcal{L} \subseteq \mathcal{I}} \quad H(\mathcal{L}, \mathcal{H}) \tag{1a}$$

$$\text{s.t. } \chi(\mathcal{L}, \mathcal{H}) \leq b$$

$$H(\mathcal{L}, \mathcal{H}) = \max_{\mathcal{O} \subseteq \mathcal{L} \setminus \mathcal{H}} \quad S(\mathcal{L}, \mathcal{O}) \tag{1b}$$

$$\text{s.t. } |\mathcal{O}| \leq r$$

$$S(\mathcal{L}, \mathcal{O}) = \min_{\Phi} \max_{j} \phi_{\Phi(j),j} \tag{1c}$$

$$\text{s.t. } \Phi(j) \in \mathcal{L} \setminus \mathcal{O} \quad \forall j \in \mathcal{J}.$$

This is the *integrated MFLHP model*. The defender's problem (1a) is to locate a set of facilities $\mathcal{L}$ from a set of candidate locations $\mathcal{I}$ and harden a subset, $\mathcal{H}$, of the located facilities to minimize the post-disruption radius $H(\mathcal{L}, \mathcal{H})$. The cost of location and hardening must be within a budget, b. The interdictor's problem (1b) is to destroy a subset, $\mathcal{O}$, of the located, unhardened facilities in order to maximize the post-disruption radius, $S(\mathcal{L}, \mathcal{O})$. The interdictor can only destroy r facilities. The operator's problem (1c) is to assign demand points to facilities. Since the operator's problem is uncapacitated, it is optimal to assign every demand point to its closest facility. Thus, the operator's problem can be represented by taking the maximum distance from any demand point to its closest non-disrupted facility.

2.1 Describing the Most Disruptive Facility Disruption Scenario

Rather than trying to solve the three-stage formulation, the MFLHP's structure can be exploited to formulate a single-stage model. In particular, because of the minimax distance objective of the integrated MFLHP model, the interdictor's optimal solution can be described as a simple expression.

Because of the structure of the interdictor's problem, we can determine exactly which facilities the interdictor would want to destroy to optimize his objective without having to solve an optimization problem. After the sets $\mathcal{L}$ and $\mathcal{H}$ have been chosen, the interdictor's problem is

$$U^{(r)*} = \max_{\mathcal{O} \subseteq \mathcal{L} \setminus \mathcal{H}} S(\mathcal{L}, \mathcal{O}), \tag{2}$$

such that $\mathcal{O}$ consists of no more than r facilities. Let $\mathcal{L}_j^i$ be the set of facilities closer to demand point j than facility i is. Let $\mathcal{O}^*(\mathcal{L}, \mathcal{H})$ be the set of facilities that optimizes the interdictor's problem (maximize the post-disruption maximum radius) given a location-hardening solution $(\mathcal{L}, \mathcal{H})$.

Definition 1. Let $U^{(r)*}$ be the optimal post-disruption radius. Demand point j' and facility i' are a *post-disruption bottleneck pair* if $U^{(r)*} = \phi_{i'j'}$. In this case, j' is called a *post-disruption bottleneck demand point* and i' is called a *post-disruption bottleneck facility.*

Theorem 1. *Let $D_j(\mathcal{L}) = \min_{i \in \mathcal{L}} \phi_{ij}$, let $\bar{D}_j(\mathcal{L}) = \max_{i \in \mathcal{L}} \phi_{ij}$, and let $\mathcal{L}_j^k \subseteq \mathcal{L}$ be the set of the k closest located facilities to j. The interdictor's optimal strategy is the following. First, choose the bottleneck demand point as*

$$j' = \arg \max_{j \in \mathcal{J}} \min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\ D_j(\mathcal{H})\}.$$

Next, depending on which facilities are located and hardened, take one of the following actions. i) If the r facilities closest to j' are unhardened, then interdict the r closest facilities to j'. ii) If the k^{th} ($k \leq r$) closest facility to j' is hardened, then interdict the $(k-1)$ closest facilities to j'.

Proof. The interdictor is able to force the post-disruption assignment distance (PDAD) for a demand point $j' \in \mathcal{J}$ to equal $\min\{\bar{D}_{j'}(\mathcal{L}_{j'}^{r+1}),\ D_{j'}(\mathcal{H})\}$ by choosing actions i) or ii) depending on which facilities are located and hardened. Action i) causes j's PDAD to equal $\bar{D}_{j'}(\mathcal{L}_{j'}^{r+1})$ and Action ii) causes it to equal $D_{j'}(\mathcal{H})$. Since the interdictor can only destroy r facilities and a demand point is always assigned to its closest operating facility after a disruption, the quantity $\min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\ D_j(\mathcal{H})\}$ is an upper bound for the PDAD for any demand point $j \in \mathcal{J}$. Since the interdictor is a maximizer and the problem has the bottleneck property (i.e., the interdictor's objective value depends on only one demand point), he or she will focus on the demand point with the largest PDAD upper bound. □

The following corollary, following directly from Theorem 1, provides a simple expression of the interdictor's optimal objective value, which will help simplify the three-level model.

Corollary 1. *The interdictor's optimal objective value is equal to*

$$\max_{j \in \mathcal{J}} \min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\ D_j(\mathcal{H})\}. \tag{3}$$

2.2 Single-Level Model

By Corollary 1, the interdictor's maximization problem (given a location-hardening solution $(\mathcal{L}, \mathcal{H})$) can be reformulated as the problem of simply choosing a demand point, j', with the largest post-disruption assignment distance to be the post-disruption bottleneck demand point, i.e.:

$$H(\mathcal{L}, \mathcal{H}) = \max_{j \in \mathcal{J}} \min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\ D_j(\mathcal{H})\}, \tag{4}$$

Substituting (4) for (1b), the three-level problem (1) converts to the following minimax problem:

$$U^{(r)*} = \min_{(\mathcal{L},\mathcal{H}) \subseteq \mathcal{K}} \max_{j \in \mathcal{J}} \min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\ D_j(\mathcal{H})\}, \tag{5}$$

where $\mathcal{K} = \{(\mathcal{L}, \mathcal{H}) :\ \mathcal{L} \subseteq \mathcal{I},\ \mathcal{H} \subseteq \mathcal{L},\ \chi(\mathcal{L}, \mathcal{H}) \leq b\}$ is the set of feasible location-hardening solutions.

2.3 MIP Model

Model (5) has the minimax structure of a bottleneck problem [14] and can thus be reformulated as a single-level MIP. Let W_{ij} be a variable that is 1 if ϕ_{ij} is an upper bound on the post-disruption assignment distance (PDAD) for demand point j and 0 otherwise. Let X_i be a binary variable that is 1 if a facility is located at i and 0 otherwise and Y_i be a variable that is 1 if a facility at i is hardened and 0 otherwise. The cost of locating a facility at i is f_i, and the cost of hardening a facility at i is g_i.

A MIP formulation of the integrated MFLHP model is:

$$\min \quad U^{(r)} \tag{6a}$$

$$\text{s.t.} \quad U^{(r)} \geq \phi_{ij} W_{ij} \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, \tag{6b}$$

$$(r+1)W_{ij} \leq (r+1)Y_i + \sum_{i' \in \mathcal{I}: \phi_{i'j} \leq \phi_{ij}} X_{i'} \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, \tag{6c}$$

$$\sum_{i \in \mathcal{I}} W_{ij} = 1 \quad \forall j \in \mathcal{J}, \tag{6d}$$

$$W_{ij} \leq X_i \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, \tag{6e}$$

$$Y_i \leq X_i \quad \forall i \in \mathcal{I} \tag{6f}$$

$$\sum_{i \in \mathcal{I}} f_i X_i + \sum_{i \in \mathcal{I}} g_i Y_i \leq b, \tag{6g}$$

$$X_i, Y_i \in \{0, 1\} \quad \forall i \in \mathcal{I}, \tag{6h}$$

$$W_{ij} \in \{0, 1\} \quad \forall i \in \mathcal{I}, j \in \mathcal{J}. \tag{6i}$$

The objective (6a) is to minimize the post-disruption radius, which is the defender's objective. Constraints (6b) require that the interdictor's objective is equal to the maximum of all of the PDAD upper bounds, obeying Proposition 1. Constraints (6c) ensure that the PDAD upper bound for a demand point j is equal to $\min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\, D_j(\mathcal{H})\}$. If a facility i is hardened, i.e., $Y_i = 1$, then $\phi_{ij} = D_j(\mathcal{H})$ in an optimal solution. If r facilities closer to j than i are located and i is located, i.e., $\sum_{i' \in \mathcal{I}: \phi_{i'j} \le \phi_{ij}} X_{i'} = r + 1$, then $\phi_{ij} = \bar{D}_j(\mathcal{L}_j^{r+1})$ in an optimal solution. Constraints (6d) require every demand point to have an upper bound. Constraints (6e), although not necessary because of the presence of Constraints (6c), tighten the LP relaxation. Constraints (6f) prevent non-located facilities from becoming hardened. Constraint (6g) is a budget on the location and hardening costs. Constraints (6h)–(6i) require that the variables be binary.

Theorem 2. *The single-level MIP formulation* (6) *is equivalent to the three-level model* (1).

