Better Delivery/Pick Up Routes in the Presence of Uncertainty

Draft Final Report

Metrans Project 06-11

August 2007

Fernando Ordóñez and Maged Dessouky

Graduate Student: Ilgaz Sungur

Industrial and Systems Engineering

University of Southern California

Los Angeles, CA 90089

METRANS
Transportation Center
USC / CSULB

1 Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the Department of Transportation, University Transportation Centers Program, and California Department of Transportation in the interest of information exchange. The U.S. Government and California Department of Transportation assume no liability for the contents or use thereof. The contents do not necessarily reflect the official views or policies of the State of California or the Department of Transportation. This report does not constitute a standard, specification, or regulation.

2 Abstract

We consider the Courier Delivery Problem, a variant of the Vehicle Routing Problem with time windows in which customers appear probabilistically and their service times are uncertain. We use scenario-based stochastic optimization with recourse for uncertainty in customers and robust optimization for uncertainty in service times. Our proposed model generates a master plan and daily schedules by maximizing the coverage of customers and the similarity of routes in each scenario while minimizing the total time spent by the couriers and the total earliness and lateness penalty. For the large scale problem instance, we develop a two-phase insertion based approximate solution procedure, MADS, to balance these different objectives. We provide simulation results and show that our heuristic improves the similarity of routes and the lateness penalty at the expense of increased total time spent when compared to a solution by independently scheduling each day. Our experimental results also show improvements over a real-life solution.

Contents

List of Tables

List of Figures

3   Disclosure

Project was funded in entirety under this contract to California Department of Transportation.

4   Acknowledgments

5   Introduction

In this study, we consider the Courier Delivery Problem (CDP) which is a variant of the Vehicle Routing Problem with Time Windows (VRPTW). We focus on courier delivery/pick-up for an urban area where travel distances are relatively small, and therefore we assume has no uncertainty. On the other hand, the number of potential customers who can make a delivery request each day is very large. Couriers, therefore, are to make frequent stops with short travel times. Each delivery to a customer site is associated with a service time and it must be done within a specific interval of time given by time windows which can be either hard constraints or soft constraints allowing penalty for violations. The problem could be either capacitated or uncapacitated depending on the type of the goods delivered. Uncertainty arises most of the time due to the unknown service times and to the probabilistic nature of the customers, i.e. daily delivery requests from potential customers are not known beforehand but they usually become available in the morning.

When the customer locations and service times are revealed at the beginning of each day, a simple solution of routing for a planning horizon is to repeat the routing solution procedure each day based on that day's requirements which would yield independent daily schedules, see (Savelsbergh and Goetschalckx 1995; Beasley and Christofides 1997). However, the challenge many times in industry is to create similar daily schedules which necessitate a master plan for the planning horizon. The master plan is generated based on historical data and then used as a basis for daily schedules of couriers to cover all the delivery requests at minimum cost. Thus, the master plan helps eliminate significant re-planning in each day, and increases the similarity between daily schedules and the familiarity of drivers with the daily routes. The master plan also constitutes a basis for the estimate of resource and cost allocations of the service provider for the planning horizon.

To formulate the CDP, we adapt a VRPTW formulation and use combination of robust optimization and stochastic programs with recourse to address the uncertainty in service times and customers. Robust optimization considers all possible realizations of uncertainty and results in solutions that exhibit little sensitivity to data variations and are said to

be *immunized* to this uncertainty. Thus, robust solutions are *good* for all possible data uncertainty. Stochastic programs with recourse however give a flexible and adaptive solution which can be different for each realization of the uncertainty via a recourse action at the expense of additional assumptions and increased complexity in the formulation. Yet, in certain applications recourse actions are inevitable, i.e. generating daily schedules based on a master plan. We consider each day in the planning horizon as a scenario and based on historical data we generate a robust master plan for customers who are most likely to occur. We also incorporate recourse actions in the problem formulation to provide flexible daily schedules based on the robust master plan. The objective of the problem is, for each scenario, to maximize the coverage of customers, to minimize the total time spent by the couriers and the total earliness and lateness penalty, and to maximize the similarity of the daily routes with the robust master plan.

