Improving supply system reliability against random disruptions: Strategic protection investment

Stefano Starita[a] and Maria Paola Scaparra[b]

[a]Sasin School of Management, Chulalongkorn University, Bangkok 10330, Thailand; [b]Centre for Logistics and Heuristic Optimisation (CLHO), Kent Business School, University of Kent, Canterbury CT2 7FS, UK

ARTICLE HISTORY
Compiled May 18, 2021

ABSTRACT
Supply chains, as vital systems to the well-being of countries and economies, require systematic approaches to reduce their vulnerability. In this paper, we propose a non linear optimisation model to determine an effective distribution of protective resources among facilities in service and supply systems so as to reduce the probability of failure to which facilities are exposed in case of external disruptions. The failure probability of protected assets depends on the level of protection investments and the ultimate goal is to minimize the expected facility-customer transport or travel costs to provide goods and services. A linear version of the model is obtained by exploiting a specialized network flow structure. Furthermore, an efficient GRASP solution algorithm is developed to benchmark the linearised model and resolve numerical difficulties. The applicability of the proposed model is demonstrated using the Toronto hospital network. Protection measures within this context correspond to capacity expansion investments and reduce the likelihood that hospitals are unable to satisfy patient demand during periods of high hospitalization (e.g., during a pandemic). Managerial insights on the protection resource distribution are discussed and a comparison between probabilistic and worst-case disruptions is provided.

KEYWORDS
Disruption management; resource allocation; fortification; healthcare; Covid-19

1. Introduction

In the wake of the 21st century's terrorist attacks and numerous natural disasters, the need of protecting supply chains and critical infrastructure systems has been widely recognized in the scientific community, among practitioners and government agencies alike. Risks to critical lifeline networks result from a complex mix of naturally occurring and human threats, including earthquakes, hurricanes, flooding, terrorist attacks, and political unrest. More recently, the ongoing Covid-19 pandemic is disrupting healthcare systems worldwide. Shortages of intensive care units (ICUs) and medical equipment are taking a severe toll, dramatically highlighting the necessity of putting in place mitigation strategies to address disruptive scenarios. Other less catastrophic contingencies, such as component and technological failures or temporary shortages of capacity, may still have enduring severe effects on a supply chain operation continuity.

CONTACT Stefano Starita. Email: stefano.starita@sasin.edu

With so many possible threats and in view of the intrinsic fragility of today complex supply-chains, proactive measures must be taken to lessen the risks and minimize the impact that disruptions to the provision of essential services (such as energy, food, water, transport, communications, health and banking) may cause on businesses and communities. Vulnerability reduction can be attained by protecting key critical components and assets. Protecting and ensuring the resiliency of critical infrastructures has become an imperative mandate in the agenda of governments and organizations and, in the last decade, government and private sector expenditures devoted to security and protection efforts have increased significantly (European Commission, 2019).

Different protection strategies can be implemented to increase systems resiliency. These depend on the nature of the hazard they address and the type of infrastructure under consideration. As an example, site fortification, monitoring devices, intelligence and surveillance may be used to protect nuclear and chemical facilities from terrorist attacks. In a broader sense, protection measures include capacity expansion or relocation of critical facilities. For example, Phua et al. (2020) recommend increasing capacity of health infrastructure, supply and staff as the key to deal with the Covid-19 pandemic. Several examples can be found of countries increasing ICU beds by rationalising hospital resources and spaces (BBC, 2020b) as well as building brand new temporary and permanent facilities (BBC, 2020a). Finally, supplier decentralization and diversification strategies can be used to protect supply-chains. This also fits the Covid-19 discourse, with companies experiencing sourcing issues and being required to proactively find alternatives to mitigate economic disruptions (Avetta, 2020).

This article focuses on the protection of service or supply facilities in median systems, where service is provided to customers by their closest facility and efficiency is measured in terms of expected travel costs or distances. This kind of efficiency measure has been largely used in application settings where maximizing consumer access to supply centers is the primary objective, as in the case of facilities providing essential services or goods (e.g., fire stations, hospitals, test centers, vaccines). Throughout the remainder of paper, the terms protection and fortification are used interchangeably. We assume that each facility can fail independently with a given probability and that customers have prior information on the operational status of the facilities. Given a budget constraint, the objective is to allocate protective resources so as to minimize the expected demand-weighted transportation costs (or travel times or distances). Our modeling approach can handle site-specific failure probabilities, the case where fortification only reduces, but does not eliminate, the probability of failure, and the case where the failure probability reduction depends on the level of protection investments. The resulting formulation is non linear. A custom GRASP algorithm is introduced to solve the model. The heuristic is also used to highlight numerical difficulties and fine tune an efficient linearisation approach which exploits a network flow structure to replace the multiplicative terms in the objective function of the non-linear model.

This paper contributes to the disruption management body of research by: introducing the Probabilistic Median Fortification Problem (PMFP) as a non-linear, mixed-integer model and providing a probability-chain linearisation; proposing a GRASP procedure exploiting efficient data structures, which can be used to benchmark the linearisation and redress numerical issues typical of probability-chain based formulations; demonstrating how the model can be applied in practice to a hospital network so as to inform decisions about hospital capacity expansion during periods of high hospitalization demand (e.g., during a pandemic); proposing a hybrid protection approach which produces robust solutions against both probabilistic and worst-case disruptions.

The reminder of this paper is organized as follows. A comprehensive review of

related research on supply chain disruption management is provided in Section 2. In Section 3, the problem is formally introduced and formulated as a non-linear model. A linear formulation is then presented in Section 4. The GRASP heuristic is described in Section 5. Some computational experiments are presented in Section 6, whereas the application of PMFP to a real case study is discussed in Section 7. Some conclusive remarks are provided in Section 8.

2. Background and related work

There is a vast body of OR publications studying problems related to supply systems and their design. Since the seminal works by Hakimi (1964) and Hakimi (1965), the p-median problem has been widely used to study the design of demand-weighted supply networks. It has been customised and applied to a wide range of contexts, for example healthcare (Daskin and Dean, 2005), waste management (Dantrakul et al., 2014) and air transport (Sasaki et al., 1999). Moreover, given its combinatorial nature, extensive efforts in devising efficient solution algorithms have been made (Mladenović et al., 2007). The reader is referred to Melo et al. (2009) and Puerto et al. (2018) for comprehensive reviews on the topic. Given the focus of this work, the following review is restricted to papers that consider disruption of service and supply systems. This body of research has increased substantially since the beginning of the millennium to denote a growing involvement of the OR community in tackling emergency logistics challenges. Emergency management efforts can be broadly casted as pre-event and post-event actions (Tufekci and Wallace, 1998). Pre-event actions include mitigation and preparedness operations which are implemented before a disaster occurs. Post-event actions incorporate response and recovery operations which take place in the aftermath of a disaster (e.g., relief distribution, evacuation and casualty transportation). Overall the four programmatic phases (mitigation, preparedness, response, and recovery) constitute the disaster operations management lifecycle (Altay and Green III, 2006). Examples of decision tools to support humanitarian logistics considering different disruption stages can be found in Çelik et al. (2012). As facility protection pertains to the mitigation stage, we only briefly discuss papers which focus on mitigation measures. An overview of the academic OR literature on post-event operations can be found in recent surveys (Çelik, 2016; Esposito Amideo et al., 2019).

2.1. Designing Reliable Supply Systems

The impact of catastrophic events and random failures can be mitigated by designing supply chains and infrastructure networks which are inherently reliable to external disruptions. Increased reliability clearly comes with a cost. However, ad-hoc optimization models can be devised to identify system configurations which are both cost-efficient and reliable (Snyder and Daskin, 2005). Building upon the seminal papers by Daskin (1983) and Drezner (1987), who consider the location of unreliable facilities in coverage and median systems respectively, a number of reliability location models have appeared in the literature in recent years (Alcaraz et al., 2012; Berman and Krass, 2011; Berman et al., 2007, 2009; Cui et al., 2010; Lei and Tong, 2013; Li and Ouyang, 2010; O'Hanley et al., 2013; Ozkan et al., 2020; Shen et al., 2011). All these models aim at identifying the optimal location of facilities which can fail with some probability so as to minimize the expected costs in normal and failure scenarios. Lim et al. (2010), Li et al. (2013) and Li and Savachkin (2013) consider the problem of locating both reliable and

unreliable facilities. Peng et al. (2011) extend previous location reliability models by using a p-robustness measure to design reliable multi-echelon supply chain networks. Jabbarzadeh et al. (2018) incorporate sustainability issues into a classic reliable design problem. In a slightly different context, O'Hanley and Church (2011) consider the design of robust coverage-type service networks to hedge against worst-case facility losses. Aksen and Aras (2012) propose bilevel fixed charge facility location problems to counter antagonist attacks. Yun et al. (2015) propose a modeling framework for reliable supply chain network design when customers are unaware of the disruption status of a facility. Zarrinpoor et al. (2017) study a location problem to design a robust hierarchical network for health service. Yu et al. (2017) propose a risk-adverse approach to the problem by incorporating conditional value-at-risk measures in the design process. Aslan and Çelik (2019) propose a model to design a three-echelons humanitarian response network and its pre-disaster inventory under probabilistic disruptions, while accounting for road restoration and relief transportation as well.

2.2. Protecting Supply Systems from Worst-Case Disruptions

Although the impact of facility disruptions can be mitigated in the initial design, redesigning an entire system may be infeasible or prohibitively expensive. Additionally, relocating facilities may lead to more robust supply chains in terms of efficiency and transportation costs but may not eliminate other risks due to external disruptions, such as loss of human lives, property damage and environmental damage. As an alternative, the reliability of existing infrastructure can be enhanced through efficient investments in protection measures. Different modeling frameworks have been adopted in the literature to model protection resource allocation to counter both strategic and probabilistic risks (Golany et al., 2009). The majority of protection models consider strategic risks (i.e., intentional attacks) and are formulated as bilevel programming models where the upper level problem of optimally distributing protection resources embeds a lower-level optimization problem which endogenously generates worst-case scenario losses (Brown et al., 2006). Examples of these defender-attacker models can be found in Qiao et al. (2007) for water supply systems, Holmgren et al. (2007) for electric power grids, Cappanera and Scaparra (2011), Starita and Scaparra (2016), Starita and Scaparra (2017) for transportation networks. Several different works consider facility protection investments against worst-case losses in median supply systems (Church and Scaparra, 2007; Liberatore et al., 2012, 2011; Losada et al., 2012b; Scaparra and Church, 2012). These models have introduced increasingly complex issues such as facility capacities, uncertainty in the number of losses, correlation of disruptive events, disaster propagation effects, recovery times, disruption frequency and investment levels. Details of defender-attacker models for other systems, such as hub networks (Ramamoorthy et al., 2018) and decentralised systems (Zhang et al., 2017), can be found in Scaparra and Church (2019).