Proof. Let $\mathcal{L} = \{i : X_i = 1\}$ and $\mathcal{H} = \{i : Y_i = 1\}$ be the set of located and hardened facilities corresponding to a solution to (6). (i) A feasible solution to (6) obeys the requirement $\mathcal{H} \subseteq \mathcal{L} \subseteq \mathcal{I}$ in (1) because Constraints (6h) enforce $\mathcal{L} \subseteq \mathcal{I}$, and the coefficients $(f_i + g_i)$ for each Y_i in (6g) ensure that every hardened facility is also located, enforcing $\mathcal{H} \subseteq \mathcal{L}$. (ii) In addition, Constraint (6g) is simply a concrete instance of the abstract constraint $\chi(\mathcal{L}, \mathcal{H}) \le b$ in (1). (iii) Finally, we will demonstrate that minimizing $U^{(r)}$ in (6) subject to Constraints (6b)–(6i) is equivalent to minimizing $H(\mathcal{L}, \mathcal{H})$ in (1). Because of Constraints (6b) and (6d), $U^{(r)} = \max_{j \in \mathcal{J}} \phi_{\Phi(j),j}$, where $\Phi(j)$ is the facility i for which $W_{ij} = 1$. Because of the bottleneck structure inherent in (6a) and (6b), there exists a demand point j' such that $U^{(r)} = \phi_{\Phi(j'),j'}$; i.e., j' is a bottleneck demand point. Because of Constraints (6c), the feasible values for $\Phi(j')$ include all i for which $Y_i = 1$ and all i for which $\sum_{i' \in \mathcal{I}: \phi_{i'j'} \le \phi_{ij'}} X_{i'} = r + 1$. Because of the minimization in (6a), $\Phi(j')$ will be chosen in a way that minimizes $\phi_{\Phi(j'),j'}$. Thus,

$$
\begin{aligned}
U^{(r)} &= \phi_{\Phi(j'),j'} = \min\{ \min_{i: \sum_{i' \in \mathcal{I}: \phi_{i'j'} \le \phi_{ij'}} X_{i'} = r+1} \phi_{i,j'},\, \min_{i: Y_i = 1} \phi_{i,j'} \} \\
&= \min\{\bar{D}_{j'}(\mathcal{L}_{j'}^{r+1}), D_{j'}(\mathcal{H})\} = \max_{j \in \mathcal{J}} \min\{\bar{D}_j(\mathcal{L}_j^{r+1}),\, D_j(\mathcal{H})\},
\end{aligned}
$$

which by Corollary 1 is equal to $H(\mathcal{L}, \mathcal{H})$. Therefore, since the single-level MIP formulation (6) is equivalent to minimizing $H(\mathcal{L}, \mathcal{H})$ subject to the Constraints $\mathcal{H} \subseteq \mathcal{L} \subseteq \mathcal{I}$ and $\chi(\mathcal{L}, \mathcal{H}) \le b$, the theorem holds. □

2.4 Alternate Formulation

We also tested an alternate formulation of Model (6) which involves redefining the X and Y variables. Let X_i be a binary variable that is 1 if a facility is located but not hardened at i and 0 otherwise and Y_i be a variable that is 1 if a facility at i is located and hardened and 0 otherwise.

The alternate formulation is as follows:

$$\text{(MIP-Alt)} \quad \min \quad U^{(r)} \tag{7a}$$

$$\text{s.t.} \quad (6b) - (6d) \tag{7b}$$

$$W_{ij} \leq X_i + Y_i \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, \tag{7c}$$

$$\sum_{i \in \mathcal{I}} f_i X_i + \sum_{i \in \mathcal{I}} (f_i + g_i) Y_i \leq b, \tag{7d}$$

$$(6h) - (6i). \tag{7e}$$

Constraints (7c) and (7d) are revised versions of Constraints (6e) and (6g), respectively. Preliminary tests showed that the alternate model (7) solved faster than Model (6).

3 Solution Methodology

Model (6) can be solved using an off-the-shelf MIP optimizer such as CPLEX. However, preliminary experiments indicated that this approach would be computationally prohibitive. Because of the bottleneck structure of the integrated MFLHP model, we chose to use a binary search (BS) algorithm instead. Hochbaum and Shmoys [14] showed that all bottleneck problems can be solved by solving a series of auxiliary problems within a binary search algorithm that searches over values in the set of all possible radii. These auxiliary problems can be thought of as inverses of their corresponding bottleneck problem. Specifically, this auxiliary problem takes a radius value as an input and outputs the cost of covering all objects within that radius.

Empirical evidence has shown that a binary search algorithm works well for the p-center problem, which is also a bottleneck problem [11]. The p-center problem is to find a set of p located facilities that minimize the radius. The auxiliary problem for the p-center problem is the set-cover problem with unitary costs, which is still an NP-hard problem [15]. If some radius δ is given as an input to the set-cover problem, the set-cover problem outputs how many facilities must be located, i.e., the cost, so that all demand points are covered within δ. Let $U^{(0)*}$ be the optimal radius and $p^*(\delta)$ be the optimal number of facilities needed to cover all demand points within δ. If $p^*(\delta)$ is greater than or equal to p, then $U^{(0)*}$ is at least δ and δ is a new lower bound. If $p^*(\delta)$ is less than p, then $U^{(0)*}$ is at most δ and δ is a new upper bound. Thus, a binary search can be performed over all values of δ to find the optimal radius, $U^{(0)*}$. Binary search has been shown to be an effective solution method for the p-center problem because the set-cover problem with unitary costs is easier to solve than the p-center problem itself.

Similar to the solution approach for the p-center problem, in this article we solve the integrated MFLHP using a binary search algorithm. However, we use an auxiliary problem that is tailored for the MFLHP along with new upper and lower bounds and a new heuristic.

3.1 Auxiliary Problem

To use a binary search algorithm for the integrated MFLHP, the auxiliary problem must first be described. Define $U^{(r)}$ as the radius for the auxiliary problem. (Note that $U^{(r)}$ is now a parameter and not a variable, as it was in Model (6).) To evaluate whether a particular $U^{(r)}$ is above or below the optimal post-disruption radius, the set-cover problem with location and hardening (SCP-LH) is used:

$$\text{SCP-LH}(U^{(r)}, r) \quad \min \quad \sum_{i \in \mathcal{I}} f_i X_i + \sum_{i \in \mathcal{I}} Y_i \tag{8a}$$

$$\text{s.t.} \quad (r+1) \sum_{i: \phi_{ij} \leq U^{(r)}} Y_i + \sum_{i: \phi_{ij} \leq U^{(r)}} X_i \geq r + 1 \quad \forall j \in \mathcal{J}, \tag{8b}$$

$$Y_i \leq X_i \quad \forall i \in \mathcal{I}, \tag{8c}$$

$$X_i, Y_i \in \{0, 1\} \quad \forall i \in \mathcal{I}. \tag{8d}$$

The SCP-LH minimizes the cost required for every demand point to have a post-disruption assignment distance less than or equal to $U^{(r)}$. The objective (8a) is to minimize the total cost of location and hardening. Constraints (8b) require that for each demand point j, either $r+1$ facilities within $U^{(r)}$ of j must be located or at least one facility within $U^{(r)}$ of j must be hardened. Since the set-cover problem is a special case of SCP-LH($U^{(r)}, r$) (set $r = 0$ and set $f_i = 1, Y_i = 0$ for all i), SCP-LH($U^{(r)}, r$) is also NP-hard.

Note that an alternate version of SCP-LH can be constructed by redefining the X and Y variables as in Section 2.4. In this alternate model the objective (8a) changes to $\sum_{i \in \mathcal{I}} f_i X_i + \sum_{i \in \mathcal{I}} (f_i + g_i) Y_i$, but Constraints (8b) remain unchanged. We refer to this alternate model as SCP-LH-Alt. Because SCP-LH-Alt outperformed SCP-LH in preliminary testing, we use SCP-LH-Alt throughout the rest of this paper.

3.2 Binary Search Algorithm

The binary search algorithm for the integrated MFLHP is similar to the binary search algorithm for the p-center problem. In addition to the modified auxiliary problem, we also add a heuristic for the auxiliary problem and use a polynomial algorithm to obtain bounds for the integrated MFLHP. The binary search algorithm is described in Medal [21].

Before starting the binary search, we attempt to find good upper and lower bounds for the integrated MFLHP to reduce the search space of the algorithm. In particular, we apply the binary search algorithm to the linear-programming relaxation of the auxiliary problem (8). Let the binary search algorithm that uses the linear-programming relaxation of the auxiliary problem be called the *relaxed binary search algorithm*. This idea was first employed by Elloumi et al. [11] for the p-center problem.

We compute a binary search (BS) lower bound for the MFHLP by applying the relaxed binary search algorithm. The optimal radius returned by the relaxed binary search algorithm is a lower bound to the integrated MFHLP. We add the following step to the relaxed binary search algorithm so that it also returns an upper bound. Each iteration of the relaxed binary search algorithm produces a new midpoint index; let D_{index} be the distance corresponding to this index. To obtain a heuristic upper bound, solve SCP-LH-Alt(D_{index}) using a greedy algorithm (see Medal [21]). This greedy algorithm extends an algorithm by Balas and Ho [3] for the set cover problem and sequentially adds facilities to be located and hardened based on their costs of location and hardening and how many additional demand points they can cover. If the total cost of this heuristic solution to SCP-LH-Alt(D_{index}) is less than the budget (b), then D_{index} is a new incumbent lowest upper bound.

There are other ways to obtain bounds for the integrated MFLHP. A linear-programming lower bound (LP) can be obtained by solving the linear-programming relaxation of the MIP model (6). A partial relaxation (PR) lower bound can be obtained by solving the MIP model (6) with only X_i and Y_i relaxed for all i. In our experimentation, we found that the BS lower bound required much less run time than the PR lower bound and yet the BS lower bounds were reasonably close (16% relative error, on average) to the PR lower bounds.

3.3 Multiobjective Optimization: Pre-Disruption Radius Vs. Post-Disruption Radius

Although the MFLHP optimizes the post-disruption radius (PostDR), decision-makers may also be concerned about the pre-disruption radius (PreDR), especially if disruptions are rare. Because these two objectives conflict, an effective solution approach is to generate a set of Pareto-efficient solutions, allowing a decision-maker to choose from among them based on his or her preferences.

