A Two-stage Vehicle Routing Model for Large-scale Bioterrorism Emergencies

Zhihong Shen
shenz@usc.edu

Maged Dessouky
maged@usc.edu

Fernando Ordóñez
fordon@usc.edu

January, 2009

Abstract

In this paper, we are interested in routing vehicles to service a large-scale bioterrorism emergency. We describe the specifics of routing vehicles in such a large-scale emergency and decompose the problem into two stages: a planning stage and an operational stage. In the planning stage we generate the routes well in advance of any emergency. In the operational stage, we take into account the planned routes and the information revealed at the time of the emergency, to decide the delivery quantity and any adjustments to the routes. We propose mathematical formulations and solution approaches for both stages. Lastly, we demonstrate the effectiveness of our formulations and solution procedures in developing robust routes through numerical experiments.

1 Introduction

According to the U.S. Centers for Disease Control and Prevention (CDC), "a bioterrorism attack is the deliberate release of viruses, bacteria, or other germs (agents) used to cause illness or death in people, animals, or plants." Such an attack creates a large-scale emergency when the event overwhelms local emergency responders and has the potential to cause substantial casualties and property damage. A key ingredient in an effective response to a bioterrorism emergency is the prompt availability of necessary medical supplies at emergency sites. Given the challenges of delivering massive supplies in a short time period to dispersed demand areas, operations research models could play an important role. Larson (2005) and Larson et al. (2006) conducted a detailed analysis based on well-known and recent large-scale emergencies.

In a bioterrorism emergency, medication or antidotes must be applied within a specified time limit from the occurrence of the event to maximize their effectiveness and, in some situations, save lives. However, local caches of medicine are likely to be insufficient and traditional pharmaceutical supply chains would not be prepared to push a huge volume of medical supplies to the affected population in a short time. In a large-scale emergency, the bulk of the medical supplies would have to come from national repositories, which in the United States is the CDC maintained Strategic National Stockpile (SNS). It is each state's responsibility to have plans to receive and distribute SNS medicine and medical supplies to local communities as quickly as possible. This work is concerned with the distribution process after the SNS supplies have been received at the local level. For this problem, a vehicle routing model provides a method to plan the disbursement of the medical supplies ahead of an attack and, in the event an attack occurs, a methodology to quickly adjust the routing solution to provide an efficient response.

The objective of this routing problem is to minimize the unmet demand. The reason for this is because the unmet demand in an emergency situation can result in loss of life, an impact that outweighs other commonly used VRP objectives such as travel time or number of vehicles used. An additional important aspect in a large-scale bioterrorism emergency is how long it takes to service a demand. We only consider a demand satisfied if the vehicle arrives within a certain time window, since delivering medicine too late can make an appreciable health difference and even lead to loss of life.

Because of the unpredictable nature of large-scale bioterrorism emergencies, we consider uncertain demand and travel times. First, at a given demand point (e.g., a neighborhood block), the quantity of required emergency supplies (antidotes, protective equipments, medication, etc.) is often proportional to the unknown size of the population and/or numbers of casualties. Second, the casualties exposures, or demands among "worried well" are hard to accurately predict. In addition, in an emergency situation, the traffic condition can become highly uncertain, due to unpredictability in people's behavior.

When we consider the problem of routing inventory in response to large-scale bioterrorism emergencies, both routes and delivery quantities are important in creating an effective dispatching plan. From a planning perspective effective routes can be used by agencies in mock trial runs of an emergency event to pre-design the routing policies and provide training opportunities. From an operational perspective, the models can be used to update the pre-designed routes and promptly determine exact delivery quantities to dispensing or treatment sites to generate real-time near-optimal dispatching solutions after an emergency situation has happened.

In this work, we propose a two-stage model for a specific bioterrorism emergency scenario, e.g., an anthrax attack. We believe that this type of model can also be applied to other emergency scenarios by further investigating and validating several key problem parameters (such as the service deadline, demand distribution etc.). The two-stage model has a planning stage and an operational stage. The purpose of the first planning stage is to generate pre-planned routes well in advance of any possible emergency for mock trial runs and training. We formulate this problem as a mixed integer programming (MIP) model to meet the special requirements for the distribution process of large-scale bioterrorism emergencies. In the second operational stage, at the time of the attack when more information is revealed, we need to quickly respond to the event and generate the delivery requirements with the planned routes as well as making adjustments on the routes if necessary. We refer to the second stage model as the recourse strategy. In this paper, we propose three different recourse strategies.

The main contributions of this article are: from the modeling perspective, we analyze the problem of routing vehicles in response to a large-scale bioterrorism emergency and decompose the problem into two stages (pre-planning and operational stages), then formulate mathematical models for both stages; from the solution approach perspective, we propose a tabu heuristic for the planning stage model and an approximation heuristic for one of the proposed recourse strategies, then demonstrate the effectiveness of the model and solution approach in developing robust routes on numerical experiments. This work is an extension on our previous work (Shen et al., to appear). The previous work emphasized the first-stage planning model (a stochastic vehicle routing model) and compared the chance-constraint programming technique with a deterministic model and a robust approach. In this work, we focus on the two-stage modeling framework and the second-stage recourse strategies. We briefly present the first-stage model and highlight the addition of allowing for split delivery to the formulation. For more details of the first-stage routing model, please refer to our previous paper.

The rest of this paper is organized as follows. In Section 2, we review the relevant literature. Section 3 proposes the mathematical model for the planning stage and presents three recourse strategies for the operational stage. Section 4 presents the solution approaches for both stages. We present some illustrative computational results in section 5, and finally, conclude in Section 6.

2 Literature Review

The Vehicle Routing Problem is defined on a given topological graph. The classical standard VRP generates a set of routes which visit each customer exactly once. It aims to minimize the total travel time and/or the operational cost. When some uncertainty in the parameters is introduced into the problem, it is called a Stochastic VRP (SVRP). A comprehensive overview of the Vehicle Routing Problem and its stochastic variant can be found in Toth and Vigo (2002) and Gendreau et al. (1996a). Routing problems with profits are a generalization of the vehicle routing problems, where it is not necessary to visit all customers. A profit is associated with each demand node and a uniform deadline is applied to every node. The problem aims to maximize the total collected profits subject to the time restriction (constrained by the travel cost). This is called the *Orienteering Problem* (OP) when only a single vehicle is present and is motivated by an outdoor sport. The multi-vehicle version of the OP is referred to as the *Team Orienteering Problem* (TOP) in the literature. In comparison to the classical routing problems, the routing problems with profits have an added element of complexity. They require selecting a subset of nodes to visit while determining an order on the tour. If all the customers can be visited within the time constraints, then it becomes the classical routing problem and the extra complexity vanishes. We review these two well-known bodies of literature in the following two subsections and follow with a review on the split delivery and tabu search heuristics applied to VRPs.

2.1 Stochastic Vehicle Routing Problems

The Stochastic Vehicle Routing Problem can be broadly classified based on the following three criteria: (1) where the uncertainty lies in the problem, e.g., the presence of the customers (Jézéquel 1985; Jaillet 1987; Bertsimas 1988; Bertsimas 1992) the demand level (Bertsimas and Simchi-Levi 1996) or the travel time (Kao 1978; Laporte et al. 1992; Jula et al. 2006) and the service time at customer sites (Hadjiconstantinou and Roberts 2002); (2) how to model the problem, e.g., by stochastic programming technique (Stewart and Golden 1983; Bertsimas 1992), by Markov decision process (Dror and Trudeau 1986; Dror 1993; Dror et al. 1989; Secomandi 2001) or by robust optimization methodology (Sungur et al. 2008); (3) how to solve the model, which heavily depends on the modeling method and can be broadly classified into two categories: exact methods (branch and cut, integer L-shape method (Gendreau et al. 1995) and generalized dynamic programming (Carraway et al. 1989)) and heuristic methods such as saving algorithms (Clarke and Wright 1964), sweep algorithms, Genetic Algorithms, tabu search (Gendreau et al. 1996b) to name a few. For a more detailed discussion based on this classification, please refer to Shen et al. (to appear).

Here we will further expand the discussion of the Stochastic VRPs modeled through the stochastic programming framework: either as a *chance constrained program* (CCP) or as a *stochastic program with recourse* (SPR). A CCP enforces the probability of satisfying the constraints with stochastic parameters to be above a given threshold; however, it does not consider the corrective actions and their cost into the formulation. On the contrary, SPR aims to incorporate the expected cost of the proposed second-stage recourse actions into the objective function in the first-stage model. This brings considerable difficulties but is more meaningful. The research in modeling SVRP with SPR has flourished over the last two decades.

In the early works (Stewart and Golden, 1983; Laporte et al. 1992), the expected penalty incurred in the first-stage objective is derived from how the failure occurred. When the demand exceeds the capacity, each extra unit of demand is penalized; or when the travel time exceeds the deadline, each extra unit of time is penalized. These works do not propose explicit recourse actions. Bertsimas (1993) applies two straightforward recourse actions (which are called "traditional recourse strategies" in later literature) and penalizes on a different dimension as where the failure occurred. When the demand exceeds the capacity, the penalty is on the expected extra travel cost to cover the failed customers. Later on, more sophisticated recourse strategies have been designed and studied (Yang et al., 2000; Ak and Erera, 2006; Novoa et al., 2006), and they follow the "traditional" strategies to penalize on a different dimension other than the failure occurring dimension. They all share the assumption that there are no deadlines on the temporal perspective, no total supply quantity constraints, and the stochastic elements only appear in the demand. Therefore, they have the luxury to return a vehicle to the depot whenever a failure occurs (or is expected to occur) and reload to continue serving customers following the designed strategies without worrying about violating time constraints or the total supply level at the depot constraint. Our work differs from these by considering a problem with uncertainty both in demand and travel time, and constraints both on capacity and deadlines.