In formulating a master plan, a variety of approaches could be used in addressing uncertainty in service times and customers. However, since one of the primary goals in generating the master plan is to improve similarities in daily schedules which are different for each outcome of the uncertainty, an intuitive approach is to use a method which would result in a master plan which would stay *good* for all possible realizations of the uncertainty and thus which will require the least modifications in the resulting daily schedules. This suggests using a robust optimization approach since using a master plan generated by an expected value approach, for instance, may require significant modifications in generating daily schedules for certain realizations of the uncertainty resulting in a poor performance of similarity. By a similar and intuitive reasoning for the uncertainty in customers, we can say that considering a subset of customers with high probability of occurrence in constructing the master plan should improve the similarity of daily schedules since the master plan would be tailored towards these most likely customers who appear in most of the days.

To solve the CDP, we develop a two-phase approximate solution procedure, Master And Daily Scheduler (MADS), using insertion as an efficient heuristic subroutine to address the large scale real-life problem. The first phase is the construction phase where a master plan

and initial daily schedules are generated from the given scenarios. The second phase is iterative. At each iteration, first the master plan is modified based on the feedback given by the daily schedules and second the daily schedules are updated based on the new master plan. The second phase and thus the two-phase algorithm end when there is no further improvement in the objective function at the end of an iteration reaching a local optimum solution. As a result, a master plan is generated based on historical data, which is to be used as a basis in constructing future daily schedules based on a recourse action during the planning horizon.

The structure of this report follows. We discuss the relevant literature in the next section. In Section 7, we present the CDP formulation for problems with service time uncertainty and probabilistic customers. In Section 8, we propose the two-phase heuristic for master and daily schedules, MADS. We present our computational results in Section 9. These include a comparison of MADS to a solution obtained by independently scheduling each day without a master plan, and to a real-life solution. We present conclusions and recommendations in Section 10 and finish with a description of implementation issues in Section 11.

6   Literature Review

Problems where a given set of vehicles with finite capacity have to be routed to satisfy a geographically dispersed demand at minimum cost are known as Vehicle Routing Problems (VRP). This class of problems was introduced by Dantzig and Ramser (1959) and since has lead to a considerable amount of research on the VRP itself and its numerous extensions and applications. General surveys of vehicle routing research can be found in Toth and Vigo (2002), Fisher (1995), and Laporte and Osman (1995). The VRP is known to be NP-Hard (Lenstra and Rinnooy Kan 1981), but nevertheless, there is considerable work on developing exact solution procedures, see for instance (Lysgaard et al. 2004; Baldacci et al. 2004; Ralphs et al. 2003; Fukasawa et al. 2006).

Cordeau et al. (2002) and Laporte et al. (2000) provide an excellent survey of the different

approximate solution approaches to the VRP and the VRP with Time Windows (VRPTW). There is a variety of different heuristics that have been applied to the VRP and to its variants. One popular heuristic is insertion and recent examples are Quadrifoglio and Dessouky (2007), and Lu and Dessouky (2006). Solomon (1987) compares different classical insertion methods with other heuristics on well-known benchmark problems. In a more recent study, Cordeau et al. (2002) use a similar method to compare several heuristics on large problem instances based on four measures: accuracy, speed, simplicity and flexibility.

The most studied areas in the stochastic vehicle routing problem literature have been the VRP with stochastic demands (VRPSD), and with stochastic customers (VRPSC). A major contribution to VRPSD comes from Bertsimas (1992), where a-priori solutions use different recourse policies to solve the VRPSD and bounds, asymptotic results and other theoretical properties are derived. Bertsimas and Simchi-Levi (1996) surveys work on VRPSD with an emphasis on the insights gained and on the algorithms proposed. Different solution algorithms are presented in Dror et al. (1989) and Dror (1993), including conventional stochastic programming and Markov decision processes for single and multi-stage stochastic models. Secomandi (2001) introduces a re-optimization type routing policy for the VRPSD.

The VRPSC, where fixed demand customers have a probability p_i of being present, and the VRP with stochastic customers and demands (VRPSCD), which combines VRPSC and VRPSD, first appeared in Jézéquel (1985), Jaillet (1987) and Jaillet and Odoni (1988). Bertsimas (1988) gives a systematic analysis and presents several properties, bounds and heuristics. Gendreau et al. (1995) proposes the first exact solution, an L-shaped method and a tabu search meta-heuristic for the VRPSCD.

