



# Facility location under demand uncertainty: Response to a large-scale bio-terror attack

Pavankumar Murali<sup>a,\*</sup>, Fernando Ordóñez<sup>b</sup>, Maged M. Dessouky<sup>b</sup>

<sup>a</sup>IBM T. J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532, United States

<sup>b</sup>Department of Industrial & Systems Engineering, University of Southern California, 3715 McClintock Avenue, Los Angeles, CA 90089, United States

## ARTICLE INFO

### Article history:

Available online 22 September 2011

### Keywords:

Capacitated facility location  
Distance-dependent coverage  
Demand uncertainty  
Emergency response  
Locate-allocate heuristic

## ABSTRACT

In the event of a catastrophic bio-terror attack, major urban centers need to efficiently distribute large amounts of medicine to the population. In this paper, we consider a facility location problem to determine the points in a large city where medicine should be handed out to the population. We consider locating capacitated facilities in order to maximize coverage, taking into account a distance-dependent coverage function and demand uncertainty. We formulate a special case of the maximal covering location problem (MCLP) with a loss function, to account for the distance-sensitive demand, and chance-constraints to address the demand uncertainty. This model decides the locations to open, and the supplies and demand assigned to each location. We solve this problem with a locate-allocate heuristic. We illustrate the use of the model by solving a case study of locating facilities to address a large-scale emergency of a hypothetical anthrax attack in Los Angeles County.

© 2011 Elsevier Ltd. All rights reserved.

## 1. Introduction

Large-scale emergency events such as a bio-terrorist attack, or a natural calamity such as an earthquake, strike with little or no warning. Such situations can lead to a big surge in the demand for medical supplies. Implementing an efficient disbursement of the, possibly limited, medical supplies needed to satisfy this demand is critical in reducing morbidity and mortality. The United States maintains stockpiles of medical supplies, the Strategic National Stockpile (SNS), to meet the extraordinary needs that can arise in such large-scale emergencies. When the government declares the need to use the SNS, the initial supplies, in the form of medical push packages, would be delivered at the affected location within 24 h. The local authority is then responsible for developing an efficient plan for distributing the supplies to the population.

A disbursement plan considered by local authorities consists of setting up points of disbursement (POD) where the population would go to receive medical supplies or medical attention. This form of distributing supplies is particularly useful if it is necessary to further screen the population or have a trained team administer the medicine. The key decisions in setting up such a disbursement plan consist of the locations of the facilities (or PODs) to be opened

and the amounts of supplies to stock at each of these facilities. Covering models are a classic solution approach for facility location problems in emergency-related scenarios [1,2]. A demand point is treated as covered only if a facility, or a set of facilities, is available to provide the required service to the demand point within a required distance or time. An important additional consideration when planning a response to a large-scale emergency is that there is a large degree of uncertainty associated with the location of the emergency and the number of people affected.

In this work we propose a capacitated facility location model to decide which facilities to open as PODs and how many supplies to make available at each POD in order to maximize the coverage of an uncertain demand in the event of a bio-terror attack. We make the following assumptions about this large-scale emergency: 1) The fraction of the population at demand point  $i$  that can be serviced by facility  $j$  decreases as distance  $d_{ij}$  increases. 2) The PODs are capacitated facilities. 3) There is a total amount of supply  $S$  available to be distributed among the facilities. 4) The demand is optimally assigned to facilities so that the demand covered is maximized. The facilities are capacitated since the speed at which people are serviced in a POD and the physical dimensions of these locations impose a restriction on the maximum rate of service. Total supply constraints can be significant in an emergency setting for a number of reasons, such as: the delay in delivery of the SNS, difficulty in estimating the emergency demand, and concurrent demands on the SNS at other sites.

\* Corresponding author.

E-mail address: [pmurali@usc.edu](mailto:pmurali@usc.edu) (P. Murali).

The assumption that an optimal assignment of demand points to facilities can be implemented is a critical assumption of this model. When distributing medical supplies from PODs, each individual has the freedom to decide whether to go to the assigned POD or to deviate from this recommendation to obtain a better service, suggesting that a congestion model would be a more representative model. We opted to consider a simple optimal assignment model of demand to facilities, as opposed to representing the congestion effects of having individuals select facilities, since due to the uncertainty present in planning a response to an emergency, the significance of a more representative demand model is not clear. Such an optimal assignment model however is realistic when facilities send a mobile unit to serve the population and it also has been a common practice to plan resources at facilities in mandatory evacuation orders during hurricanes where people are advised to go to particular relief centers [3]. We therefore assume that in a large-scale emergency, local authorities aim for an efficient distribution of medical supplies by advising the affected population to visit a particular set of facilities, based on their residential location, medical conditions and supplies available at the facilities.

The question we seek to answer in this work is whether a facility location model with these characteristics can obtain an efficient distribution of scarce medical supplies. This possibility stems from the idea that a critical mistake would be to place the scarce supplies at locations where they are not consumed due to the uncertainty in demand. We can avoid this by placing the supplies at locations where they are more likely to be consumed. If we take into account that a fraction of the demand can be covered by facilities that are further away, then the idea is to select facilities at locations that could receive demand from more demand points and are thus more likely to experience a stable demand. We develop efficient solution methods for this optimization problem that allows us to solve real sized problems representing a large urban area. We present computational results on a possible bio-terror attack in Los Angeles County and simulated random demands to evaluate the solutions found.

The rest of the paper is organized as follows. In Section 2, we provide a brief review of prior work on covering problems and stochastic facility location. In Section 3, we present the proposed capacitated facility location model with chance-constraints, and in Section 4 we present our solution algorithm to solve this optimization problem. We present results from simulation experiments which verify the performance of our model and solution algorithm in Section 5. Finally, in Section 6, we present our conclusions and some future work.

## 2. Literature review

Given that our objective is to design an effective response strategy to a large-scale emergency to reduce casualties, a maximal covering location problem (MCLP) [4] is appropriate for our purpose. In the event of a large-scale emergency, a very important decision would be to locate facilities in tune with the location and intensity (that is, number of affected people) of the attack. One of the key assumptions of the MCLP is that a demand point is assumed to be fully covered if located within a distance  $r$  from the facility and not covered if it is farther than  $r$  away from the facility. However, while planning for an emergency scenario, due to the gravity of the situation, the possible damage to the transportation network and individual decisions, it would be next to impossible to precisely predict whether a person would be able/willing to travel to the recommended facility to receive medicines, making it unrealistic to enforce the binary coverage assumptions of the MCLP. Instead, it is more realistic to assume that the further away the facility is, the

smaller the fraction of the demand it can cover. In the generalized MCLP (GMCLP) as defined by [5], each demand point  $i$  has multiple sets of coverage levels, with corresponding coverage radii. That is, if a facility is located within a distance  $r_i$  from  $i$ , then the coverage level is  $a_i(r_i)$ . The coverage levels can be thought of as a decreasing step function of the distance between a demand point and an open facility. This work was extended by [6] to the gradual covering decay model where they considered general forms of the coverage decay function. Berman et al. (2009) [7] considered the variable radius covering problem, where the decision-maker needs to determine coverage radii for the facilities, in addition to the numbers and locations of facilities, to cover a discrete number of demand points with minimal cost. In our work, we adopt the idea of multiple coverage levels proposed by [5].