2.3. Protecting Supply Systems from Probabilistic Disruptions

In the last decade, researchers started focusing on protection models for systems facing probabilistic disruptions. The early OR papers discussing protection models in a probabilistic environment consider investment decisions to strengthen infrastructure links in transportation networks. As an example, Liu et al. (2009) proposes a two-stage stochastic programming problem to minimize the total expected physical and

social losses caused by potential disruptions. Fan and Liu (2010) analyse the problem of allocating limited retrofit resources to highway bridges so as to reduce structural losses and travel delays. Both these models assume that the probability of damage of protected components is reduced to zero, to avoid the computational difficulties associated with modeling non-zero reduction in the failure likelihood. Peeta et al. (2010) consider the problem where the survival probabilities of the protected components depend on the level of investment. The model aims at maximizing the post-disaster connectivity and minimizing traversal costs between origin and destination nodes in highway networks. Sawik (2013) introduce a model where binary fortification of suppliers and emergency inventory are used to mitigate the impact of disruption scenarios. Qin et al. (2013) introduce a two-stage stochastic fortification model where the planner trades off between fortifying facilities and placing emergency inventory resources. Zhang et al. (2014) propose a protection model for p-median systems which accounts for worst-case and random attacks. They assume that protected facilities are immune to disruptions. Medal et al. (2016) relax binary assumptions by assuming that different levels of investments can be made to decrease the disruption likelihood of facilities. Their model relies on pre-defining a set of disruptive scenarios and aims at maximising the expected utility gained from successful shipments to costumers.

2.4. Combining Design with Protection

A recent line of research has also investigated the combination of protection and design models. For example, Li et al. (2013) consider facilities facing probabilistic disruptions and model decisions on locations, main and back-up assignments and binary fortifications. This work is further extended in Li and Savachkin (2016) by considering continuous fortification investments. Each open facility has a failure probability that can be reduced through investments. However, no simultaneous disruptions are considered. Jabbarzadeh et al. (2016) propose a robust formulation for the capatted fixed-charge location problem where the disruption probability of unreliable facilities decreases with the level of protection investments. Unlike our model, the authors assume the existence of completely reliable facilities which serve the customers when an unreliable facility fails. Rohaninejad et al. (2018) introduce a scenario-based formulation to design multi-echelon supply chains with facility hardening options.

Overall, the current literature on supply chain protection has mostly focused on worst-case disruptions and network design problems. Stochastic disruptions have mainly been investigated for the fixed-charge location problem, assuming either individual disruptions or requiring to pre-define a set of disruption scenarios. This paper aims to widen the body of research on probabilistic disruptions on pre-existing supply chains. A novel formulation is introduced, modelling simultaneous disruptions and incorporating investment level decisions aimed at reducing the probability of failure of facilities in p-median systems. Unlike assumptions made in previous models, this formulation does not rely on pre-defining a set of disruptive scenarios and computes the expected transportation costs considering the probabilities of failure of all facilities, without assuming the presence of reliable facilities which never fail.

3. The p-Median Fortification Problem (PMFP)

The p-median fortification problem assumes the existence of a supply system with p operating facilities, hence no location decisions are considered. Each facility has unlimited capacity and the system users receive service from their nearest facility. Let J, indexed by j, be the set of p existing facilities, and I, indexed by i, the set of N customers. Each customer $i \in I$ has a demand of h_i units and each unit shipped to customer i from facility j incurs a transportation cost of d_{ij}. We assume that transportation costs are directly related to distances or travel times. Therefore, in the following, we will refer to d_{ij} as cost, distance or travel time interchangeably. Also, we denote by i_k the kth closest facility to customer i (ties can be broken arbitrarily for equidistant facilities).

Each facility j may fail or be disrupted with a fixed probability q_j. Failures are assumed to be independent. A disrupted facility becomes inoperable, so that the customers served by it must be reassigned to their closest (cheapest) non-disrupted facility. We assume that each customer i accepts service only from facilities within an admissible area K_i. For example, K_i can be defined as the set of the K closest facilities to customer i, with $|K| \leq p$, or as the set of facilities within a predefined radius ρ_i, i.e. $K_i = \{j \in J | d_{ij} \leq \rho_i\}$. If all the facilities in K_i fail, the demand of customer i is lost and the system incurs a penalty, or lost-sale cost, θ_i per unit of demand. In some application settings, customers may be willing to receive service by any facility, independently from the cost involved. This can be easily captured in the model by setting $K = p$ if the admissible area is defined in terms of number of closest facilities, or by setting ρ_i as large as the cost of serving any customer from any location in the network, if the admissible area is defined by a maximum radius. If customers can receive service by any facility, the penalty is incurred only when all the facilities fail.

A limited amount of protection resources, also referred to as protection budget, b, is available for protecting (or hardening or fortifying) some of the facilities. Different amounts of protective resources can be invested on a single facility. The failure probability after protection depends on the level of protection investments. Let λ be the number of possible investment levels and c_{jl} be the cost for protecting facility j at level l, with $c_{jl} > c_{jl-1}$ for each $j \in J$ and each $l = 0, ..., \lambda$. A security investment at level l reduces j's failure probability from q_j to q_{jl}, with $q_{jl} < q_{jl-1}$. We assume $q_{j0} = q_j$ and $c_{j0} = 0$, i.e., a protection at level zero is equivalent to no protection. In the following, we denote by $\bar{q}_j = 1 - q_j$ the probability that facility j is operational if not protected, and by $\bar{q}_{jl} = 1 - q_{jl}$ the probability that facility j is operational if protected at level l.

The objective of the PMFP is to identify the protection investment strategy that minimizes the expected transportation costs. Protection decisions are modeled through the following set of decision variables:

$$z_{jl} = \begin{cases} 1, & \text{if facility } j \text{ is protected at level } l \\ 0, & \text{otherwise} \end{cases}$$

The expected cost for one unit of demand of customer i can be computed as a function

of the vector of protection variables z as follows:

$$\begin{aligned} d_i(z) &= \sum_{m=1}^{|K_i|} \left[\prod_{k=1}^{m-1} \sum_{l=0}^{L} q_{i_k l} z_{i_k l} \right] \left[\sum_{l=0}^{L} \bar{q}_{i_m l} z_{i_m l} \right] d_{i i_m} \\ &+ \left[\prod_{k=1}^{|K_i|} \sum_{l=0}^{L} q_{i_k l} z_{i_k l} \right] \theta_i \end{aligned} \tag{1}$$

The first term in the right-hand side of expression (1) represents the expected cost for serving customer i by the facilities in the service area K_i. Specifically, customer i is served by its mth closest facility at a cost d_{ii_m} only if all the closest facilities i_k, with $k < m$, are not operational, which occurs with probability $\prod_{k=1}^{m-1} \sum_{l=0}^{L} q_{i_k l} z_{i_k l}$, but facility i_m is operational, which occurs with probability $\sum_{l=0}^{L} \bar{q}_{i_m l} z_{i_m l}$. The second term in the right-hand side of the expression represents the expected penalty which is incurred when all the facilities in the set K_i fails. The probability of this event is $\prod_{k=1}^{|K_i|} \sum_{l=0}^{L} q_{i_k l} z_{i_k l}$.

The PMFP can then be stated as the following combinatorial optimization problem:

$$\min \quad D = \sum_{i \in I} h_i d_i(z) \tag{2}$$

$$\text{s.t.} \quad \sum_{j \in J} \sum_{l=0}^{\lambda} c_{jl} z_{jl} \leq b \tag{3}$$

$$\sum_{l=0}^{L} z_{jl} = 1, \quad \forall j \in J \tag{4}$$

$$z_{jl} \in \{0, 1\}, \quad \forall j \in J, l = 0, ..., \lambda \tag{5}$$

The objective function (2) minimizes the demand weighted sum of expected transportation costs. Constraint (3) specifies that the total protection investments cannot exceed the available budget. Constraints (4) state that only one level of protection can be chosen for a facility. Finally, constraints (5) represent the integrality requirements of the decision variables.

4. A linear formulation for PMFP

It is obvious that model (2)-(5) is not amenable to solution by general purpose optimization solvers due to the complex, highly non-linear structure of its objective function. This section presents a linearization of PMFP, based on the use of a specialized network flow structure to evaluate the compound probability terms in the expected costs (1) of the objective function (2). This type of linearization technique was first used by Morton et al. (2007) in the context of nuclear smuggling interdiction. Losada et al. (2012a) used a similar network structure to solve a stochastic interdiction problem in median systems. O'Hanley et al. (2013) and Tran et al. (2016) proposed different network representations, referred to as *probability chain* and *probability lattice*.

In our formulation, we use a specialized network flow structure which, for each

customer i, computes the correct probability that i is allocated to each facility j while taking into account that the probability of failure of each facility can be reduced by different amounts, depending on the level of protection investment.

To reformulate PMFP in linear form, referred to as PMFP-L, we define the following additional variables:

y_{ijl} : when facility j is protected at level l, this represents the probability that all the facilities which are closer to customer i than j are not operational. If facility j is protected at any other level, y_{ijl} is zero.

u_i: this is the probability that all the facilities in the service region of customer i are not operational.

The variables y_{ijl} are defined for each customer i, each facility $j \in K_i$, and each protection level l. Note that when $j = i_1$ (i.e., j is the closest facility to customer i), the variable y_{ii_1l} is one if i_1 is protected at level l, zero otherwise.
PMFP-L is formulated as follows:

$$\min \quad \sum_{i \in I} h_i \left[\sum_{k=1}^{|K_i|} d_{ii_k} \sum_{l=0}^{\lambda} \bar{q}_{i_k l} y_{ii_k l} + \theta_i u_i \right] \tag{6}$$

$$\text{s.t.} \quad \sum_{l=0}^{L} y_{ii_k l} = 1 \qquad \forall i \in I, k = 1 \tag{7}$$

$$\sum_{l=0}^{L} q_{i_k l} y_{ii_k l} = \sum_{l=0}^{L} y_{ii_{k+1} l} \qquad \forall i \in I, k = 1, \ldots, |K_i| - 1 \tag{8}$$

$$\sum_{l=0}^{L} q_{i_k l} y_{ii_k l} = u_i \qquad \forall i \in I, k = |K_i| \tag{9}$$

$$y_{ijl} \leq z_{jl} \qquad \forall i \in I, k = 1, ..., |K_i|, l = 0, ..., \lambda \tag{10}$$

$$(3) - (5)$$

$$y_{ijl} \in [0, 1] \qquad \forall i \in I, k = 1, ..., |K_i|, l = 0, ..., \lambda \tag{11}$$

$$u_i \in [0, 1] \qquad \forall i \in I \tag{12}$$

The objective function (6) minimizes the expected transportation cost defined in terms of the probability variables y_{ijl} and u_i. Constraints (7)-(9) enforce these variables to take the correct probability value. These can be seen as flow balance constraints which propagate the failure probabilities from facility i_k to facility i_{k+1} according to the protections that are made. Constraints (10) force the variables y_{ijl} to be zero if facility j is not protected at level l. Finally, constraints (11)-(12) force the probability variables to be continuous variables in the interval $[0, 1]$.

If customers can receive service from any facility in the system, this model has $p\lambda$ integer variables, $p\lambda n + n$ continuous variables and $p\lambda n + np + n + p + 1$ constraints. It is noteworthy that this linear construct requires the same number of integer variables as the non-linear formulation, i.e., all the linearization variables are continuous.

5. A GRASP procedure for PMFP

GRASP is an iterative, two-phase metaheuristic originally developed by Feo and Resende (1995). The first phase consists of a randomized greedy construction algorithm which produces a feasible solution to the problem. In the second phase, a local search is applied to the random greedy solution to obtain a local optimum in its neighborhood (Festa and Resende, 2011). As pointed out in Mannino et al. (2007), there are two tasks which affect the performance of a GRASP algorithm: the evaluation of the greedy function and the optimization over the neighborhood. In the following, we describe how the computational complexity of these tasks can be reduced by using appropriate data structures.