A number of models and solution procedures for VRPSD and VRPSC allow recourse actions, see Gendreau et al. (1996) for a good survey. Such models allow actions to adjust an a-priori solution after the uncertainty is revealed, and therefore can consider a broader set of possible routes than a method that does not allow recourse. Different recourse actions have been proposed in the literature, such as skipping non-occurring customers, returning to the depot when the capacity is exceeded, or complete reschedule for occurring customers

(Jaillet 1988; Bertsimas et al. 1990; Waters 1989). Recent work by Morales (2006) uses robust optimization for the VRPSD with recourse. It considers that vehicles replenish at the depot, computes the worst-case value for the recourse action by finding the longest path on an augmented network, and solves the problem with a tabu search heuristic.

Compared with stochastic customers and demands, the VRP with stochastic service and travel times (VRPSSTT) has received less attention. Laporte et al. (1992) proposed three models for VRPSSTT: chance constrained model, 3-index recourse model and 2-index recourse model, and presented a general branch-and-cut algorithm for all three models. The VRPSSTT model was applied to a banking context and an adaptation of the savings algorithm was used in the work of Lambert et al. (1993). Jula et al. (2006) developed a procedure to estimate the arrival time to the nodes in the presence of hard time windows. In addition, they used these estimates embedded in a dynamic programming algorithm to determine the optimal routes.

In this work, we use robust optimization for the VRP with stochastic service times. We follow the robust optimization methodology as introduced by Ben-Tal and Nemirovski (1998, 1999) and El-Ghaoui et al. (1998) for linear, quadratic, and general convex programs, which is recently extended to integer programming by Bertsimas and Sim (2003). The general approach of robust optimization is to optimize against the worst instance that might arise due to data uncertainty by using a min-max objective. This typically results in solutions that exhibit little sensitivity to data variations and are said to be *immunized* to this uncertainty. Robust solutions have the potential to be efficient solutions in practice, since they tend not to be far from the optimal solution of the deterministic problem and significantly outperform the deterministic optimal solution in the worst-case (Goldfarb and Iyengar 2003; Bertsimas and Sim 2004).

The robust optimization methodology assumes the uncertain parameters belong to a given bounded uncertainty set. For fairly general uncertainty sets, the resulting robust counterpart can have a comparable complexity to the original problem. For example, a linear program with uncertain parameters belonging to a polyhedral uncertainty set has a robust problem

which is an LP whose size is polynomial in the size of the original problem (Ben-Tal and Nemirovski 1999). This nice complexity result however does not carry over to robust models of problems with recourse, where LPs with polyhedral uncertainty can result in NP hard problems, see (Ben-Tal et al. 2003). An important question, therefore, is how to formulate a robust problem that is not more difficult to solve than its deterministic counterpart. In particular, Sungur (2007) and Sungur et al. (2006) showed that obtaining robust solutions for VRP with demand and travel time uncertainty is not more difficult than obtaining the deterministic solutions.

The particular real-life Courier Delivery Problem (CDP) we consider in this study is an uncapacitated VRPTW with uncertainty in service times together with probabilistic customers. The delivery requests become available each morning and it is desired to construct a master plan for the planning horizon which is to be used as a basis for daily schedules to minimize the cost and to maximize the coverage of customers as well as the similarity of daily routes with the master plan. Some of the recent findings in the literature on the CDP and on its variants follow. For the deterministic pickup and delivery problem systems, Hall (1996) studied the optimal route designs for overnight carriers using continuous space approximation models and demonstrated how these constraints affect the designs. Haughton and Stenger (1998) modeled customer service for fixed route delivery systems under stochastic demand. Later, Haughton (2000) developed a framework for quantifying the benefits of route re-optimization, again under stochastic customer demands. Research that considers the design of large scale logistic systems for uncertain environments comes from Erera (2000). His work adapts a continuum approximation methodology developed for deterministic problems for analysis of large scale vehicle dispatching problems. The methodology is to find expected cost approximations that allow near-optimal configuration of such systems under stochastic customer locations and demand. A recent work by Zhong et al. (2007) proposes the concepts of "cell", "core area", and "flex-zone" to develop a two-stage model which uses the capacity that is not assigned in the first stage to adapt to the demand uncertainty in the second stage. They also improve familiarity of drivers with the daily routes over the

planning horizon by increasing the similarity of each day's schedule. One important distinction between our model and the work of Zhong et al. (2007) is that we explicitly account for the travel distances generated by the routes whereas they use continuous approximation techniques to account for travel time.