Previous works have considered a number of models of how demand is assigned to facilities. Spatial interactions models, also known as “gravity models”, have been used by researchers to assign demand to facilities as a function of parameters that could lead to attraction between demand points and facilities. Spatial interaction models such as those proposed by [8] propose a model based on the distance minimization approach to represent how people assign themselves to facilities. Berman and Krass (1998) [9] focus on competitive facility location assuming that people decide which facility to visit based on the facility’s attractiveness and the travel distance. Aboolian et al. (2007) [10] study a location problem wherein facilities compete for customer demand based on the service utility that they provide. In these prior papers demand is assigned to facilities under competition or congestion and every customer chooses a facility with an objective to maximize its individual utility leading to a game. Although such selfish behavior could occur in a large-scale emergency, in this work we assume that demand is assigned optimally in order to maximize the demand covered.

There exists a fair amount of literature on facility location models dealing with response to an emergency. One of the earliest models in this area was developed by [11] where they developed the location problem as a set covering problem with equal costs in the objective. The sets were composed of the potential facility points within a certain distance or time from each demand point. They solved this problem using linear programming techniques. Rawls and Turnquist (2010) [12] presented a two-stage stochastic optimization model to locate facilities and assign supply to them under an emergency scenario. They develop a mixed integer program to address uncertainty in the demands and in the capacity of the transportation network. Berman and Gavius (2007) [13] presented competitive location models to locate facilities that contain resources required for response to a terrorist attack. They consider the worst-case scenario where the terrorist has knowledge of the location of the facilities and the State needs to take this into account. Jia et al. (2006) [2] presented an uncapacitated version of the covering model to locate staging areas in the event of a large-scale emergency. The location of the facilities and the allocation of the demand points to the open facilities are primarily based on distance constraints. In our paper, we extend this model to a capacitated facility location model. Given the uncertainty associated with a large-scale emergency, it is important to accurately determine the quantity of supplies to be stocked at each potential facility site so that the coverage provided is maximized. To address this issue, we consider the available supply at each facility to be a variable. Yi and Özdamar (2007) [14] propose a mixed integer multi-commodity network flow model for logistics planning in emergency scenarios. They address the aspects of distributing supplies to temporary emergency centers and the evacuation of wounded people to emergency units. The location problem is implicitly handled by allocating optimal

service rates to emergency centers according to which wounded patients are discharged from the system. Görmez et al. (2010) [15] study the problem of locating disaster response facilities to serve as storage and distribution points. Under an emergency, supplies will flow from these facilities to local dispensing sites where they will be distributed to people. They decompose the problem into a two-stage approach – in the first stage, they decide the locations of the local dispensing sites, and in the second stage they treat the local dispensing sites as demand points and decide the locations of the response facilities. The model is solved for a worst-case earthquake scenario in Istanbul. Balcik and Beamon (2008) [16] develop a deterministic model that decides the number and locations of stocking points in a humanitarian relief network and the type and amount of supplies to be pre-positioned in these locations under budgetary and capacity constraints. Chang et al. (2007) [17] present a decision-making tool that could be used by disaster prevention and rescue agencies for planning flood emergency logistics preparation. Two stochastic programming models are developed to determine the structure of rescue organization, the location of rescue resource storehouses, the allocation of rescue resources within capacity restrictions and the distribution of rescue resources.

Facility location models aid decisions that are expensive and difficult to change. Hence, facility location models should consider the uncertainty associated with the demand, supply and distance parameters over a time horizon. In the past, researchers have utilized stochastic and robust optimization approaches to model uncertainty in facility location problems. Snyder (2006) [18] gives an in-depth review of the work done using these two approaches. Mean-outcome models minimize the expected travel cost or maximize the expected profit. The mean-outcome models introduced by [19–21], minimized expected costs or distance. Balachandran and Jain (1976) [22] presented a capacitated facility location model with an objective to minimize expected cost of location, production and transportation. Berman and Odoni (1982) [23] considered travel-times to be scenario-based, with transitions between states or scenarios occurring according to a discrete-time Markov process. The objective was to minimize travel-times and facility relocation costs. Weaver and Church (1983), Mirchandani et al. (1985), Louveaux (1986) and Louveaux and Peeters (1992) [24–27] presented stochastic versions of the P-median problem to choose facilities and allocate demand points. Berman and Drezner (2008) [28] presented a P-median problem that handles uncertainty by minimizing the expected cost of serving all demand nodes in the future. Mean–variance models address the variability in performance and the decision-maker’s aversion toward risk. Such models include [29–31]. Yet another method to model uncertainty is chance-constrained programming. In this procedure, the parameters that are unknown at the time of planning are assumed to follow certain probability distributions. A chance-constraint requires the probability of a certain constraint, involving the uncertain parameter(s), holding to be sufficiently high. Carbone (1974) formulated a P-median model to minimize the distance traveled by a number of users to fixed public facilities. The uncertainty in the number of users at each demand node is handled using chance-constraints. The model seeks to minimize a threshold and ensure that the total travel distance is within the threshold with a probability  $\alpha$ . Interested readers are referred to [18,32] for a review of papers dealing with robust facility location.

While there exists a large amount of literature in the area of capacitated facility location, there have not been many papers on maximal covering models that use chance-constraints to deal with demand uncertainty. In an application like the one presented in this paper, where the number of people affected by a large-

scale emergency and its location are unknown well in advance, facility location modeling under uncertainty is vital. Our model assigns the supply to be stored at each facility by considering it as a decision variable that depends on an unknown demand. Previous papers that have supply as a variable [12,26] do not consider the relation between supply and a random unknown demand. Since the supply at each facility depends on a demand unknown to us a-priori, we model this as a capacitated facility location problem and use chance-constraints to handle the demand uncertainty.

### 3. Facility location model for large-scale emergencies

In this section, we present the capacitated covering model with chance-constraints to handle demand uncertainty. As explained earlier, our objective is to maximize the percentage of the affected population that successfully receives the medication. That is, our goal is to maximize coverage or minimize unmet demand.

#### 3.1. Coverage bound function

In the work presented in this paper, we adopt the idea of multiple coverage levels introduced by [5]. We assume that the fraction of demand at point  $i$  that is assigned by planners to any facility  $j$  to get service decreases with  $d_{ij}$ , the distance between  $i$  and  $j$ . We assume this decrease follows a step function. That is, given positive distances  $0 = \delta_0 < \delta_1 < \delta_2 < \dots < \delta_K$  we let  $f_k D_i$  be the total demand from point  $i$  that can assigned to the group of facilities located from  $\delta_{k-1}$  to  $\delta_k$ , where  $D_i$  is the demand of point  $i$  and  $K$  is the number of coverage levels. The quantities  $f_k$  denote the fraction of the demand at  $i$  that can be satisfied in the interval of distance  $(\delta_{k-1}, \delta_k)$  from  $i$ . These fractions satisfy  $1 = f_1 > f_2 > \dots > f_K > 0$ . Since  $\sum_{k=1}^K f_k \geq 1$  is always valid, we enforce the condition that  $\sum_{j \in J} t_{ij} \leq D_i$  in the deterministic and chance-constrained models presented in the following sections, where  $t_{ij}$  is the amount of medical supplies allocated to demand point  $i$  by facility  $j$ . Note that we can have  $\delta_K$  large enough to represent the maximum distance local authorities would be willing to consider while assigning the affected population to open facilities.

