A Routing and Scheduling Approach for Planning Medication Distribution

Jeffrey W. Herrmann
Sara Lu
Kristen Schalliol
Institute for Systems Research
University of Maryland, College Park, MD 20742, USA

Abstract

This paper presents a two-stage approach for solving the medication distribution problem. The problem addresses a critical issue in emergency preparedness. Public health officials must plan the logistics for distributing medication to points of dispensing (PODs), which will give medication to the public in case of a bioterrorist attack such as anthrax. We consider the problem at the state and local levels. Our approach separates the problem into two subproblems: (1) the "routing problem" assigns PODs to vehicles and creates routes for each vehicle, and (2) the "scheduling problem" determines when the vehicles should start these routes and how much material should be delivered on each trip. This paper describes the approach and presents the results of using this approach to construct solutions for two scenarios.

Keywords
Logistics and inventory, vehicle routing, scheduling, homeland security

1. Introduction
Improving emergency preparedness requires planning responses to bioterrorist attacks. In the case of a large scale bioterrorist event, such as the release of anthrax, public health officials may decide that mass dispensing of medication is needed. According to the Centers for Disease Control and Prevention, large cities and metropolitan areas need to dispense antibiotics to their entire identified population within 48 hours of the decision to do so [1]. Cities in every state are developing plans for opening points of dispensing (PODs) to give prophylactic medication to persons who are currently healthy but may have been exposed to a pathogen. PODs may be setup in schools, recreation centers, churches, and other non-medical facilities. Other modes of dispensing medication are being considered, but PODs are the primary focus of planning activities.

The proposed research is motivated by work with county public health departments in the state of Maryland who must plan the logistics for distributing medication to the PODs from a central location. We consider the problem at the state and local levels (not the national level). After the decision for mass dispensing is made, county public health departments will begin preparing to open multiple PODs simultaneously at a designated time. The state will request medication from the federal government, who will deliver an initial but limited supply of medication to a state receipt, storage, and stage (RSS) facility (which we call the "depot"). Contractors will deliver more medication to the depot, but the state will begin shipping medication from the depot to the PODs before everything arrives from the contractors. The deliveries to the depot arrive in batches that we call "waves."

Poor medication distribution plans will delay the time that some PODs receive medication. This can delay the opening of these PODs, and some residents may not get their medication in a timely manner, which increases their risk of death or illness. Clearly, there are many uncertainties in medication distribution, including the timing of shipments to the depot, the time needed to load and unload trucks, travel times, and the demand for medication at each POD. For this reason, planners need a robust plan. In particular, it is better if the plan calls for delivering medication to PODs much earlier than it is needed. This improves the likelihood that the PODs will open on-time, will not run out of medication during operations, and will dispense medication to the largest number of people in a timely manner.

The operations of firefighters, emergency medical services, and police departments have motivated research into location models [2-4] and dynamic vehicle routing models [5-7]. However, these models are not relevant to the medication distribution problem, which is more closely related to the inventory routing problem [8-11] and the production-distribution scheduling problem [12]. Still, the models used for those problem are also not directly relevant.

This paper addresses the single-product, deterministic problem. Inventory is treated as a continuous variable, but the number of pallets must be an integer. We measure the medication with the number of regimens. In mass dispensing, each person will get one predetermined regimen, which is a bottle with a specific number of pills. All PODs have the same hours of operation, and loading and unloading times are independent of the quantity. We are ignoring other resources such as the loading docks at the depot, the available drivers, and the number of available pallets. The paper formulates the problem, presents a two-stage approach for constructing solutions, and discusses the results of applying this approach to two scenarios. More details about the problem and the approach can be found in [13].

2. Problem Formulation

A problem instance specifies the following information. Without loss of generality we let time $t = 0$ correspond to the first instant that the depot has medication. PODs will begin operating at time $t = T_1$ and continue to operate until time $t = T_2$. In practice, these times may be on the order of 12 to 48 hours.

There are n PODs (sites). Each site ($k = 1, \ldots, n$) has a dispensing rate of L_i regimens per time unit. This is the rate at which the site consumes medication. The site needs a total of $(T_2 - T_1) L_i$ regimens. There is a depot ($k = 0$) that has a supply of medication. Let $I(t)$ be the cumulative amount of medication delivered to the depot at time t. $I(t)$ is a discontinuous, non-decreasing function due to the batch deliveries that are made there.