Because the binary search algorithm presented in Section 3.2 efficiently computes the optimal solution to the MFLHP, it can also be used to efficiently compute the Pareto-efficient set for multiple objectives. To describe how to compute the Pareto-efficient set for the pre-disruption and post-disruption radii, let us introduce two new optimization problems. Let MFLHP($\cdot, U^{(r)}$) denote the problem "minimize the pre-disruption radius such that the post-disruption radius is no greater than $U^{(r)}$, and MFLHP($U^{(0)}, \cdot$) denote the problem "minimize the post-disruption radius such that the pre-disruption radius is no greater than $U^{(0)}$. Both of these problems can be solved using a modification of the binary search algorithm from Section 3.2. Specifically, a single-cover constraint is added to the SCP-LH($U^{(r)}, r$) auxiliary problem, resulting in the distance-constrained set-cover problem with location and hardening (DC-SCP-LH):

$$\text{DC-SCP-LH}(U^{(0)}, U^{(r)}, r) \quad \min \quad \sum_{i\in\mathcal{I}} f_i X_i + \sum_{i\in\mathcal{I}} g_i Y_i \tag{9a}$$

$$\text{s.t.} \quad \sum_{i:\phi_{ij}\le U^{(0)}} X_i \ge 1 \quad \forall j \in \mathcal{J}, \tag{9b}$$

$$Y_i \le X_i \quad \forall i \in \mathcal{I} \tag{9c}$$

$$(r+1) \sum_{i:\phi_{ij}\le U^{(r)}} Y_i + \sum_{i:\phi_{ij}\le U^{(r)}} X_i \ge r+1 \quad \forall j \in \mathcal{J}, \tag{9d}$$

which minimizes the cost of location and hardening (9a) subject to the requirements that i) the pre-disruption radius be at most $U^{(0)}$ (9b) and ii) the post-disruption radius be at most $U^{(r)}$ (9d).

The DC-SCP-LH$(U^{(0)}, U^{(r)}, r)$ can be used to solve MFLHP$(\cdot, U^{(r)})$ and MFLHP$(U^{(0)}, \cdot)$ via binary search:

- To solve MFLHP$(\cdot, U^{(r)})$, fix $U^{(r)}$ in Constraints (9d) and vary $U^{(0)}$ in Constraints (9b) via binary search (see Section 3.2) in order to find the optimal pre-disruption radius, $U^{(0)*}$.

- To solve MFLHP$(U^{(0)}, \cdot)$, fix $U^{(0)}$ in Constraints (9b) and vary $U^{(r)}$ in Constraints (9d) within a binary search algorithm (see Section 3.2).

Given the problems MFLHP$(\cdot, U^{(r)})$ and MFLHP$(U^{(0)}, \cdot)$, we can use the ϵ-constraint method of multiobjective optimization [6] to generate the Pareto-efficient set. The ϵ-constraint method was used for the p-center problem with facility disruptions in Medal et al. [22] and is modified in this article to include facility hardening. Specifically, the auxiliary problem DC-SCP-LH$(U^{(0)}, U^{(r)}, r)$ is used in place of the auxiliary problem employed in Medal et al. [22]. Algorithm 1 describes this ϵ-constraint method. On line 4, an upper bound is obtained by taking the maximum distance from a demand point j to its $(r+1)^{\text{st}}$ closest facility site, denoted as i_j^{r+1}. On line 6, MFLHP$(U_k^{(0)}, \cdot)$ is solved, returning the post-disruption radius for the k^{th} member of the Pareto set, and on line 8, MFLHP$(\cdot, U_k^{(r)} - \epsilon)$ is solved, returning the pre-disruption radius. The value ϵ is subtracted from $U_k^{(r)}$ to ensure that the pre-disruption radius corresponding to Pareto point $(k+1)$ is greater than the pre-disruption radius corresponding to Pareto point k.

Algorithm 1 Constructing the Pareto-efficient set for pre- and post-disruption radius objectives.

function GenerateParetoEfficientSet
  $k \leftarrow 0$; $\epsilon \leftarrow$a small number
  Set of Pareto-efficient points $\mathcal{S} \leftarrow \emptyset$
  Solve MFLHP$(\cdot, \max_{ij}\{\phi_{i_j^{r+1} j}\})$, returning minimum PreDR $U_k^{(0)}$.
  loop
    Solve MFLHP$(U_k^{(0)}, \cdot)$, returning $U_k^{(r)}$. ▷ Min. PostDR s.t. PreDR constraint.
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{(U_k^{(r)}, U_k^{(0)})\}$
    Solve MFLHP$(\cdot,\ U_k^{(r)} - \epsilon)$, returning $U_{k+1}^{(0)}$. ▷ Min. PreDR s.t. PostDR constraint.
    if MFLHP$(\cdot,\ U_k^{(r)} - \epsilon)$ is infeasible, **break**.
    $k \leftarrow k + 1$
  return $\mathcal{S}$

4 Illustrative Example

We will first demonstrate our research on a simplified example, using a dataset from Daskin [8] that consists of 150 nodes representing the most populous cities in the United States. This example attempts to locate seven facilities to service each of the 150 demand points. To make solutions easier to visualize, each of the demand points are given a unit weight. Without considering facility disruptions, the design objective is to minimize the maximum distance from a demand point to its closest facility. Thus, this is the classic p-center problem [12]. Because the p-center problem is known to produce many optimal solutions, we utilize the total distance as a secondary objective in all of the results presented in this section and the next. See Appendix A for a model that optimizes the total distance subject to constraints on the pre-disruption and post-disruption radii.

Figure 1a shows the optimal solution to this 7-center problem. Seven facilities are located in the optimal solution, and the (pre-disruption) radius for this solution is 420 miles. One limitation of the p-center model is that it may produce solutions that are vulnerable to facility disruptions. Figure 1b shows how the solution is impacted by the worst-case facility disruption scenario consisting of the disruption of $r = 3$ facilities. Specifically, the demand points assigned to the facilities at Albequerque, NM; Boise City, ID; and Glendale, CA must be reassigned, causing the radius to increase. The new radius, the post-disruption radius, is 1,624 miles which is a 286% increase over the pre-disruption radius.

A decision-maker is likely to be concerned about this 286% increase in the radius. To reduce the potential increase in the radius the MFLHP model can be used to locate and harden facilities. However, it is likely that the decision-maker is interested in both the pre-disruption and post-disruption radii. In many cases a decision-maker may not know his or her preferred tradeoff between the pre-disruption and post-disruption radii. Thus, it would be helpful if the decision-maker had a

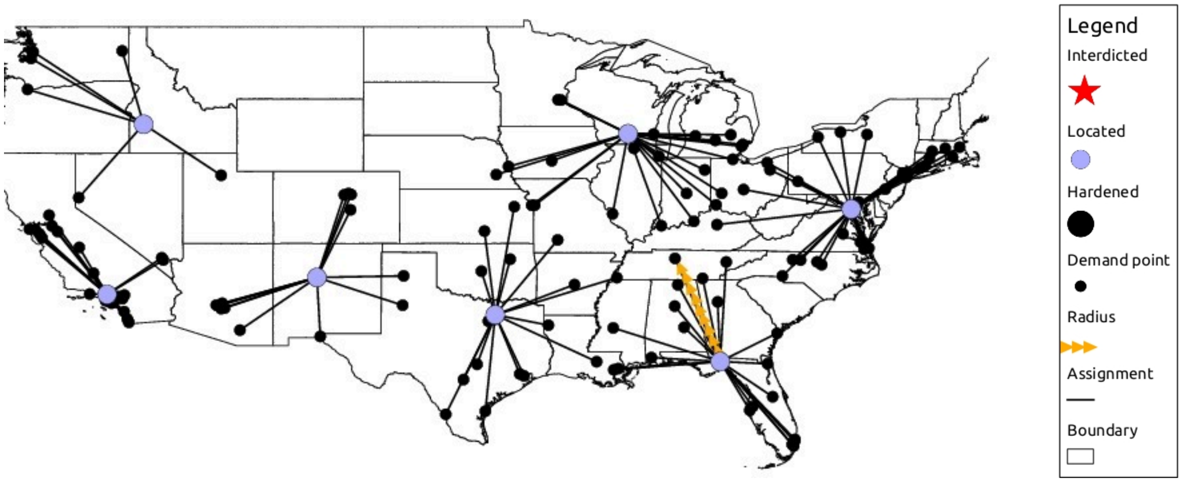

(a) Before disruptions

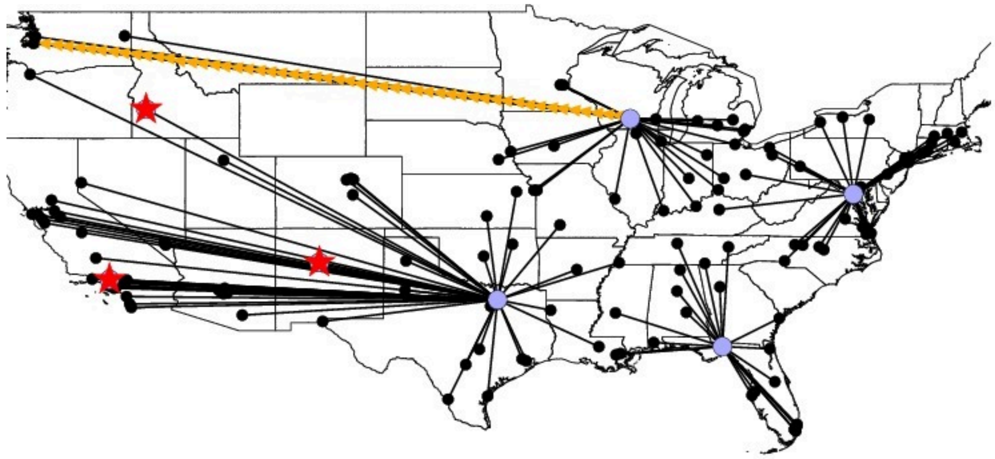

(b) After disruptions

Figure 1: Optimal solution to 7-center problem

set of Pareto-efficient solutions to choose from. Fortunately, Algorithm 1 can be used to compute the Pareto-efficient set for the pre-disruption and post-disruption radii.