We illustrate the coverage bound function in Fig. 1. Here, demand points are denoted by stars and facilities by triangles. There are three coverage levels, shown by the concentric circles around demand points. Facilities that fall beyond the third concentric circle of a demand point, say  $i$ , are assumed to be too far to satisfy any of the demand of  $i$ . Demand point 1 (DP1), with a demand of  $D_1$ , has zero facilities in the first coverage level, two in the second coverage level and zero in the third level. Following our coverage bound function, facilities F1 and F3 can cover at most  $f_2 D_1$

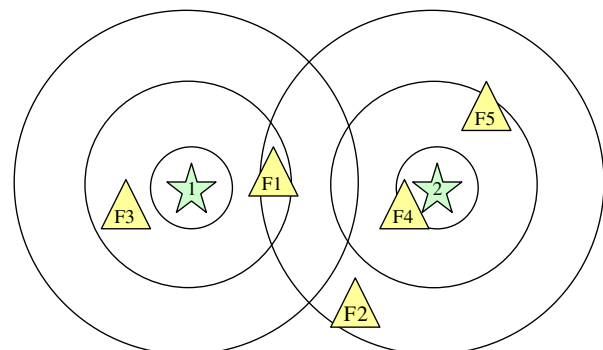


Fig. 1. Demand loss function.

of the demand of  $DP1$ . Since these are the only two facilities in the area of coverage of  $DP1$ , the *upper bound* on the coverage  $DP1$  can obtain is the *lowest* of three quantities:  $f_2D_1$ ,  $D_1$  and the total supply stored at facilities 1 and 3. Since  $F1$  and  $F3$  might also serve other demand points, the actual coverage  $DP1$  will obtain could be lower than this upper bound. The other demand point,  $DP2$ , has one facility in the first coverage level, one in the second level, and two facilities in the third level. Hence, an upper bound on the coverage of  $DP2$  is the minimum of three quantities:  $f_1D_2 + f_2D_2 + f_3D_2$ , the demand  $D_2$ , and the supply stored at  $F1$ ,  $F2$ ,  $F4$  and  $F5$ . When the sum of the fractions  $f_1$ ,  $f_2$  and  $f_3$  exceeds 1, then the maximum possible coverage is limited by  $D_2$ . Again, the actual coverage of the demand at  $DP2$  could be lower than this upper bound because  $F1$ ,  $F2$ ,  $F4$  and  $F5$  might need to serve other demand points as well.

Note that the upper bound on the coverage at a certain coverage level is a property of the demand and is not related to the number of facilities opened at that level. For example, the bound on the coverage of  $DP2$  at the third coverage level is at most  $f_3D_2$  regardless of whether both  $F1$  and  $F2$  are available or only one of them is open.

The covering model we present allows distributing the demand of one point between multiple facilities located possibly at different coverage levels in order to maximize the amount of demand serviced. For example, even if the facilities located in the first coverage level  $[0, \delta_1]$  have sufficient supply to satisfy all the demand, it might be convenient to distribute some of this demand to facilities at the second, or later, coverage levels, since the resources placed at each facility need to satisfy the demand from other points also. In other words, the proposed model maximizes the population covered by deciding which facilities to open, the quantity of supplies to stock at each open facility, and how each demand point is serviced.

### 3.2. Deterministic model

In the model presented below, we determine which of a set of pre-specified facilities need to be opened when a large-scale emergency occurs. We consider a set  $I$  of demand points and a set  $J$  of facility locations. We also consider that each demand point has  $K$  levels of coverage. The model considers the following parameters and decision variables.

#### Parameters

- $S$ : total supply available during an emergency
- $N$ : total number of facilities that need to be opened
- $\beta_j$ : capacity of facility  $j \in J$
- $D_i$ : demand for medical supplies from demand point  $i \in I$
- $f_k$ : a fraction,  $f_k D_i$  is the  $k$ -th level coverage bound on the demand of point  $i$
- $\delta_k$ : radius of the  $k$ -th coverage level from a demand point
- $d_{ij}$ : distance between demand point  $i$  and facility  $j$

#### Decision variables

- $x_j$ : takes a value of 1 if facility  $j \in J$  is open and 0 otherwise
- $s_j$ : supply to be assigned to facility  $j \in J$
- $t_{ij}$ : amount of medical supplies allocated to demand point  $i$  by facility  $j$

In the deterministic model the location and the intensity (size of the affected population) is known ahead of time. That is, the demand at each demand point is known. Then, the problem at hand is to identify the locations of open facilities and their respective supplies,  $s_j$ , out of the available amount  $S$  so as to maximize coverage. The deterministic coverage model is given below:

$$\mathbf{DM} : \max \sum_{i \in I, j \in J} t_{ij} \tag{1}$$

$$\text{s.t. } \sum_{j \in J} x_j = N$$

$$\sum_{i \in I} t_{ij} \leq s_j \quad \forall j \in J \tag{2}$$

$$s_j \leq \beta_j x_j \quad \forall j \in J \tag{3}$$

$$\sum_{j \in J} s_j \leq S \tag{4}$$

$$\sum_{j | \delta_{k-1} < d_{ij} \leq \delta_k} t_{ij} \leq f_k D_i \quad \forall i \in I, k = 1, \dots, K \tag{5}$$

$$\sum_{j \in J} t_{ij} \leq D_i \quad \forall i \in I \tag{6}$$

$$\sum_{j | d_{ij} > \delta_k} t_{ij} = 0 \quad \forall i \in I \tag{7}$$

$$x_j \in \{0, 1\}; \quad s_j, t_{ij} \geq 0$$

The objective of the above integer programming model is to maximize the number of people who receive medication. The first constraint ensures that exactly  $N$  facilities are opened. The second constraint ensures that the supplies distributed for all demand points from  $j$  cannot exceed the available supply at  $j$ . Constraints (3) and (4) make sure supplies are only assigned to open facilities and that these supplies satisfy the facility capacities and total supplies available. The coverage bound function is enforced by constraint (5), where the amount of demand that can be assigned to all the facilities within  $(\delta_{k-1}, \delta_k]$  of demand point  $i$  is bounded by  $f_k D_i$ . The sixth constraint ensures that the amount of supplies assigned to demand point  $i$  from all facilities (at all coverage levels) are not more than the demand at  $i$ ,  $D_i$ . This way even if the  $\sum_{k=1}^K f_k \geq 1$ , supplies sent to  $i$  do not exceed its demand.

### 3.3. Chance-constrained model

We now consider the case when the demand for medical supplies at each demand point is not known a-priori. This means that constraints 5 and 6 of **DM** have some degree of uncertainty around them. We assume that the possible demand values due to an emergency event follow a random variable for which we know the probability distribution at each demand point. One solution in this case is to ignore the uncertainty, replace the random variable  $D_i$  with its expected value  $\mathbb{E}[D_i]$  and use the deterministic model **DM** to get a solution. However, ignoring the uncertainty can be risky in emergency situations.

The proposed chance-constrained model instead requires that the uncertain constraints be satisfied with high probability. That is, if we let  $\xi^i$  represent the random demand at  $i$ , the fifth and sixth constraints of **DM** has to be satisfied except for a small probability  $\epsilon \in (0, 1)$ .