5.1. Construction phase

The greedy randomized phase for PMFP builds a solution starting from the trivial solution $z_{j0} = 1, \forall j \in J$ (i.e., all the facilities are unprotected). At each iteration, an unprotected facility j is protected at some level l. The pair (j, l) is selected at random from a restricted candidate list (RCL). The construction of the RCL is guided by an adaptive greedy function which computes the unit reduction in the objective function (2) obtained by protecting an unprotected facility j at some level l. Namely, let z be the partial solution at the current iteration. For each unprotected facility j and for each protection level $l = 1, ..., \lambda$, the greedy function is given by:

$$b_{jl} = \frac{\Delta_{jl}}{c_{jl}} = \frac{\sum_{i \in I} h_i (d_i(z) - d_i(z'))}{c_{jl}} \tag{13}$$

where $z' = z$ except for the component z_{jl} (i.e., $z_{jl} = 0$ and $z'_{jl} = 1$). Note that the evaluation of the greedy function requires computing the new expected transportation cost $d_i(z')$ for each customer i according to equation (1). Therefore, the evaluation of (13) requires $O(np^2\lambda)$ time.

A pair (j, l) is considered and its greedy function b_{jl} evaluated only if the cost of protecting facility j at level l, c_{jl}, does not exceed the protection budget still available at that iteration. A pair (j, l) is inserted in the RCL if its greedy function value is in the range $[b_{max} - \alpha(b_{max} - b_{min}), b_{max}]$, where b_{max} and b_{min} are the greedy function values of the best and the worst insertion at that construction iteration. The restricted candidate parameter α varies between 0 and 1. Values of α close to 0 increase the greediness of the approach, whereas values close to 1 increase the amount of randomness. In our computational experiments we used a self-calibrating mechanism for automating the choice of α known as Reactive Grasp (Prais and Ribeiro, 2000) and implemented as in Scaparra and Church (2005). The construction phase terminates when the residual protection budget is insufficient to protect any unprotected facility or all facilities have already been protected at some level $l \neq 0$.

5.2. Local search phase

In the local search phase, we simply explore a *swap* neighborhood. Namely, the neighbors of the current solution are all the solutions that can be obtained by removing the protection of a protected facility and protecting un unprotected facility at some level

$l \neq 0$. The neighborhood only includes feasible solutions with respect to the budget constraint. As the evaluation of each move profitability is computationally expensive, we use a first improvement strategy during the search. Further details on the actual implementation of swap moves are provided in the next section.

5.3. Efficient evaluation of move profitability

To reduce the time for evaluating every possible insertion in the RCL at each step of the constructive phase and every swap in the local optimization phase, we maintain an expected distance matrix $W = [w_{ik}]$, where each entry stores the expected transportation cost for customer i when only its first k closest facilities are considered. For each customer i and each $k = 1, ..., |K_i|$, the corresponding entry in the matrix is initialized as follows:

$$w_{ik} = \sum_{m=1}^{k} \prod_{n=1}^{m-1} q_{i_n} \bar{q}_{i_m} d_{ii_m} \tag{14}$$

Additionally, when $k = |K_i| + 1$, the matrix entry for customer i is:

$$w_{ik} = \sum_{m=1}^{|K_i|} \prod_{n=1}^{m-1} q_{i_n} \bar{q}_{i_m} d_{ii_m} + \prod_{k=1}^{|K_i|} q_{i_k} \theta_i. \tag{15}$$

By using this matrix, the expected transportation cost for each customer i deriving from adding or dropping the protection of a facility can be computed in constant time as follows.

Let z be the current solution with $z_{h0} = 1$, and z' the solution obtained from z by protecting facility h at level l for some $l \in \{1, ..., \lambda\}$, i.e., $z' = z$ except for $z'_{h0} = 0$ and $z'_{hl} = 1$. Assume that facility h is the mth closest facility to customer i (i.e., $h = i_m$). The expected transportation cost for customer i, $d_i(z')$, after the protection of facility h at level l is simply:

$$d_i(z') = w_{im-1} + (w_{im} - w_{im-1})\bar{q}_{hl}/\bar{q}_h + (d_i(z) - w_{im})q_{hl}/q_h. \tag{16}$$

Similarly, let z' be the solution obtained from z by dropping the protection of facility u at level t, with $t \in \{1, ..., \lambda\}$. Let u be the mth closest facility to customer i (i.e., $u = i_m$). The expected transportation cost for customer i, $d_i(z')$, after the removal of the protection of facility u at level t is:

$$d_i(z') = w_{im-1} + (w_{im} - w_{im-1})\bar{q}_u/\bar{q}_{ut} + (d_i(z) - w_{im})q_u/q_{ut}. \tag{17}$$

For both add and drop moves, the overall objective function (2) of the new solution z' can be evaluated in $O(N)$ time by summing the demand weighted $d_i(z')$ over all the customers, i.e.:

$$D(z') = \sum_{i \in I} h_i d_i(z') \tag{18}$$

When a move is actually performed, the matrix W must be updated. This can be done as follows. Assume than an add move sets $z'_{hl} = 1$ ($l \neq 0$) and h is the m-th

closest facility to customer i. Then each new matrix entry, w'_{ik}, is obtained from the previous entry according to the formula:

$$w'_{ik} = \begin{cases} w_{ik}, & \text{if } k < m, \\ (w_{ik} - w_{ik-1})\bar{q}_{hl}/\bar{q}_h + w'_{ik-1}, & \text{if } k = m, \\ (w_{ik} - w_{ik-1})q_{hl}/q_h + w'_{ik-1}, & \text{if } k > m. \end{cases} \tag{19}$$

Similarly, each new matrix entry, w'_{ik}, after a drop move which sets $z'_{ut} = 0$ ($t \neq 0$), where u is the m-th closest facility to customer i, is:

$$w'_{ik} = \begin{cases} w_{ik}, & \text{if } k < m, \\ (w_{ik} - w_{ik-1})\bar{q}_u/\bar{q}_{ut} + w'_{ik-1}, & \text{if } k = m, \\ (w_{ik} - w_{ik-1})q_u/q_{ut} + w'_{ik-1}, & \text{if } k > m. \end{cases} \tag{20}$$

The update of the matrix W can hence be performed in $O(NP)$ time.

In the construction phase of the GRASP procedure, formulae (16) and (18) are used to evaluate the greedy function for each unprotected facility h and level l. When a pair (h, l) is then randomly selected from the RCL, formula (19) is used to update the expected cost matrix W.

In the local search phase, each swap move is evaluated as the combination of a drop and an add move. A scheme for the neighborhood optimization phase is given below. Let z be the solution obtained at the end of the construction phase and rb the residual protection budget. L_{in} and L_{out} are respectively the sets of protected and unprotected facilities in the current solution.

Procedure LocalSearch(z)
0. Set $z' = z$, $L_{in} = \{j \in J | z_{j0} = 0\}$ and $L_{out} = \{j \in J | z_{j0} = 1\}$;
1. **While** z not locally optimal **do**
2. Select u from L_{in}. Let t be the protection level such that $z_{ut} = 1$.
Set $z'_{ut} = 0$, $z'_{u0} = 1$ and update matrix W according to (20) (*drop move*)
3. **For each** h in L_{out} **do**
4. Identify the largest protection level l such that $rb + c_{ut} - c_{hl} \geq 0$.
5. Set $z'_{hl} = 1$, $z'_{h0} = 0$ and compute $D(z')$ by using formulae (16) and (18) (*add move*).
6. If $D(z') < D(z)$, (*improving swap detected*)
7. $rb := rb - c_{ut} + c_{hl}$, $L_{in} := L_{in} \setminus \{u\} \cup \{h\}$, $L_{out} := L_{out} \setminus \{h\} \cup \{u\}$
Set $z = z'$, update matrix W according to (19). Go to step 2.
8. Else $z'_{hl} = 0$, $z'_{h0} = 1$ (*restore previous solution and try next h*)
9. **End for**
10. Set $z'_{ut} = 1$, $z'_{u0} = 0$ and restore matrix W according to formula (19) with $h = u$ and $l = t$.
Go to step 1. (*No improving swap found for facility u. Undo drop and try next u*).
11. **End while**
12. **Return** z

Note that in the implementation of the local search procedure above, the facilities in the lists L_{in} and L_{out} are scanned in a circular fashion and a locally optimal solution is found after a full pass of the two lists without discovering a cost-decreasing move.

6. Computational results

Computational tests are conducted on three different data sets that have been frequently used in related literature: the US dataset with 263 demand points (Snyder and Daskin, 2005), the UK dataset with 250 demand points (Liberatore and Scaparra, 2011) and the London Ontario (LN) dataset with 150 demand points (Goodchild and Noronha, 1983). The GRASP approach is implemented in C++ and compiled using Microsoft Visual Studio 2017. The linear formulation is solved using CPLEX version 12.5 Callable Libraries. Experiments are performed on the same AMD Ryzen 7-1700 CPU @ 3.0 GHz workstation with 8GB of RAM, running on Microsoft Windows 10.

With the aim of providing a thorough analysis, we consider three cases to model the probability of failures. Specifically, we look at economic concepts such as production function and return of scale to formulate how failure probabilities depend on protection investments. In economy, a return to scale is said to be increasing (decreasing) when the output grows more (less) than proportionally with the input. With this in mind, we introduce a convex function ($q_{jl} = \frac{q_{j0}}{2^l}$) showing an increasing return to scale whereby the rate of improvement (i.e., reduction of the probability of failure) decreases as the investment increases. This models cases where significant benefits can be attained with small initial investments. Conversely, the concave function ($q_{jl} = q_{j0}(\frac{\lambda - l}{\lambda})^{0.6}$) is used to model cases where big investments are required to obtain significant reductions of the risk of failure. In this case, the rate of failure probability reduction increases as more resources are employed. We also consider a linear function ($q_{jl} = q_{j0} - (\frac{q_{j0} - 0.5q_{j0}/2^{\lambda}}{\lambda})l$) where the reduction of failure probability is directly proportional to the level of investment. Furthermore, we consider two risk scenarios by setting q_{j0} equal to 0.1 and 0.3, respectively (Figures 1). These capture two cases where unprotected facilities have 10% and 30% chances of being disrupted. Note that we also run experiments with site-specific failure probabilities q_{j0} drawn in the interval [0.1, 0.3]. However, the results were very similar and therefore are not reported for the sake of brevity.