2.2 Routing Problem with Profits

Hayes and Norman (1984) first introduced the orienteering problem with real data for the 1981 Karrimor Mountain Marathon and proposed a dynamic programming solution approach. Tsiligirides (1984) deduced a common mathematical model, proposed heuristics and ran the numerical tests on up to 33 locations. Later on, a number of heuristics have been proposed by various groups of researchers, e.g., Golden et al. (1987, 1988), Ramesh et al. (1991), and Chao et al. (1996a). All these heuristics were tested using the small size benchmark problems provided in (Tsiligirides 1984) with improving solution quality. More recently, many types of sophisticated meta-heuristic techniques have been applied on the OP, for example: artificial neural networks (Wang et al. 1995), tabu search (Gendreau et al. 1998b), genetic algorithms (Tasgetiren and Smith 2000), and ant colony optimization (Liang and Smith 2001). These meta-heuristics can provide comparable results on much larger scale problems (up to 300 nodes); among them, the tabu search has been shown to provide an optimality gap that is less than 1% in a number of experiments. Another focus of the research on the OP are exact solution methods. Laporte and Martello (1990) used a branch-and-bound approach with a knapsack bound; while Ramesh et al. (1992) applied a shortest spanning tree relaxation style bounding within the same framework. Leifer and Rosenwein (1994) discussed an LP-based bounding procedure. Later on, Fischetti et al. (1997, 1998) and Gendreau et al. (1998a) proposed a branch-and-cut method to solve instances with up to 500 nodes. Feillet et al. (2005) provided a more detailed review on this topic. Different variations of the standard OP have attracted many groups of researchers in recent years. For example, Geem et al. (2005) applied a harmony search mimicking the improvisation process of music players to study a generalized OP with multiple objectives; Mak and Thomadsen (2006) investigated a quadratic OP with edge-reward and cardinality constraint through its polyhedral properties; and Jozefowiez et al. (2008) introduced a hybrid meta-heuristic for OP which combines an ejection chain local search and a multi-objective evolutionary algorithm.

Two works have considered stochastic elements in the OP. Tang and Miller-Hooks (2005a) addressed an orienteering problem with stochastic service time. The problem was formulated as a chance-constrained stochastic programming problem. An exact branch-and-cut algorithm and a "construct-and-adjust" heuristic were proposed to solve it. Ilhan et al. (2008) focused on the uncertainty in profits in the OP and maximized the probability of collecting more than a pre-defined target level of the total return. They proposed an exact solution method based on a parametric formulation and developed a bi-objective GA heuristic.

The routing problem with profits was extended to a multi-vehicle context by Butt and Cavalier (1994). They were concerned with the athletes recruitment problem and proposed a greedy algorithm for solving it. Wu (1992) and Dunn (1992) faced an underway replenishment of a dispersed carrier battlegroup. Dunn (1992) applied dynamic programming to obtain a quick and useful schedule for this military operational situation. Chao et al. (1996b) used a straightforward extension of their heuristic for the OP (Chao et al. 1996a) to the TOP context and solved instances with up to 102 nodes. Later on, Butt and Ryan (1999) proposed a branch-and-price exact solution procedure based on a set-partitioning formulation, and made efficient use of constraint branching and column generation. In recent years, Tang and Miller-Hooks (2005b) introduced tabu heuristics for the TOP; Archetti et al. (2007) proposed two variants of a generalized tabu search algorithm and a variable neighborhood search algorithm and their computational experiments on benchmark problems (Chao et al. 1996b) generated best solutions by that time. Ke et al. (2008) introduced Ant Colony Optimization (ACO) and Vansteenwegen et al. (to appear) used a guided local search (GLS) meta-heuristic for the TOP. ACO outperforms Archetti et al. (2007)'s result in terms of the solution quality for some benchmark problems and the GLS produces an equivalent quality solution with better efficiency (less running time). Boussier et al. (2007) first proposed an exact solution approach, branch and price algorithm, to the TOP and solves instances with up to 100 customers. All the works mentioned above for the TOP consider the constraint on the temporal dimension only and assume an uncapacitated fleet. Goel and Gruhn (2005) studied a more complicated problem arising in air-cargo transport. They were concerned with a multiple pickup and delivery problem with a heterogeneous fleet. Load acceptance, compatibility constraints, time windows constraints, capacity constraints, as well as a timely manner of handling dynamic requests were under consideration at the same time. A variable neighborhood search method (Mladenovic and Hansen 1997) was developed to solve the problem on real-life scale test cases.

2.3 Split Delivery and Tabu Search Heuristics for VRPs

Split delivery allows multiple visits to a demand node. This is an extension of the classical VRPs and has also been investigated by a number of researchers (Dror and Trudeau 1989; Dror and Trudeau 1990; Dror et al. 1994; Archetti et al. 2006). Dror and Trudeau (1989, 1990) have analyzed the saving generated by allowing split deliveries in a VRP. Archetti et al. (2006) demonstrated a tabu search heuristic to solve problems with up to 199 nodes; most instances were solved in less than 10 minutes and a few cases used around 3 hours.

Tabu search is a local search procedure which iteratively moves from a solution to its best neighbor until some stopping criteria are satisfied. A comprehensive review on this technique and its applications can be found in Glover and Laguna (1997). Tabu search was first introduced for solving the VRP by Willard (1989). Later on, different groups of researchers designed a variety of neighborhood/moves and adopted some problem-specific mechanism to significantly improve the performance (Osman 1993; Gendreau et al. 1994; Xu and Kelly 1996; Toth and Vigo 2003). Tabu search has also been applied to major variants of VRP, e.g., VRP with time windows (Taillard et al. 1997), VRP with split delivery (Archetti et al. 2006), the pick up and delivery problem (Bianchessi and Righini 2006), as well as the stochastic VRP (Gendreau et al. 1996b). It has been shown that tabu search generally yields very good results on a set of benchmark problems and some larger instances (Gendreau et al. 2002).

3 Mathematical Models

In this section, we propose a model for the problem of routing for large-scale bioterrorism emergencies. The primary objective is to minimize the unmet demand (to maximize the life-saving). This can be seen as a case of routing with profits, since minimizing the unmet demand corresponds to maximizing the collected profits when we set the profit of each node to its demand level. However, compared with the traditional routing problem with profits, whose uncollected demand nodes are only due to the deadline constraint, our problem is further complicated by potentially insufficient supply at the depot or limited vehicle capacity.

In our problem setting, the large-scale bioterrorism emergency makes both the demand level and the traffic condition highly unpredictable; hence uncertainty exists both in demand and in travel time. The overwhelming demand with limited resources and the urgency of the timely delivery make the problem constrained by both the vehicle capacity (in demand dimension) and the service deadline (in temporal dimension). Since our model can be applied to a single specific type of emergency scenario (e.g., anthrax attack), the service deadline can be obtained from CDC regulations well before the attack happens at the planning stage. We formulate the route planning stage as a stochastic programming problem where we quantify all the unmet demand, due to either insufficient capacity, supply, or late delivery. We model this problem with a single depot by considering the fact that, in most emergency scenarios, the SNS supplies are usually shipped to one regional central warehouse (e.g., LAX airport in the Los Angeles Metro Area) for further distribution at the local level. We use chance constraints to represent the uncertainty in demand and travel times.

In addition to minimizing the unmet demand, the planning problem also considers a secondary objective to guide the construction of a *complete route*. Instead of ignoring the nodes not selected by the team orienteering problem due to tight deadline, insufficient supply, or limited vehicle capacity, we include them at the end of each route according to this secondary objective. These complete routes that visit all demand nodes even after the deadline or with an empty truck, provide a blueprint for recourse actions in the operational stage. The secondary objective is the arrival time at each node. We use a significantly small coefficient κ to weigh this secondary objective so it will not interfere with the route generation for nodes that can be satisfied within constraints, which is guided by the minimum unmet demand primary objective.

We illustrate this model for the planning stage with the following figures. Figure 1-(a) shows the routes generated by a classical VRP, which visit all the demand nodes and aim to minimize the total travel time. Figure 1-(b) demonstrates the routes obtained from a routing problem with profits on the same topological graph, which visit the nearby high-profit nodes within the given deadline while ignoring those far-away low-profit ones. Figure 1-(c) illustrates the result by the routing problem to minimize the unmet demand and then complete the route as we just proposed. It generates exactly the same result as the routing problem with profits before the deadline. After the deadline, it visits the remaining far-away nodes guided by the secondary objective.

For the operational stage we consider models that will adapt the preplanned complete routes to satisfy the realized demand in the large scale emergency. In this paper we present three recourse operational models that differ in how much re-optimization we can do in the operational stage. We consider an LP recourse strategy in which we can only adjust the quantity sent on each vehicle, a knapsack recourse strategy which, in addition to setting the quantity, allows to modify the routes by skipping customers in each route, and for comparison purposes we also consider a re-planning strategy in which we look for the optimal routing solution for the realized demand. The operational stage problem should balance the need for familiarity with a planned strategy, which can be used by responders for training purposes, with efficiency in the solution and a quick solution time for a rapid response.