The time spent at site i (to load or unload a vehicle) is p_i for $i = 0, \ldots, n$. The time to go from site i to site j is c_{ij}. There are V vehicles. Vehicle v has a capacity of C_v pallets of material. At each site, a vehicle will deliver one or more pallets. A pallet can hold at most P regimens.

Given a problem instance, a solution specifies one or more routes for each vehicle and the quantity delivered at each site. The key decision variables are the sequences, start times, and delivery quantities. Let r_v be the number of routes that vehicle v makes. For the j-th route for vehicle v, $m(v, j)$ is the number of sites on the route, $\sigma_{vj} = \{i_1, ..., i_{m(v,j)}\}$ is the sequence of sites that the vehicle visits, and t_{vj} is the start time at which the vehicle begins loading at the depot. Finally, q_{vjk} is the quantity delivered to each site $k \in \sigma_{vj}$. Note that $q_{vjk} = 0$ if and only if $k \notin \sigma_{vj}$.

Given a solution, we can evaluate its feasibility as follows. Let y_{vj} be the total duration of a route. Let w_{vjk} be the duration between the start of the route and the time that the delivery at site $k \in \sigma_{vj}$ is complete. Let h_{vjk} be the (integer) number of pallets required to deliver q_{vjk} regimens to site $k \in \sigma_{vj}$.

$$y_{vj} = p_0 + c_{0i_1} + p_{i_1} + c_{i_1 i_2} + \cdots + p_{i_{m(vj)}} + c_{i_{m(v,j)} 0} \quad (1)$$

$$w_{vjk} = p_0 + c_{0i_1} + p_{i_1} + c_{i_1 i_2} + \cdots + p_k \quad (2)$$

Certain constraints must be satisfied for the solution to be feasible. The quantity shipped from the depot cannot exceed the amount delivered to the depot (Equation 3). A vehicle cannot begin a new route until it returns to the depot (Equation 4). The number of pallets used must be sufficient (Equation 5). The vehicle capacity cannot be exceeded on any route (Equation 6). Each and every site must receive all needed medication (Equation 7). All route start times must be non-negative (Equation 8).

$$\sum_{(a,b):t_{ab} \le t_{vj}} \sum_{k \in \sigma_{ab}} q_{abk} \le I\left(t_{vj}\right) \quad v = 1, \ldots, V;\ j = 1, \ldots, r_v \quad (3)$$

$$t_{vj} \ge t_{v,j-1} + y_{v,j-1} \quad v = 1, \ldots, V;\ j = 2, \ldots, r_v \quad (4)$$

$$0 \le q_{vjk} \le Ph_{vjk} \quad v = 1,\ldots,V\ ;\ j = 1,\ldots,r_v \tag{5}$$

$$\sum_{k\in\sigma_{vj}} h_{vjk} \le C_v \quad v = 1,\ldots,V\ ;\ j = 1,\ldots,r_v \tag{6}$$

$$\sum_{v=1}^{V}\sum_{j=1}^{r_v} q_{vjk} = \left(T_2 - T_1\right) L_k \quad k = 1,\ldots,n \tag{7}$$

$$t_{vj} \ge 0 \quad v = 1,\ldots,V\ ;\ j = 1,\ldots,r_v \tag{8}$$

The problem is to find a feasible solution with the largest amount of minimum slack. Given a feasible solution, evaluating its minimum slack requires measuring the slack of each route. For each site $k \in \sigma_{vj}$, let Q_{vjk} be the total quantity previously delivered to that site on previous routes. This depends upon the set E_{vjk} of routes (*a*, *b*) such that $k \in \sigma_{ab}$ and $t_{ab} + w_{abk} \le t_{vj} + w_{vjk}$. Note that E_{vjk} does not include the route (*v*, *j*).

$$Q_{vjk} = \sum_{(a,b)\in E_{vjk}} q_{abk} \tag{9}$$

The expected time at which that site runs out of medication (if this delivery is delayed) is $T_1 + Q_{vjk} / L_k$. Let s_{vj} be the slack of route (*v*, *j*). That is, if the start of the route were delayed more than s_{vj} time units and no more medication were delivered to the sites $k \in \sigma_{vj}$, at least one of these sites would run out of medication. The minimum slack *S* of a solution is the minimum slack over all vehicles and routes.