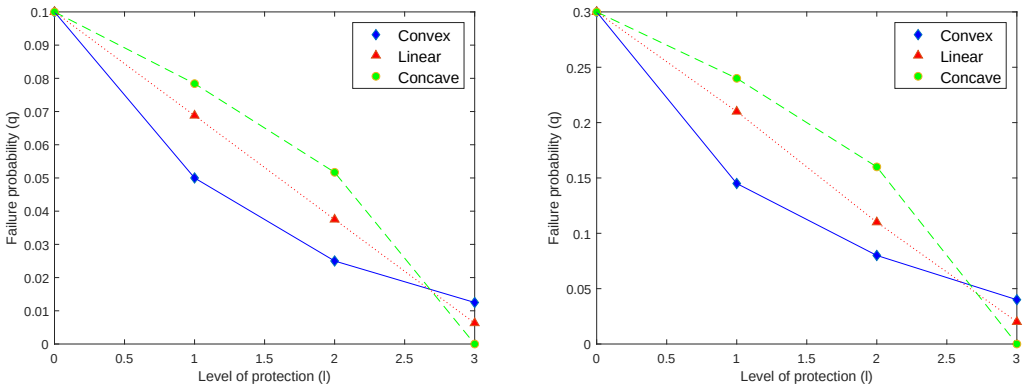

(a) $q_{j0} = 0.1$ (b) $q_{j0} = 0.3$

Figure 1.: Failure probabilities VS protection levels. Convex, linear and concave functions.

Similarly to Berman et al. (2007) and O'Hanley et al. (2013), preliminary tests highlight numerical difficulties when using Cplex to solve PMFP-L. These numerical issues are due to the presence in the model objective of terms which differ by several orders of magnitude (obtained as a product of probabilities). As a results, the solver

sometimes converges to solutions which are proven to be sub-optimal since they are worse than the ones obtained heuristically. To address this issue, we rescale the objective and constraints by multiplying them by a suitable scale factor. The value of the scale factor is computed by trial and error, using the GRASP solutions as benchmarks. Its value is set to 10,000. We refer to this scaled version of the model as PMFP-LS.

Let $\bar{b} = \lambda p$ be the budget needed to protect the whole system at the highest level. In the experiments we define b as a percentage of $\bar{b}$ and let it range in $\{5\%, 15\%, 25\%\}$. Finally, all datasets are tested with 30, 40 and 50 open facilities. Considering larger numbers of facilities within this context seems unnecessary given that supply systems with a very large number of assets are inherently resilient to external disruption (Scaparra and Church, 2008b).

Tables 1, 2 and 3 summarize the computational results on the UK, US and LN datasets, for the three probability functions. The first three columns show the risk scenario, the number of facilities and the available budget, respectively. The best objectives found by PFMP-LS are listed in the *obj* columns. Columns *diff* show the relative difference between obj and the solutions found by PMFP-L and GRASP. Finally, columns t report the computing times in seconds for each method.

With the concave probability function (Table 1), all three methods return the same solutions for every scenario in the US and UK datasets, without significant differences in computing times. In the LN dataset, CPLEX fails to identify the optimal solutions while solving PFMP-L when $q_{j0} = 0.1$. This phenomenon is due to the numerical errors previously discussed. Specifically, PMFP-L returns the naïve solution (highlighted with the star superscript) without any protection implemented. This results in large objective gaps that increase with increasing values of the budget b.

			UK						US						LN					
			PFMP-LS		PMFP-L		GRASP		PFMP-LS		PMFP-L		GRASP		PFMP-LS		PMFP-L		GRASP	
q_{j0}	p	b%	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec
0.1	30	5	415429	1.2	0.0	0.8	0.0	0.1	465979	1.4	0.0	1.2	0.0	0.1	103288	0.5	1.4*	0.3	0.0	0.1
		15	397111	1.2	0.0	0.9	0.0	0.4	431735	1.5	0.0	1.0	0.0	0.3	100563	0.4	4.2*	0.2	0.1	0.2
		25	384635	1.2	0.0	0.8	0.0	0.5	407606	1.3	0.0	0.9	0.0	0.5	98640	0.4	6.2*	0.2	0.0	0.3
	40	5	325422	0.9	0.0	0.7	0.0	0.2	328848	1.3	0.0	0.8	0.0	0.2	78781	0.4	2.2*	0.2	0.0	0.1
		15	308444	1.0	0.0	0.6	0.0	0.5	301666	1.1	0.0	0.7	0.0	0.6	76748	0.4	4.9*	0.2	0.0	0.3
		25	296867	0.9	0.0	0.6	0.0	0.8	287830	1.2	0.0	0.7	0.0	0.9	74909	0.4	7.5*	0.3	0.0	0.4
	50	5	267178	1.3	0.0	0.9	0.0	0.3	251895	1.4	0.0	0.9	0.0	0.3	63434	0.4	2.0*	0.2	0.0	0.1
		15	248334	1.1	0.0	0.7	0.0	0.9	233739	1.5	0.0	0.9	0.0	1.0	61152	0.4	5.8*	0.2	0.0	0.4
		25	238987	1.1	0.0	0.8	0.0	1.2	221845	1.3	0.0	0.8	0.0	1.5	59236	0.3	9.3*	0.2	0.0	0.5
0.3	30	5	604323	1.6	0.0	1.8	0.0	0.1	746660	2.9	0.0	1.7	0.0	0.1	138245	0.9	0.0	0.9	0.0	0.1
		15	535772	2.8	0.0	1.5	0.0	0.4	625624	1.8	0.0	1.7	0.0	0.3	128101	0.8	0.0	0.7	0.0	0.2
		25	492473	3.4	0.0	1.5	0.0	0.6	545696	2.8	0.0	1.9	0.0	0.5	121135	1.3	0.0	1.2	0.0	0.3
	40	5	485594	1.8	0.0	1.7	0.0	0.2	542957	1.6	0.0	1.4	0.0	0.2	109325	0.5	0.0	0.5	0.0	0.1
		15	424739	1.7	0.0	1.3	0.0	0.6	452617	1.4	0.0	1.4	0.0	0.7	102020	0.9	0.0	1.0	0.0	0.3
		25	386546	2.1	0.0	1.5	0.0	0.9	398320	2.4	0.0	1.6	0.0	0.9	95427	0.6	0.0	0.7	0.0	0.4
	50	5	412671	3.9	0.0	1.7	0.0	0.3	434233	2.3	0.0	2.3	0.0	0.3	92316	0.8	0.0	1.1	0.0	0.1
		15	348182	3.6	0.0	1.5	0.0	0.9	364506	3.9	0.0	2.1	0.0	1.1	84679	1.4	0.0	1.1	0.0	0.4
		25	315727	4.1	0.0	1.6	0.0	1.3	319778	4.4	0.0	1.9	0.0	1.4	77948	0.7	0.0	0.7	0.0	0.7
Average				**1.9**	**0.0**	**1.2**	**0.0**	**0.6**		**2.0**	**0.0**	**1.3**	**0.0**	**0.6**		**0.6**	**2.4**	**0.6**	**0.0**	**0.3**

Table 1.: Computational results with the concave probability function

The Linear and Convex cases show more serious numerical difficulties across all instances when $q_{j0} = 0.1$ (see Tables 2 and 3). Only naïve solutions are obtained solving PFMP-L, for the US and UK datasets. As for the LN dataset, occasionally the solver converges to a solution which is non-naïve but sub-optimal because the budget is only used partially. This leads to gaps that despite being smaller than those observed with the naïve solution, are still significant. Overall, PFMP-LS, returns the same or

better solutions than PMFP-L and never finds solutions worse than the ones found by GRASP. This accuracy comes at a slight computational cost. However, given the modest computing times for solving these realistic-size problems, the increased time is negligible.

			UK					US					LN				
			PFMP-LS	PMFP-L		GRASP		PFMP-LS	PMFP-L		GRASP		PFMP-LS	PMFP-L		GRASP	
q_{j0}	p	b%	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec
0.1	30	5	409689	1.2	8.2*	1.2	0.0	0.1	456195	1.0	7.0*	1.0	0.0	0.1	102507	0.4	0.0	0.4	0.0	0.1
		15	390119	1.4	13.6*	1.4	0.1	0.3	416425	1.2	17.2*	1.2	0.6	0.3	99398	0.4	0.2	0.5	0.0	0.2
		25	374125	1.1	18.4*	1.1	0.6	0.5	392364	1.2	24.4*	1.2	1.2	0.5	96857	0.5	0.0	0.4	0.0	0.2
	40	5	322056	0.9	10.4*	0.9	0.0	0.2	321876	1.0	9.1*	1.0	0.0	0.2	78339	0.3	0.0	0.3	0.0	0.1
		15	301316	1.0	18.0*	1.0	0.8	0.5	292986	1.2	19.9*	1.2	0.7	0.6	75552	0.3	0.0	0.4	0.0	0.3
		25	289657	1.4	22.8*	1.4	1.1	0.7	278595	1.5	26.0*	1.5	1.8	0.9	73086	0.4	0.0	0.4	0.3	0.3
	50	5	263390	0.9	12.1*	0.9	0.0	0.3	247342	0.9	8.1*	0.9	0.0	0.3	62950	0.3	0.0	0.4	0.0	0.1
		15	243059	0.9	21.5*	0.9	1.1	0.9	225894	1.0	18.3*	1.0	0.7	1.0	59882	0.4	0.0	0.4	0.0	0.4
		25	232893	1.1	26.8*	1.1	0.8	1.2	213038	2.9	25.5*	2.9	1.0	1.3	57532	0.3	0.2	0.3	0.5	0.4
0.3	30	5	610193	2.2	0.0	2.3	0.0	0.1	749334	3.1	0.0	2.5	0.0	0.1	138327	0.7	0.0	0.8	0.0	0.1
		15	549991	3.6	0.0	2.5	0.0	0.3	641521	3.8	0.0	2.2	0.0	0.3	129359	1.1	0.0	1.3	0.0	0.2
		25	510842	6.6	0.0	3.0	0.0	0.5	570419	4.1	0.0	2.4	0.0	0.5	123094	2.8	0.0	1.6	0.0	0.2
	40	5	495355	2.0	0.0	1.8	0.0	0.2	550852	1.5	0.0	1.7	0.0	0.2	109920	0.6	0.0	0.5	0.0	0.1
		15	440210	2.3	0.0	1.7	0.0	0.5	469524	2.3	0.0	1.9	0.0	0.6	103345	1.6	0.0	1.0	0.0	0.2
		25	405801	7.5	0.0	2.7	0.0	0.7	419913	9.4	0.0	2.2	0.0	0.8	97396	4.2	0.0	1.2	0.0	0.3
	50	5	420129	1.9	0.0	1.9	0.0	0.2	438086	2.7	0.0	2.1	0.0	0.3	92589	0.9	0.0	0.7	0.0	0.1
		15	362821	2.5	0.0	2.1	0.0	0.7	376112	5.5	0.0	2.5	0.0	1.0	85722	1.4	0.0	1.1	0.0	0.4
		25	333401	7.8	0.0	2.2	0.0	1.3	335959	11.0	0.0	2.8	0.0	1.3	79657	3.0	0.0	0.9	0.0	0.4
Average				**2.6**	**8.4**	**1.7**	**0.2**	**0.5**		**3.1**	**8.6**	**2.8**	**0.3**	**0.6**		**1.1**	**0.0**	**0.7**	**0.0**	**0.2**