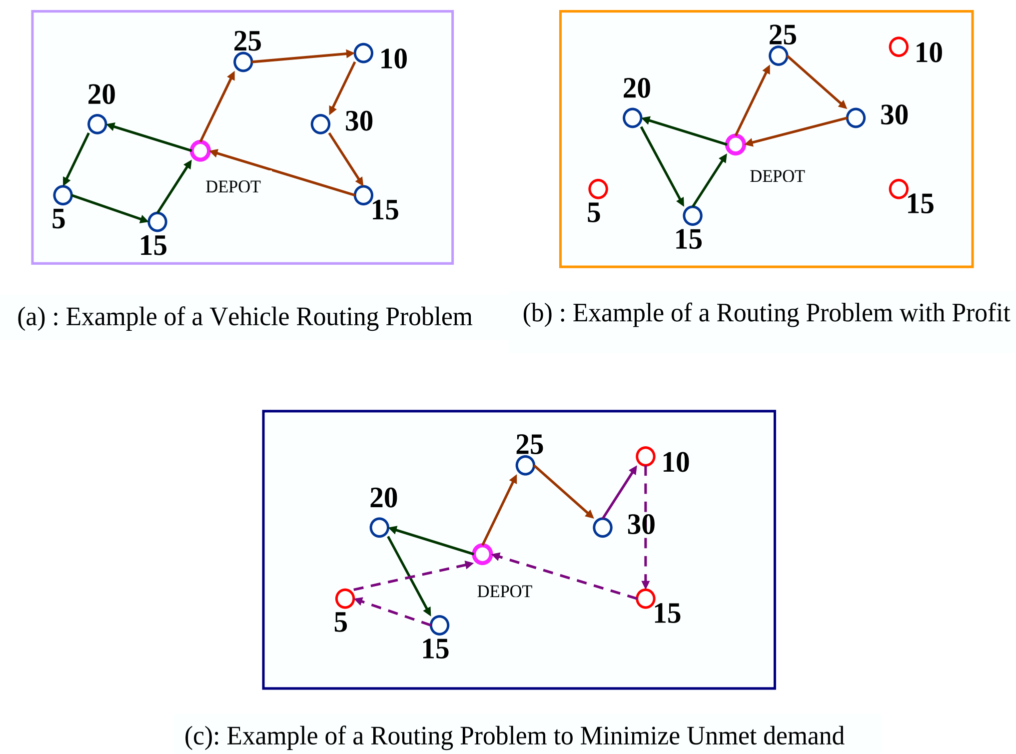

(a) : Example of a Vehicle Routing Problem

(b) : Example of a Routing Problem with Profit

(c): Example of a Routing Problem to Minimize Unmet demand

Figure 1: A priori route model at the planning stage.

3.1 Notation

We first introduce the notation used in both the planning and operational stages.

We consider a set K of vehicles and a set D of demand nodes. We identify an additional node, node 0, as the supply node (depot) and let $C = D \cup \{\text{node } 0\}$ represent the set of all nodes. Indexed on sets K and C, we define the following *deterministic parameters* for the planning model:

n:	initial number of vehicles at the supply node (depot)
s :	amount of supply at the supply node (depot)
c_k:	load capacity of vehicle k
dl_i:	service deadline at demand node i.

We use a large constant M to express nonlinear relationships through linear constraints and a small weight κ to balance the primary and secondary objectives. We also consider the following two parameters to represent the uncertain travel time and demand, respectively, in the planning stage

$\tau_{i,j,k}$: time required to traverse arc (i,j) for vehicle k
ζ_i : amount of demand needed at node i.

Finally, we define the binary and non-negative decision variables as follows for the planning model, indexed on sets K, C:
Binary:

$X_{i,j,k}$: flow variables, equal to 1 if (i,j) is traversed by vehicle k and 0 otherwise
$S_{i,k}$: service variables, equal to 1 if node i can be serviced by vehicle k

Non-negative:

$Y_{i,j,k}$: amount of commodity traversing arc (i,j) using vehicle k
U_i: amount of unsatisfied demand of commodity at node i
$T_{i,k}$: visit time at node i of vehicle k
$\delta_{i,k}$: delay incurred by vehicle k in arriving at i.

The variable $\delta_{i,k}$ represents the delay of the visit time if a vehicle reaches the node after the deadline. It is set to zero if it arrives before the deadline.

In the second operational stage, the uncertain parameters have been revealed and become deterministic as defined below:

$t_{i,j,k}$: actual time required to traverse arc (i,j) for vehicle k
d_i : actual amount of demand needed at node i.

Since the planned route is assumed as given at the operational stage, the vehicle flows $X_{i,j,k}$ become parameters. In our first recourse strategy, the value of $S_{i,k}$ is uniquely determined once $X_{i,j,k}$ are given. Hence, they also are no longer decision variables. However, in our second recourse strategy, we introduce a new binary variable

$Z_{i,j,k}$: flow variables, equal to 1 if (i,j) is traversed by vehicle k and 0 otherwise

and establish its relationship with $S_{i,k}$. Hence, $S_{i,k}$ remain as variables.

3.2 The Planning Stage – A Chance-constrained Programming Model

We follow the same modeling framework as we proposed in our previous work (Shen et al. to appear) and relax the model to allow split delivery. The chance-constrained programming technique is used to model the constraints with uncertain travel time and demand. We assume the distribution of these stochastic elements is known at the planning stage. The CCP model then can be readily transformed to a deterministic counterpart. We present the complete model here and refer the reader to the previous paper for a detailed explanation of the model and its derivation.

$$\textbf{CCP}: \quad \begin{array}{ll} \text{minimize} & \sum_{i \in D} U_i + \kappa \sum_{i \in D, k \in K} T_{i,k} \\ \text{subject to} & \text{constraints } (1) - (17)\ , \end{array}$$

Route Feasibility:

$$\sum_{i \in D} \sum_{k \in K} X_{0,i,k} \leq n \tag{1}$$

$$\sum_{i \in D} \sum_{k \in K} X_{i,0,k} \leq n \tag{2}$$

$$\sum_{j \in D} X_{0,j,k} = \sum_{j \in D} X_{j,0,k} = 1 \qquad (\forall k \in K) \tag{3}$$

$$\sum_{j \in C} \sum_{k \in K} X_{i,j,k} \geq 1 \qquad (\forall i \in D) \tag{4}$$

$$\sum_{j \in C} \sum_{k \in K} X_{j,i,k} \geq 1 \qquad (\forall i \in D) \tag{5}$$

$$\sum_{j \in C} X_{i,j,k} = \sum_{j \in C} X_{j,i,k} \qquad (\forall i \in D \quad k \in K) \tag{6}$$

Time feasibility:

$$T_{0,k} = 0 \qquad (\forall k \in K) \tag{7}$$

$$P\left\{\tau | (T_{i,k} + \tau_{i,j,k} - T_{j,k}) \leq (1 - X_{i,j,k})M\right\} \geq 1 - \alpha_T \qquad (\forall i, j \in C \quad k \in K) \tag{8}$$

$$\sum_{j \in C} X_{i,j,k} M \geq T_{i,k} \geq 0 \qquad (\forall i \in D \quad k \in K) \tag{9}$$

$$dl_i \sum_{j \in C} X_{i,j,k} \geq T_{i,k} - \delta_{i,k} \geq 0 \qquad (\forall i \in D \quad k \in K) \tag{10}$$

Node service constraints

$$(1 - S_{i,k})M \geq \delta_{i,k} \qquad (\forall i \in D \quad k \in K) \tag{11}$$

$$\sum_{j \in C} X_{i,j,k} \geq S_{i,k} \qquad (\forall i \in D,\, k \in K) \tag{12}$$

$$\sum_{j \in C} Y_{j,i,k} - \sum_{j \in C} Y_{i,j,k} \leq S_{i,k} M \qquad (\forall i \in D,\, k \in K) \tag{13}$$

Demand flows:

$$s - \sum_{k \in K} \left[\sum_{j \in C} Y_{0,j,k} - \sum_{j \in C} Y_{j,0,k} \right] \geq 0 \tag{14}$$

$$P \left\{ \zeta | \sum_{k \in K} \left[\sum_{j \in C} Y_{j,i,k} - \sum_{j \in C} Y_{i,j,k} \right] + U_i - \zeta_i \geq 0 \right\} \geq 1 - \alpha_D \qquad (\forall i \in D) \tag{15}$$

$$X_{i,j,k} c_k \geq Y_{i,j,k} \qquad (\forall i, j \in C,\, k \in K) \tag{16}$$

Binary and non-negativity properties:

$$X_{i,j,k}, S_{i,j} \text{ binary}; \qquad Y_{i,j,k} \geq 0; \qquad U_i \geq 0; \qquad T_{i,k} \geq 0; \qquad \delta_{i,k} \geq 0. \tag{17}$$

We only highlight the two constraints (4) and (5) which we relax for split delivery. These two constraints state that each demand node must be visited at least once, which allows split delivery, e.g., multiple visits to a demand node. In a large-scale emergency it is likely that the demand at points which are seriously affected will not be able to be satisfied by a single truckload. Hence we consider split delivery, which also provides greater solution flexibility to potentially service more demand nodes before the deadline.