Figure 2 shows the Pareto-efficient set for the pre-disruption and post-disruption radii. (For this example, we assume that the cost of hardening a facility is the same as the cost of locating a facility (e.g., $g_i = f_i$).) The left and right endpoints of the Pareto set are obtained by optimizing the post-disruption and pre-disruption radii, respectively, in isolation. The Pareto set is relatively flat when the post-disruption radius is between 902 and 1,381. This means the post-disruption radius can be decreased from 1,381 to 902 (35% reduction) while only increasing the pre-disruption radius from 543 to 612 (13% increase). Overall, the slope of the curve is gradual, indicating that a large reduction in the post-disruption radius can be obtained with a modest increase in the pre-disruption radius. More analysis on the tradeoff between the pre-disruption and post-disruption radii is given in Section 5.1.

Figure 3 shows three of the solutions corresponding to Pareto-efficient points in Figure 2. Figures

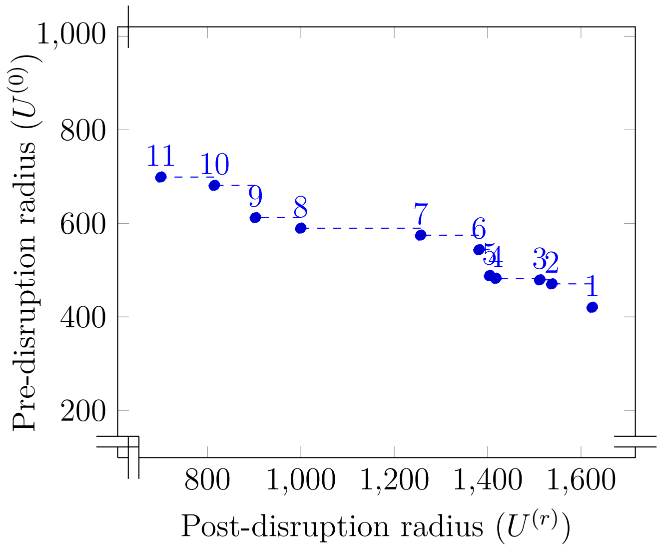

Figure 2: Pre-disruption radius vs. post-disruption radius: Pareto-efficient set

3a, 3c, and 3e show the solutions and assignments without disruptions for solutions 6, 9, and 11 of the Pareto set in Figure 2. Figures 3b, 3d, and 3f show the post-disruption assignments. (Solution 1 is shown in Figure 1.) As the Figures show, as the post-disruption radius receives more emphasis, less of the budget is spent on location and more is spent on hardening. A description of each of the solutions in Figures 3a–3f is given in the supplemental material.

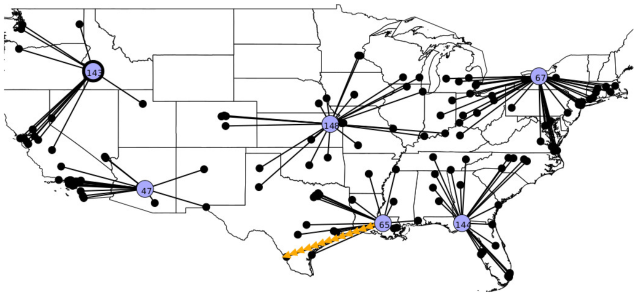

(a) Soln. 6: Pre-disruption radius favored; no disruption

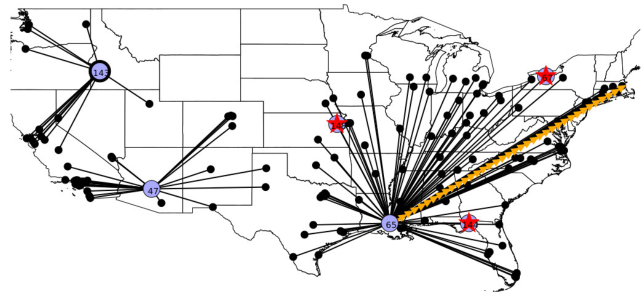

(b) Soln. 6: Pre-disruption radius favored; after disruption

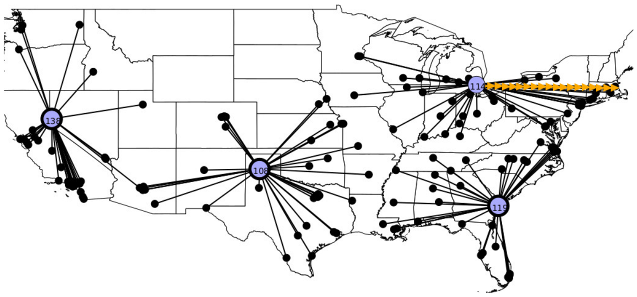

(c) Soln. 9: Post-disruption radius favored; no disruption

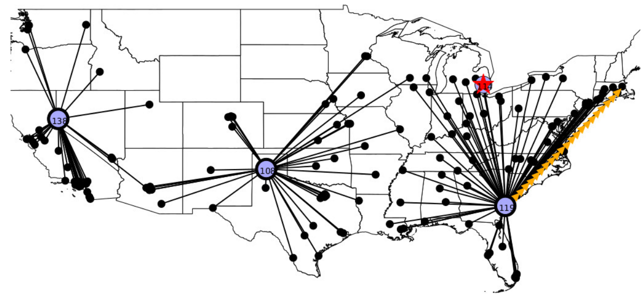

(d) Soln. 9: Post-disruption radius favored; after disruption

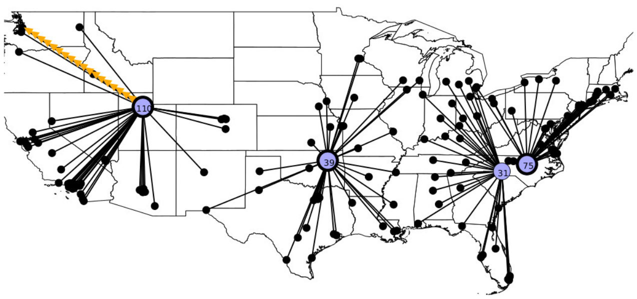

(e) Soln. 11: Post-disruption radius optimized in isolation; no disruption

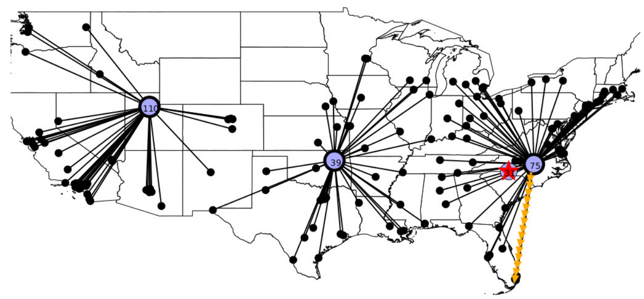

(f) Soln. 11: Post-disruption radius optimized in isolation; after disruption

Figure 3: Optimal facility location and hardening solutions

5 Experimental Results

This section reports the results of computational experimentation with the MFLHP. Section 5.1 analyzes the tradeoff between the pre-disruption radius and post-disruption radius. Section 5.2 analyzes the benefit of including facility hardening when locating facilities that are subject to disruptions.

All experiments were run on a compute node on the Arkansas High Performance Computing Cluster. The node has 2 Xeon X5670 Intel processors, which each have 6 cores and a clock speed of 2.93GHz. The total of 12 cores share 24GB of memory. Computations were done on a 64-bit

Linux operating system. All of the MIP and LP models, including the set-cover problems, were solved using CPLEX v12 Parallel MIP Optimizer with 12 parallel threads and default settings. The binary search algorithm was programmed in Java using CPLEX Concert technology.

Table 1 describes the datasets used in the experimentation. These datasets, taken from the facility location literature, vary in size (49–818 nodes), distance metric used (e.g., Euclidean, great circle, etc.), and whether demand weights and facility location costs are homogeneous or non-homogeneous. Column 2 shows the number of nodes in the dataset. Columns 3 and 4 indicate whether the demand weights and location costs are homogeneous (H) or non-homogeneous (NH). When a dataset has demand weights, then ϕ_{ij} is the weighted distance; otherwise it is the unweighted distance. Some of the datasets are based on geographical data and some are not. Column 6 indicates which distance metric is used: Euclidean, Great Circle, or road distances measured from real data.

Table 1: Datasets used in experimentation

Name (1)	$\lvert\mathcal{I}\rvert = \lvert J\rvert$ (2)	Demand weights (3)	Location costs (4)	Source of data (5)	Distance measure (6)	Ref. (7)
d49	49	NH	NH	49 US state capitals and Washington, D.C.	Great circle	Daskin [8]
d88	88	NH	NH	cities in US	Great circle	[8]
d150	150	NH	NH	cities in US	Great circle	[8]
sw55	55	NH	1.0	population centers in Washington, D.C.	Euclidean	Swain [36]
lor100	100	NH	1.0	population centers in San Jose Dos Campos, Brazil (SJDC)	Road	Lorena and Senne [19]
lon150	150	NH	1.0	population centers in London, Ontario	Road	Alp et al. [2]
lor200	200	NH	1.0	SJDC	Road	[19]
lor300a	300	NH	1.0	SJDC	Road	[19]
lor402a	402	NH	1.0	SJDC	Road	[19]
lor818	818	H	1.0	SJDC	Road	[19]
rl1323	1323	H	1.0	drilling problem TSP	Euclidean	Reinelt [28]

Table 2 gives a list of all of the parameters that were varied in our experimentation. First, the budget (b) for a particular instance is a percentage of the cost of locating every facility. Therefore, letting B be the budget multiplier, the budget (b) equals $\lceil B \sum_{i\in\mathcal{I}} f_i \rceil$. Second, the cost of hardening a facility located at i is a multiple of the cost of locating a facility at i. That is, $g_i = H f_i$, where H

is the hardening cost multiplier. The number of facility disruptions, r, was also varied. To facilitate comparison across datasets, $r = 1$ and $r = 2$ were considered for all datasets. We also considered the number of disruptions to be a percentage of the number of facility locations.