$$s_{vj} = \min_{k\in\sigma_{vj}} \left\{ T_1 + Q_{vjk} / L_k - \left(t_{vj} + w_{vjk}\right)\right\} \tag{10}$$

$$S = \min_{v=1,\ldots,V\,;\, j=1,\ldots,r_v} \left\{ s_{vj} \right\} \tag{11}$$

3. Example

Consider a two-site, one-vehicle problem instance. T_1 = 24 hours = 1440 minutes. T_2 = 48 hours. L_1 = 10,000 regimens per hour, and L_2 = 5,000 regimens per hour. P = 10,000 regimens per pallet. C_1 = 10 pallets. $p_0 = p_1 = p_2 = 15$ minutes. The travel times (in minutes) are given in Table 1. The depot will receive three waves: 100,000 regimens at *t* = 0, 125,000 regimens at *t* = 4 hours, and 135,000 regimens at *t* = 8 hours.

Table 2 describes a feasible solution in which the vehicle travels the same sequence for five routes: σ_{1j} = {1, 2} for *j* = 1, …, 5. Then, w_{1j1} = 15 + 10 + 15 = 40 minutes, and w_{1j2} = 40 + 25 + 15 = 80 minutes. The total route duration is y_{1j} = 80 + 30 = 110 minutes. Table 3 shows the slack calculations. In this simple example, the minimum slack is 1360 minutes (22.67 hours), which is quite large.

Table 1. Travel times (in minutes) for example.

From \ To	Depot	Site 1	Site 2
Depot	-	10	30
Site 1	10	-	25
Site 2	30	25	-

Table 2. Feasible solution for example. All times in minutes.

Route	t_{1j}	q_{1j1}	h_{1j1}	q_{1j2}	h_{1j2}
1	0	70,000	7	30,000	3
2	240	70,000	7	30,000	3
3	345	10,000	1	15,000	2
4	480	70,000	7	30,000	3
5	585	20,000	2	15,000	2

Table 3. Slack calculations for example. All times in minutes.

Route j	Site 1			Site 2			Route slack
	Q_{1j1}	Run out time	Slack for site	Q_{1j2}	Run out time	Slack for site	
1	0	1440	1400	0	1440	1360	1360
2	70,000	1860	1580	30,000	1800	1480	1480
3	140,000	2280	1895	60,000	2160	1735	1735
4	150,000	2340	1820	75,000	2340	1780	1780
5	220,000	2760	2135	105,000	2700	2035	2035

4. Solution Approach

Instead of attempting to solve the problem as a large integer program, we adopt a two-stage solution approach that separates the problem into two subproblems: (1) the “routing problem” assigns PODs (sites) to vehicles and creates routes for each vehicle, and (2) the “scheduling problem” determines when the vehicles should start these routes and how much material should be delivered to each site on each trip.

In this approach, each available vehicle will have exactly one route (sequence of sites). A vehicle may perform that route more than once with different delivery quantities each time.

4.1 The Routing Problem

The routing problem is formulated as a capacitated vehicle routing problem (CVRP), which has been studied extensively [14]. The quantity to be delivered to each site is that site’s fair share of the largest wave. That is, if one site has a demand that is two times another site’s demand, the quantity delivered to the first site will be two times the quantity delivered to the second site. If we consider our example, site 1 will get a delivery of (10,000/15,000)(135,000) = 90,000 regimens, and site 2 will get a delivery of 45,000 regimens.

To solve the routing problem, we used the TourSolver route optimization software [15], which searches for a solution that minimizes the total cost, which includes the full day cost of operating vehicles, the cost of unused vehicles, the driving cost, and the overtime cost. Including the cost of unused vehicles and overtime penalizes solutions that do not use the given vehicles completely and evenly as possible. We solved the problem multiple times and changed the length of a “day” each time until we found a solution in which the route durations were nearly the same. Solving the routing problem generates a feasible route for each vehicle, and these routes cover all of the sites. We call these “single-wave routes.” The number of sites per route can vary greatly.

We then double all of the quantities to be delivered and solve the CVRP again. We call these the “double-wave routes.” Due to the fixed vehicle capacity, the larger delivery quantities limit our ability to find routes with nearly equal durations, so this leads to more variability in the route durations.

4.2 The Scheduling Problem

Given a set of routes, the scheduling problem determines how many times each vehicle should perform its route, when it should leave the depot, and how much should be delivered to each site on its route. The objective is to maximize the minimum slack of a solution.