3.3 The Operational Stage

After an emergency occurs we can use the observed outcome of the uncertain parameters to make a quick adjustment to the preplanned routing solution to respond to the event. This response needs to provide the delivery requirements (how much to load on each vehicle) based on the planning result as well as adjustments to the routes to make the delivery more effective. This second stage model must be able to be solved quickly so that the adjustments do not delay the emergency response. We present three different recourse strategies and their corresponding mathematical models in the following subsections.

3.3.1 The LP Recourse Strategy

In the first recourse strategy, we strictly follow the pre-planned route obtained in the first-stage but we can adjust the supply carried by each vehicle and the amount to be dropped at different demand nodes according to the revealed demand level and travel time. In this approach, the commodity flow variables become the only amount to be determined. This is the easiest strategy with least flexibility among the three we propose in this paper. Since it is a linear programming problem, it can be solved very efficiently by a CPLEX solver.

Model **DLP**:

$$\text{minimize} \sum_{i\in D} U_i + \kappa \sum_{i\in D, k\in K} T_{i,k}$$

subject to: constraints (13)-(14), (16), (18)-(19).

$$\sum_{k\in K}\left[\sum_{j\in C} Y_{j,i,k} - \sum_{j\in C} Y_{i,j,k}\right] + U_i - d_i \;\geq\; 0 \qquad (\forall i \in D) \tag{18}$$

$$Y_{i,j,k} \geq 0; \qquad U_i \geq 0. \tag{19}$$

3.3.2 The Knapsack Recourse Strategy

The alternative recourse strategy is inspired by the classical recourse strategy B in Bertsimas (1992), which assumes the demand level will be revealed before the fleet departs from the depot. Hence, the zero demand customers will be skipped. In our strategy, we share the same assumption that the actual demand and travel time needed will become known before the vehicles leave the depot. Besides deciding the amount carried by each vehicle, the recourse allows for skipping low demand nodes when a direct visit to the next high demand node could result in a less total unmet demand level. This recourse strategy can be modeled into an MIP as below.

We use the same set of parameters and decision variables as in the first-stage planning model except that $X_{i,j,k}$ becomes a parameter (its value will be given from the pre-planning solution) and we add a new set of vehicle flow variables $Z_{i,j,k}$. We keep the same objective function and replace $X_{i,j,k}$ with $Z_{i,j,k}$ in the time constraints, node service constraints and demand flow constraints (constraint (8) - (17)). In the route feasibility constraints group, we will establish the relationship between the new vehicle flow variable $Z_{i,j,k}$ and the pre-planned routes $X_{i,j,k}$.

$$\mathbf{DP}: \quad \begin{array}{ll} \text{minimize} & \sum_{i\in D} U_i + \kappa \sum_{i\in D, k\in K} T_{i,k} \\ \text{subject to} & \text{constraints } (20)-(35)\ , \end{array}$$

The following constraints (20)-(24) enforce the route feasibility.

$$\sum_{i\in D}\sum_{k\in K} Z_{0,i,k} \;=\; \sum_{i\in D}\sum_{k\in K} X_{0,i,k} \tag{20}$$

$$\sum_{i \in D} \sum_{k \in K} Z_{i,0,k} = \sum_{i \in D} \sum_{k \in K} X_{i,0,k} \tag{21}$$

$$\sum_{j \in D} Z_{0,j,k} = \sum_{j \in D} Z_{j,0,k} = 1 \qquad (\forall k \in K) \tag{22}$$

$$\sum_{i,j \in R_k, i \text{before} j} Z_{i,j,k} \leq \sum_{i,j \in R_k, i \text{before} j} X_{i,j,k} \qquad (\forall k \in K) \tag{23}$$

$$\sum_{j \in C} Z_{i,j,k} = \sum_{j \in C} Z_{j,i,k} \qquad (\forall i \in D \quad k \in K) \tag{24}$$

Constraints (20)-(22) require the number of vehicles dispatched in the second-stage to be the same as the solution obtained from the first-stage pre-planned solution. R_k in constraint (23) represents the subset of nodes visited by the pre-planned route k. It is defined as $R_k = \{i : i \in D | X_{i,j,k} = 1\}$. The expression i before j states the sequence of visiting. The right hand side of constraint (23) represents the number of edges (except for the edge from the last visited demand node back to the depot) in the pre-planned route for vehicle k; the left hand side represents the same meaning in the to-be-decided operational stage route for vehicle k, which is bounded by the given pre-planned route. Hence it allows skipping of some nodes in a given route. Constraint (24) guarantees the connectivity of the vehicle flow.

Constraints (25)-(28) enforce schedule feasibility with respect to time considerations.

$$T_{0,k} = 0 \qquad (\forall k \in K) \tag{25}$$

$$(T_{i,k} + t_{i,j,k} - T_{j,k}) \leq (1 - Z_{i,j,k})M \qquad (\forall i, j \in C \quad k \in K) \tag{26}$$

$$\sum_{j \in C} Z_{i,j,k} M \geq T_{i,k} \geq 0 \qquad (\forall i \in D \quad k \in K) \tag{27}$$

$$dl_i \sum_{j \in C} Z_{i,j,k} \geq T_{i,k} - \delta_{i,k} \geq 0 \qquad (\forall i \in D \quad k \in K) \tag{28}$$

Constraints (29)-(31) address node service constraints.

$$(1 - S_{i,k})M \geq \delta_{i,k} \qquad (\forall i \in D \quad k \in K) \tag{29}$$

$$\sum_{j \in C} Z_{i,j,k} \geq S_{i,k} \qquad (\forall i \in D, \, k \in K) \tag{30}$$

$$\sum_{j \in C} Y_{j,i,k} - \sum_{j \in C} Y_{i,j,k} \leq S_{i,k} M \qquad (\forall i \in D, \, k \in K) \tag{31}$$

Constraints (32)-(34) state the construction on the demand flows.

$$s - \sum_{k \in K} \left[\sum_{j \in C} Y_{0,j,k} - \sum_{j \in C} Y_{j,0,k} \right] \geq 0 \tag{32}$$

$$\sum_{k\in K}\left[\sum_{j\in C}Y_{j,i,k}-\sum_{j\in C}Y_{i,j,k}\right]+U_i-d_i \geq 0 \qquad (\forall i \in D) \tag{33}$$

$$Z_{i,j,k}c_k \geq Y_{i,j,k} \qquad (\forall i,j \subseteq C,\, k\in K) \tag{34}$$

$$Z_{i,j,k}, S_{i,j} \text{ binary}; \qquad Y_{i,j,k}\geq 0; \qquad U_i \geq 0; \qquad T_{i,k}\geq 0; \qquad \delta_{i,k}\geq 0. \tag{35}$$

Constraint (35) states the binary and non-negativity properties of the decision variables.

The second recourse strategy will lead to a result at least as good as the first one by providing extra freedom on executing the route which allows some skipping on the low demand level nodes while still preserving the advantages of customer-vehicle familiarity and the readiness for training purposes as well. The trade-off is that it requires additional computational effort since the vehicle flow variables are introduced into the model again. However, the flexibility on the vehicle flow is restricted within the fixed route, which dramatically decreases the admissible search space compared with the first-stage model. Hence, it can be solved efficiently. We call it the knapsack recourse strategy because we propose an approximation solution approach by solving a knapsack problem. The problem can be solved efficiently by this approximation algorithm and the details of the approach will be discussed in the next section.

3.3.3 The Re-planning Strategy

The last proposed recourse strategy is to solve a deterministic routing problem to minimize the unmet demand used for the first planning stage by plugging in the actual travel time and demand value. The advantage of this strategy lies in that it provides the largest search space and most flexibility regarding the current emergency scenario. But it loses the pre-knowledge of the planned routes which makes it difficult for training and driver learning on the routes. Since familiarity of the routes and training are important aspects of reliable solutions, it is impractical to apply this re-plan strategy in real life. However, this strategy can provide us with a bound to compare and evaluate the previous two strategies in the experimental section below.

4 Solution Approaches

4.1 The Planning Stage – Tabu Search Heuristics

Tabu search is a widely used metaheuristic algorithm. The search keeps a tabu list which prohibits revisiting a recently explored solution unless some aspiration criteria are met to avoid cyclic movement. It allows a solution to temporarily move to a worse position to escape a local optimum. Tabu search has been successfully used in solving hard optimization problems in many fields. Applying tabu search to a particular problem requires a fair amount of problem-specific knowledge. Given the success of tabu heuristic in the classical VRP and its variants, and the similar structure of our model, we believe this approach holds much promise in solving our problem. The algorithm we propose here uses some ideas from the standard VRP, and incorporates new features taking into account the unmet demand objective of our problem.

