Proceedings of the 6th INFORMS Workshop on Data Mining and Health Informatics (DM-HI 2011)
P. Qian, Y. Zhou, C. Rudin, eds.

A Planning Model for Public Health Emergency Response Logistics

Kathleen King
School of Operations Research and Information Engineering
Cornell University
Ithaca, NY
kak59@cornell.edu

John Muckstadt
School of Operations Research and Information Engineering
Cornell University
Ithaca, NY

Abstract

The United States' response plan for many types of public health emergencies requires the large-scale distribution of medical supplies to the affected populations. We discuss how this system would operate in response to an aerosol anthrax attack, and we present a model of the planned three echelon distribution network. One feature of the distribution network that makes it unique from commercial systems is the existence of service capacity constraints at the lower echelon clinics where individuals receive treatment. We describe several methods for making medical supply allocation decisions and discuss how the model can be used by public health officials to assist in policy decision-making.

Keywords: emergency preparedness, emergency response, anthrax, supply chain logistics

Introduction

Public health emergencies, such as bioterrorist attacks and pandemics, require carefully coordinated logistical responses at the federal, state, and local levels to ensure that the affected populations receive care as quickly as possible. Many types of emergencies require large-scale distribution of vaccines, antibiotics, or antiviral medication for prophylaxis or treatment to minimize potential mortality and morbidity. Planning and managing distribution networks in emergency settings are challenging tasks due to the high levels of uncertainty in the timing, scale, and location of such events. To help policy makers address these problems, we have developed models of both the planning and operational phases of public health emergencies which give quantitative estimates of the resources and staff required by various response plans; show the probable consequences of dispensing network designs; and provide tools for decision-making with limited information. In this paper we will present a planning model of the United States' mass dispensing emergency response network and discuss how the model may be used to guide policy decisions for response to an aerosol anthrax attack, which has been identified as one of the most significant bioterrorist threats [8].

The United States Department of Health and Human Services and the Centers for Disease Control and Prevention (CDC) created the Strategic National Stockpile (SNS), formerly the National Pharmaceutical Stockpile in 1999 to ensure that essential pharmaceuticals and medical supplies will be available to the American population during emergencies [5]. Once an emergency such as an anthrax attack has been detected, the SNS will be activated. Individuals contract aerosol anthrax upon breathing in anthrax spores, and mortality rates for patients who inhale large quantities of spores increase drastically after about 48 hours [11]. There is no fast method for testing whether an asymptomatic individual has inhaled enough spores to contract anthrax, but individuals who receive antibiotic prophylaxis before the onset of symptoms have a high likelihood of survival [11]. Thus, the CDC's goal is to ensure that all individuals exposed to anthrax spores begin a course of antibiotic prophylaxis within 48 hours of the decisions to deploy the SNS [4].

The SNS stores its inventory in multiple warehouses around the country to ensure that it can deliver an initial shipment of antibiotics and medical supplies to any affected state within twelve hours [5]. Subsequent shipments may be sent over time, depending on the number of people affected by the attack. Each state has at least one Receiving Storing and Staging (RSS) warehouse, which will be opened to accept materials from the SNS and send them to the appropriate dispensing locations. These Points of Dispensing (PODs) may be ad hoc clinics set up by state or local health departments to distribute antibiotics to the general public, or they may be "closed" locations that serve well-defined populations such as hospital patients, nursing home residents, or prison inmates [6].

Given the potentially severe consequences of failing to provide antibiotics quickly to exposed individuals, it is essential that the distribution network described above operates efficiently, but managing the logistics of the response effort is difficult. In the early phase of the emergency, supplies of antibiotics, transportation capacity, and clinic staffing levels may be limited, so all of these resources must be used carefully. As part of the long term planning process, our goals are to assist policy-makers in determining the effectiveness of potential distribution network design plans under various patient arrival patterns; understanding the value of increasing various resources; and recognizing the value of sharing information across the network. When an emergency actually occurs, we need to help planners determine how to make decisions quickly with limited information. Public health officials need tools that will help them adjust to uncertainty over time and make the best decisions possible given limited information.