7 CDP Formulation

In this section, we formulate the CDP which is a variant of the VRPTW with uncertain service times and probabilistic customers. The delivery requests arrive daily from potential customers with known time windows but uncertain service times at the beginning of each day. The locations of the customers are known but it is not known a-priori whether a particular customer requests a delivery at a given day. There are limited number of couriers to route for a limited time period each day. The first goal is to construct an a-priori master plan for the planning horizon to be used in constructing daily schedules by making modifications every day according to daily customer requests. The reason is twofold. First, it is desired to have similarity in daily schedules to minimize the cost of re-planning every day. In particular, similarity in daily schedules allows for better familiarity in the routes by the drivers, thereby potentially improving the driver productivity. Second, for a given planning horizon, the company will need to plan its resources and budget beforehand. The master plan constitutes a basis for the estimate of such resource and cost allocations. The second goal is to construct a daily schedule of couriers to serve as many customers as possible by meeting the deadlines. The vehicles are allowed to wait at a customer site if they arrive before the earliest start time which might cause an earliness penalty; and if they arrive after the latest start time then a lateness penalty might be incurred. Therefore, the additional challenges in constructing daily schedules is to minimize these penalties as well as minimizing the total time spent by the vehicles in each day, which accounts for travel, waiting, and service times.

Given historical data, total of D days, we consider each past day (scenario) as a realization of uncertainty and we construct scenario-based uncertainty sets for service times and

customer occurrences. In constructing the master plan, we address the probabilistic nature of the customers by attempting to serve only customers with high probability of occurrence and we use the worst-case service times due to robust optimization; both ideas pertain to improving similarity of daily schedules with the master plan. Hence, we obtain robust a-priori routes for the master plan which are then adjusted in each day by following the recourse action of partial rescheduling of routes. That is, in each day, robust routes are used as a basis in constructing daily routes in the light of observed realizations of uncertainty by increasing the similarities with the master plan. This method uses robust optimization to generate an a-priori master plan and stochastic optimization with recourse to generate adjusted daily schedules. We can say that the robust master plan is *trained* by the past data to generate future daily schedules. The implicit assumption here is that the following days are similar to past realizations, i.e. there is no seasonal trends or irregular outcomes.

To formulate the CDP, as a general methodology we adapt a nominal VRPTW. We formulate the master plan to determine the a-priori routes for customers with high probability of occurrence with worst-case service times. We indicate the variables, uncertain parameters, and related sets with superscript 0. Then for each scenario d out of D days, based on historical data we formulate a similar VRPTW with variables, uncertain parameters, and related sets superscripted with d. Lastly, we combine all these *separate* models by relating the recourse variables to a-priori variables based on the recourse action of partial rescheduling of routes in order to increase the similarity of daily schedules with the master plan. Thus, the size of the CDP is $D + 1$ times the size of the nominal VRPTW. In this formulation, the master plan can be in fact considered as a *special* day with selected set of customers and worst-case service times.

The general definition of the nominal VRPTW follows. Given a network of customers and the location of the depot, the objective is to route the fleet of vehicles to serve the maximum number of customers based on their service times and time windows. We consider the soft time windows in this study. That is, if a vehicle arrives to a customer before its earliest start time, there will be a waiting penalty; and similarly if it arrives after its latest

8

start time, there will be a lateness penalty. The vehicles start and end their routes at the depot. The length of a route is composed of travel, waiting, and service times, which is also the total time spent. There is a common due date for vehicles to return to the depot at the end of the day which is a hard constraint. A customer can be served by only a single vehicle and there is no capacity considerations in the problem.