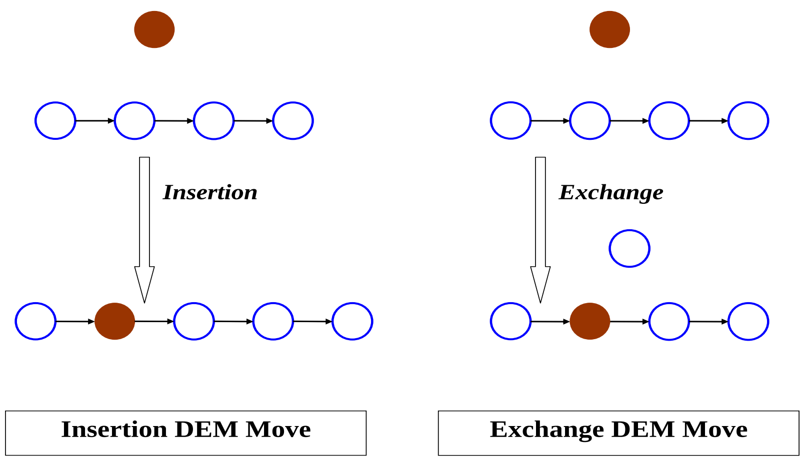

Figure 2: Illustration of the DEM Moves.

A key element in designing a tabu heuristic is the definition of the neighborhood of a solution or equivalently the possible moves from a given solution. Beside adopting the standard 2-opt exchange move and the λ-interchange move, we implement a new DEM-move to try to reduce the unmet demand. The DEM-move inserts an unassigned demand node into a current route or exchanges an unassigned node with some (one to three) consecutive node(s) in a route by abiding both deadline and capacity constraints. Figure 2 illustrates some examples of the DEM move. The upper left graph represents an unassigned node and an existing route before the DEM move. The lower left graph gives the result after an insertion DEM move. The unassigned node has been inserted between the first and second nodes of the existing route. The lower left graph is one of the upper left graph's neighbors. This unassigned node could also be inserted between the second and third nodes or between the third and fourth nodes as long as the deadline and capacity constraints can be met after the insertion. The two graphs at the right hand side show an example of the exchange DEM move. The λ-interchange move exchanges a subsequence of at most λ consecutive customers on a route with a different sequence also of up to λ customers. This move includes a combination of vertex reassignments to a different route and vertex interchanges between two routes. With $\lambda = 3$, the reassigning/interchanging segment is up to three consecutive vertices in a route. (See figure 3 for the illustration with $\lambda = 1$ insertion at the left and $\lambda = 2$ exchange at the right.)

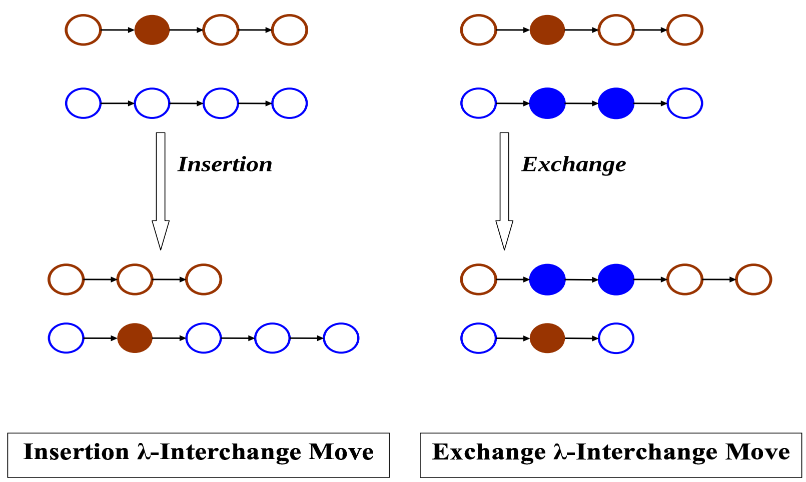

Figure 3: Illustration of the λ-interchange Moves.

Because the problem has insufficient supplies at the depot in most scenarios, we use an *UnassignNodeManager* list of nodes to keep track of all the nodes with unmet demand. We also use a *RouteManager* list of solutions to keep track of all the incomplete (may not include all the nodes) but feasible (meet both capacity and time constraints) routes. We initiate a solution by using a visit-nearest-node heuristic and put all the unvisited nodes into the *UnassignNodeManager* list. We also maintain a *MoveManager* list to record all possible moves from the current solution. After a solution moves to its neighbor, instead of reconstructing the whole neighborhood of the new solution, we only update the non-overlapping neighbors. That is, we eliminate those moves that are relevant (sharing the same route or sharing the same unassigned node, etc.) to the move that just has been executed, and generate new feasible moves that are relevant to the updated route(s) only. This significantly reduces the computational effort of exploring the neighborhood.

After we obtain the tabu search solution, we complete the route by adopting a variation of a "next-earliest-node" heuristic to insert all unassigned nodes after the deadline of each route (keeping the before-deadline part intact). In this post-processing heuristic, each unassigned node picks a route to locate itself after the deadline, where the summation of its arrival time and the arrival time to the next node is minimized at the time of assignment.

The quality of this tabu heuristic is evaluated in the next section.

4.2 The Operational Stage

The first LP recourse strategy can be solved by CPLEX, while the third re-plan strategy shares the same model as the planning stage, hence can be solved by the tabu heuristic we presented in the previous subsection. Therefore, in this subsection, we focus on presenting an approximation algorithm for the second knapsack recourse strategy.

After the actual demand and travel time have been revealed at the beginning of the second stage, the knapsack recourse strategy follows the first-stage pre-planned routes and allows skipping low demand level nodes. Hence, there are two decisions to make: (1) which nodes to be skipped, and (2) how much supply to load on each vehicle under the given total supply and deadline, to minimize the total unmet demand. We propose the following solution procedure:

For each pre-planned route, we enumerate all the feasible routes by skipping some nodes and visiting the rest in the original sequence. A feasible route must be finished within the pre-specified deadline. This is equivalent to assigning a binary digit to each node in the pre-planned route, 1 if selected, 0 otherwise. Therefore, a vehicle that visits q demand nodes could have up to 2^q feasible routes, which means there are at most 2^n feasible set of routes for the entire fleet. We use L_k to denote the set of feasible routes for vehicle k. For every feasible route l for vehicle k in L_k, an associated demand level $a_{k,l}$ is the summation of the demand of the nodes in the route. With a known total supply quantity at the depot, we can solve an integer programming problem to decide the route to be used for each vehicle.

Parameters:

s : amount of supplies at the supply node (depot)
$a_{k,l}$: the total demand on route l of vehicle k.

Binary Decision variables:

$W_{k,l}$: equals to 1 if route l is used by vehicle k, 0 otherwise.

$$
\mathbf{IP}: \quad \begin{array}{ll} \text{maximize} & \sum_{k \in K, l \in L_k} a_{k,l} W_{k,l} \\ \text{subject to} & \text{constraints } (36) - (37)\ , \end{array}
$$

$$
\sum_{l \in L_k} W_{k,l} = 1 \qquad (\forall l \in K) \tag{36}
$$

$$
\sum_{k \in K, l \in L_k} a_{k,l} W_{k,l} \leq s \tag{37}
$$

The IP aims to maximize the demand that can be serviced, which is the same as minimizing the total unmet demand. Constraint (36) states that only one route can be selected for each vehicle. Constraint (37) enforces the total demand to be serviced will not exceed the supply available at the depot. This problem is called a multiple-choice knapsack problem, which is a 0-1 knapsack problem with the addition of the disjoint multiple-choice constraints. Pisinger (1995) presented a simple partitioning algorithm and incorporated it into a dynamic programming framework to solve the problem efficiently. Problems with 1000 different classes of items and each class with 100 items can be solved within 200 seconds.

The solution of this IP will specify which route calculated in the first enumeration process should be used (hence we know which node(s) would be skipped); and the associated parameter $a_{k,l}$ states the quantity of supply loaded on vehicle k. The constraints from the temporal perspective are considered in the first enumeration procedure (the definition of a feasible route); and the ones from the demand/supply perspective are stated in the second IP model.

5 Computational Experiments

In this section, we demonstrate how the proposed models and two-stage solution approach presented in this paper are used.

We define a *problem setting* as a randomly generated network with a single depot and a set of dispersed demand nodes whose mean demand quantity is also randomly generated. The size of a *problem setting* is defined as the number of the demand nodes. We conduct two sets of experiments. The first is to evaluate the performance of the tabu heuristic for the first-stage model on different size problems by comparing the heuristic results with CPLEX bounds. The second is to perform simulations on 10 different settings with size 50 and average the results. The purpose of the second experiment is to compare the effect of different recourse strategies through simulations.

We first describe how to generate the input parameters, then demonstrate and discuss the experimental results.

5.1 Data Generation of Input Parameters

We test our tabu heuristic on 3 different problem sizes: with 10, 20 and 50 demand nodes. We generate 5 different *problem settings* for each size. For each setting we uniformly generate the coordinates of demand nodes and 1 depot node in a 200 by 200 square domain; and a mean demand quantity for each demand node ranging from 5 to 15. We set the uncapacitated fleet size as 3, 4 and 10 for problem size 10, 20 and 50 respectively. We use the Euclidean distances between any two nodes and assume a symmetric complete graph topology. The mean travel time between any two nodes is proportional to the distance.