Table 2: Parameter values for experiments

Parameter	Description	Values used
$\lvert\mathcal{I}\rvert$	number of facility locations	depends on dataset
$\lvert J\rvert$	number of demand points	depends on dataset
B	budget multiplier	$0.1, 0.2, 0.3$
H	hardening cost multiplier	$0.1, 0.25, 0.5, 1$
r	number of facility disruptions	1,2,$\lceil 0.05B\lvert\mathcal{I}\rvert\rceil$,$\lceil 0.1B\lvert\mathcal{I}\rvert\rceil$,$\lceil 0.15B\lvert\mathcal{I}\rvert\rceil$

5.1 Multiobjective Analysis: Pre-Disruption Radius vs. Post-Disruption Radius

As discussed in Section 3.3, a decision-maker may be interested in the tradeoff between the pre-disruption radius and the post-disruption radius. In this section, this tradeoff is analyzed from two perspectives. First, the Pareto-efficient set is displayed in Section 5.1.1. Second, Section 5.1.2 addresses the question: If the post-disruption radius is optimized in isolation, what is the effect on the pre-disruption radius, and vice versa? Finally, the computation time for the multi-objective algorithm is reported in Section 5.1.3.

5.1.1 Pareto-Efficient Set

Figure 4 shows Pareto-efficient sets for the pre-disruption and post-disruption radii for several instances of the d49 and lon150 datasets with the cost of hardening a facility equal to the cost of locating a facility (e.g., $g_i = f_i$). Each Pareto set was generated using Algorithm 1. The point in the southwest corner of the plot is the solution for zero interdictions; thus, the pre- and post-disruption radii are equal for this point.

Two main observations can be made from Figure 4.

1. *Tradeoff ratio.* The plots become steeper as the post-disruption radius decreases in Figures 4b, 4c, and 4d. This indicates that as the post-disruption decreases, the tradeoff ratio (i.e., the marginal increase in the pre-disruption radius per unit decrease in the post-disruption radius) increases. On the other hand, for larger values of the post-disruption radius, the tradeoff ratio is smaller. This indicates that significant decreases in the post-disruption radius can

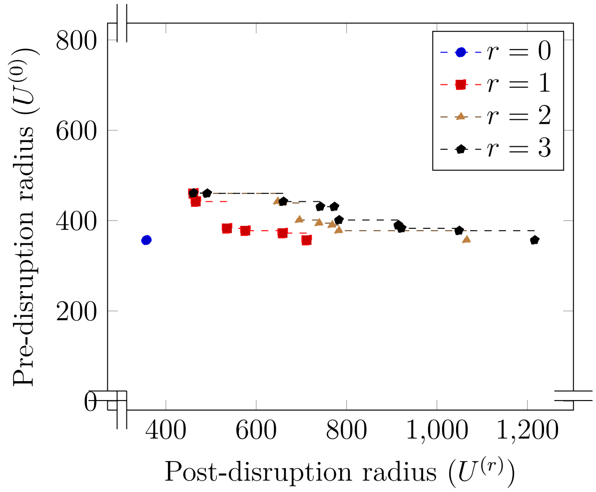

(a) Unweighted Daskin 49-node dataset (d49), $B = 0.2$

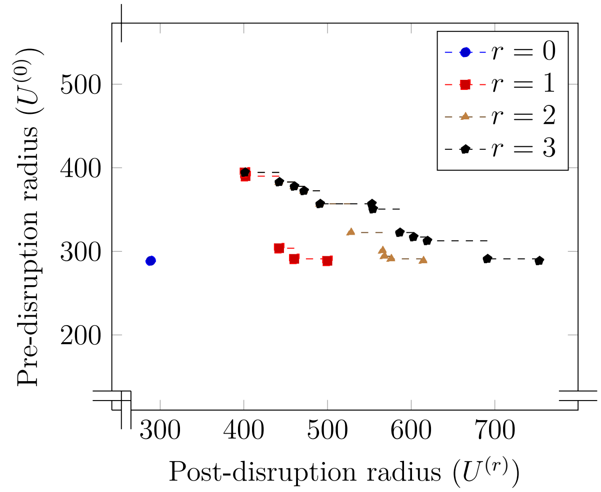

(b) Unweighted Daskin 49-node dataset (d49), $B = 0.3$

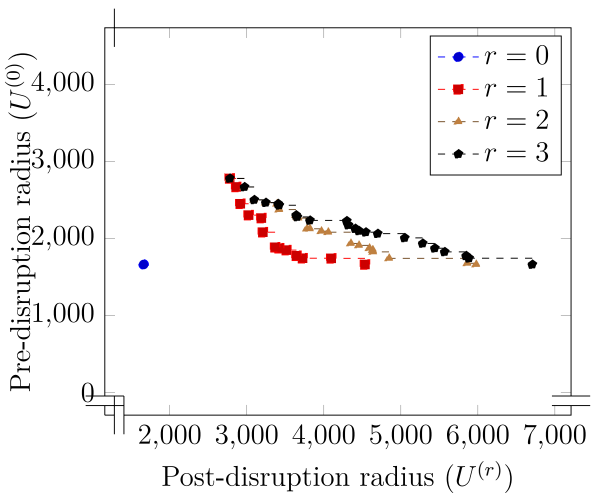

(c) London 150-node dataset (lon150), $B = 0.2$

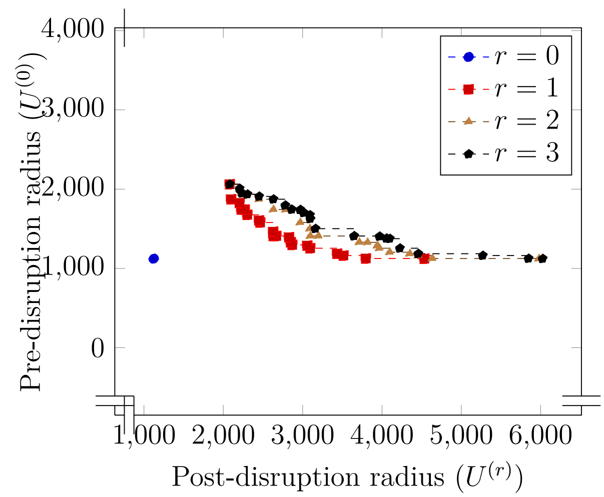

(d) London 150-node dataset (lon150), $B = 0.3$

Figure 4: PreDR vs. PostDR: Pareto efficient sets for several problem instances

be attained with small increases in the pre-disruption radius, a result that is consistent with Snyder and Daskin [34].

2. *Flat regions.* Most of the plots with more interdictions ($r = 2, 3$) have more flat regions in which the pre-disruption radius stays constant as the post-disruption radius decreases (e.g., d49, $B = 0.2$, $r = 3$). Thus, it is important for a decision-maker to be able to view the Pareto-efficient set, especially when the number of interdictions (r) is large (> 2). For, if the decision-maker optimized the pre-disruption radius subject to an arbitrary constraint on the post-disruption radius, they may identify a non-Pareto solution, i.e., a solution on the flat region. Such a solution is undesirable because the post-disruption radius can be decreased without increasing the pre-disruption radius.

5.1.2 Relative Objective Function Increase for Only Considering a Single Objective

In this section we examine the penalty incurred for optimizing either the pre- or post-disruption radius in isolation, rather than using a multi-objective approach. For example, if the pre-disruption radius is optimized in isolation, the post-disruption radius may be much higher than its optimal value. We analyzed this penalty by recording the relative objective function increases for optimizing either of these two objectives in isolation.

For our analysis we used two models to solve various problem instances: the integrated MFLHP model and the b-center model, which is the budget-constrained version of the classic p-center problem [12] and minimizes the pre-disruption radius. For each model, we recorded two values: the post-disruption radius (PostDR) and the pre-disruption radius (PreDR). We then used these values to compute two types of relative objective function increases. The Type I relative objective function increase occurs when the decision-maker plans for a disruption to occur (e.g., optimizes the post-disruption radius) and yet no disruption occurs. Type II occurs when the decision-maker does not plan for a disruption and yet a disruption occurs. (These metrics are analogous to Type I and II errors in hypothesis testing.) Let $XY^*_{(r)}$ be the optimal location and hardening solution generated from the integrated MFLHP model, and let $f_{(0)}(XY^*_{(r)})$ and $f_{(r)}(XY^*_{(r)})$ be the pre-disruption and post-disruption radii, respectively, for this solution. Let $X^*_{(0)}$ be the optimal location solution generated by the b-center problem, and let $f_{(0)}(X^*_{(0)})$ and $f_{(r)}(X^*_{(0)})$ be the pre-disruption and post-disruption radii, respectively, for this solution. The *type I relative objective function increase* is

$$\frac{(\text{PreDR} \mid \text{disruptions}) - (\text{PreDR} \mid \text{no disruption})}{(\text{PreDR} \mid \text{no disruption})} = \frac{f_{(0)}(XY^*_{(r)}) - f_{(0)}(X^*_{(0)})}{f_{(0)}(X^*_{(0)})}, \tag{10}$$

and the *type II relative objective function increase* is

$$\frac{(\text{PostDR} \mid \text{no disruption}) - (\text{PostDR} \mid \text{disruptions})}{(\text{PostDR} \mid \text{disruptions})} = \frac{f_{(r)}(X^*_{(0)}) - f_{(r)}(XY^*_{(r)})}{f_{(r)}(XY^*_{(r)})}. \tag{11}$$

Table 3 shows the type I and type II relative objective function increases for instances of the d49 and sw55 datasets. The table shows that for the d49 dataset, the average Type I relative objective function increase was much lower than the average Type II relative objective function increase. Conversely, the opposite behavior occurred for the sw55 dataset. It is unclear why the behavior for the d49 dataset is different because two other datasets also had the same behavior as the sw55 dataset: lor100 (Avg. Type I = 3.55, Avg. Type II = 2.23) and lon150 (Avg. Type I = 6.85, Avg. Type II = 2.73).