Table 2.: Computational results with the linear probability function

			UK					US					LN				
			PFMP-LS	PMFP-L		GRASP		PFMP-LS	PMFP-L		GRASP		PFMP-LS	PMFP-L		GRASP	
q_{j0}	p	b%	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec	obj	sec	diff %	sec	diff %	sec
0.1	30	5	416096	1.2	5.6*	1.2	0.0	0.1	461001	1.0	3.9*	1.0	0.0	0.1	102741	0.3	1.0	0.2	0.0	0.1
		15	396920	1.1	10.5*	1.1	0.2	0.4	429181	1.3	11.6*	1.3	0.2	0.3	100034	0.4	2.6	0.2	0.0	0.2
		25	385230	1.3	13.9*	1.3	0.3	0.5	411095	1.4	16.5*	1.4	0.5	0.5	97776	0.4	4.6	0.2	0.1	0.2
	40	5	326653	0.9	7.3*	0.9	0.0	0.2	325813	1.0	7.8*	1.0	0.0	0.2	78613	0.3	2.4*	0.2	0.0	0.1
		15	310189	1.2	12.5*	1.2	0.2	0.5	303961	1.1	15.5*	1.1	0.2	0.7	75991	0.3	6.0*	0.2	0.1	0.3
		25	299511	1.1	16.6*	1.1	0.4	0.8	291035	1.4	20.7*	1.4	0.1	1.0	73918	0.3	8.9*	0.2	0.2	0.4
	50	5	266817	0.8	8.4*	0.8	0.0	0.3	250190	0.8	4.6*	0.8	0.0	0.3	62976	0.3	2.8*	0.2	0.0	0.1
		15	251078	1.0	14.9*	1.0	0.1	0.9	233287	1.1	9.9*	1.1	0.2	1.1	60408	0.3	7.1*	0.2	0.0	0.4
		25	242201	1.2	19.1*	1.2	0.4	1.2	222031	1.2	14.8*	1.2	0.2	1.4	58631	0.4	10.4*	0.2	0.2	0.5
0.3	30	5	606080	2.4	0.0	2.5	0.0	0.1	727017	1.9	0.0	2.2	0.0	0.1	135988	0.7	0.0	0.8	0.0	0.1
		15	536842	2.6	0.0	2.1	0.2	0.4	617371	2.6	0.0	2.4	0.3	0.3	126101	1.5	0.0	1.4	0.1	0.2
		25	495940	5.3	0.0	3.0	0.4	0.5	556119	6.7	0.0	2.8	0.3	0.5	118158	2.2	0.0	1.0	0.3	0.2
	40	5	488250	1.5	0.0	1.8	0.0	0.2	535404	1.6	0.0	1.7	0.0	0.2	108630	0.8	0.0	0.7	0.0	0.1
		15	429674	5.9	0.1	2.8	0.5	0.6	455900	5.1	0.0	2.5	0.4	0.6	99257	1.1	0.0	1.1	0.1	0.3
		25	391658	6.6	0.0	2.6	0.4	0.8	407515	8.7	0.0	3.1	0.9	0.9	92196	1.4	0.0	0.8	0.4	0.3
	50	5	410609	1.4	0.0	1.7	0.0	0.3	425865	2.0	0.0	2.1	0.0	0.3	90859	0.8	0.0	0.7	0.0	0.1
		15	355371	3.1	0.0	2.0	0.3	0.9	362392	10.0	0.0	2.7	0.5	1.0	81766	1.1	0.1	0.9	0.1	0.4
		25	323116	6.5	0.0	2.3	0.9	1.2	320602	14.1	0.1	3.4	0.7	1.3	75568	2.1	0.0	1.0	0.3	0.5
Average				**2.5**	**6.0**	**1.7**	**0.2**	**0.5**		**3.5**	**5.8**	**1.9**	**0.3**	**0.6**		**0.8**	**2.6**	**0.6**	**0.1**	**0.2**

Table 3.: Computational results with the convex probability function

As expected, scenarios with high p and b values lead to slightly higher computing times for both PFMP methods. The q_{j0} values also have some impact on the solution time of the PMFP models, suggesting that riskier scenarios (higher values for q_{j0}) are somewhat more difficult to solve. GRASP, on the other hand, is neither sensitive to the probability function nor to the value of q_{j0}.

In Table 4, we evaluate how protection resources are invested. The purpose of the analysis is to provide insights on how the failure probability function affects protection policies. As in the previous tables, the first three columns identify the risk scenario,

number of facilities and budget available. The other columns show the number of times facilities are protected at a given level, for each of the probability functions. Similar results (not reported) can be observed on the US and LN datasets.

			Protection level l								
			1	2	3	1	2	3	1	2	3
q_{j0}	p	$b\%$	Concave			Linear			Convex		
0.1	30	5	1	0	1	0	2	0	1	0	1
	30	15	1	0	4	1	6	0	8	1	1
	30	25	1	0	7	0	11	0	10	3	2
	40	5	0	0	2	0	3	0	3	0	1
	40	15	0	0	6	0	9	0	9	3	1
	40	25	0	0	10	0	15	0	13	7	1
	50	5	1	0	2	1	3	0	4	0	1
	50	15	1	0	7	0	11	0	9	5	1
	50	25	1	0	12	1	18	0	17	4	4
	Total		**6**	**0**	**51**	**3**	**78**	**0**	**74**	**23**	**13**
0.3	30	5	1	0	1	1	0	1	2	1	0
	30	15	1	0	4	1	0	4	8	1	1
	30	25	1	0	7	1	0	7	11	4	1
	40	5	0	0	2	0	0	2	3	0	1
	40	15	0	0	6	0	0	6	9	3	1
	40	25	0	0	10	0	0	10	17	5	1
	50	5	1	0	2	1	0	2	4	0	1
	50	15	1	0	7	1	0	7	9	5	1
	50	25	1	0	12	1	0	12	17	4	4
	Total		**6**	**0**	**51**	**6**	**0**	**51**	**80**	**23**	**11**

Table 4.: Number of fortified facilities for each protection level. UK dataset.

The Concave and Convex cases show opposite results for the two risk scenarios. Specifically, with a concave probability function, most of the resources are used to implement few high-level protections ($l = 3$). Conversely, with the convex probability function, the protection focus is mostly on low-level protections ($l = 1$). The linear case is more sensitive to the value of q_{j0}. In fact, when $q_{j0} = 0.1$, the solutions appear to be a trade-off between the Concave and Convex cases, with most of the protections enforced at the mid-level ($l = 2$). However, when $q_{j0} = 0.3$, the behaviour changes significantly, with the linear case generating exactly the same protection plans as the concave case and protecting most of the facilities at the highest possible level. These results highlight how choosing a suitable probability function is paramount to identifying efficient fortification strategies. Estimating this function is largely problem dependent and can be driven by historical data and experts' judgment.

7. A case study

This section provides an application of PMFP to a real life problem. The supply network is the Toronto general hospital network (Berman et al., 2007). The dataset is generated by dividing Toronto into 96 areas whose demand is estimated based on the number of residents and business dwellings. Each area is represented by its centroid. The resulting 96-node network is shown in figure 2. The nodes with an identification number represent the hospital locations. Of these, 16 are simple locations and 3 are co-locations (i.e., more than one hospital located in the same node).
This case study is chosen to highlight how the model can be used to support planning

decisions in the healthcare context, a topic that is sadly popular due to the ongoing Covid-19 pandemic. The failure probability q is interpreted as the likelihood that hospitals do not have enough resources (e.g., number of available beds) to satisfy the required demand of patients, leading to the redirection of ambulances toward other facilities. Accordingly, what we referred to as protection levels can now be interpreted as investment levels to increase hospitals' capacities. Small levels represent relatively inexpensive plans (e.g., converting small available spaces, rationalising resources etc.), whereas higher levels are used to model large and more expensive plans (e.g., building a new ward). The protection budget ranges from 0% to 25%. Similarly to the computational analysis, we consider two risk scenarios with q_{j0} equal to 0.1 and 0.3. For example, when $q_{j0} = 0.3$, hospital j has a 30% chance that the emergency room is full and ambulances have to be re-directed to other facilities, unless investments are made. These probabilities are purposely set to high values to model scenarios where health systems struggle to cope with high patient demand. This is particularly true at peak periods throughout the year or during a pandemic. For example, during the winter 2018 in the UK, one in six patients waited at least 30 minutes in the ambulance before being admitted for treatment and, during the Christmas week, 39 patients had to be diverted because hospitals had no resources available (The Guardian, 2018). At the worst point of the Covid-19 pandemic in England, up to 40% of intensive care patients had to be transferred to hospitals hundreds of miles away due to bed shortages (BBC, 2021). Canada faces similar problems, with hospitals often congested and increasingly long wait times (CIHI, 2019). In Canada as well, the ongoing Covid-19 pandemic further exacerbated the situation (CBC, 2020).

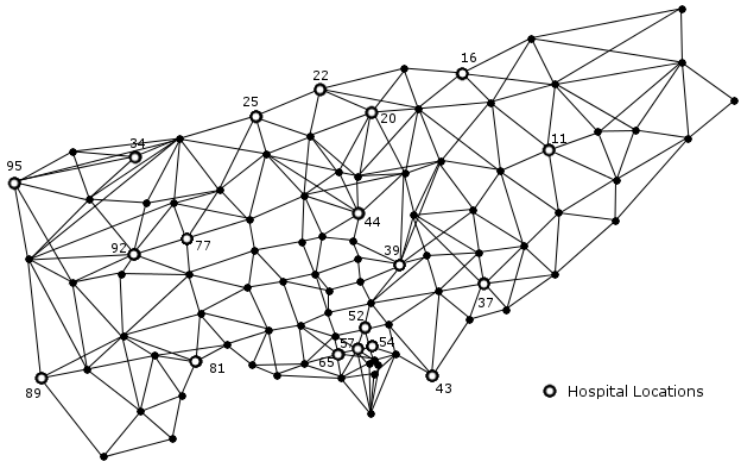

Figure 2.: Location of Toronto Hospitals

7.1. Solution analysis

This analysis aims at studying the impact of increasing the protection budget on the total expected transport cost (or travel distance). Solutions are obtained by solving PFMP-LS with Cplex. Figures 3a and 3b show the percentage reduction in the expected cost against the protection budget levels considered.

In the lower risk scenario, as the budget increases, the linear function results in the largest cost reduction, with a 6.4% decrement when $b = 25\%$. The concave and convex

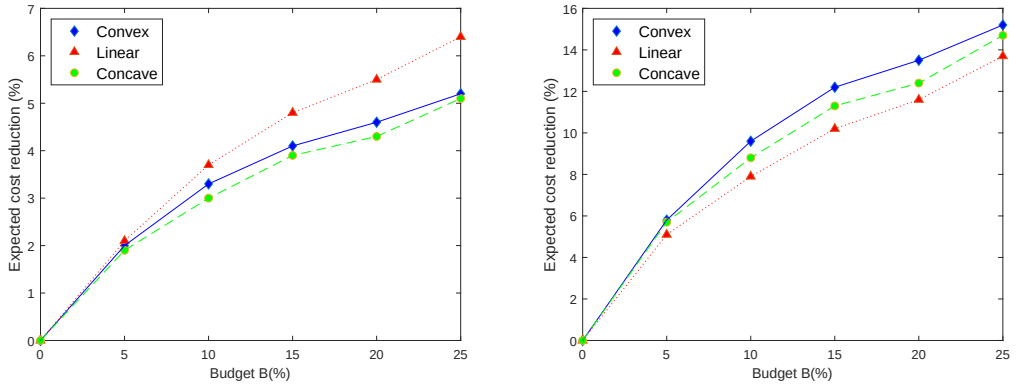

(a) Low risk ($q_{j0} = 0.1$) (b) High Risk ($q_{j0} = 0.3$)

Figure 3.: Expected cost reduction for different protection budget levels

functions show a similar behaviour in both risk scenarios. The cost reductions due to protection are more significant in the high risk scenario. This is expected as often larger risk also means larger room for improvements.