In the CCP model, we use a lognormal distribution with a randomly generated mean value; the standard deviation is set to be proportional to the mean value of demand (20% of the mean value) and given by the following relationship for travel time uncertainty $\sigma = \frac{UB-\mu}{100}\mu$, where UB is an upper bound on the mean travel time. This relationship creates arcs with small uncertainty and large mean travel times. We set the confidence level as 85% for the chance constrained model, which sets the values κ_D and κ_T to 1.04. Since the routes generated from the model are sensitive to both the deadline and the total supply at the depot, we vary these two parameters and observe the results. We use 70%, 80%, 90%, 100% and 120% of the *base quantity* as the available supply quantity. The *base quantity* of the total supply at the depot is defined to be the summation of the mean demand quantity at all the demand nodes, which is 100, 200, and 500 on average for the three different size problem settings respectively. The deadline is set to 40%, 50%, 60%, 80%, 100% and 120% of the *base route length.* The *base route length* is defined as the average length of all the edges in the graph times the average number of served nodes per vehicle. For example, for the size 50 problem, 50 demand nodes are served by 10 vehicles; so on average, each vehicle serves 5 demand nodes; hence the average number of served nodes per vehicle is 5. We call one combination of the deadline and the total supply parameters a *test case.* In the experiment to evaluate the quality of the tabu heuristic, we run each problem setting on 9 *test cases*, which are identified by the different combinations of the deadline on the demand nodes (3 types: 40%, 80%, 120%) and the total supply at the depot (3 types: 70%, 90%, 120%). In the simulation to compare different recourse strategies, we have 30 *test cases* for each problem setting, which are the combinations of 6 types of deadline (40%, 50%, 60%, 80%, 100%, 120%) and 5 types of total supply (70%, 80%, 90%, 100%, 120%).

5.2 Quality of Tabu Heuristic

To evaluate the quality of our tabu heuristic, we run this search process over 5 random *problem settings* over each problem size (10, 20 and 50) as described above. We also run these problems using CPLEX 9.0 for one hour and record the lower bound and upper bound obtained from the CPLEX solver to compare with the tabu result. The results for size 10, 20 and 50 are presented in Tables 1, 2 and 3 respectively.

From these tables, we can see that for smaller size problems (size 10), the CPLEX solver can obtain a feasible solution within 1 hour for most of the instances (40 out of the 45 instances). The Tabu heuristic results obtained within 200 seconds generate better solutions than the CPLEX results for about half of the instances (17 out of the 40 instances). The CPLEX solver cannot obtain a feasible solution within 1 hour for all larger size problems (size 20, 50). However, Tabu search can generate a solution within 200 seconds close to the lower bound obtained from the CPLEX. The lower bound provided by the CPLEX solver primarily reflects the shortage of the supply at the depot compared with the total demand. These tables also show that there are bigger gaps between the CPLEX lower bound and the tabu result for the tight deadline and low total supply level combination. We suspect this is due to the tight deadline which might prevent the timely delivery even with enough supplies. A tighter lower bound with respect to a short deadline is a subject for future research. However, in general, we conclude that the proposed tabu search method is very effective in minimizing the unmet demand, which is our primary concern in this model.

		10 Nodes	3 Vehs			CPLEX =	3600 sec				
Deadline			150	- (40%)		250	- (80%)		400	- (120%)	
		CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu	
Supply											
70 - (70%)	case 1	14633	N/A	14792	14633	N/A	14792	14633	14664	14792	
	case 2	30523	30537	33002	30522	30538	30913	30522	30537	30913	
	case 3	40162	40176	41774	40162	40176	41774	40162	40180	41774	
	case 4	22043	22059	24996	22044	22059	22275	22043	22059	22275	
	case 5	34762	41121	37764	34762	34784	37764	34763	34781	35964	
90 - (90%)	case 1	4.11	N/A	14792	3.08	N/A	11.86	3.39	32.01	11.86	
	case 2	10522	18801	18789	10522	10538	11008	10523	10539	11008	
	case 3	20162	22984	23133	20162	20179	21912	20162	20181	21912	
	case 4	2044	29233	17952	2043	2060	6582	2044	2063	6582	
	case 5	14763	30712	21961	14763	14782	16301	14763	14781	15811	
120 - (120%)	case 1	3.11	34074	14792	3.30	5964	11.86	3.25	27.60	11.86	
	case 2	2.30	10536	18789	2.14	31.45	8.16	2.67	16.20	8.16	
	case 3	1.70	8980	16539	1.85	21.94	8.24	2.10	20.24	8.24	
	case 4	3.29	N/A	17952	3.15	21.35	11.37	3.65	20.67	11.37	
	case 5	2.47	37776	21961	2.57	27.70	10.70	3.02	26.46	10.70	

Table 1: Tabu heuristic result with CPLEX bounding for 10 customers and 3 vehicles

		20 Nodes	4 Vehs		CPLEX =	3600 sec				
Deadline			200	- (40%)		400	- (80%)		600	- (120%)
		CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu
Supply										
140 - (70%)	case 1	36461	N/A	36499	36461	N/A	36498	36461	N/A	36498
	case 2	70551	N/A	70851	70552	N/A	70753	70552	N/A	70753
	case 3	75171	N/A	75400	75171	N/A	75344	75171	N/A	75344
	case 4	56362	N/A	60026	56362	N/A	58375	56363	N/A	58375
	case 5	57491	N/A	58152	57492	N/A	59671	57492	N/A	59671
180 - (90%)	case 1	1.15	N/A	14822	1.08	N/A	20.78	1.21	N/A	20.78
	case 2	30558	N/A	30978	30552	N/A	30846	30552	N/A	30846
	case 3	35171	N/A	36337	35171	N/A	35360	35171	N/A	35360
	case 4	16362	N/A	53166	16362	N/A	20168	16362	N/A	20168
	case 5	17492	N/A	28996	17491	N/A	19945	17491	N/A	19945
240 - (120%)	case 1	1.15	N/A	14822	1.15	N/A	20.78	1.21	N/A	20.78
	case 2	1.77	N/A	7800	1.83	N/A	20.39	1.69	N/A	20.39
	case 3	0.98	N/A	15332	0.96	N/A	21.89	0.99	N/A	21.89
	case 4	2.53	N/A	53166	2.53	N/A	22.97	2.51	N/A	22.97
	case 5	1.58	N/A	28996	1.42	N/A	23.19	1.59	N/A	23.19

Table 2: Tabu heuristic result with CPLEX bounding for 20 customers and 4 vehicles

50 Nodes 10 Vehs CPLEX = 3600 sec

Deadline			200	- (40%)		400	- (80%)		600	- (120%)
		CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu	CPLEX LB	CPLEX UB	Tabu
Supply										
350 - (70%)	case 1	139752	N/A	140100	139752	N/A	139810	139752	N/A	139810
	case 2	154132	N/A	154190	154132	N/A	154190	154132	N/A	154190
	case 3	137861	N/A	137960	137861	N/A	137960	137861	N/A	137960
	case 4	147863	N/A	147930	147863	N/A	147940	147863	N/A	147940
	case 5	150453	N/A	150710	150453	N/A	150550	150453	N/A	150550
450 - (90%)	case 1	39751	N/A	40569	39751	N/A	40125	39751	N/A	40125
	case 2	54132	N/A	59333	54132	N/A	54195	54132	N/A	54195
	case 3	37861	N/A	38950	37861	N/A	37949	37861	N/A	37949
	case 4	47863	N/A	48825	47863	N/A	47924	47863	N/A	47924
	case 5	50453	N/A	50766	50453	N/A	50536	50453	N/A	50536
600 - (120%)	case 1	1.49	N/A	40569	1.49	N/A	47.23	1.49	N/A	47.23
	case 2	1.63	N/A	40470	1.63	N/A	47.77	1.63	N/A	47.77
	case 3	1.22	N/A	38950	1.22	N/A	47.55	1.22	N/A	47.55
	case 4	3.36	N/A	30781	3.36	N/A	48.34	3.36	N/A	48.34
	case 5	2.70	N/A	29562	2.70	N/A	58.99	2.70	N/A	58.99

Table 3: Tabu heuristic result with CPLEX bounding for 50 customers and 10 vehicles

5.3 Simulation and Analysis

In this simulation experiment, we evaluate and compare the effectiveness of different recourse strategies we discussed in section 4. We present the process of the simulation as a flowchart in figure 4. First, with different given deadline and total supply combination and other input parameters (include the stochastic travel time and demand parameters), we use the tabu search heuristic to solve the first-stage chance-constrained model to obtain the set of routes and commodity flow values. However, at the pre-emergency stage, only the preplanned routes are of interest since they provide the guide for training and preparation purposes; the commodity flow solution which determines how much to load on each vehicle is not very meaningful at the planning stage since the actual demand may vary significantly. Hence we discard the commodity flow solution and only keep the preplanned routes. Then we randomly generate the realization of those stochastic parameters (travel time and demand) following the distribution we used to model the first-stage CCP formulation, which is a lognormal distribution. The lognormal distribution is widely used to simulate travel time uncertainties in the transportation literature and we also use it to eliminate the negative sample value. For simplicity, we apply the same distribution for the demand uncertainty as well. The literature which addresses the distribution process of large-scale emergencies is very limited and it is hard to collect and analyze real-life data. A more precise probability distribution for the demand and travel times may be the subject of further studies. For each *test case*, 100 parameter realizations have been generated to simulate different settings. The simulated data and the preplanned routes are sent to the three different recourse strategy models to obtain the new commodity flow solution as well as the actual total unmet demand. In summary, we average each *test case* (a supply and deadline combination) over 10 different *problem settings* with each setting over 100 *replications*. Hence the final result is the average value of 1000 raw data points.