A detailed review of emergency preparedness and response literature is presented by [2], but we will briefly explain where our model fits in to this framework. Much of the anthrax emergency response literature has focused on long term planning questions about the cost-effectiveness of various prevention and mitigation measures [1, 3, 20, 7, 10]. Another stream of of literature has worked to optimize POD operations [9, 16]. Some papers have considered logistical planning problems like facility location or the large-scale estimation of resource and dispensing capacity needs [17, 19, 21]. In contrast, our work focuses on more detailed logistics of the emergency response network, particularly the supply chain operations. To the best of our knowledge, the only other study that carefully models emergency supply chains is [18], but their model is for more general emergencies and does not include some key features of the SNS supply chain. Our model closely represents the structure and operations of the SNS supply chain; we have developed approximate inventory allocation methodologies; and we describe how the model can assist in public health policy decision-making.

In the following sections, we formulate the planning model as an inventory allocation dynamic program and briefly describe several solution techniques. We then discuss how modeling various network structures, decision-making policies, and resource capacities can guide public health authorities in determining emergency preparedness policy.

Inventory Allocation Planning Model

The antibiotic distribution network planned by the CDC and state departments of health is a three echelon supply chain. We model the SNS as a single location, called location 0, since the SNS inventory is centrally managed, and it is likely that during an anthrax attack only one SNS warehouse would respond. The network also contains M RSSs (locations $1, ..., M$) and N PODs (locations $M+1, ..., M+N$). Let $\mathcal{R} = \{1, ..., M\}$ be the set of RSSs, $\mathcal{P} = \{M+1, ..., M+N\}$ be the set of all PODs, and $\mathcal{L} = \{0, ..., M+N\}$ be the set of all locations. The indicator variable u_{mn} is 1 if RSS m serves POD n and 0 otherwise, and $u_{0m} = 1$ for all of the RSSs, $m \in \mathcal{R}$, since the SNS serves each RSS. Time is divided into discrete periods. The SNS may be resupplied by external vendors over the course of the emergency; the lead time from the vendors to the SNS is τ_0 and the maximum inventory that may be shipped to the SNS in period t is p_t^0. The lead time from the SNS to RSS m is τ_m, where $\tau_m = \tau_1$ for all $m \in \mathcal{R}$, and the shipping capacity from the SNS to all of the RSSs in period t is p_{0t}. The shipping capacity from RSS m to its PODs in period t is p_{mt}, and the lead time from an RSS to a POD is assumed to be less than one time period. There is an inventory processing delay at the SNS and RSSs, so inventory that arrives to any of these locations in period t may not be shipped out until time $t+1$. However, inventory that arrives at the PODs in period t may be used immediately to serve patients. The number of patients that can be served at POD n in period t is limited by the available inventory, the number of patients present, and the service capacity, a_{nt}. Because of the service capacity constraint, a POD may have both inventory and unserved patients remaining at the end of a time period.

We will let x_{nt} represent the echelon on-hand inventories at each location n at the beginning of period t. At the PODs, x_{nt} is the on-hand inventory. At an RSS m, x_{mt} includes the inventory on-hand at RSS m plus all inventory on-hand at PODs served by n plus inventory currently in-transit to RSS m. At the SNS, x_{0t} is the sum of all inventory

on-hand or in-transit to any location in the distribution network. After the inventory allocation decisions are made and inventory is shipped in period t, the echelon on-hand inventory position increases to y_{nt}; no returns are allowed, so we always have $y_{nt} \geq x_{nt}$. Note that vectors are indicated by bold font and suppressing subscripts; for example, $\mathbf{x_t} = (x_{0t}, x_{1t}, ..., x_{M+N,t})$. The number of patients waiting in a queue at POD n at the beginning of time period t who arrived in previous periods but have not yet been served is q_{nt}. The number of echelon unserved patients at each RSS m is $q_{mt} = \sum_{n \in \mathcal{P}} u_{mt} q_{nt}$, and the SNS echelon unserved patients is $q_{0t} = \sum_{n \in \mathcal{P}} q_{nt}$. The number of patients who arrive to POD n in period t is a random variable D_{nt} with density function $f_{nt}(\cdot)$, and the echelon demands for $m \in \{0\} \cup \mathcal{R}$ are $D_{mt} = \sum_{n \in \mathcal{L}} u_{mn} D_{nt}$. For this planning model, we assume that these demands are independent across time and locations, but we relax this assumption in our operational models which consider costs incurred over shorter time horizons [13]. In time period t', the demands for periods $t = 1, ..., t' - 1$ have been realized, and we represent these values by $d_{n1}, ..., d_{n,t'-1}$. The cumulative demand at location n for periods $t_1, ..., t_2$ is $d_{n,t_1,t_2} = \sum_{t'=t_1}^{t_2} d_{nt'}$.