The chance-constrained model **CCM** is formulated as follows:

$$\mathbf{CCM} : \max \sum_{i \in I, j \in J} t_{ij} \tag{8}$$

$$\text{s.t. } \sum_{j \in J} x_j = N$$

$$\sum_{i \in I} t_{ij} \leq s_j \quad \forall j \in J \tag{9}$$

$$\begin{aligned} s_j &\leq \beta_j x_j \quad \forall j \in J \\ \sum_{j \in J} s_j &\leq S \end{aligned} \tag{10}$$

$$\Pr \left( \sum_{j | \delta_{k-1} < d_{ij} \leq \delta_k} t_{ij} \leq f_k \xi^i \right) \geq 1 - \varepsilon \quad \forall i \in I, k \in 1, \dots, K \tag{12}$$

$$\Pr \left( \sum_{j \in J} t_{ij} \leq \xi^i \right) \geq 1 - \varepsilon \quad \forall i \in I \tag{13}$$

$$\sum_{j | d_{ij} > \delta_K} t_{ij} = 0 \quad \forall i \in I \tag{14}$$

$$x_i \in \{0, 1\}; \quad s_j, t_{ij} \geq 0$$

The right hand side of constraints (12) and (13) ( $f_k \xi^i$  and  $\xi^i$ ) and of the corresponding constraints on **DM** represent an upper bound on the amount of demand from  $i$  that could be satisfied by different groups of facilities. The chance-constraint aims to identify an assignment of demand to facilities  $t_{ij}$  that could meet the demand at the facilities with high probability. In other words, regardless of how the demand  $\xi^i$  changes, the assignment to facilities ensures that there is supply to satisfy some of this demand. We note that if the demand  $\xi^i$  turns out to be higher than expected then the assignment to the supply in facilities  $t_{ij}$  would satisfy only part of this demand. However, since we are maximizing  $\sum_{i \in I, j \in J} t_{ij}$ , the amount of demand that is satisfied is as large as possible, given the facility capacity and supply constraints and distribution of demand. In practice, once facilities are opened and supply levels assigned, when the emergency occurs and actual demand levels observed, a linear optimization problem can be solved to determine optimal demand assignments.

While maximizing coverage, the **CCM** model assigns supply  $s_j$  to facilities in  $J$  such that facilities further away from demand points with  $\xi^i > 0$  are allotted lower to little supply compared to facilities closer to those demand points. This ensures that wastage of supplies is minimized. We evaluate the suitability of the proposed approach in the experimental section, where we perform simulations with random demand to observe how much of the actual demand our chance-constrained model is able to satisfy.

In our work, we assume that the demand generated from the demand points follows a log-normal distribution with mean  $\mu'$  and standard deviation  $\sigma'$ . We make this assumption because demand generated at a demand point cannot be negative and the log-normal distribution has a positive support. The relationship between the parameters of the log-normal and the normal distributions are given by  $\mu' = \log \mu - \frac{1}{2} \sigma'^2$ ,  $\sigma'^2 = \log(\mu^2 + \sigma^2) / (\mu^2)$ . We define  $\kappa$  to denote the  $Z$  value of the normal distribution corresponding to the confidence level  $\varepsilon$ .  $\kappa$  is called the *safety factor*, where  $\phi(\kappa) = 1 - \varepsilon$ . Using this relation, we linearize the chance-constraint (12) as follows:

$$\sum_{j | \delta_{k-1} < d_{ij} \leq \delta_k} t_{ij} \leq f_k e^{\mu_i - \kappa \sigma_i} \quad \forall i \in I, k \in 1, \dots, K. \tag{15}$$

Chance-constraint (13) can be linearized similarly.

The objective function for the chance-constrained model remains unchanged from the deterministic model. We note that these linearized versions of the chance-constraints (12) and (13) have the same structure as constraints (5) and (6) on **DM**, with the demand value  $D_i$  replaced by  $e^{\mu_i - \kappa \sigma_i}$ . While in **DM** we replace  $D_i$  by its expected value, in **CCM** we replace it with the value given by the chance-constrained expression.

#### 4. A locate-allocate heuristic

The locate-allocate heuristic was introduced by [19] and, for instance, used in [33] to solve facility location problems. Jia et al. (2007) [34] show that the locate-allocate heuristic outperforms a genetic algorithm procedure and is nearly as good as a Lagrangean-relaxation heuristic for solving an uncapacitated facility location problem of locating medical supplies for a large-scale emergency. In terms of computational time, they showed that the locate-allocate heuristic is much faster than a genetic algorithm and a Lagrangean-relaxation heuristic. The efficiency shown for the uncapacitated version of the problem suggest the use of this type of heuristic to solve the chance-constrained model described previously. To describe the locate-allocate heuristic, we use the deterministic model **DM**. To solve this heuristic for **CCM**, we need to simply replace  $D_i$  with the value given by the chance-constrained expression.

The idea behind this heuristic is very simple. The first step is to choose an initial location for the  $N$  facilities to be opened. This can be done by using a simple greedy approach. In our case, we solve an integer program with an objective to maximize the amount of supplies transported between facilities and demand points. At this stage, we do not consider the demand loss function, but simply specify that a facility is allowed to transport supplies to demand points that lie within a radius  $\delta_K$  from it. This mixed integer program **IP**<sub>1</sub> is as follows:

$$\begin{aligned} \mathbf{IP}_1 : \quad & \max \sum_{i \in I, j \in J} t_{ij} \\ & \text{s.t.} \sum_{j \in J} x_j = N \end{aligned} \tag{16}$$

$$\sum_{i | d_{ij} \leq \delta_k} t_{ij} \leq \frac{1}{N} \sum_{i \in I} D_i x_j \quad \forall j \in J \tag{17}$$

$$x_j \in \{0, 1\}; \quad t_{ij} \geq 0$$

The initial facilities opened, using **IP**<sub>1</sub> above, consider the demand balanced over all open facilities with the help of the second constraint. The second step of the heuristic is to allocate demand points to the facilities that were opened, with an objective to maximize coverage. In our case, we solve the following linear program **LP**<sub>A</sub> that does this allocation of demand points to open facilities. We define a set  $J'$  to denote the set of facilities that have been opened and  $\xi^i_\varepsilon$  as the demand realized at the confidence level of  $\varepsilon$ .

$$\begin{aligned} \mathbf{LP}_A : \quad & \max \sum_{i \in I, j \in J'} t_{ij} \\ & \text{s.t.} \sum_{i \in I} t_{ij} \leq s_j \quad \forall j \in J' \end{aligned} \tag{18}$$

$$s_j \leq \beta_j \quad \forall j \in J' \tag{19}$$

$$\sum_{j \in J'} s_j \leq S \tag{20}$$

$$\sum_{j \in J' | \delta_{k-1} < d_{ij} \leq \delta_k} t_{ij} \leq f_k \xi_e^i \quad \forall i \in I, k \in 1, \dots, K \tag{21}$$

$$\sum_{j \in J' | d_{ij} > \delta_k} t_{ij} = 0 \quad \forall i \in I \tag{22}$$

$$\begin{aligned} \sum_{j \in J} t_{ij} &\leq \xi_e^i \quad \forall i \in I \\ s_j, t_{ij} &\geq 0 \end{aligned} \tag{23}$$