Figures 4a and 4b show marginal cost reduction contributions of each budget increment. This is a useful information that decision makers can use to decide whether a given budget increment is beneficial. The two risk cases do not show notable differences. Overall, as often happens in fortification problems, the impact of the first increments is the most significant. In fact, in both cases a 5% budget is responsible for more than one third of the total benefits. The benefit of subsequent increments steadily decreases until $b = 20\%$, where minor benefits are obtained. This trend changes over the last budget increment and seems to indicate a positive impact in increasing the protection budget from 20% to 25%.

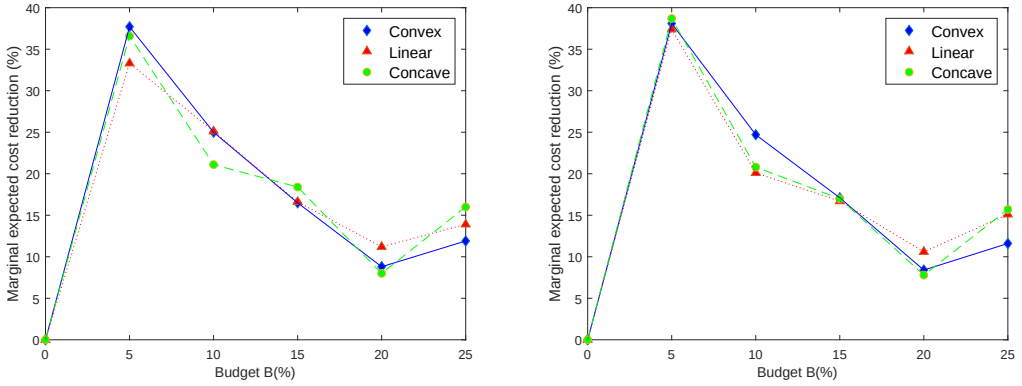

(a) Low risk ($q_{j0} = 0.1$) (b) High risk ($q_{j0} = 0.3$)

Figure 4.: Marginal cost reduction for increasing protection budget levels

Figures 5 and 6, show the optimal investment plans obtained with a budget equal to 5% and 25%. The concentric circles represent the three levels of investment.

Similarly to the computational experiments on the UK, US and LN datasets, the plans with linear (Figures 6c and 6d) and concave (Figures 6e and 6f) probabilities are the same under the higher risk scenario. Moreover, investment plans for the concave

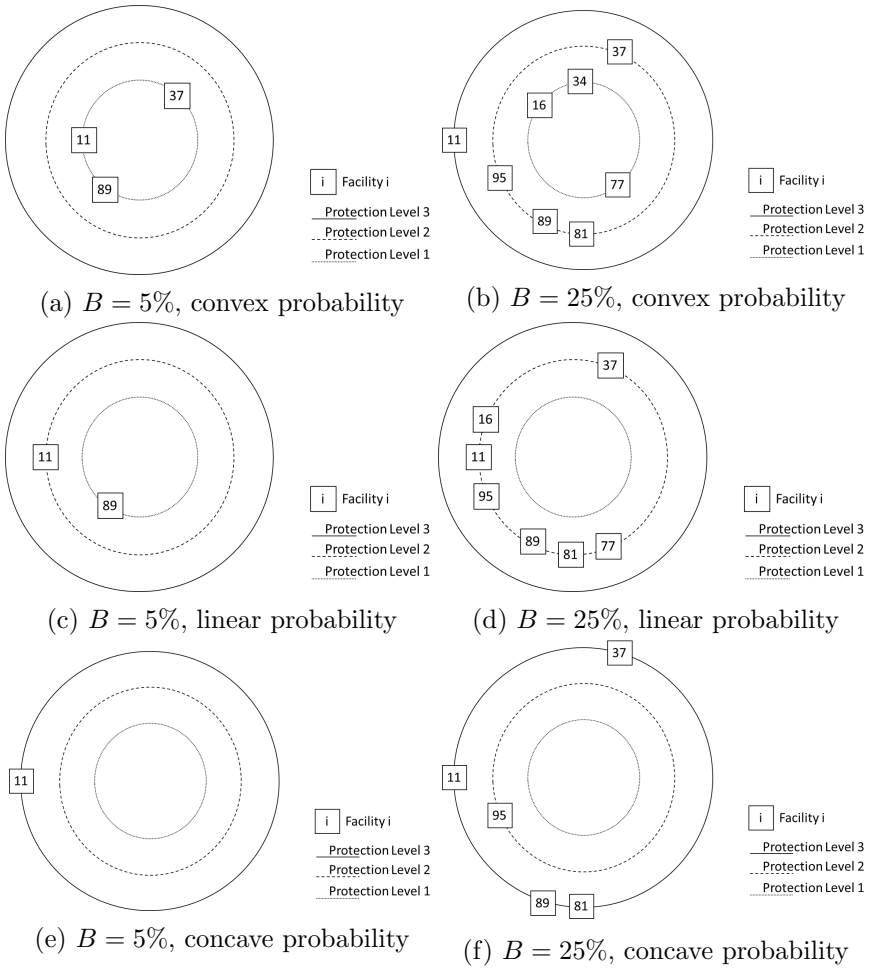

(a) $B = 5\%$, convex probability

(b) $B = 25\%$, convex probability

(c) $B = 5\%$, linear probability

(d) $B = 25\%$, linear probability

(e) $B = 5\%$, concave probability

(f) $B = 25\%$, concave probability

Figure 5.: Investment plans (Low risk)

probability function do not change across risk scenarios (Figures 5e, 5f, 6e and 6f). On the other hand, plans with the convex probability slightly change. For example, under lower risk and $B = 5\%$ (Figure 5a), lowest level investments are selected for 3 hospitals (11, 37, and 89). Increasing the risk results in hospital 37 being dropped from the investment plan and hospital 11 being upgraded to level 2 (Figure 6a). More noticeable are the differences in the investment plans with the linear probability function. In fact, under lower risk, the plans result in medium-low level investments on a large number of hospitals (7 when $B = 25\%$, Figure 5d). However, the higher risk scenario leads to plans with fewer facilities protected but higher level investments (Figure 6d).

Overall, independently on the risk scenario or probability function, the figures suggest that there are a number of hospitals which are clearly more critical than others. An example is hospital 11, which appears in every investment plan.

7.2. Random VS worst-case disruptions

In the following analysis, we compare the solutions obtained with PMFP to hedge against random disruptions against the solutions computed for worst-case disruptions. For the latter, we use the RIMF model introduced by Scaparra and Church (2008a).

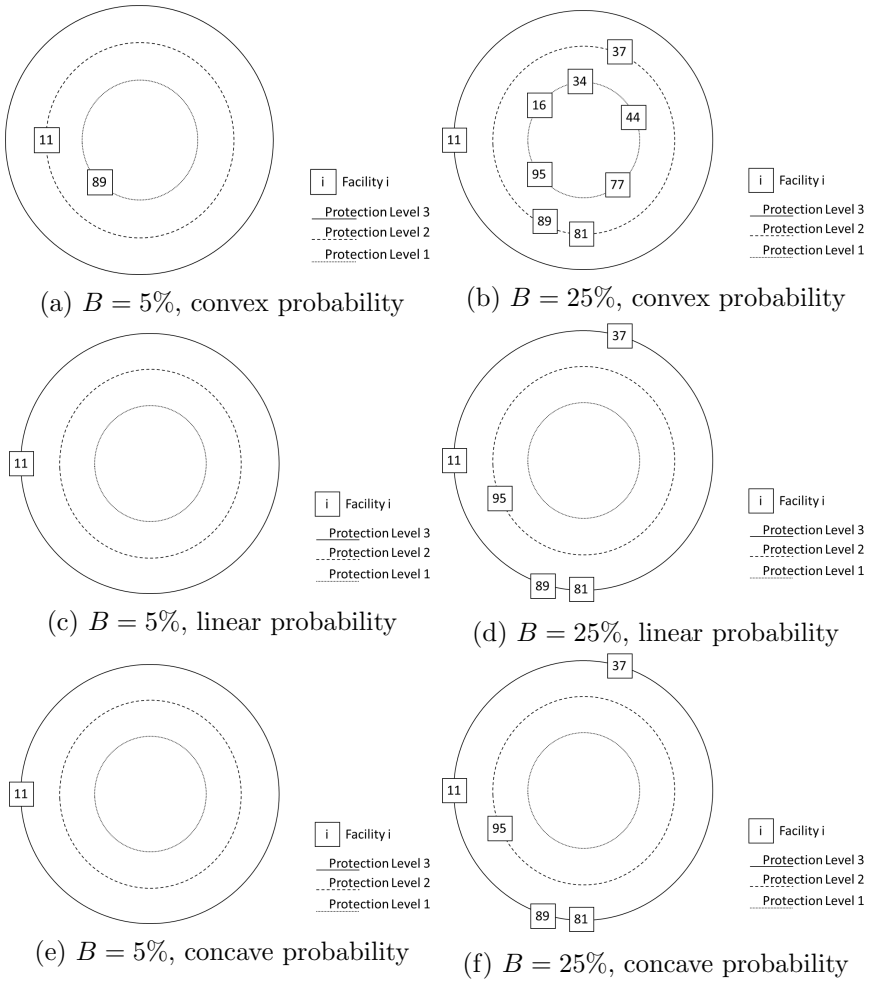

(a) $B = 5\%$, convex probability

(b) $B = 25\%$, convex probability

(c) $B = 5\%$, linear probability

(d) $B = 25\%$, linear probability

(e) $B = 5\%$, concave probability

(f) $B = 25\%$, concave probability

Figure 6.: Investment plans (High risk)

This model identifies which facilities to protect so as to minimise the impact of the worst-case loss of r facilities. Both interdictions and protections are binary. Consequently, in order to provide a meaningful comparison, we restrict PMFP to two investment levels (i.e., $l = 0, 1$). We let q_{j0} range in $[0.10, 0.15, 0.20, 0.25, 0.30]$, whereas q_{j1} is always set to 0. In RIMF, we set $r = [1, 2, 3, 4, 5]$. We test the same budget levels (from 5% to 25%), assuming the same cost for investments across the different models. Overall, 25 scenarios are considered for both PMFP and RIMF. The goal of this analysis is not to compare costs directly, as they are obtained under significantly different assumptions, but to analyse to what extent the investment plans differ and how decision makers can combine the findings from both models.

Table 5 aggregates the results of the 25 scenarios considered, providing insights about the most critical hospitals according to PMFP and RIMF. The first column shows the hospital index. Columns 2 to 5 report the number and percentages of times a specific hospital appears in the optimal investment plans for PMFP and RIMF, respectively. The table indicates that there are some similarities between the optimal investments obtained assuming random and worst-case disruptions. In fact, both models clearly suggest that hospital 11 plays a critical role and requires investments independently on the type of risk scenario considered. Four additional hospitals are

selected by PMFP, five by RIMF. However, how frequently these hospitals are selected changes between the two cases. For example, hospital 37 is selected 60% of the time by PMFP, only 28% by RIMF. Another noticeable difference is the presence of hospitals 16, 77 and 92 in the RIMF solutions Interestingly, these 3 hospitals are selected by RIMF for large values of r (i.e., 4, 5). This seems to indicate that the differences between PMFP and RIMF are more significant in risky scenarios (i.e., for high values of r).