Inventory allocation decisions $\mathbf{y_t}$ must be made before the new patient demands are observed in each time period. These decisions are constrained by the available on-hand inventory at SNS and RSSs at the beginning of each period and by the shipping capacities. The on-hand inventory at location $m \in \{0\} \cup \mathcal{R}$ in period t is given by the inventory that was shipped to location m one lead time ago, less the patients that were served in the interim plus the patients still waiting at the PODs downstream from m, less the inventory that is currently on-hand downstream:

$$\text{On-hand inventory at } m \text{ at the beginning of } t \quad = \quad x_{m,t-\tau_m} - q_{m,t-\tau_m} - d_{m,t-\tau_m,t-1} + q_{mt} - \sum_{n \in \mathcal{L}} u_{mn} x_{nt}. \quad (1)$$

Notice that calculating the on-hand inventory level requires knowing past values for the echelon on-hand inventories, unserved patients, and patient demands. Define $\mathbf{x_t^{past}} = (\mathbf{x_{t-\tau_0}}, ..., \mathbf{x_t})$ to be the vector of past on-hand echelon inventory values and define $\mathbf{q_t^{past}}$ similarly. The vector of past demands is defined similarly, but it does not include $\mathbf{D_t}$, since this value is not yet known at the beginning of period t. Using these variables, we can state the set of feasible inventory allocations in time period t:

$$\begin{aligned} \mathcal{U}_t(\mathbf{x_t^{past}}, \mathbf{q_t^{past}}, \mathbf{d_t^{past}}) \quad = \quad & \Big\{ \mathbf{y_t} : y_{nt} \geq x_{nt} \quad \text{for } n \in \mathcal{L}; \\ & \sum_{n \in \mathcal{L}} u_{mn}(y_{nt} - q_{nt}) \leq x_{m,t-\tau_m} - q_{m,t-\tau_m} - d_{m,t-\tau_m,t-1} \text{ for } m \in \{0\} \cup \mathcal{R}; \\ & \sum_{n \in \mathcal{L}} u_{mn}(y_{nt} - x_{nt}) \leq p_{mt} \quad \text{for } m \in \{0\} \cup \mathcal{R}; \\ & y_{0t} - x_{0t} \leq p_t^0 \Big\}. \end{aligned} \quad (2)$$

The model considers two types of costs: inventory holding costs and patient delay penalties. The inventory holding costs are linear in the number of unused units of inventory on-hand at each location at the end of each time period. However, linear costs are not appropriate for patient delay penalties. A few extra patients waiting to be served at the beginning of the next period are not problematic for POD operations, but as the number waiting grows larger, crowd management could become a serious challenge. Hence, we use patient delay penalties that are convex and increasing in the number of unserved patients at each POD at the end of each period. We let h_{nt} be the per unit holding cost charged at location n at the end of period t and $B_{nt}(\cdot)$ be the penalty function for unserved (or "backordered") patients. At the end of time period t, the inventory on-hand at POD n is y_{nt} minus the number of patients served during the period, and the number of unserved patients is given by q_{nt} plus the new demands less the number served. The number of people served in period t at POD n is

$$S_{nt}(y_{nt}, q_{nt}, D_{nt}) = \min(y_{nt}, q_{nt} + D_{nt}, a_{nt}). \quad (3)$$

At the SNS and the RSSs, the remaining inventory that is either on-hand or in-transit is given by expression (1) plus the new allocations less the inventory sent to lower echelons. So the total expected cost expressions for the end of period t are

$$\overline{C}_{nt}(y_{nt}, q_{nt}) = E\Big[h_{nt}\Big(y_{nt} - S_{nt}(y_{nt}, q_{nt}, D_{nt})\Big) + B_{nt}\Big(q_{nt} + D_{nt} - S_{nt}(y_{nt}, q_{nt}, D_{nt})\Big)\Big] \text{ for } n \in \mathcal{P}; \quad (4)$$

$$\overline{C}_{mt}(y_{m,t-\tau_m}, \mathbf{y_t}) = h_{mt}\Big(y_{m,t-\tau_m} - q_{m,t-\tau_m} - E[D_{m,t-\tau_m,t-1}] - \sum_{n\in\mathcal{L}} u_{mn}(y_{nt} - q_{nt})\Big) \quad \text{for } m \in \{0\} \cup \mathcal{R}. \quad (5)$$