Table 3: Relative objective function increases associated with optimizing only pre-disruption radius and post-disruption radius, respectively

(a) d49 dataset

No.	B	H	r	Type I	Type II
1	0.2	0.25	1	0.26	5.03
2	0.2	0.25	2	0.26	10.84
3	0.2	1	1	0.73	3.20
4	0.2	1	2	0.73	7.24
5	0.3	0.25	1	0.27	8.60
6	0.3	0.25	2	0.27	12.51
7	0.3	0.25	3	0.27	18.36
8	0.3	1	1	1.19	4.82
9	0.3	1	2	1.19	7.20
10	0.3	1	3	1.19	10.74
Average				0.64	8.85

(b) sw55 dataset

No.	b	H	r	Type I	Type II
1	11	0.25	1	0.41	1.21
2	11	0.25	2	0.41	2.13
3	11	1	1	2.94	0.79
4	11	1	2	9.32	1.46
5	16	0.25	1	4.72	2.29
6	16	0.25	2	4.72	2.47
7	16	0.25	3	4.72	4.49
8	16	1	1	0.79	1.10
9	16	1	2	0.79	1.21
10	16	1	3	0.79	2.50
Average				2.96	1.97

5.1.3 Computational Times

The single-objective binary-search algorithm, described in Section 3.2, was tested over a set of datasets and problem instances, the largest being the rl1323 dataset. The results showed that the instances solved in a reasonable amount of time (see Appendix B). In fact, the run time of the algorithm for most of the rl1323 instances was between 412 and 3389 seconds.

Experimental results also showed that the bi-objective instances solved in a reasonable amount of time for small- to medium-sized datasets (see Appendix C) even though the entire Pareto-efficient frontier must be computed. For problem instances with up to nodes, the run time was never more than 5.8 hours. However, some instances of larger datasets (300–818 nodes) did not solve within 24 hours.

5.2 Benefit of Hardening

In this section we measure the benefit of considering facility hardening when locating facilities subject to disruptions. We measure this benefit using two objectives: the pre-disruption radius and the post-disruption radius. The most interesting finding in our results is that while considering facility hardening in a model did produce solutions that have a lower post-disruption radius, it did not affect the pre-disruption radius.

To measure the benefit of hardening, we compare the integrated MFLHP model, which includes facility hardening (MFLHP), with the *location-only-with-disruptions* (LOWD) problem, which does not include hardening. The LOWD problem locates facilities subject to a budget (b) in order to minimize the post-disruption radius. To solve the LOWD, a modification of the binary search algorithm for the MFLHP can be used; this modified algorithm uses the multi-level location set-covering model [4] as its auxiliary sub-problem (see Medal et al. [22]). The binary search algorithm for the LOWD can also be used to generate Pareto-efficient sets for the pre-disruption and post-disruption radii by using an algorithm similar to Algorithm 1 (see Medal et al. [22]). In the remainder of this section we compare the MFLHP to the LOWD method based on the pre-disruption radius vs. post-disruption radius Pareto sets that they produce.

Figure 5 displays the Pareto sets for the LOWD and MFLHP methods for the unweighted d49 dataset and the lon150 dataset with different values of the budget (b), number of interdictions (r), and hardening cost multiplier (H).

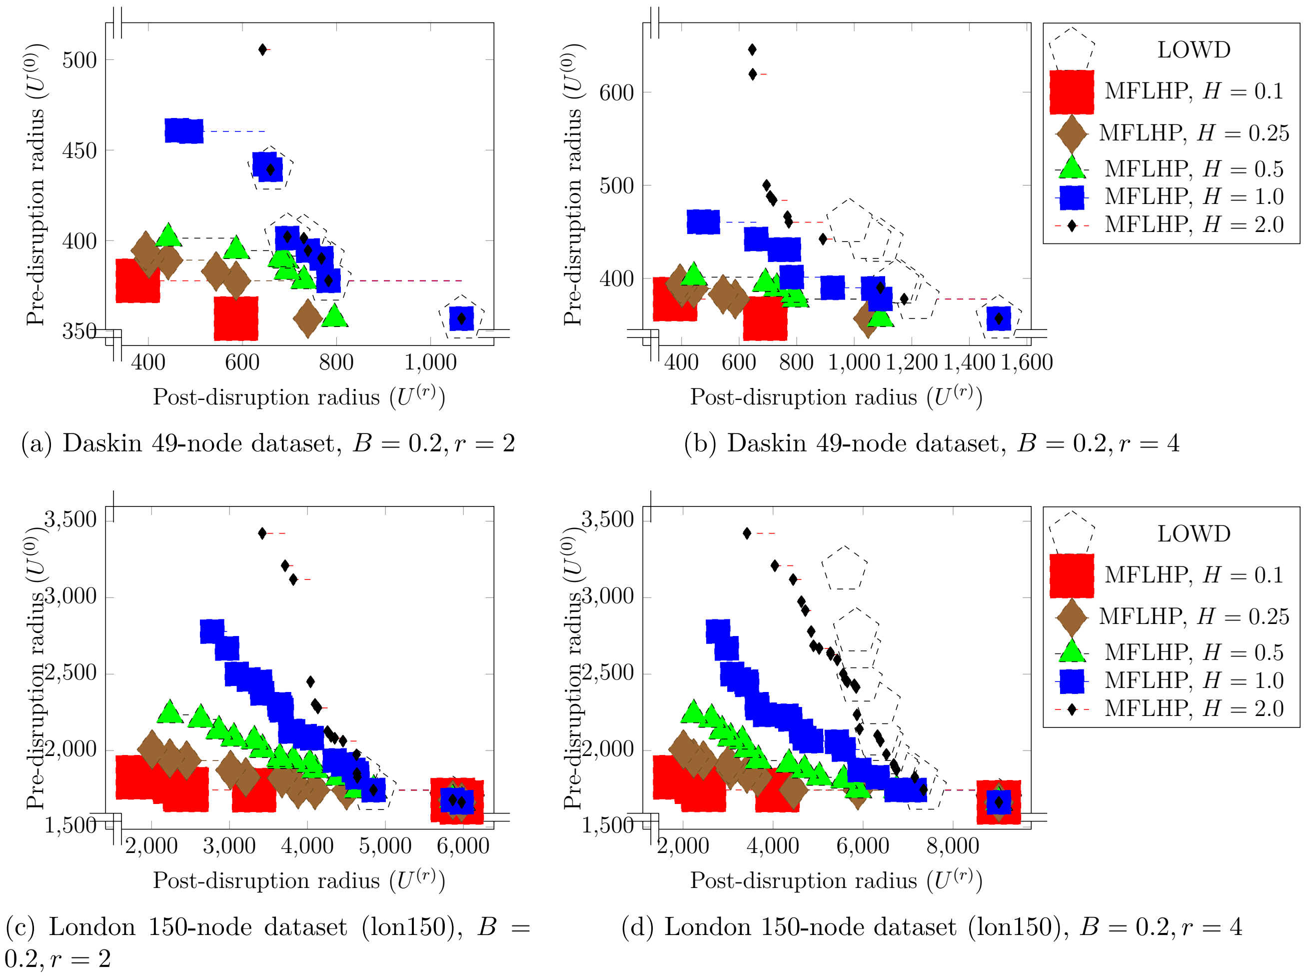

(a) Daskin 49-node dataset, $B = 0.2, r = 2$

(b) Daskin 49-node dataset, $B = 0.2, r = 4$

(c) London 150-node dataset (lon150), $B = 0.2, r = 2$

(d) London 150-node dataset (lon150), $B = 0.2, r = 4$

Figure 5: Pre-disruption radius vs. post-disruption radius for MFLHP and LOWD methods

Several interesting observations can be made from Figure 5.

1. The tradeoff ratio increases as the hardening cost multiplier (H) increases. This is because when hardening is cheap, the post-disruption radius can be reduced by hardening more facilities, which does not increase the pre-disruption radius. However, when the hardening cost is

high, the budget does not allow more facilities to be hardened; thus, facilities must be located differently, which increases the pre-disruption radius.

2. The Pareto sets become further apart as the post-disruption radius ($U^{(r)}$) decreases: when the post-disruption radius is large (because the pre-disruption radius is required to be small), hardening is not as beneficial, and the MFLHP solution will be similar to the LOWD solution. However, when the post-disruption radius is required to be small, hardening is more beneficial and the MFLHP will produce solutions with a lower post-disruption radius than the LOWD model.

3. As the hardening costs increase (especially when $H > 1$), the MFLHP Pareto sets approach the LOWD sets. This indicates that the integrated location-hardening model is most likely to give a significantly better solution than a location-only model when the hardening costs are less than the location costs.

Table 4 shows the number of facilities located and hardened for each of the Pareto points for Figures 5a and 5b; the Pareto points are listed right to left. The results support observation 3 above: when the hardening costs increase the number of hardened facilities decreases. In addition, more facilities are hardened under high hardening costs ($H \geq 1$) when there are more facility failures.