The third step of the locate-allocate heuristic is to create clusters of demand points served by each open facility. Then, for each cluster, we try to relocate the open facility from its current site to another one such that the total travel distance between the facility and the demand points in the cluster is minimized. In our approach, we define clusters  $C_1, \dots, C_N$  to denote the  $N$  clusters formed by the  $N$  open facilities. The relocate integer program is:

$$\begin{aligned} \text{IP}_R: \quad \min \quad & \sum_{C_y} \left[ \theta_{C_y, j} \sum_{i | i \in C_y} d_{ij} \right] \\ \text{s.t.} \quad & \sum_{j \in J} x_j = N \end{aligned} \tag{24}$$

$$\sum_j \theta_{C_y, j} = 1 \quad \forall y \in 1, \dots, N \tag{25}$$

$$\begin{aligned} \sum_{C_y} \theta_{C_y, j} &= x_j \quad \forall j \in J \\ x_j &\in \{0, 1\} \quad \forall j \in J \end{aligned} \tag{26}$$

Here,  $\theta_{C_y, j}$  is set to 1 if facility  $j$  is assigned to cluster  $C_y$  and 0 otherwise. The first constraint ensures that every cluster is assigned exactly one facility, and the second constraint enforces the condition that an open facility is assigned to only one cluster, and a closed facility is assigned to none. On solving the above integer

program, if a facility different from the one used in the allocate step of the heuristic is found to minimize the total travel distance for any particular cluster, then the new facility is assigned the same supply as the previous facility. The allocate and relocate steps of the locate-allocate heuristic are repeated until no further relocation occurs. At this stage, the heuristic is terminated. The favorable performance of the location-allocation heuristic in solving facility location models has been presented by [33,35–37], among others. In our implementation of this heuristic, we use integer programs to solve the relocate and allocate procedures. This reduces the number of iterations required to find the best set of locations and allocations. The heuristic stops when there is no further reduction in the objective function value of the allocate step. While the convergence of the heuristic cannot be proven for a general case, it did converge for all the instances in our experiments. Its convergence for specific cases has also been discussed in the aforementioned papers. In the following section, we test the performance of our heuristic and compare it with a simulated annealing heuristic presented by Berman and Drezner (2006) which also considers a distance-dependent demand.

### 5. Experimental analysis

In this section, we present experiments to show how the deterministic and the chance-constrained models and the locate-allocate heuristic can be used to locate staging areas for mass distribution of the medical supplies with an objective to maximize coverage in the event of an anthrax attack on Los Angeles County. Anthrax is a deadly disease that requires vaccines and antibiotics to treat the affected persons and immunize the high-risk population. During an anthrax emergency, the strategic national stockpiles (SNS) are mobilized for local emergency medical services. We use the centroid of each census tract, representing the aggregated population in that tract as the demand points. Using this procedure, we have 1939 demand points in Los Angeles

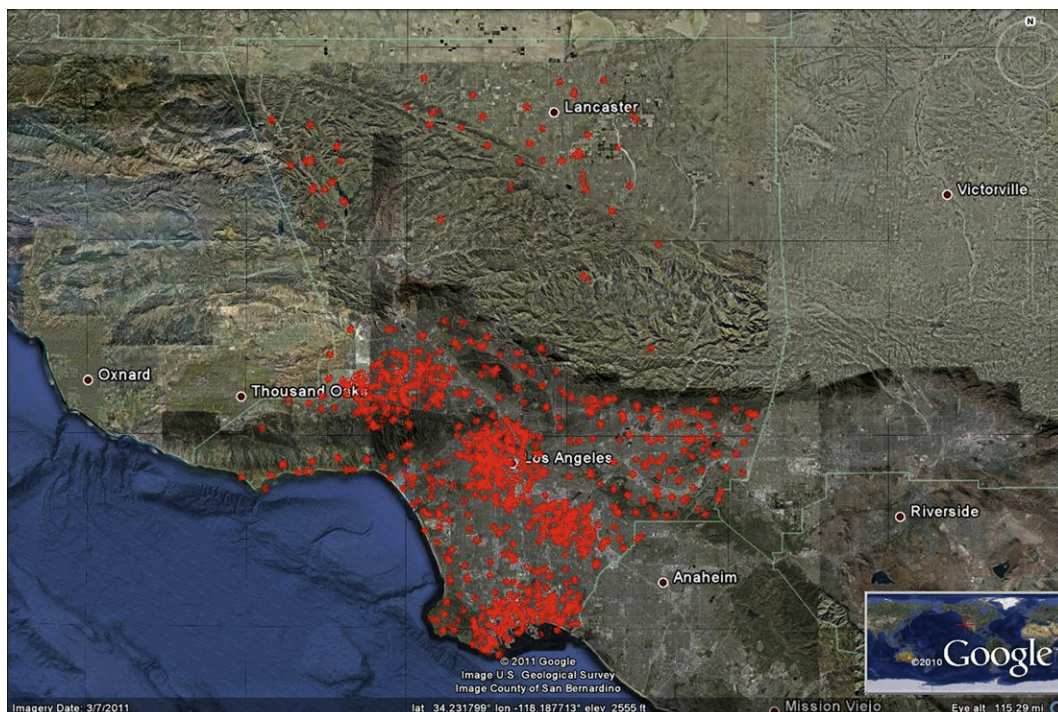


Fig. 2. Locations of demand points within Los Angeles County.

County. Under a large-scale emergency scenario, the demand arising from each demand point could be anywhere between 0 and the aggregate population represented by it. The mean demand for the 1939 demand points were provided to us. We were also provided with approximately 200 eligible facility sites. To protect the confidentiality of the exact site locations, we use exactly 200 of these sites with a slight perturbation in their geographical coordinates. Fig. 2 shows the distribution of the demand points in the County.

For the loss function, we consider three coverage levels of radii 4 mi, 8 mi and 12 mi respectively, for every open facility. Corresponding to these coverage levels, we assume  $f_1 = 100\%$ ,  $f_2 = 65\%$  and  $f_3 = 30\%$ . Additionally, as presented by [38], we assume the supply of vaccines in each facility ( $\beta_j$ ) to be at least equal to 140,000 units per week, with a response time period of 4 weeks. This means that the facilities have a supply of 560,000 units. In the experiments below we take the total supply  $S$  to be some fraction of the expected total demand. That is, we take  $S = \gamma \sum_{i \in I} \mathbb{E}[D_i]$  for some service level  $\gamma \in (0, 1]$ . We do this because the difficult situation is when there is a lack of supplies.

We coded the simulated annealing heuristic presented by [39] in C++ and ran it using Microsoft Visual Studio 5.0 on a Pentium IV computer with 512 MB RAM. We maintained the same settings as above, and compared its performance with our locate-allocate heuristic. The simulated annealing heuristic uses an approach similar to that of Vogel's approximation method [40] to allocate demand points to facilities. Since their heuristic does not consider supply as a variable, we assumed the capacity of every open facility to be equal to the  $\min((\gamma \sum_{i \in I} \mathbb{E}[D_i])/N, \beta_j)$ . The demand experienced by a facility is discounted as per the loss function. Facilities are sorted based on their opportunity cost for which we used distance as a measure. Demand points are allocated to facilities as long as the total demand assigned to a facility does not exceed its capacity.

The locate-allocate heuristic was coded in C++ using ILOG Concert Technology. For all the experiments provided below, CPLEX and the locate-allocate heuristic were performed on a Dell Precision 670 computer with a 3.2 GHz Intel Xeon Processor and 2 GB RAM running CPLEX 9.0. The heuristics executed rather quickly in about 10–15 min. To contrast, we compared the results obtained against solving the whole facility location problem with CPLEX for 2.0 CPU hours keeping the best solutions found so far. The convergence of the CPLEX solution slowed down considerably after 2.0 h.