The last allocation decision that affects the cost incurred at location n in period $t + \tau_n$ is made in period t. Costs incurred at location n before time $\tau_n + 1$ are unavoidable fixed costs. Since there are no decisions that we can make to mitigate the fixed costs, we will define a slightly different cost function that does not include these costs. Our modified cost function $C(\cdot)$ will evaluate the consequences that a set of decisions $\mathbf{y_t}$ made in time t will have on the future. We define $C(\cdot)$ as follows:

$$C_{0t}(y_{0t}, q_{0t}) = h_{0,t+\tau_0}(y_{0t} - q_{0t}) - h_{0,t+\tau_0}E[D_{0,t,t+\tau_0-1}]; \quad (6)$$

$$C_{mt}(y_{mt}, q_{mt}) = (h_{m,t+\tau_1} - h_{0t})(y_{mt} - q_{mt}) - h_{m,t+\tau_m}E[D_{m,t,t+\tau_m-1}] \quad \text{for } m \in \mathcal{R}; \quad (7)$$

$$C_{nt}(y_{nt}, q_{nt}) = \overline{C}_{nt}(y_{nt}, q_{nt}) \quad \text{for } n \in \mathcal{P}; \text{ and} \quad (8)$$

$$C_t(\mathbf{y_t}, \mathbf{q_t}) = \sum_{n=0}^{M+N} C_{nt}(y_{nt}, q_{nt}), \quad (9)$$

where $h_{nt} = 0$ and $D_{nt} = 0$ for $t > T$ and $n \in \mathcal{L}$.

During an emergency scenario, the goal of public health officials in time period t would be to make allocation decisions that minimize the holding costs and unserved patient penalties for the remainder of the time horizon. We model that decision process as a dynamic program whose value function is

$$V_t(\mathbf{x_t^{past}}, \mathbf{q_t^{past}}, \mathbf{d_t^{past}}) = \min_{\mathbf{y_t}\in\mathcal{U}_t(\mathbf{x_t^{past}},\mathbf{q_t^{past}},\mathbf{d_t^{past}})} C_t(\mathbf{y_t}, \mathbf{q_t}) + E\big[V_{t+1}(\mathbf{x_{t+1}^{past}}, \mathbf{q_{t+1}^{past}}, \mathbf{d_{t+1}^{past}})\big]\Big). \quad (10)$$

Solving problem (10) would allow us to make the best possible allocation decisions in period t. We can then run computer simulations in which we make these decisions to explore the consequences of various anthrax scenarios and different constraints on the antibiotic dispensing network. However, due to the curse of dimensionality, we cannot solve problem (10) exactly for realistically large data sets. Instead, we consider three methods for making allocation decisions, which we describe briefly below; for more details see [12].

1. **Truncated Cumulative Approximation (TCA) Method.** Problem (10) is easier to solve if we consider only a truncated time horizon rather than considering every future time period through T. In time period t, we can make decisions that minimize only the expected costs for periods $t, t+1, ..., t+k$. If we consider too few time periods (k is too small), then we will fail to capture the impact that inventory ordering decisions at the upper echelons will eventually have on service levels and unserved patient penalty costs at the PODs. To ensure that our decisions are reasonable, we must have $k \geq \tau_0 + \tau_1 + 2$. We further reduce the complexity of the problem by estimating the cumulative number of people served over time using the minimum of the cumulative available inventory, service capacity, and demand. This gives an upper bound on the actual number of people served during that time period, and this approximation simplifies the truncated problem sufficiently so that we can solve it in each time period.