	PMFP		RIMF	
Hospital	No. times	Perc	No. times	Perc
11	25	100%	25	100%
89	20	80%	15	60%
37	15	60%	7	28%
81	10	40%	11	44%
95	5	20%	8	32%
77	0	0%	**5**	20%
16	0	0%	**3**	12%
92	0	0%	**1**	4%
20	0	0%	0	0%
22	0	0%	0	0%
25	0	0%	0	0%
34	0	0%	0	0%
39	0	0%	0	0%
43	0	0%	0	0%
44	0	0%	0	0%
52	0	0%	0	0%
54	0	0%	0	0%
57	0	0%	0	0%
65	0	0%	0	0%

Table 5.: Hospitals' frequency in the optimal plans obtained by PMFP and RIMF

The following analysis aims at testing the robustness of the solutions generated by the two models. Specifically, we test how well a plan obtained by RIMF performs under probabilistic disruptions. Table 6 shows the relative cost increment generated by using the RIMF optimal investment plans within the PMFP model.

		r				
		1	2	3	4	5
	0.1	0.0%	0.0%	0.3%	0.3%	0.7%
	0.15	0.0%	0.0%	0.5%	0.5%	0.9%
q_{j0}	0.2	0.0%	0.0%	0.5%	0.6%	1.2%
	0.25	0.0%	0.0%	0.6%	0.8%	1.4%
	0.3	0.0%	0.0%	0.6%	0.9%	1.5%

Table 6.: Cost increment using RIMF optimal investment plans

Results highlight that cost increments are in general small. The first two columns report no cost increments as the solutions obtained by RIMF with $r = 1, 2$ are equal to the ones obtained by PFMP. On the other hand, the RIMF solutions are different for larger values of r. In other words, decisions generated by RIMF considering large disruption scenarios tend to work poorly within PMFP as they select new hospitals for protection (i.e., 16, 77 and 92) which are not deemed critical by PFMP.

Overall, this analysis indicates that even under simplified assumptions (i.e., $\lambda = 1$), PMFP and RIMF show differences that can lead to sub-optimal decisions if investment plans are to face disruption scenarios different from the ones they are designed to

tackle. We further investigate this issue by suggesting and testing decision policies which incorporate RIMF plans into the PMFP model. For this analysis, we return to the more generic assumptions on investment levels ($\lambda = 3$), while considering concave, linear and convex probability functions and both risk scenarios. Let J^{RIMF} be the set of facilities fortified by a RIMF optimal solution. The following policy, referred to as $\pi(\bar{l})$, enforces that any hospital selected by $RIMF$, should receive an investment of at least $\bar{l}$. Formally, this is achieved by adding to PMFP the following set of constraints:

$$\sum_{l \geq \bar{l}} z_{jl} \geq 1 \quad \forall j \in J^{RIMF} \tag{21}$$

		Lower Risk			Higher Risk		
	r	Concave	Linear	Convex	Concave	Linear	Convex
$\pi(1)$	1	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	2	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	3	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	4	0.1%	0.0%	0.0%	**0.3%**	0.2%	0.0%
	5	0.1%	0.0%	0.0%	**0.2%**	0.1%	0.0%
$\pi(2)$	1	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	2	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	3	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	4	0.2%	0.0%	0.1%	0.3%	**0.4%**	0.3%
	5	0.2%	0.0%	0.0%	**0.5%**	0.4%	0.2%
$\pi(3)$	1	0.0%	**1.3%**	0.3%	0.0%	0.0%	**1.3%**
	2	0.0%	**1.3%**	0.3%	0.0%	0.0%	**1.3%**
	3	0.1%	1.3%	0.4%	0.2%	0.1%	**1.4%**
	4	0.4%	1.6%	0.6%	0.7%	0.6%	**1.9%**
	5	0.4%	1.6%	0.5%	0.9%	0.7%	**1.7%**

Table 7.: PMFP cost increments under $\pi(l)$ policy

Table 7 shows the relative cost increment when policy π is used assuming a budget of 25%. Results show that some decisions from RIMF can be incorporated with small trade-offs on costs. Protecting the facilities selected by RIMF at the lowest level (i.e., $\pi(1)$) results in cost increments of at-most 0.3%. Marginally higher cost increments are observed with $\pi(2)$. However, more significant cost increments appear with $\pi(3)$. This is the consequence of forcing protection investments on facilities selected by RIMF. Generally, the policy's performance decreases as the risk increases (i.e., higher values of q_{j0}). This could seem counterintuitive as one would expect RIMF to be a model better suited to tackle high risk scenarios. However, results clearly indicate that even within high risk scenarios there are differences between randomised and strategic disruptions. In summary, this analysis suggests that a synthesis between the two models for this case study could be achieved via policy π at a relatively low cost.

8. Discussion and Conclusions

In this paper, we introduce an optimisation model to identify protection investment strategies which minimise the expected transportation cost of a median service system under probabilistic disruptions. We consider different levels of investment to reduce the probability of failure of an asset. The resulting model is non-linear. A linearisation

of the model, referred to as PFMP-L, is achieved via probability chains. As noted in the literature, the numerical issues induced by formulations with a similar structure to PFMP-L often cause general purpose solvers to terminate with sub-optimal solutions. The proposed GRASP heuristic helps identify such sub-optimal solutions and calibrate suitable scale factors to redress the numerical problems. The heuristic also proves effective at identifying optimal and near-optimal solutions. Finally, a case study on the Toronto hospitals' network offers insights on how the proposed model can be used to improve healthcare systems so that they are better prepared to cope with high patient demand, for example during epidemic or pandemic events. The case study also shows how robust solutions can be obtained by combining the findings of our model and the finding of a worst-case scenario model in the decision making process.

Future research direction should look into the temporal components of the problem. For instance, the duration of the disruption and the recovery time can be linked to the protection investments. In the network reliability literature the assumptions of independent failure is often made. This is realistic in some applications settings. However, when considering natural disasters affecting regions or large areas, using correlated failure probabilities is more realistic and should be considered in future studies. This is true also during pandemics as all health care centers in hotspot areas may face similar disruptions due to sudden surges in demand for hospitalization. Adding location decisions as well as introducing capacities and linking them to disruption and investment efforts is also a realistic and promising research direction. Moreover, it will be of interest to adapt and apply the model to specific industries, using application-driven probability functions and different numbers of protection levels. All the directions mentioned above are bound to make mathematical models and solution algorithms more complex. Extending the current linearisation approach to a capacitated network is a challenging problem. On the other hand, the GRASP heuristic could be modified to account for capacities with relatively minor changes.

Finally, the Covid-19 pandemic provides a strong motivation to develop more specialized complex models for healthcare systems which combine facility protection measures (e.g., capacity expansion as in our case study) with the optimal re-configuration of the system (e.g., creation of dedicated Covid-19 wards/hospitals and repurposed areas). The development of such models is crucial to ensure that hospitals can deal with sudden patient surge while also maintaining the provision of other essential clinical services and minimizing the risk of infections for both patients and health care personnel. Within this context, another interesting direction is to combine facility protection models with models which optimize staffing and supply of equipment (e.g., personal protection equipment). This is motivated by the direct impact that the lack of equipment can have on staffing shortages and, hence, on service provision (CDC, 2020). Holistic optimisation models incorporating all these aspects should be developed and possibly hybridized with simulation methods to model the complexity and uncertain factors involved in pandemics (Currie et al., 2020).

Acknowledgements

The authors are very grateful to Oded Berman, Dmitry Krass and Mozart B. C. Menezes for providing the Toronto hospitals' network case study.

References

Aksen, D. and Aras, N. (2012). A bilevel fixed charge location model for facilities under imminent attack. *Computers & Operations Research*, 39(7):1364–1381.

Alcaraz, J., Landete, M., and Monge, J. F. (2012). Design and analysis of hybrid metaheuristics for the reliability p-median problem. *European Journal of Operational Research*, 222:54–64.

Altay, N. and Green III, W. G. (2006). OR/MS research in disaster operations management. *European Journal of Operational Research*, 175(1):475–493.

Aslan, E. and Çelik, M. (2019). Pre-positioning of relief items under road/facility vulnerability with concurrent restoration and relief transportation. *IISE Transactions*, 51(8):847–868.

Avetta (2020). Supply chain lessons learned from the coronavirus and sars outbreaks - supply chain 24/7 paper. `https://bit.ly/3hp5ctt`. (Accessed on 07/03/2020).

BBC (2020a). Coronavirus: How NHS Nightingale was built in just nine days - BBC News. `https://www.bbc.com/news/health-52125059`. (Accessed on 07/07/2020).

BBC (2020b). Coronavirus: We can cope with peak, hospital bosses say - BBC News. `https://www.bbc.com/news/health-52293762`. (Accessed on 07/07/2020).

BBC (2021). Covid: Intensive-care patients moved hundreds of miles - BBC News. `https://www.bbc.com/news/health-56483444`. (Accessed on 03/26/2021).

Berman, O. and Krass, D. (2011). On n-facility median problem with facilities subject to failure facing uniform demand. *Discrete Applied Mathematics*, 159:420–432.

Berman, O., Krass, D., and Menezes, M. B. C. (2007). Facility reliability issues in network p-median problems: Strategic centralization and co-location effects. *Operations Research*, 55(2):332–350.

Berman, O., Krass, D., and Menezes, M. B. C. (2009). Locating facilities in the presence of disruptions and incomplete information. *Decision Science*, 40(4):845–868.

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

Cappanera, P. and Scaparra, M. P. (2011). Optimal allocation of protective resources in shortest-path networks. *Transportation Science*, 45:64–80.

CBC (2020). Ontario's hospital logjam brings record wait times, before arrival of covid-19 — cbc news. `https://bit.ly/2ON7INL`. (Accessed on 07/03/2020).

CDC (2020). Strategies to mitigate healthcare personnel staffing shortages — cdc. `https://bit.ly/30B0OAL`. (Accessed on 07/03/2020).

Çelik, M. (2016). Network restoration and recovery in humanitarian operations: Framework, literature review, and research directions. *Surveys in Operations Research and Management Science*, 21(2):47–61.

Çelik, M., Ergun, Ö., Johnson, B., Keskinocak, P., Lorca, Á., Pekgün, P., and Swann, J. (2012). Humanitarian logistics. In *New directions in informatics, optimization, logistics, and production*, pages 18–49. INFORMS.

Church, R. L. and Scaparra, M. P. (2007). Protecting critical assets: The r-interdiction median problem with fortification. *Geographycal Analysis*, 39:129–146.

CIHI (2019). Emergency department wait times in canada continuing to rise. Available at: https://www.cihi.ca/en/. (Accessed on 10/30/2019).

Cui, T., Ouyang, Y., and Shen, Z. J. M. (2010). Reliable facility location design under the risk of disruptions. *Operations Research*, 58(4):998–1011.

Currie, C. S., Fowler, J. W., Kotiadis, K., Monks, T., Onggo, B. S., Robertson, D. A., and Tako, A. A. (2020). How simulation modelling can help reduce the impact of covid-19. *Journal of Simulation*, pages 1–15.