Table 4: Number of facilities located and hardened for each Pareto point in Figures 5a and 5b

r	H	Number located and (hardened)
2	0.1	11 (4); 10 (10)
2	0.25	11 (2); 10 (5); 10 (5); 9 (7); 9 (8); 9 (9)
2	0.5	11 (1); 10 (3); 11 (2); 10 (2); 10 (3); 9 (5); 8 (7)
2	1.0	11 (0); 11 (0); 11 (0); 10 (1); 10 (1); 10 (1); 10 (2); 7 (5); 6 (6)
2	2.0	11 (0); 11 (0); 11 (0); 12 (0); 12 (0); 12 (0); 12 (0); 11 (1)
4	0.1	11 (4); 10 (10)
4	0.25	11 (2); 10 (5); 10 (5); 9 (7); 9 (8); 9 (9)
4	0.5	11 (1); 10 (3); 10 (3); 9 (4); 9 (4); 8 (7)
4	1.0	11 (0); 11 (1); 9 (2); 9 (2); 8 (3); 9 (3); 8 (3); 8 (4); 7 (5); 6 (6)
4	2.0	11 (0); 10 (1); 10 (1); 8 (2); 7 (3); 6 (3); 7 (3); 7 (3); 5 (4); 5 (3); 4 (4)

6 Summary and Future Work

6.1 Summary

In this article we studied the integration of facility location and hardening decisions. In order to perform our analysis, we developed an efficient method to model and solve the integrated minimax facility location-hardening problem (MFLHP) model and used our approach to discover new insights about location and hardening. To model the problem, we first demonstrated that, because of the problem's structure, the three-stage location-interdiction-distribution problem can be modeled as a single-stage problem. We then formulated this single-stage problem as a mixed-integer program (MIP). Rather than solving the MIP directly, we decomposed it into set-cover-like auxiliary problems and used a binary search algorithm to solve it.

Experiments with the integrated MFLHP model provided interesting decision-making insights. The Pareto-efficient sets for the pre-disruption radius (PreDR) and post-disruption radius (PostDR) objectives showed that significant decreases in the post-disruption radius can be attained with small increases in the pre-disruption radius. Further experimental results showed that the average relative objective function increase for only optimizing the PreDR was higher than the relative increase for only optimizing the PostDR for all of the datasets used except one. Finally, the results in Section 5.2 indicated that an increased preference towards the PostDR objective should lead the decision-maker to harden more facilities when hardening cost is less than the location cost.

The experimental results also showed that the integrated method is attractive from a computational standpoint: the integrated model solved problems with up to 1323 nodes on a 12-core machine within 3389 seconds. The bi-objective solution algorithm is able to solve small- to medium-sized datasets within a reasonable amount of time: problem instances with up to 200 nodes were solved to optimality within 5.8 hours.

6.2 Future Work

Several assumptions in our model would be useful and interesting to relax. First, in our model a hardened facility cannot fail, and a facility is either hardened or not. In reality, a facility can never be completely immune to disruptions, and investing more protection resources into a facility makes it more resilient. Second, in our model the facilities are uncapacitated. This may be a good approximation in some contexts, but in others it may be important to model facility capacity. Third, we do not include the facility failure *duration* in our model. In reality, facilities may have different recovery times, which could affect the location and hardening decisions (see Losada et al. [20]).

Acknowledgements

We would like to thank two anonymous referees for their valuable comments, which helped improve this paper. We also wish to thank the US Department of Homeland Security (DHS) for sponsoring this work through the Mack-Blackwell Rural Transportation Centre at the University of Arkansas through Grant Number DHS-1101. However, the views expressed in this paper do not represent those of DHS, but rather those of the authors.

References

[1] Aksen, D., Aras, N., Piyade, N., 2013. A bilevel P-Median model for the planning and protection of critical facilities. Journal of Heuristics 19 (2), 1–26.

[2] Alp, O., Erkut, E., Drezner, Z., 2003. An efficient genetic algorithm for the P-Median problem. Annals of Operations Research 122 (1), 21–42.

[3] Balas, E., Ho, A., 1980. Set covering algorithms using cutting planes, heuristics, and subgradient optimization: a computational study. In: Padberg, M. W. (Ed.), Combinatorial Optimization. Vol. 12 of Mathematical Programming Studies. Elsevier North-Holland, The Netherlands, pp. 37–60.

[4] Church, R. L., Gerrard, R. A., 2003. The Multi-Level location set covering model. Geographical Analysis 35 (4), 277–290.

[5] Church, R. L., Scaparra, M. P., 2007. Protecting critical assets: the R-Interdiction median problem with fortification. Geographical Analysis 39 (2), 129–146.

[6] Cohon, J. L., 2004. Multiobjective programming and planning. Vol. 140. Courier Dover Publications.

[7] Cui, T., Ouyang, Y., Shen, Z.-J. M., 2011. Reliable facility location under the risk of disruptions. Operations Research 58 (4-Part-1), 998–1011.

[8] Daskin, M. S., 1995. Network and Discrete Location: Models, Algorithms, and Applications. John Wiley and Sons, Hoboken, New Jersey.

[9] Daskin, M. S., 2000. A new approach to solving the vertex p-center problem to optimality: Algorithm and computational results. Communications of the Operations Research Society of Japan 45 (9), 428–436.

[10] Drezner, Z., 1987. Heuristic solution methods for two location problems with unreliable facilities. Journal of The Operational Research Society 38 (6), 509–514.

[11] Elloumi, S., Labbé, M., Pochet, Y., 2004. A new formulation and resolution method for the P-Center problem. INFORMS Journal on Computing 16 (1), 84–94.

[12] Hakimi, S. L., 1965. Optimum distribution of switching centers in a communication network and some related graph theoretic problems. Operations Research 13 (3), 462–475.

[13] Hochbaum, D. S., Pathria, A., 1997. Generalized p-center problems: complexity results and approximation algorithms. European Journal of Operational Research 100 (3), 594–607.

[14] Hochbaum, D. S., Shmoys, D. B., 1986. A unified approach to approximation algorithms for bottleneck problems. Journal of the ACM 33 (3), 533–550.

[15] Johnson, D. S., Dec. 1974. Approximation algorithms for combinatorial problems. Journal of Computer and System Sciences 9 (3), 256–278.

[16] Kim, C.-R., Mar. 2012. Toyota: supply chain will be ready by autumn for next big quake. Reuters.

[17] Li, Q., Zeng, B., Savachkin, A., 2013. Reliable facility location design under disruptions. Computers and Operations Research 40 (4), 901–909.

[18] Lim, M., Daskin, M., Chopra, S., Bassamboo, A., 2010. A facility reliability problem: Formulation, properties, and algorithm. Naval Research Logistics 57 (1), 58–70.

[19] Lorena, L. A. N., Senne, E. L. F., 2004. A column generation approach to capacitated P-Median problems. Computers and Operations Research 31 (6), 863–876.

[20] Losada, C., Scaparra, M. P., O'Hanley, J. R., Mar. 2012. Optimizing system resilience: A facility protection model with recovery time. European Journal of Operational Research 217 (3), 519–530.

[21] Medal, H., 2012. Locating and protecting facilities subject to random disruptions and attacks. Ph.D. thesis, University of Arkansas, Fayetteville, Arkansas.

[22] Medal, H., Rainwater, C., Pohl, E., Rossetti, M., 2013. A bi-objective analysis of the R-All-neighbor P-Center problem. Tech. rep., Mississippi State University, Starkville, MS.

[23] Mladenovic, N., Labbe, M., Hansen, P., 2003. Solving the p-center problem with tabu search and variable neighborhood search. Networks 42 (1), 48–64.

[24] O'Hanley, J. R., Church, R., Gilless, J. K., Jan. 2007. Locating and protecting critical reserve sites to minimize expected and Worst-Case losses. Biological Conservation 134 (1), 130–141.

[25] O'Hanley, J. R., Church, R. L., 2011. Designing robust coverage networks to hedge against Worst-Case facility losses. European Journal of Operational Research 209 (1), 23–36.
[26] O'Hanley, J. R., Church, R. L., Gilless, 2007. The importance of *In Situ* site loss in nature reserve selection: Balancing notions of complementarity and robustness. Biological Conservation 2 (1), 170–180.
[27] O'Hanley, J. R., Paola Scaparra, M., Garcia, S., Mar. 2013. Probability chains: A general linearization technique for modeling reliability in facility location and related problems. European Journal of Operational Research 230 (1), 63–75.
[28] Reinelt, G., 1991. TSPLIB- a traveling salesman problem library. ORSA Journal of Computing 3 (4), 376–384.
[29] Salmeron, J., Wood, K., Baldick, R., 2009. Worst-Case Interdiction Analysis of Large-Scale Electric Power Grids. IEEE Transactions on Power Systems 24 (1), 96–104.
[30] Scaparra, M. P., Church, R. L., 2008. A bilevel Mixed-Integer program for critical infrastructure protection planning. Computers and Operations Research 35 (6), 1905–1923.
[31] Scaparra, M. P., Church, R. L., 2008. An exact solution approach for the interdiction median problem with fortification. European Journal of Operational Research 189 (1), 76–92.
[32] Scaparra, M. P., Pallottino, S., Scutella, M. G., 2004. Large-scale local search heuristics for the capacitated vertex p-center problem. Networks 43 (4), 241–55.
[33] Smith, J. C., 2011. Basic interdiction models. In: Cochran, J., Cox, L. A., Keskinocak, P., Kharoufeh, J. P., Smith, J. C. (Eds.), Wiley Encyclopedia of Operations Research and Management Science. Wiley, Hoboken, NJ.
[34] Snyder, L. V., Daskin, M. S., 2005. Reliability models for facility location: The expected failure cost case. Transportation Science 39 (3), 400–416.
[35] Suzuki, A., Drezner, Z., 1996. The p-center location problem in an area. Location Science 4 (1-2), 69–82.
[36] Swain, R. W., 1971. A decomposition algorithm for a class of facility location problems. Ph.D. thesis, Cornell University.