2. **Lagrangian Relaxation and Decomposition (LRD) Method.** Instead of truncating problem (10), we can use Lagrangian relaxation to simplify the structure of the problem. The new problem can be decomposed by location into $M + N + 1$ single variable dynamic programs, each of which is straightforward to solve. Solving the dynamic programs for a wide range of possible state variables provides an estimate for the value function for every possible state. Then, in each time period t, we can replace the expected future cost with the estimated value function. This method is an extension of the methods presented in [14, 15].

3. **Fair Share Method.** The current method that public health officials plan to use for inventory allocation is called the "fair share" method. In this method, inventory is dispensed based on population. In each time period, the available inventory at the upper echelons is pushed out based on the population that each region is expected to serve. The inventory allocation does not depend on the actual demands that are observed over time; decisions only depend on the total population that is expected to arrive over the time horizon.

For various simulated disaster scenarios and response networks, we make decisions using these methods. The simulations are straightforward; in each time period we make allocation decisions based on the state of the system, then randomly generate patient demands and calculate the resulting costs and updated state variables. We continue this over the entire time horizon and then repeat the whole process multiple times, with different random demand drawn from the same distributions to obtain a confidence interval for the cost of using a particular allocation method for a particular set of emergency parameters and a given dispensing network. The relative success of these three methods is described in detail in [12], but we note that the costs incurred by the Fair Share Method are at least twice the costs incurred by the first two methods (TCA and LRD) in approximately 80% of demand scenarios considered. The TCA and LRD methods provide comparable performance (costs within 20%) under about half of the scenarios run. In the remaining scenarios, the LRD method performs better when demand varies significantly and the TCA performs better in scenarios when the demand distributions are closer to being stationary. In the following section, we will describe a couple examples of how the model and solution methods described can be used to assist policy-makers.

Policy Implications of the Model

As we have mentioned earlier, this model is intended to be a tool that allows public health officials to explore the performance of different distribution network structures and the relative value of different resources. Policy-makers can test the relative value of investing limited funds in stockpiling inventory, increasing staffing levels, or adding transportation capacity by changing these parameters and running the simple simulation we described above to determine the costs associated with each scenario. They can investigate the relative improvements in patient delays given increases in each of these capacities. They can also obtain confidence intervals for the total inventory, staff time, and transportation capacity that may be required under a variety of anthrax attack scenarios.

The model can be used to address questions about the topology of the dispensing network. There is significant political pressure on policy-makers to open more PODs to ensure that all neighborhoods are equally well-served and the distribution network appears "fair." We can use the model to help policy makers explore the potential consequences, in terms of patient delays and inventory costs, of constructing fewer larger PODs compared to many smaller PODs. It is essential to ensure that all people are able to reach the PODs in a manner that is not overly burdensome, but opening excessive numbers of PODs would certainly decrease the efficiency of the distribution network since staff and initial shipments of antibiotics are limited. Our model provides quantitative evidence to support the need for a moderate number of larger PODs in an anthrax response scenario, since the number of people to be served will be large and the length of the time horizon will be very short. Requiring individuals to travel slightly farther for service in some cases could save lives.

We offer significant reasons for public health authorities to invest in command and control infrastructure. A strong command and control system would allow the New York City area to share staff between PODs, so that staffing levels can vary over time in response to patient demand. We can show that the same total staffing capacity in the system can provide much better service if the staffing levels at each POD can be adjusted periodically. Also, a command and control system that supports detailed inventory information sharing allows for much more intelligent inventory allocation. We show that the Fair Share Method described above wastes large quantities of inventory and creates unnecessarily large patient delays compared to the TCA and LRD methods. Since there is also political pressure to divide inventory by population rather than need, the quantitative evidence provided by our model can help support the decision to use resources more efficiently and show that this will provide much better service for all patients.

References

[1] Baccam, P. and M. Boechler, 2007, "Public Health Response to an Anthrax Attack: An Evaluation of Vaccination Policy Options," Biosecurity and Bioterrorism, 5(1), 2634.

[2] Brandeau, M.L., J.H. McCoy, N. Hupert, J.E. Holty, and D.M. Bravata, 2009, "Recommendations for Modeling Disaster Responses in Public Health and Medicine: A Position Paper for the Society for Medical Decision Making," American Journal of Disaster Medicine, 29(4), 438-460.

[3] Brookmeyer, R., E. Johnson, and R. Bollinger, 2004, "Public Health Vaccination Policies for Containing an Anthrax Outbreak," Nature, 432, 901904.