Dantrakul, S., Likasiri, C., and Pongvuthithum, R. (2014). Applied p-median and p-center algorithms for facility location problems. *Expert Systems with Applications*, 41(8):3596–3604.

Daskin, M. (1983). The maximal expected covering location model: formulation, properties and heuristic solution. *Transportation Science*, 17:48–70.

Daskin, M. S. and Dean, L. K. (2005). Location of health care facilities. In *Operations research*

and health care, pages 43–76. Springer.

Drezner, Z. (1987). Heuristic solution methods for two location problems with unreliable facilities. *Journal of the Operational Research Society*, 38(6):509–514.

Esposito Amideo, A., Scaparra, M. P., and Kotiadis, K. (2019). Optimising shelter location and evacuation routing operations: The critical issues. *European Journal of Operational Research*, 279(2):279–295.

European Commission (2019). Eu to invest nearly 225 million in security research. Available at: https://ec.europa.eu/info/news/. (Accessed on 29/10/2019).

Fan, Y. and Liu, C. Z. (2010). Solving stochastic transportation network protection problems using the progressive hedging-based method. *Networks and Spatial Economics*, 10(2):193–208.

Feo, T. A. and Resende, M. G. C. (1995). Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133.

Festa, P. and Resende, M. (2011). GRASP: Basic components and enhancements. *Telecommunication Systems*, 46(3):253–271.

Golany, B., Kaplan, E. H., Marmur, A., and Rothblum, U. G. (2009). Nature plays with dice - terrorists do not: Allocating resources to counter strategic versus probabilistic risks. *European Journal of Operational Research*, 192(1):198–208.

Goodchild, M. F. and Noronha, V. T. (1983). *Location-allocation for small computers*, volume 8. Department of Geography, University of Iowa.

Hakimi, S. L. (1964). Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations research*, 12(3):450–459.

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.

Holmgren, A. J., Jenelius, E., and Westin, J. (2007). Evaluating strategies for defending electric power networks against antagonistic attacks. *IEEE Transactions on Power Systems*, 22(1):76–84.

Jabbarzadeh, A., Fahimnia, B., and Sabouhi, F. (2018). Resilient and sustainable supply chain design: sustainability analysis under disruption risks. *International Journal of Production Research*, pages 1–24.

Jabbarzadeh, A., Fahimnia, B., Sheu, J. B., and Moghadam, H. S. (2016). Designing a supply chain resilient to major disruptions and supply/demand interruptions. *Transportation Research Part B: Methodological*, 94:121–149.

Lei, T. L. and Tong, D. (2013). Hedging against service disruptions: an expected median location problem with site-dependent failure probabilities. *Journal of Geographical Systems*, 15:491–512.

Li, Q. and Savachkin, A. (2013). A heuristic approach to the design of fortified distribution networks. *Transportation Research Part E*, 50:138–148.

Li, Q. and Savachkin, A. (2016). Reliable distribution networks design with nonlinear fortification function. *International Journal of Systems Science*, 47(4):805–813.

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

Li, X. and Ouyang, Y. (2010). A continuum approximation approach to reliable facility location design under correlated probabilistic disruptions. *Transportation Research Part B*, 44(4):535–548.

Liberatore, F. and Scaparra, M. P. (2011). Optimizing protection strategies for supply chains: comparing classic decision making criteria in an uncertain environment. *Annals of the Association of American Geographers*, 101(5):1–17.

Liberatore, F., Scaparra, M. P., and Daskin, M. (2012). Optimization methods for hedging against disruptions with ripple effects in location analysis. *Omega*, 40(1):21–30.

Liberatore, F., Scaparra, M. P., and Daskin, M. S. (2011). Analysis of facility protection strategies against an uncertain number of attacks: The stochastic r-interdiction median problem with fortification. *Computers & Operations Research*, 38(1):357–366.

Lim, M., Daskin, M. S., Bassamboo, A., and Chopra, S. (2010). A facility reliability problem:

Formulation, properties, and algorithm. *Naval Research Logistics*, 57:58–70.

Liu, C. Z., Fan, Y., and Ordonez, F. (2009). A two-stage stochastic programming model for transportation network protection. *Computers & Operations Research*, 36(5):1582–1590.

Losada, C., Scaparra, M. P., Church, R. L., and Daskin, M. S. (2012a). The stochastic interdiction median problem with disruption intensity levels. *Annals of Operations Research*, 201:345–365.

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

Mannino, C., Rossi, F., and Smriglio, S. (2007). The network packing problem in terrestrial forecasting. *Operations Research*, 54(4):611–626.

Medal, H. R., Pohl, E. A., and Rossetti, M. D. (2016). Allocating protection resources to facilities when the effect of protection is uncertain. *IIE Transactions*, 48(3):220–234.

Melo, M. T., Nickel, S., and Saldanha-Da-Gama, F. (2009). Facility location and supply chain management–a review. *European Journal of Operational Research*, 196(2):401–412.

Mladenović, N., Brimberg, J., Hansen, P., and Moreno-Pérez, J. A. (2007). The p-median problem: A survey of metaheuristic approaches. *European Journal of Operational Research*, 179(3):927–939.

Morton, D. P., Pan, F., and Saeger, K. J. (2007). Models for nuclear smuggling interdiction. *IIE Transactions*, 39(1):3–14.

O'Hanley, J. and Church, R. L. (2011). Designing robust coverage networks to hedge against worst-case facility losses. *European Journal of Operational Research*, 209:23–36.

O'Hanley, J., Scaparra, M. P., and Garcia, S. (2013). Probability chains: A linearization technique for modeling reliability in facility location and related problems. *European Journal of Operational Research*, 230(1):63–75.

Ozkan, O., Ermis, M., and Bekmezci, I. (2020). Reliable communication network design: The hybridisation of metaheuristics with the branch and bound method. *Journal of the Operational Research Society*, 71(5):784–799.

Peeta, S., Salman, F. S., Gunnec, D., and Viswanath, K. (2010). Pre-disaster investment decisions for strengthening a highway network. *Computers & Operations Research*, 37:1708–1719.

Peng, P., Snyder, L. V., Lim, A., and Liu, Z. (2011). Reliable logistics networks design with facility disruptions. *Transportation Research Part B*, 45:1190–1211.

Phua, J., Weng, L., Ling, L., Egi, M., Lim, C.-M., Divatia, J. V., Shrestha, B. R., Arabi, Y. M., Ng, J., Gomersall, C. D., et al. (2020). Intensive care management of coronavirus disease 2019 (covid-19): challenges and recommendations. *The Lancet Respiratory Medicine*.

Prais, M. and Ribeiro, C. (2000). Reactive GRASP: An application to a matrix decomposition problem in TDMA traffic assignment. *INFORMS Journal on Computing*, 12(3):164–176.

Puerto, J., Ricca, F., and Scozzari, A. (2018). Extensive facility location problems on networks: an updated review. *TOP*, 26(2):187–226.

Qiao, J., Jeong, D., Lawley, M., Richard, J. P., Abraham, D. M., and Yih, Y. (2007). Allocating security resources to a water supply network. *IIE Transactions*, 39:95–109.

Qin, X., Liu, X., and Tang, L. (2013). A two-stage stochastic mixed-integer program for the capacitated logistics fortification planning under accidental disruptions. *Computers & Industrial Engineering*, 65(4):614–623.

Ramamoorthy, P., Jayaswal, S., Sinha, A., and Vidyarthi, N. (2018). Multiple allocation hub interdiction and protection problems: Model formulations and solution approaches. *European Journal of Operational Research*, 270(1):230–245.

Rohaninejad, M., Sahraeian, R., and Tavakkoli-Moghaddam, R. (2018). Multi-echelon supply chain design considering unreliable facilities with facility hardening possibility. *Applied Mathematical Modelling*, 62:321–337.

Sasaki, M., Suzuki, A., and Drezner, Z. (1999). On the selection of hub airports for an airline hub-and-spoke system. *Computers & operations research*, 26(14):1411–1422.

Sawik, T. (2013). Selection and protection of suppliers in a supply chain with disruption risks.

International Journal of Logistics Systems and Management, 15(2-3):143–159.

Scaparra, M. and Church, R. (2005). A GRASP and path relinking heuristic for rural road network development. *Journal of Heuristics*, 11:89–108.

Scaparra, M. P. and Church, R. L. (2008a). A bilevel mixed integer program for critical infrastructure protection planning. *Computers & Operations Research*, 35:1905–1923.

Scaparra, M. P. and Church, R. L. (2008b). An exact solution approach for the interdiction median problem with fortification. *European Journal of Operational Research*, 189:76–92.

Scaparra, M. P. and Church, R. L. (2012). Protecting supply systems to mitigate potential disaster: a model to fortify capacitated facilities. *International Regional Science Review*, 35:188–210.

Scaparra, M. P. and Church, R. L. (2019). Location problems under disaster events. In Laporte, G., Nickel, S., and da Gama, F. S., editors, *Location Science*, pages 631–656. Springer-Verlag.

Shen, Z. J. M., Zhan, R. L., and Zhang, J. (2011). The reliable facility location problem: Formulations, heuristics, and approximation algorithms. *INFORMS Journal on Computing*, 23(3):470–482.

Snyder, L. V. and Daskin, M. S. (2005). Reliability models for facility location: The expected failure cost case. *Transportation Science*, 39(3):400–416.

Starita, S. and Scaparra, M. P. (2016). Optimizing dynamic investment decisions for railway systems protection. *European Journal of Operational Research*, 248(2):543–557.

Starita, S. and Scaparra, M. P. (2017). Passenger railway network protection: a model with variable post-disruption demand service. *Journal of the Operational Research Society*, 69(4):603–618.

The Guardian (2018). 16,900 people in a week kept in NHS ambulances waiting for hospital care. Available at: https://bit.ly/30FHVfW. (Accessed on 22/07/2020).

Tran, T. H., O'Hanley, J. R., and Scaparra, M. P. (2016). Reliable hub network design: Formulation and solution techniques. *Transportation Science*, 51(1):358–375.

Tufekci, S. and Wallace, W. A. (1998). The emerging area of emergency management and engineering. *IEEE Transactions on Engineering Management*, 45(2):103–105.

Yu, G., Haskell, W. B., and Liu, Y. (2017). Resilient facility location against the risk of disruptions. *Transportation Research Part B: Methodological*, 104:82–105.

Yun, L., Qin, Y., Fan, H., Ji, C., Li, X., and Jia, L. (2015). A reliability model for facility location design under imperfect information. *Transportation Research Part B: Methodological*, 81:596–615.

Zarrinpoor, N., Fallahnezhad, M. S., and Pishvaee, M. S. (2017). Design of a reliable hierarchical location-allocation model under disruptions for health service networks: A two-stage robust approach. *Computers & Industrial Engineering*, 109:130–150.

Zhang, X., Zheng, Z., Zhu, Y., and Cai, K. Y. (2014). Protection issues for supply systems involving random attacks. *Computers & Operations Research*, 43:137–156.

Zhang, X. Y., Zheng, Z., and Cai, K. Y. (2017). A fortification model for decentralized supply systems and its solution algorithms. *IEEE Transactions on Reliability*, 67(1):381–400.