To solve this problem, we developed a variety of heuristic techniques to construct a feasible solution. We will discuss three here (others are discussed in [13]). Note, however, that these scheduling heuristics are much different from dispatching, which maintains a queue of vehicle waiting to start their routes, uses simple policies to prioritize the vehicles in the queue, and starts the highest priority vehicle as soon as sufficient material is available at the depot. Previous studies have shown that such dispatching is highly myopic and cannot generate high-quality solutions because it ignores the pattern of deliveries to the depot.

We assign each route to one vehicle. Determining route start times depends upon the duration of its route and the time between deliveries to the depot (the “waves”). Ideally, each wave should be followed by vehicles leaving the depot to take the newly-arrived material from the depot to the sites. However, if the route duration is long, the vehicle may not be finished with its route when the next wave arrives at the depot.

We considered three types of schedules. (1) The *single wave* schedule uses the single-wave routes, and each vehicle leaves the depot as soon as possible after each wave. The delivery quantity to each POD is a fraction of that wave. (2) The *double wave* schedule uses the double-wave routes, and each vehicle leaves the depot after every other wave. The first delivery quantity to each POD equals a fraction of the first wave. Each remaining delivery quantity to a POD is a fraction of two waves. If the number of waves is even, one extra delivery is done after the last wave, and the delivery quantity to each POD equals a fraction of the last wave. (3) The *modified double wave* schedule modifies the double wave schedule by allowing vehicles with shorter duration routes to leave the depot after the even-numbered waves (instead of waiting to the next odd-numbered wave). In all cases, the delivery quantities are the same as in the double wave schedule.

5. Results

We considered two scenarios for testing the approaches. Due to space limitations, we can only briefly describe the scenarios and the results.

5.1 One-county Scenario

The first scenario used 50 PODs from the TourSolver example [15]. We assumed that each POD will open 12 hours after medication distribution begins and will operate for 12 hours. The total number of regimens dispensed is given by the example. The PODs dispense between 10,000 and 53,000 regimens. The total number of regimens is 950,389. medication will arrive at the depot in five waves every three hours. Each one of the first four waves will supply 200,000 regimens, and the fifth will supply 150,389 regimens. The nine trucks provided in the example have different capacities from 5 to 40 pallets, and each pallet can hold at most 11,200 regimens.

First, we created a CVRP instance in which just over 21% of each POD's total quantity is delivered. (The total number of regimens was 199,982.) We first ran TourSolver using a 3 hour run time for each truck. This yielded a solution that used all nine trucks, but some trucks were busy for all of the three hours, and others had much shorter routes. Using routes with such duration variability will yield a poor solution. Because the average time that a truck was busy was approximately 2:20, we ran TourSolver again using a 2:20 run time for each truck. In the resulting solution, six trucks were busy for nearly the entire run time. The three other trucks required some overtime (5, 10, and 40 minutes). Thus, the routes were more nearly equal in duration.

We constructed a single wave schedule in which each route starts after each wave. The delivery quantities were the same for the first four waves and smaller for the fifth wave. The minimum slack was 354 minutes, and this occurred in the fifth wave. We also found double-wave routes by doubling the delivery quantities. Both the double wave schedule and the modified schedule had a minimum slack of 240 minutes, which occurred after the fifth wave. The slack decreased with each wave because the depot receives only 2.4 hours' worth of medication every 3 hours.

Finally, we considered a simple dispatching scheme in which the entire quantity for a POD is delivered in two shipments. We reused the single wave routes and prioritized them by the travel time to the POD so that the routes with the longest travel times were done first. The minimum slack was 181 minutes, which occurs in the fifth wave.

5.2 Three-county Scenario

We considered a realistic scenario with three counties in the state of Maryland. Medication arrives at the state RSS (depot) in seven waves, one every two hours, with roughly the same amount of medication in each wave. A total of 189 PODs will dispense medication from T_1 = 24 hours to T_2 = 48 hours. First, we created the single-wave and double-wave routes and then the single-wave, double-wave, and modified double-wave schedules. The minimum slack of the double-wave and modified double-wave schedules was 509 minutes, which occurred after the third wave. The slack increased with each wave because the deliveries brought over six hours worth of medication every four hours. The minimum slack of the single-wave schedules was 360 minutes, which occurred after the seventh wave. The slack of some vehicles' routes decreased with each wave because their route durations exceeded the time value of the deliveries (in the single wave case, each delivery supplied about 3.4 hours worth of medication).