[4] Centers for Disease Control and Prevention, 2004, "Cities Readiness Initiative Guidance, Appendix 3," http://www.bt.cdc.gov/planning/guidance05/pdf/appendix3.pdf [Accessed September 2, 2011].

[5] Centers for Disease Control and Prevention, 2011, "Strategic National Stockpile," http://www.bt.cdc.gov/stockpile/ [accessed September 2, 2011].

[6] Coordinating Office For Terrorism Preparedness And Emergency Response, Division Of Strategic National Stockpile, 2008, "Point of Dispensing (POD) Standards," http://health.mo.gov/emergencies/sns/pdf/12-PODStandards.pdf, [Accessed September 2, 2011].

[7] Fowler, R. A. and S. Shafazand, 2011, "Anthrax Bioterrorism: Prevention, Diagnosis and Management Strategies," Journal of Bioterrorism and Defense, 2(2), doi = 10.4172/2157-2526.1000107.

[8] Henderson, D. A., 1999, "The Looming Threat of Bioterrorism," Science, 283, 1279-1282.

[9] Hupert, N., W. Xiong, K.A. King, M. Castorena, C. Hawkins, C. Wu, and J.A. Muckstadt, 2009, "Uncertainty and Operational Considerations in Mass Prophylaxis Workforce Planning," Disaster Medicine and Public Health Preparedness, 3(2), S121.

[10] Hupert, N., D. Wattson, J. Cuomo, E. Hollingsworth, K. Neukermans, and W. Xiong, 2009, "Predicting Hospital Surge after a Large-Scale Anthrax Attack: A Model-Based Analysis of CDC's Cities Readiness Initiative Prophylaxis Recommendations," Medical Decision Making, 29, 424-437.

[11] Inglesby, T. V. et al., "Anthrax as a Biological Weapon, 2002: Updated Recommendations for Management," Journal of the American Medical Association, 287(17), 2236-2252.

[12] King, K. A. and J. A. Muckstadt, [in prep], "Planning Models for Public Health Emergency Response."

[13] King, K. A. and J. A. Muckstadt, [in prep], "Models for Operational Decision-Making During Public Health Emergencies."

[14] Kunnumkal, S. and H. Topaloglu, 2008, "A Duality-Based Relaxation and Decomposition Approach for Inventory Distribution Systems," Naval Research Logistics Quarterly, 55(7), 612-631.

[15] Kunnumkal, S. and H. Topaloglu, 2011, "Linear Programming Based Decomposition Methods for Inventory Distribution Systems," European Journal of Operational Research, 211(2), 282-297.

[16] Lee, E.K., S. Maheshwary, J. Mason, and W. Glisson, 2006, "Decision Support System for Mass Dispensing of Medications for Infectious Disease Outbreaks and Bioterrorist Attacks," Annals of Operations Research, 148(1), 25-53.

[17] Lee, E.K., H.K. Smalley, Y. Zhang, F. Pietz, and B. Benecke, 2009, "Facility Location and Multi-Modality Mass Dispensing Strategies and Emergency Response For Biodefence and Infectious Disease Outbreaks," International Jourtnal of Risk Assessment and Management, 12(2-4), 311-351.

[18] Lee, Y. M., 2008, "Analyzing Dispensing Plan for Emergency Medical Supplies in the Event of Bioterrorism," Proceedings of the 2008 Winter Simulation Conference, Edited by S. J. Mason et al., December 7-10, Miami, FL.

[19] Savachkin, A. and A. Uribe, 2011, "Dynamic Redistribution of Mitigation Resources During Influenza Pandemics," Socio-Economic Planning Sciences, doi = 10.1016/j.seps.2011.05.001.

[20] Wein, L.M. and D.L. Craft, 2005, "Evaluation of Public Health Interventions for Anthrax: A Report to the Secretary's Council on Public Health Preparedness," Biosecurity and Bioterrorism, 3(4), 348-356.

[21] Zaric, G.S., D.M. Bravata, J.C. Holty, K.M. McDonald, D.K. Owens, and M.L. Brandeau, 2008, "Modeling the Logistics of Response to Anthrax Bioterrorism," Medical Decision Making, 28, 332-250.