5.1. Deterministic model

In this case, we use the mean demand for the 1939 demand points as the actual demand for the respective demand points. We present below the coverage that could be obtained from the deterministic model using the locate-allocate heuristic. We test the sensitivity of the algorithm for varying number of facilities to be opened  $N$  and varying service levels  $\gamma$ .

We compare the above results with that obtained by using the simulated annealing procedure, presented by Berman and Drezner (2006), to solve the deterministic facility location problem. The results from the simulated annealing procedure are presented below.

In the figure below we plot the performance of the locate-allocate (L-A) heuristic and the simulated annealing (SA) procedure for  $N = 40$ , and compare them to the best lower bound (LB) and upper bound (UB) of the integer program **DM** obtained using CPLEX after 2.0 CPU hours, also for  $N = 40$ . Fig. 3 shows that for the settings of our problem, the locate-allocate heuristic outperforms the simulated annealing procedure in locating facilities to

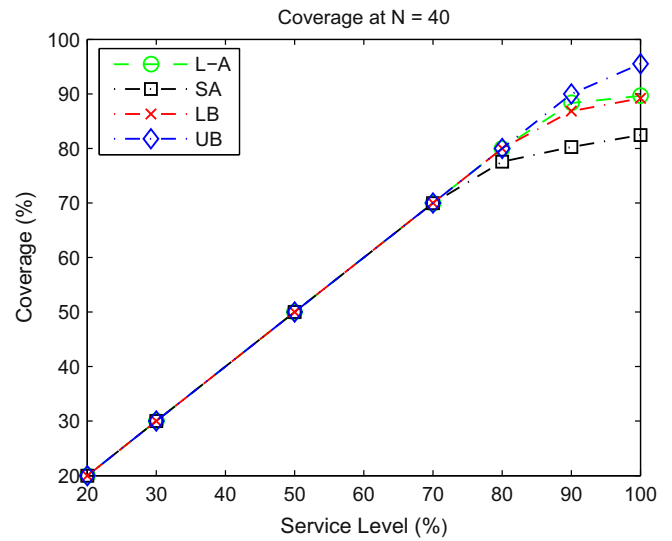


Fig. 3. Comparison of locate-allocate and simulated annealing for  $N = 30$ .

maximize coverage. In addition, while the coverage achieved by the locate-allocate heuristic is mostly greater than the CPLEX lower bound, the coverage achieved by simulated annealing is usually slightly below this lower bound for  $\gamma$  values of 100%, 90%, and 80%. This trend in both the heuristics holds true for all the values of  $N$  tested, as shown in Tables 1 and 2 above.

To further evaluate the performance of our algorithm in comparison to that of simulated annealing, we design an experiment that consists of 500 demand points and 50 potential facility sites. The demand points and facility sites were chosen randomly from the entire set for Los Angeles that was presented earlier. As before,  $\beta_j = 140,000$  units per week and  $\gamma = [0,1]$  and the response time is 4 weeks. We solve the deterministic problem **DM** using L-A and SA for different coverage levels radii ( $r$ ) and coverage fractions ( $f$ ):  $r^1 = 4$  mi, 8 mi, 12 mi;  $r^2 = 3$  mi, 6 mi, 12 mi;  $f^1 = 100\%$ , 65%, 30%;  $f^2 = 100\%$ , 75%, 50%. The results of these experiments are presented in Tables 3–8 below.

On comparing results in Tables 3 and 4, we conclude that L-A outperforms SA under these settings. We also note that for both heuristics the coverage drops when the number of open facilities is 5. This could be due to the fact that 5 facilities may not be sufficient to provide reasonable coverage to the spatially distributed demand points. There is a slight improvement in coverage when the coverage fractions increase (Table 5). This could mean that the total supply,  $S$ , stored at the open facilities in the experiments in Table 3 was less than  $4N\beta_j$ , where  $N$  is the number of open facilities and the response time is 4 weeks. On increasing the coverage fractions, some of the remaining supply can now be allotted to demand points in all three coverage levels, thereby increasing overall coverage. However, the improvement in coverage is negligible with 5 open facilities probably because  $S$  was very close to  $4N\beta_j$  and very few of the 500 demand points are within 12 miles from the open facilities. If  $S \approx 4N\beta_j$  under  $f^1$ , then little supply is left over for

Table 1 Coverage from the deterministic model using locate-allocate heuristic.

$\gamma$	100%	90%	80%	70%	50%	30%	20%
$N = 50$	94.19	89.99	80	70	50	30	20
$N = 40$	89.66	88.32	80	70	50	30	20
$N = 30$	84.96	83.68	77.01	70	50	30	20
$N = 20$	74.56	72.26	72.09	70	50	30	20

**Table 2**  
Coverage from the deterministic model using the simulated annealing heuristic.

$\gamma$	100%	90%	80%	70%	50%	30%	20%
$N = 50$	90.86	84.71	79.96	70	50	30	20
$N = 40$	82.45	80.24	77.58	69.97	50	30	20
$N = 30$	79.67	76.44	74.30	69.89	50	30	20
$N = 20$	71.34	69.91	68.83	68.67	50	30	20

**Table 3**  
Coverage from DM using L-A heuristic, under  $r^1$  and  $f^1$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	95.62	90	80	70
$N = 10$	86.49	83.25	76.80	69.79
$N = 5$	75.44	73.37	69.64	64.76

**Table 4**  
Coverage from DM using SA heuristic, under  $r^1$  and  $f^1$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	91.41	87.31	80	70
$N = 10$	79.11	72.73	65.35	53.33
$N = 5$	70.47	61.93	53.80	43.71

**Table 5**  
Coverage from DM using L-A heuristic, under  $r^1$  and  $f^2$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	98.74	90	80	70
$N = 10$	89.56	85.67	78.25	70
$N = 5$	75.81	73.37	69.64	64.76

**Table 6**  
Coverage from DM using L-A heuristic, under  $r^2$  and  $f^1$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	89.22	84.36	77.53	70
$N = 10$	76.48	74.57	70.70	66.15
$N = 5$	69.61	69.01	67.14	63.79

**Table 7**  
Coverage from DM using SA heuristic, under  $r^2$  and  $f^1$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	82.57	76.41	73.13	68.93
$N = 10$	68.62	61.43	55.89	50.44
$N = 5$	63.13	57.10	49.60	40.57

allocation to demand points upon increasing coverage fractions. The coverage remains unchanged because the optimal solution obtained with the original coverage fractions would still remain optimal.

Similar to the previous set of experiments, we notice from Tables 6 and 7 that L-A performs better than SA. Once again coverage slightly improves when coverage fractions are increased.

**Table 8**  
Coverage from DM using L-A heuristic, under  $r^2$  and  $f^2$ .

$\gamma$	100%	90%	80%	70%
$N = 20$	91.52	85.25	77.79	70
$N = 10$	77.83	75.33	70.78	66.21
$N = 5$	69.61	69.01	67.14	63.79

However, the gain in coverage obtained in Table 8 is lower than the gain obtained in Table 5. This could mean that a significant number of demand points were located between 3 and 4 miles and between 6 and 8 miles from the open facilities. Demand points located between 3 and 4 miles could receive up to 100% coverage in the first set of experiments, but only up to 75% in the second set. Demand points located between 6 and 8 miles could receive up to 75% coverage in the first set of experiments, but up to 50% in the second set. Hence, the gain is lower in Table 8.