The final result is presented in table 4. The routes obtained in the planning stage are from the chance-constrained model. The number in each grid is the percentage of the unmet demand quantity over the total demand level. For every *test case*, the re-plan column gives the best result (lowest total unmet demand), since it regenerates both the routes and the commodity flows at the same time, which provides the largest search space. We note that the solution for the re-planned routes may not necessarily be optimal since we used the Tabu heuristic to identify them. Nevertheless, the re-planned routes outperforms the other two strategies. However, with different realization parameters, this approach may generate significantly different routes. It loses the advantage of the familiarity of preplanned routes. On the other hand, both the LP and knapsack recourse strategies serve for training and preparation purposes. Since the knapsack recourse strategy provides extra flexibility by allowing to skip low-demand and/or far-away nodes in a route, it has a larger search space than the LP strategy. The LP can be solved to optimality using CPLEX and the knapsack strategy can be solved to a near optimal solution by the approach we presented in section 5. In our numerical simulation, the knapsack recourse strategy always outperforms the

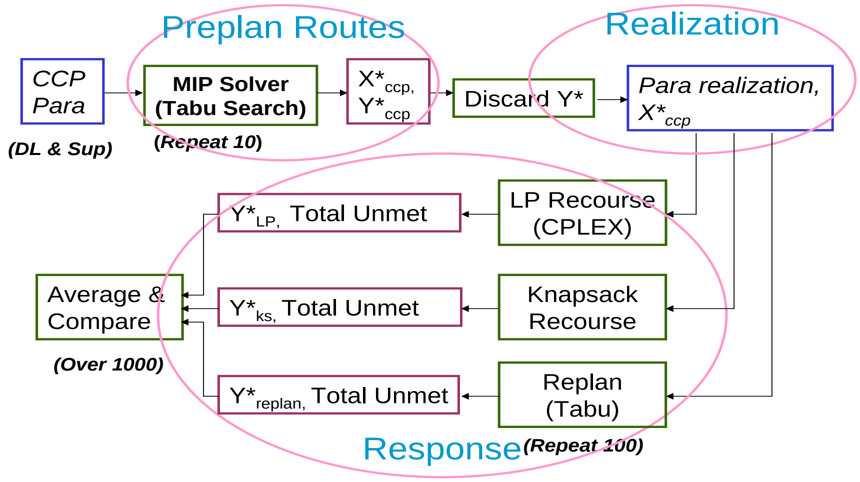

Figure 4: Flowchart of Simulation Process.

LP recourse.

	DL 200			DL 250			DL 300		
	LP	Knapsack	Replan	LP	Knapsack	Replan	LP	Knapsack	Replan
SUP 350	30.7%	30.3%	30.3%	30.5%	30.3%	30.3%	30.3%	30.3%	30.3%
SUP 400	22.1%	20.4%	20.4%	21.2%	20.4%	20.4%	20.8%	20.4%	20.4%
SUP 450	16.7%	11.3%	10.5%	13.7%	10.8%	10.5%	12.5%	10.6%	10.5%
SUP 500	15.2%	8.1%	1.9%	10.5%	5.7%	1.9%	8.0%	4.2%	1.9%
SUP 600	15.2%	6.8%	0.0%	10.3%	5.6%	0.0%	7.5%	3.2%	0.0%

	DL 400			DL 500			DL 600		
	LP	Knapsack	Replan	LP	Knapsack	Replan	LP	Knapsack	Replan
SUP 350	30.3%	30.3%	30.3%	30.3%	30.3%	30.3%	30.3%	30.3%	30.3%
SUP 400	20.4%	20.4%	20.4%	20.4%	20.4%	20.4%	20.4%	20.4%	20.4%
SUP 450	11.2%	10.5%	10.5%	10.9%	10.5%	10.5%	10.7%	10.5%	10.5%
SUP 500	5.4%	3.2%	1.9%	4.1%	2.7%	1.9%	3.5%	2.5%	1.9%
SUP 600	4.5%	1.5%	0.0%	2.9%	1.0%	0.0%	2.1%	0.8%	0.0%

Table 4: Unmet demand percentage comparison between 3 different recourse strategies for 50 customers and 10 vehicles (average over 10 networks)

When the problem has a tight total supply level these 3 strategies perform similarly, regardless of the deadline. All are very close to the trivial lower bound – the shortage of the supply compared with the total demand. This shows that these models did their best to send out all the available supply within the given deadline in such conditions. On the other hand, when there are tight deadlines, as the total supply level goes up, the gaps between the re-plan strategy and the knapsack strategy increase. This increase, but more pronounced, is also observed when comparing the re-plan and the LP strategy. This extra solution search space gives the re-plan and knapsack model more chances to reach a better result. In summary, the knapsack recourse strategy provides a nice trade-off between maintaining the familiarity of the preplanned routes and an efficient solution and quick solution times. Indeed, it is the most efficient recourse strategy that uses the preplanned routes and the proposed knapsack approximation algorithm is able to obtain solutions quickly.

Finally we note that the tabu heuristic presented in this paper performs better in practice than the tabu heuristic in our previous work (Shen et al. to appear). This means that the heuristic of this paper obtains solutions that improve the average unmet demand in the simulation analysis. The solutions obtained with the new tabu heuristic (Table 4, LP column) obtain more than an 8% drop in unmet demand from the old heuristic (Shen et al. to appear, Table 3, CCP model) for instances with tight deadlines and ample supply of medicine.

6 Conclusions

In this paper, we studied routing problems in response to large-scale bioterrorism emergencies. We analyzed the characteristics of routing problems for large-scale bioterrorism attacks and decomposed the problem into two stages: planning and operational stages. We proposed a chance constrained model for the planning stage and three different recourse strategies for the operational stage. Solution approaches for both stages were presented. Finally, we illustrated the effectiveness of the models and solution approaches by computational experiments and concluded that the first-stage chance-constrained model combined with the second-stage knapsack recourse strategy were most effective in solving our problem at hand.

We note that in classical stochastic programming models with recourse the first-stage solution is informed by the outcome of the recourse stage. We did not pursue this approach for two reasons: in a large scale emergency situation the uncertainty is so significant that it is difficult to represent it with scenarios accurately. The second reason is because such a stochastic model with recourse would be considerably more difficult to solve. Nevertheless, a comparison of our preplanned solution against the solution from a stochastic model with recourse is an interesting problem and the topic of future research.

Acknowledgment

This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE), grant number EMW-2004-GR-0112. However, any opinions, findings, and conclusions or recommendations in this document are those of the author(s) and do not necessarily reflect views of the U.S. Department of Homeland Security. Also the authors wish to thank Hongzhong Jia, Harry Bowman, Richard Larson and Terry O'Sullivan for their valuable input and comments for the improvement of this paper.

References

Archetti, C., A. Hertz, and M. Speranza (2006). A tabu search algorithm for the split delivery vehicle routing problem. *Transportation Science 40*, 64–73.

Archetti, C., A. Hertz, and M. G. Speranza (2007). Metaheurisitics for the team orienteering problem. *Journal of Heuristics 13*, 49–76.

Bertsimas, D. (1988). *Probabilistic combinational optimization problems*. Ph. D. thesis, Operation Research Center, Massachusetts Institute of Technology, Cambridge, MA.

Bertsimas, D. (1992). A vehicle routing problem with stochastic demand. *Operations Research 40*(3), 574–585.

Bertsimas, D. and D. Simchi-Levi (1996). A new generation of vehicle routing research: robust algorithms, addressing uncertainty. *Operations Research 44*(2), 286–304.

Bianchessi, N. and G. Righini (2006). Heuristic algorithms for the vehicle routing problem with simultaneous pick-up and delivery. *Computers and Operations Research 34*, 578–594.

Boussier, S., D. Feillet, and M. Gendreau (2007). An exact algorithm for team orienteering problems. *4OR 5*, 211–230.

Butt, S. E. and T. M. Cavalier (1994). A heuristic for the multiple tour maximum collection problem. *Computers and Operations Research 21*(1), 101–111.

Butt, S. E. and D. M. Ryan (1999). A optimal solution procedure for the multiple tour maximum collection problem using column generation. *Computers and Operations Research 26*(4), 427–441.

Carraway, R., T. Morin, and H. Moskowitz (1989). Generalized dynamic programming for stochastic combinatorial optimization. *Operations Research 37*(5), 819–829.

Chao, I.-M., B. Golden, and E. Wasil (1996a). A fast and effective heuristic for the orienteering problem. *European Journal of Operational Research 88*(3), 475–489.

Chao, I.-M., B. Golden, and E. Wasil (1996b). The team orienteering problem. *European Journal of Operational Research 88*(3), 464–474.

Clarke, G. and J. Wright (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research 12*, 568–581.

Dror, M. (1993). Modeling vehicle routing with uncertain demands as a stochastic program: properties of the corresponding solution. *European Journal of Operational Research 64*, 432–441.

Dror, M., G. Laporte, and P. Trudeau (1989). Vehicle routing with stochastic demands: properties and solution framework. *Transportation Science 23*(3), 166–176.

Dror, M., G. Laporte, and P. Trudeau (1994). Vehicle routing with split deliveries. *Discrete Applied Mathematics 50*, 239–254.