A Model for Minimizing Total Distance Subject to a Radius Constraint

The following model minimizes the pre-disruption total distance subject to the constraints that 1) the pre-disruption radius must be at most $U^{(0)}$ and 2) the post-disruption radius must be at most $U^{(r)}$. Redefine Y_i to equal 1 if a facility at i is hardened and zero otherwise. Let V_{ij} be a variable that is 1 if demand point j is assigned to the facility located at i as when a disruption does not occur, e.g., the pre-disruption assignment. Then, the model is as follows:

$$\min \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \phi_{ij} V_{ij} \tag{A.1a}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{I}} V_{ij} = 1 \quad \forall j \in \mathcal{J} \tag{A.1b}$$

$$V_{ij} \le X_i \quad \forall i \in \mathcal{I}, j \in \mathcal{J} \tag{A.1c}$$

$$\sum_{i:\phi_{ij}\le U^{(0)}} X_i \ge 1 \quad \forall j \in \mathcal{J}, \tag{A.1d}$$

$$(r+1) \sum_{i:\phi_{ij}\le U^{(r)}} Y_i + \sum_{i:\phi_{ij}\le U^{(r)}} X_i \ge r+1 \quad \forall j \in \mathcal{J} \tag{A.1e}$$

$$\sum_{i\in\mathcal{I}} f_i X_i + \sum_{i\in\mathcal{I}} g_i Y_i \le b \tag{A.1f}$$

$$Y_i \le X_i \quad \forall i \in \mathcal{I} \tag{A.1g}$$

$$X_i, Y_i \in \{0,1\} \quad \forall i \in \mathcal{I} \tag{A.1h}$$

$$V_{ij} \in \{0,1\} \quad \forall i \in \mathcal{I}, j \in \mathcal{J} \tag{A.1i}$$

The objective (A.1a) minimizes the sum of the distances between demand points and their assigned facilities. Every demand point must be assigned to one facility (A.1b) and that facility must be located (A.1c). Constraints (A.1d) and (A.1e) ensure that the pre-disruption and post-disruption radii are no more than $U^{(0)}$ and $U^{(r)}$, respectively. Finally, the model includes a budget constraint (A.1f), a constraint allowing a variable to be hardened only if it is located (A.1g), and variable-type constraints ((A.1h and (A.1i)).

When $r = 1$, the hardening variables, Y_i, are not needed and Constraints (A.1e) can be removed because they become identical to Constraints (A.1d).

B Computation Times for the Binary Search Algorithm

Another important consideration in deciding whether the integrated method is preferred over other methods is the computation time required to solve problem instances. Specifically, it is important to consider these two questions: 1) is the integrated method able to solve problem instances that would be of interest to decision makers? and 2) is the integrated method able to solve these problem instances within an amount of time that is satisfactory to the decision makers? The experimental results in this section show that the answer to both of these two questions is yes.

To obtain the run times analyzed in this section, we solved the integrated MFLHP model for several problem instances of the rl1323 dataset, which is motivated by a circuit board drilling problem [28]. The instances were obtained from the rl1323 dataset by varying the parameters b, g_i, and r. The instances were solved using the binary search algorithm described in Section 3.2.

In a preliminary set of experiments, the run time for the binary search algorithm was several orders of magnitude smaller than the run time for the MIP model. Thus, we do not report results for the MIP model.

Experimental results for the binary search algorithm computational performance are shown in Table 5. Each row contains the results for an instance of the rl1323 dataset, the largest dataset for which the binary search algorithm was able to solve all instances to optimality. Columns (5)–(11) show the initial percentage gap between the lower (5) and upper bound (6) and the final optimal solution, the initial total proportional gap (7), the number of times the auxiliary problem is solved to optimality (8), the presolve run time in which the initial bounds are obtained (9), the run time of the binary search algorithm (10), and the total run time (11).

Table 5: Run times (s) of binary search algorithm for rl1323 dataset

				Initial gaps				**Run time (s)**		
No. (1)	b (2)	H (3)	r (4)	LB (5)	UB (6)	LB+UB (7)	# solved to opt. (8)	Pre-solve (9)	Solve (10)	Total (11)
1	132	0.25	1	0.22	0.09	0.31	10	1110	283	1393
2	132	0.25	2	0.16	0.09	0.25	9	590	228	818
3	132	0.25	7	0.08	0.09	0.17	9	525	228	754
4	132	0.25	14	0.07	0.09	0.15	9	503	224	726
5	132	1	1	0.21	0.13	0.34	10	393	203	596
6	132	1	2	0.17	0.13	0.30	11	367	196	563
7	132	1	7	0.08	0.13	0.20	10	353	152	505
8	132	1	14	0.05	0.13	0.18	10	253	159	412
9	264	0.25	1	0.26	0.21	0.47	9	2502	887	3389
10	264	0.25	2	0.20	0.21	0.41	9	2199	781	2980
11	264	0.25	14	0.06	0.07	0.13	7	1436	600	2036
12	264	0.25	27	0.04	0.07	0.11	7	1307	580	1887
13	264	1	1	0.23	0.24	0.48	10	1110	374	1483
14	264	1	2	0.18	0.24	0.42	9	1292	366	1659
15	264	1	14	0.04	0.10	0.14	8	592	288	880
16	264	1	27	0.02	0.10	0.12	8	721	307	1028

Table 5 indicates that the binary search algorithm can solve large problems in a reasonable amount of time. The lower bound percentage gaps are usually good, ranging from 2 to 26%. The heuristic and initial bounds are effective in reducing the number of times the auxiliary problem is solved to optimality. In the worst case performance of the binary search algorithm, 21 auxiliary problems must be solved to optimality; however, these results show that the auxiliary problem needed to be solved to optimality between 7 and 11 times. Most of the problems solve in a reasonable amount of time, from 412 to 3389 seconds.

As Table 6 shows, the average run times over a set of problem instances for the remaining datasets used in this paper are mostly small, increasing with the number of nodes in the dataset. However, the lor818 dataset takes much longer than the others.

Table 6: Average run times (s) for binary search algorithm

Dataset	Average run time (s)
d49	0.18
d88	0.44
d150	0.60
sw55	0.23
lor100	0.52
lon150	0.65
lor200	1.37
lor300a	4.24
lor402a	12.25
lor818	166.8

C Computation Times for the Bi-Objective Algorithm

Table 7 contains run times for Algorithm 1 for various problem instances for the d49, sw55, d88, d150, and lon150 datasets. These run times include the time to optimize the total distance as a secondary objective for each point in the Pareto set. As the table shows, the main factors that influence run time are the number of nodes and the hardening cost multiplier.

Table 8 shows the run times for several larger problem instances. The run time is again influenced by the number of nodes and the hardening cost multiplier. The table also shows that the run time for these larger datasets is orders of magnitude larger than the run time of the smaller datasets shown in Table 7. Indeed, some of the instances did not complete within the 24 hour time limit.

Table 7: Run time for pre-disruption radius vs. post-disruption radius algorithm (Algorithm 1)

Dataset	Budget, b	Interdictions, r	Hardening cost multiplier, H	Run time
d49	10	2	0.25	6(s)
d49	10	2	1	14(s)
d49	15	3	0.25	7(s)
d49	15	3	1	16(s)
sw55	11	2	0.25	0.8(s)
sw55	11	2	1	2.9(s)
sw55	11	3	0.25	0.8(s)
sw55	11	3	1	2.4(s)
sw55	17	3	0.25	1.3(s)
sw55	17	3	1	2.4(s)
sw55	17	4	0.25	1.2(s)
sw55	17	4	1	2.3(s)
d88	18	3	0.25	10(s)
d88	18	3	1	16(s)
d88	18	4	0.25	8(s)
d88	18	4	1	22(s)
d88	27	5	0.25	10(s)
d88	27	5	1	25(s)
d88	27	6	0.25	10(s)
d88	27	6	1	23(s)
d150	30	5	0.25	14(s)
d150	30	5	1	61(s)
d150	30	6	0.25	12(s)
d150	30	6	1	42(s)
d150	45	7	0.25	12(s)
d150	45	7	1	28(s)
d150	45	9	0.25	11(s)
d150	45	9	1	61(s)
lon150	30	5	0.25	27(s)
lon150	30	5	1	68(s)
lon150	30	6	0.25	25(s)
lon150	30	6	1	71(s)
lon150	45	7	0.25	50(s)
lon150	45	7	1	57(s)
lon150	45	9	0.25	20(s)
lon150	45	9	1	48(s)

Table 8: Run time for pre-disruption radius vs. post-disruption radius algorithm (Algorithm 1)

Dataset	Budget, b	Interdictions, r	Hardening cost multiplier, H	Run time
lor200	40	6	0.25	534(s)
lor200	40	6	1	5.8(h)
lor200	40	8	0.25	451(s)
lor200	40	8	1	1.4(h)
lor200	60	9	0.25	165(s)
lor200	60	9	1	1.8(h)
lor200	60	12	0.25	123(s)
lor200	60	12	1	3131(s)
lor300	60	9	0.25	1.5(h)
lor300	60	9	1	>24(h)
lor300	60	12	0.25	1766(s)
lor300	60	12	1	23.1(h)
lor300	90	14	0.25	742(s)
lor300	90	14	1	8.8(h)
lor300	90	18	0.25	632(s)
lor300	90	18	1	1.9(h)
lor402	81	13	0.25	7.8(h)
lor402	81	13	1	389(s)
lor402	81	17	0.25	2.5(h)
lor402	81	17	1	>24(h)
lor402	121	19	0.25	101(s)
lor402	121	19	1	2(h)
lor402	121	25	0.25	1850(s)
lor402	121	25	1	12.2(h)
lor818	164	25	0.25	>24(h)
lor818	164	25	1	>24(h)
lor818	164	33	0.25	>24(h)
lor818	164	33	1	>24(h)
lor818	246	37	0.25	1.1(h)
lor818	246	37	1	1067(s)
lor818	246	50	0.25	4.7(h)
lor818	246	50	1	>24(h)