5.2. Chance-constrained model

Having verified the quality of the locate-allocate heuristic for solving the deterministic model, we now investigate the performance of the chance-constrained model and the heuristic under demand uncertainty. For all the results presented in this section, we assume that  $N = 20$  and  $\gamma = 0.8$ . We use a log-normal distribution with the same mean values as was used in the deterministic case for generating random demand for the simulations presented in this section. The standard deviations are taken to be a certain percentage (10%, 20% etc.) of the respective mean demand values. In Table 9 below, we present the results from an experiment to study how the  $\kappa$  values impact the coverage provided by the chance-constrained model.

For low values of  $\epsilon$  (corresponding to high values of  $\kappa$  and low uncertainty), the constraint that  $\sum_j t_{ij}$  should never exceed  $f_{k\epsilon}^i$  can be satisfied easily. Hence,  $\sum_j t_{ij} \leq f_{k\epsilon}^i$  is a tight constraint. This explains the low coverage values in Table 9 above. As  $\epsilon$  increases (that is, uncertainty increases),  $f_{k\epsilon}^i$  would need to be set to a higher value so that  $\sum_j t_{ij}$  never exceeds  $f_{k\epsilon}^i$ . As a result, coverage values increase.

Next, we perform two sets of simulation experiments to evaluate the quality of the locations of the open facilities and the supply stored at each open facility in terms of unmet demand through simulations. For both the simulation experiments, we fix the facility locations and supplies, that is, the  $x_j$  and  $s_j$  solution values for all the locations, for each combination of  $\kappa$  and  $\sigma$  as per Table 9. Next, we generate random demands for the 1939 demand points based on their respective mean values for each value of  $\sigma$ . We compute the following for each sample of demand: (1) the coverage of the CCM from the performance obtained by the facility sites and supply solutions under this random demand for a given  $\kappa$ - $\sigma$  combination is recorded as the coverage obtained by using the CCM for allocating demand points to facilities, (2) coverage obtained if demand were known in advance (i.e., deterministic) in which case the DM model is used for locating facilities and allocating demand points and supplies to facilities. Then, the ratios (1)/(2) are computed. Each recording in Table 10 below is the average of 20 such ratios.

The ratios compare the performance of facility locations and their respective supplies outputted by the CCM model under a random demand to the performance of the locations and their supplies outputted by the DM model optimal for that demand alone. In Table 10, the denominator values remain constant for each

**Table 9**  
Coverage from the chance-constrained model.

$\kappa$	$\epsilon$	$\sigma = 10\%$	$\sigma = 20\%$	$\sigma = 40\%$
1.96	0.025	62.52	52.57	41.62
1.64	0.050	63.47	53.16	42.11
1.44	0.075	64.08	53.38	42.65
1.28	0.100	64.66	54.09	44.83
1.15	0.125	64.94	55.13	45.54
1.04	0.150	65.18	56.69	46.73
0.93	0.175	65.27	57.21	48.44
0.84	0.200	65.78	59.32	49.23



**Table 10**

Ratio of the coverage from the CCM under a random demand to the best possible coverage had we known this demand in advance.

$\kappa$	$\epsilon$	$\sigma = 10\%$	$\sigma = 20\%$	$\sigma = 40\%$
1.96	0.025	0.8253	0.7386	0.6868
1.64	0.050	0.8352	0.7558	0.7078
1.44	0.075	0.8391	0.7734	0.7312
1.28	0.100	0.8502	0.7815	0.7457
1.15	0.125	0.8539	0.8040	0.7492
1.04	0.150	0.8606	0.8121	0.7598
0.93	0.175	0.8626	0.8200	0.7704
0.84	0.200	0.8646	0.8317	0.7844

**Table 11**

Ratio of the performance of the CCM to the DM in response to a random demand.

$\kappa$	$\epsilon$	$\sigma = 10\%$	$\sigma = 20\%$	$\sigma = 40\%$
1.96	0.025	0.9799	0.9820	1.0152
1.64	0.050	0.9917	1.0048	1.0727
1.44	0.075	0.9962	1.0282	1.1082
1.28	0.100	1.0094	1.0390	1.1301
1.15	0.125	1.0138	1.0689	1.1353
1.04	0.150	1.0218	1.0797	1.1514
0.93	0.175	1.0242	1.0902	1.1675
0.84	0.200	1.0265	1.1058	1.1887

column. Similar to Table 9, the numerator values increase with a decrease in the  $\kappa$  value and decrease with an increase in uncertainty, represented by the standard deviation. Table 10 shows that in the worst-case, the locations chosen by the locate-allocate heuristic cover approximately 69% of the demand that could be covered had we known this random demand well in advance. For the best-case scenario, this value increases to approximately 86%.

For the second simulation experiment, we compute the following: (1) the numerator is computed just as how they were computed in Table 10 above, (2) we consider the sites, the  $x_j$  and the  $s_j$  values, used by the deterministic model, presented in Table 1 with  $N = 20$  and  $\gamma = 0.8$  and allocate demand points to these open facilities under the same demands as were considered for the numerator (1). Then, the ratios (1)/(2) are computed. Each recording in Table 11 is averaged over 20 such ratios.

Table 11 shows the merit of using the chance-constrained model to locate facilities, determine their supplies and allocate demand points to the open facilities. It also shows the relative gain in the coverage, achieved by using the chance-constrained model over the deterministic model, increases with a decrease in the safety factor and an increase in the standard deviation. The first is due to an increase in the numerator values as  $\kappa$  decreases, and the second is due to a decrease in the denominator values as demand uncertainty increases. Under the best-case scenario, we see nearly a 20% gain in coverage by resorting to the locations given by the chance-constrained model.

## 6. Conclusions

In this study, we consider the problem of locating POD for medicines in response to a bio-terror attack. To address the tremendous magnitude and low frequency of large-scale emergencies we obtain a solution that maximizes the number of people serviced under such uncertain and limited resources/time conditions.

The main contribution of our work is in designing a response strategy to distribute supplies in a large-scale emergency that considers distance-sensitive coverage, in addition to demand uncertainty. In the problem we consider here, first, the facilities are

capacitated by the service rate of a POD. Second, the demand satisfied depends on the distance to the facility. This is because while planning response to a large-scale emergency scenario, it is reasonable to assume that the number of people expected to be assigned to a particular POD decreases as their distance to that POD increases. Thirdly, given the unpredictability as to when and where such an emergency scenario could occur and how many people would be affected, there is a significant uncertainty in demand values. The aim is to identify locations and a way of distributing supplies that will be effective in meeting the uncertain demand. In an emergency situation an overriding objective is to service as much of the demand as possible. In such a situation, an effective placement of supplies so that they service demand is more important than an accurate model of how demand is distributed among facilities. Finally in this work we aim to include the aggregate effect of uncertainty in the demand at different facilities and leave for future work the inclusion of more detailed demand distribution models, such as a consumer-choice demand model.

## Acknowledgments

This research was partly supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) at USC. The findings, conclusions and recommendations presented in this paper are those of the authors, and do not necessarily reflect the views of the Department of Homeland Security. The authors would like to thank Hongzhong Jia, Yen-Ming Lee, Wei Ye and Zhihong Shen for their help and valuable inputs through the course of this research work. We also thank the anonymous referees for their valuable suggestions to help us improve the paper.