Dror, M. and P. Trudeau (1986). Stochastic vehicle routing with modified savings algorithm. *European Journal of Operational Research 23*, 228–235.

Dror, M. and P. Trudeau (1989). Savings by split delivery routing. *Transportation Science 23*, 141–145.

Dror, M. and P. Trudeau (1990). Split delivery routing. *Naval Research Logistics 37*, 383–402.

Dunn, J. (1992). Scheduling underway replenishment as a generalized orienteering problem. Master's thesis, Naval Postgraduate School, Monterey, California, USA.

Feillet, D., P. Dejax, and M. Gendreau (2005). Traveling salesman problems with profits. *Transportation Science 39*(2), 188–205.

Fischetti, M., J. S. Gonzalez, and P. Toth (1997). A branch-and-cut algorithm for the generalized traveling salesman problem. *Operations Research 45*(3), 378–394.

Fischetti, M., J. S. Gonzalez, and P. Toth (1998). Solving the orienteering problem through branch-and-cut. *INFORMS Journal of Computing 10*(2), 133–148.

Geem, Z. W., C. Tseng, and Y. Park (2005). *Lecture notes in computer science*, Chapter Harmony Search for Generalized Orienteering Problem: Best Touring in China, pp. 741–750. Springer-Verlag.

Gendreau, M., A. Hertz, and G. Laporte (1994). A tabu search heuristic for the vehicle routing problem. *Management Science 40*, 1276–1290.

Gendreau, M., G. Laporte, and J. Y. Potvin (2002). *The Vehicle Routing Problem*, Chapter Metaheuristics for the capacitated VRP, pp. 129–154. SIAM Monographs on Discrete Mathematics and Applications, SIAM Publishing.

Gendreau, M., G. Laporte, and R. Seguin (1995). An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science 29*(2), 143–155.

Gendreau, M., G. Laporte, and R. Seguin (1996a). Stochastic vehicle routing. *European Journal of Operational Research 88*, 3–12.

Gendreau, M., G. Laporte, and R. Seguin (1996b). A tabu search heuristic for the vehicle routing problem with stochastic demands and customers. *Operations Research 44*(3), 469–477.

Gendreau, M., G. Laporte, and F. Semet (1998a). A branch-and-cut algorithm for the undirected selective traveling salesman problem. *Networks 32*, 263–273.

Gendreau, M., G. Laporte, and F. Semet (1998b). A tabu search heuristic for the undirected selective traveling salesman problem. *European Journal of Operational Research 106*, 539–545.

Glover, G. and M. Laguna (1997). *Tabu Search*. Kluwer, Boston, MA.

Goel, A. and V. Gruhn (2005). Large neighborhood search for rich vrp with multiple pickup and delivery locations. In *MEC-VNS: 18th Mini Euro Conference on VNS*.

Golden, B., L. Levy, and R. Vohra (1987). The orienteering problem. *Naval research logistics 34*(3), 307–318.

Golden, B., Q. Wang, and L. Liu (1988). A multifaceted heuristic for the orienteering problem. *Naval research logistics 35*, 359–366.

Hadjiconstantinou, E. and D. Roberts (2002). *The Vehicle Routing Problem*, Chapter Routing under uncertainty: an application in the scheduling of field service engineers, pp. 331–352. SIAM Monographs on Discrete Mathematics and Applications, SIAM Publishing.

Hayes, M. and J. Norman (1984). Dynamic programming in orienteering: route choice and sitting of controls. *Journal of the Operational Research Society 35*, 791–796.

Ilhan, T., S. Iravani, and M. Daskin (2008). The orienteering problem with stochastic profits. *IIE Transactions 40*, 406–421.

Jaillet, P. (1987). Stochastic routing problem. In G. Andreatta, F. Mason, and P. Serafini (Eds.), *Stochastics in Combinatorial Optimization*, pp. 178–186. World Scientific, New Jersey.

Jézéquel, A. (1985). Probabilistic vehicle routing problems. Master's thesis, Department of Civil Engineering, Massachusetts Institute of Technology.

Jozefowiez, N., F. Glover, and M. Laguna (2008). Multi-objective meta-heuristics for the traveling salesman problem with profits. *Journal of Mathematical Modelling and Algorithms 7*, 177–195.

Jula, H., M. M. Dessouky, and P. Ioannou (2006). Truck route planning in non-stationary stochastic networks with time-windows at customer locations. *IEEE Transactions on Intelligent Transportation Systems 37*, 51–63.

Kao, E. (1978). A preference order dynamic program for a stochastic travelling salesman problem. *Operations Research 26*, 1033–1045.

Ke, L., C. Archetti, and Z. Feng (2008). Ants can solve the team orienteering problem. *Computers and Industrial Engineering 54*, 648–665.

Laporte, G., F. Louveaux, and H. Mercure (1992). The vehicle routing problem with stochastic travel times. *Transportation Science 26*(3), 161–170.

Laporte, G. and S. Martello (1990). The selective traveling salesman problem. *Discrete applied mathematics 26*, 193–207.

Larson, R. C. (2005). *The McGraw-Hill Handbook of Homeland Security*, Chapter Decision Models for Emergency Response Planning, pp. 911–930. The McGraw-Hill Companies.

Larson, R. C., M. Metzger, and M. Cahn (2006). Responding to emergencies: Lessons learned and the need for analysis. *Interfaces 36*(6), 486–501.

Leifer, A. and M. Rosenwein (1994). Strong linear programming relaxations for the orienteering problem. *European Journal of Operational Research 73*, 517–523.

Liang, Y.-C. and A. Smith (2001). An ant colony approach to the orienteering problem. Technical report, Department of industrial and systems engineering, Auburn University, Auburn, AL.

Mak, V. and T. Thomadsen (2006). Polyhedral combinatorics of the cardinality constrained quadratic knapsack problem and the quadratic selective travelling salesman problem. *Journal of Combinatorial Optimization 11*, 421–434.

Mladenovic, N. and P. Hansen (1997). Variable neighborhood search. *Computers and Operations Research 24*, 1097–1100.

Novoa, C., R. Berger, J. Linderoth, and R. Storer (2006). A set-partitioning-based model for the stochastic vehicle routing problem. Technical Report 06T-008, Lehigh University, 200 W. Packer Ave., Bethlehem, PA 18015, USA.

Osman, I. H. (1993). Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research 41*, 421–451.

Pisinger, D. (1995). A minimal algorithm for the multiple-choice knapsack problem. *European Journal of Operational Research 83*, 394–410.

Ramesh, R., Y.-S. Yoon, and M. Karwan (1991). An efficient four-phase heuristic for the generalized orienteering problem. *Computers and Operations Research 18*(2), 151–165.

Ramesh, R., Y.-S. Yoon, and M. Karwan (1992). An optimal algorithm for the orienteering tour problem. *ORSA Journal of Computing 4*(2), 155–165.

Secomandi, N. (2001). A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research 49*(5), 796–802.

Shen, Z., F. Ordóñez, and M. M. Dessouky (to appear). *Springer book series on Optimization and its Applications*, Chapter The Minimum Unmet Demand Stochastic Vehicle Routing Problem. Springer.

Stewart, W. and B. Golden (1983). Stochastic vehicle routing: a comprehensive approach. *European Journal of Operational Research 14*, 371–385.

Sungur, I., F. Ordóñez, and M. M. Dessouky (2008). A robust optimization approach for the capacitated vehicle routing problem with demand uncertainty. *IIE Transactions 40*, 509–523.

Taillard, E. D., P. Badeau, M. Gendreau, F. Guertin, and J. Y. Potvin (1997). A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science 31*, 170–186.

Tang, H. and E. Miller-Hooks (2005a). Algorithms for a stochastic selective travelling salesperson problem. *Journal of the Operational Research Society 56*, 439–452.

Tang, H. and E. Miller-Hooks (2005b). A tabu search heuristic for the team orienteering problem. *Computers and Operations Research 32*, 1379–1407.

Tasgetiren, F. M. and A. E. Smith (2000). A genetic algorithm for the orienteering problem. In *Proceeding of 2000 Congress Evolutionary Computation*, pp. 1190–1195.

Toth, P. and D. Vigo (2002). *The Vehicle Routing Problem.* SIAM Monographs on Discrete Mathematics and Applications, SIAM Publishing.

Toth, P. and D. Vigo (2003). The granular tabu search and its application to the vehicle routing problem. *INFORMS Journal on Computing 15*, 333–346.

Tsiligirides, T. (1984). Heuristic methods applied to orienteering. *Journal of Operational Research Society 35*(9), 797–809.

Vansteenwegen, P., W. Souffriau, G. V. Berghe, and D. V. Oudheusden (to appear). A guided local search metaheuristic for the team orienteering problem. *European Journal of Operational Research.*

Wang, Q., X. Sun, B. Golden, and J. Jia (1995). Using artificial neural networks to solve the orienteering problem. *Annuls of Operational Research 61*, 111–120.

Willard, J. A. G. (1989). Vehicle routing using r-optimal tabu search. Master's thesis, The Management School, Imperial College, London.

Wu, T. (1992). Optimization models for underway replenishment of a dispersed carrier battle group. Master's thesis, Naval Postgraduate School, Monterey, California, USA.

Xu, J. and J. P. Kelly (1996). A network flow-based tabu search heuristic for the vehicle routing problem. *Transportation Science 30*, 379–393.