6. Summary and Conclusions

This paper introduced the medication distribution problem, an important part of planning the response to a bioterrorism attack, and presented a two-stage routing and scheduling approach for constructing solutions. Because a robust plan is desirable, our objective was to maximize the minimum slack of the solution. Instead of attempting to solve the problem as a large integer program, we adopted a two-stage solution approach that separates the problem into two subproblems. This practical separation reduces the solution effort, though it is not guaranteed to find an optimal solution. To demonstrate the approach, we applied it to two scenarios, including one for three counties in the state of Maryland.

Future work is needed to automate the routing and scheduling approach to enable a decision support tool for public health emergency preparedness planners, to develop optimization techniques for finding even better solutions, and to test these approaches on other scenarios.

Acknowledgements

The work of Sara Lu and Kristen Schalliol occurred during their participation in the NSF REU site “Introducing the Systems Engineering Paradigm to Young Researchers and Future Leaders” (NSF grant EEC 0243803). Kay Aaby, Rachel Abbey, and Kathy Wood at the Montgomery County, Maryland, Public Health Services provided excellent guidance and assistance. Cooperative Agreement Number U50/CCU302718 from the CDC to NACCHO supported this publication. Its contents are solely the responsibility of the University of Maryland and do not necessarily represent the official views of CDC or NACCHO. The discussion of related work relies in part on material prepared by Zhi-Long Chen.

References

1. Centers for Disease Control and Prevention (CDC), 2008, “Key Facts about the Cities Readiness Initiative (CRI),” http://www.bt.cdc.gov/CRI/facts.asp, accessed September 16, 2008.
2. Daskin, M.S., and Stern, E.H., 1981, “Hierarchical objective set covering model for emergency medical service vehicle deployment,” Transportation Science, 15, 137 – 152.
3. Ball, M.O., and Lin, F.L., 1993, “A reliability model applied to emergency service vehicle location,” Operations Research, 41, 18 – 36.
4. Ceyhun, A., Selim, H., and Ozkarahan, I., 2007, “A fuzzy multi-objective covering-based vehicle location model for emergency services,” Computers and Operations Research, 34, 705 – 726.
5. Sivanandan, R., Hobeika, A.G., Ardekani, S.A., and Lockwood, P.B., 1988, “Heuristic shortest-path method for emergency vehicle assignment - a study on the Mexico City network,” Transportation Research Record, no. 1168, 1988, 86 – 91.
6. Weintraub, A., Aboud, J., Fernandez, C., Laporte, G., and Ramirez, E., 1999, “An emergency vehicle dispatching system for an electric utility in Chile,” Journal of the Operational Research Society, 50, 690-696.
7. Haghani, A., Tian, Q., and Hu, H., 2004, “Simulation model for real-time emergency vehicle dispatching and routing,” Transportation Research Record, no. 1882, 176 – 183.
8. Dror, M., Ball, M., and Golden, B., 1985, “Computational comparisons of algorithms for the inventory routing problem,” Annals of Operations Research, 4, 3 – 23.
9. Campbell, A., Clarke, L., Kleywegt, A.J., and Savelsbergh, M.W.P., 1998, “The inventory routing problem,” in Fleet Management and Logistics, Crainic, T.G., and Laporte, G. (eds.), Kluwer Academic Publishers, Dordrecht, The Netherlands, Chapter 4.
10. Baita, F., Ukovich, W., Pesenti, R., and Favaretto, D., 1998, “Dynamic routing-and-inventory problems: A review,” Transportation Research, Part A, 32, 585 – 598.
11. Moin, N.H., and Salhi, S., 2007, “Inventory routing problems: A logistical overview,” Journal of the Operational Research Society, 58, 1185 – 1194.
12. Chen, Z.-L., 2008, “Integrated production and outbound distribution scheduling: Review and extensions,” To appear in Operations Research.
13. Herrmann, J.W., Lu, S., and Schalliol, K., “Routing and Scheduling for Medication Distribution Plans,” Technical Report 2008-26, Institute for Systems Research, University of Maryland, College Park.
14. Toth, P. and Vigo, D., 1998, “Exact Solution of the Vehicle Routing Problem,” in Fleet Management and Logistics, Crainic, T.G., and Laporte, G. (eds.), Kluwer Academic Publishers, Dordrecht, The Netherlands, 1-31.
15. C2Logix, 2008, “SNSTourSolver,” http://cdcstockpilerouting.c2logix.com/Citrix/AccessPlatform/auth/login.aspx, accessed July 25, 2008.