## References

- Gendreau M, Laporte G, Semet F. The maximal expected coverage relocation problem for emergency vehicles. *J Opl Res Soc* 2006;57:22–8.
- Jia H, Ordóñez F, Dessouky MM. A modeling framework for facility location of medical services for large-scale emergencies. *IIE Trans* 2006;39:41–55.
- Sherali HD, Carter TB. A location-allocation model and algorithm for evacuation planning under hurricane/flood conditions. *Transport Res Part B* 1991;25:439–52.
- Church R, ReVelle C. The maximum covering location problem. *Papers of the Reg Sci Association* 1974;32:101–18.
- Berman O, Krass D. Generalized maximum cover location problem. *Comput Oper Res* 2002;29:563–81.
- Berman O, Krass D, Drezner Z. The gradual covering decay location problem on a network. *Eur J Opl Res* 2003;151:474–80.
- Berman O, Drezner Z, Krass D, Wesolowsky GO. The variable radius covering problem. *Eur J Opl Res* 2009;196:516–25.
- Hodgson MJ, Jacobsen SK. A hierarchical location-allocation model with travel based on expected referral distances. *Ann Oper Res* 2009;176:271–86.
- Berman O, Krass D. Flow intercepting spatial interaction model: a new approach to optimal location of competitive facilities. *Location Sci* 1998;6:41–65.
- Aboolian R, Berman O, Krass D. Competitive facility location and design problem. *Eur J Opl Res* 2007;182:40–62.
- Toregas C, Swain R, ReVelle C, Bergman L. The location of emergency service facility. *Oper Res* 1971;19:1363–73.
- Rawls CG, Turnquist MA. Pre-positioning of emergency supplies for disaster response. *Transport Res Part B* 2010;44:521–34.
- Berman O, Gavious A. Location of terror response facilities. *Eur J Opl Res* 2007;177:1113–33.
- Yi W, Özdamar L. A dynamic logistics coordination model for evacuation and support in disaster response activities. *Eur J Opl Res* 2007;179:1177–93.
- Görmez N, Köksalan M, Salman FS. Locating disaster response facilities in Istanbul. *J Opl Res Soc*; 2010:1–14.
- Balcik B, Beamon BM. Facility location in humanitarian relief. *Int J Logist Res Appl* 2008;11:101–21.
- Chang MS, Tseng YL, Chen JW. A scenario planning approach for the flood emergency logistics preparation problem under uncertainty. *Transport Res Part E* 2007;43:737–54.
- Snyder L. Facility location under uncertainty: a review. *IIE Trans* 2006;38(7):537–54.
- Cooper L. A random locational equilibrium problem. *J Regional Sci* 1974;14(1):47–54.

- [20] Sheppard ES. A conceptual framework for dynamic location-allocation analysis. *Environ Plan A* 1974;6:547–64.
- [21] Mirchandani PB, Oudjit A. Localizing 2-medians on probabilistic and deterministic tree networks. *Networks* 1980;10(4):329–50.
- [22] Balachandran V, Jain S. Optimal facility location under random demand with a general cost structure. *Nav Res Logist Q* 1976;23:421–36.
- [23] Berman O, Odoni AR. Locating mobile servers on a network with Markovian properties. *Networks* 1982;12:73–86.
- [24] Weaver JR, Church RL. Stochastic p-robust location problems. *Transport Sci* 1983;17(2):168–80.
- [25] Mirchandani PB, Oudjit A, Wong RT. Multi-dimensional extensions and a nested dual approach for the m-median problem. *Eur J Opl Res* 1985;21(1):121–37.
- [26] Louveaux FV. Discrete stochastic location models. *Ann Oper Res* 1986;6:23–34.
- [27] Louveaux FV, Peeters D. A dual-based procedure for stochastic facility location. *Oper Res* 1992;40(3):564–73.
- [28] Berman O, Drezner Z. The P-median problem under uncertainty. *Eur J Opl Res* 2008;189:19–30.
- [29] Jucker JV, Carlson RC. Public facilities location under a stochastic demand. *Oper Res* 1976;24(6):1045–55.
- [30] Jucker JV, Carlson RC. The simple plant location problem under uncertainty. *Oper Res* 1976;24(6):1045–55.
- [31] Hodder JE, Dincer MC. A multifactor model for international plant location and financing under uncertainty. *Comput Oper Res* 1986;13(5):601–9.
- [32] Snyder L, Daskin MS. Stochastic p-robust location problems. *IIE Trans* 2006;38(11):971–85.
- [33] Larson RC, Brandeau M. Extending and applying the hypercube queuing model to deploy ambulances in boston. In: Swersey A, Ignall E, editors. *Delivery of urban services*; 1986.
- [34] Jia H, Ordóñez F, Dessouky MM. Solution approached for facility location of medical supplies for large-scale emergencies. *Comput Indus En* 2007;52:257–76.
- [35] Cooper L. Heuristic methods for location-allocation problems. *SIAM Rev* 1964;6:37–53.
- [36] Densham P, Rushton G. Designing and implementing strategies for solving large location-allocation problems. tech. rep. National Center for Geographic Information and Analysis; 1991.
- [37] Taillard ÉD. Heuristic methods for large centroid clustering problems. *J Heuristics* 2003;9:51–73.
- [38] Bravata DM, Zaric GS, Holty J-EC, Brandeau ML, Wilhelm ER, McDonald KM, et al. Reducing mortality from anthrax bioterrorism: costs and benefits of alternative strategies for stockpiling and dispensing medical and pharmaceutical supplies. *Biosecur Bioterror, Biodefense Strategy, Practice Sci* 2006;4:244–62.
- [39] Berman O, Drezner Z. Location of congested capacitated facilities with distance-sensitive demand. *IIE Trans* 2006;38:213–31.
- [40] Reinfield NV, Vogel WR. *Mathematical programming*. Englewood Cliffs, New Jersey: Prentice-Hall; 1958.

**Pavankumar Murali** is a Research Staff Member at the IBM T. J. Watson Research Center in New York. He obtained his Ph.D and M.S. degrees in Industrial Engineering and Operations Research from the University of Southern California, and an undergraduate degree in Mechanical Engineering from the Indian Institute of Technology – Madras. His research interests lie in the areas of integer programming, combinatorial optimization, and design and analysis of algorithms. His work includes the applications of these tools to problems in transportation, logistics and supply chain.

**Fernando Ordóñez** is an Associate Professor at the University of Southern California. He obtained his Ph.D degree in Operations Research from the Massachusetts Institute of Technology, and an undergraduate degree in Mathematical Engineering from the University of Chile. His interests include convex optimization, robust optimization, complexity of algorithms, sensitivity analysis, condition number theory, and applications of optimization to engineering and management science.

**Maged Dessouky** is a Professor at the University of Southern California. He received his Ph.D in Industrial Engineering and Operations Research from the University of California - Berkeley, and his M.S. and undergraduate degrees in Industrial Engineering from Purdue University. His research has been funded by the National Science Foundation, Society of Manufacturing Engineers, PATH, Caltrans, FTA, and Department of Defense. He was recipient of the 2007 Transportation Science & Logistics Best Paper Prize. He is a Fellow of the Institute of Industrial Engineers.