The mathematical formulation of the CDP is given in Problem 1. The depot is located at node 0. K is the number of vehicles. Let N_D be the set of integers from 0 to D to indicate scenarios including master plan at index 0. Let O^d be the n dimensional data vector indicating the occurrence of each node on a given scenario d out of $D+1$ total scenarios. In this data, if node i appears in scenario d then $O_i^d = 1$, otherwise $O_i^d = 0$. Then the general probability of occurrence of customer i is given by: $p_i = \frac{\sum_{d=1}^{D} O_i^d}{D}$. Let C^d be the set of customers that occur in a given scenario d. That is, customer $i \in C^d$ if $O_i^d = 1$. Then V^d is the set of all the nodes that occur in a given scenario d with $|V^d| = |C^d| + 1 + K$, including the depot at index 0 and K artificial nodes (identical copies of the depot). Also let set A^d be the indices of these artificial nodes: $|C^d| + 1 \ldots |C^d| + 1 + K$. In other words, $V^d = C^d \cup A^d \cup \{0\}$. The time to traverse the arc from node i to node j is given by t_{ij} and s_i^d is the service time of node i in scenario d. In particular $\forall d \in N_D$ the followings are true: $s_0^d = 0$; $O_0^d = 1$; $\forall i \in A^d$, $O_i^d = 1$ and $t_{i0} = t_{0i} = s_i^d = 0$; $\forall i \in A^d, \forall j \in C^d$, $t_{ij} = t_{0j}$ and $t_{ji} = t_{j0}$; and $\forall i, j \in A^d$, $i \neq j$, $t_{ij} = t_{ji} = 0$. a_i^d is the earliest start time and b_i^d is the latest start time to serve customer i in scenario d. The common due date for all vehicles is L. If customer i is selected to be served in the solution of scenario d, then z_i^d is one, otherwise z_i^d is zero. If the arc from node i to node j is selected in the solution of scenario d, then x_{ij}^d is one, otherwise x_{ij}^d is zero. The continuous variable y_i^d is the arrival time to node i in scenario d except the depot in which case it is the departure time, i.e. $y_0^d = 0$. In particular, y_i^d for $i \in A^d$ corresponds to the arrival time to the depot of a vehicle. Note that $y_i^d = 0$ for a customer i which is not served in the solution, i.e. when $z_i^d = 0$. The continuous variable e_i^d is the earliness penalty and l_i^d is the lateness penalty of customer i in scenario d; similarly $e_i^d = l_i^d = 0$ when $z_i^d = 0$. The binary c_{ij}^d variable is the additional

recourse variable associated with each scenario d (such that $d \neq 0$) and each arc (i, j) in the network to indicate the common arc (i, j) which is traversed both in the scenario d and in the a-priori route given by x_{ij}^0. M is a sufficiently large number which can be set to $M = 3L$.

$$\max \sum_{d \in N_D} \left(\alpha_1 \sum_{i \in C^d} z_i^d - \alpha_2 \sum_{i \in A^d} y_i^d - \alpha_3 \sum_{i \in C^d} l_i^d - \alpha_4 \sum_{i \in C^d} e_i^d \right) \tag{1.1}$$
$$+ \sum_{d \in N_D \setminus \{0\}} \alpha_5 \sum_{i \in V^d} \sum_{j \in V^d, \ j \neq i} c_{ij}^d t_{ij}$$

$$\text{s.t.} \quad \sum_{j \in V^d, \ i \neq j} x_{ji}^d = z_i^d \qquad\qquad i \in C^d, \ d \in N_D \tag{1.2}$$

$$\sum_{j \in V^d, \ j \neq i} x_{ij}^d = z_i^d \qquad\qquad i \in C^d, \ d \in N_D \tag{1.3}$$

$$\sum_{i \in V^d \setminus \{0\}} x_{0i}^d = K \qquad\qquad d \in N_D \tag{1.4}$$

$$\sum_{i \in V^d \setminus \{0\}} x_{i0}^d = K \qquad\qquad d \in N_D \tag{1.5}$$

$$x_{i0}^d = 1 \qquad\qquad i \in A^d, \ d \in N_D \tag{1.6}$$

$$y_i^d + t_{ij} + s_i^d \leq y_j^d + M(1 - x_{ij}^d) \qquad \begin{array}{l} i \in V^d, \ j \in V^d \setminus \{0\} \\ i \neq j, \ d \in N_D \end{array} \qquad (1.7) \quad (1)$$

$$a_i^d \leq y_i^d + e_i^d + M(1 - z_i^d) \qquad\qquad i \in C^d, \ d \in N_D \tag{1.8}$$

$$y_i^d \leq b_i^d + l_i^d \qquad\qquad i \in C^d, \ d \in N_D \tag{1.9}$$

$$y_i^d \leq L z_i^d \qquad\qquad i \in C^d, \ d \in N_D \tag{1.10}$$

$$x_{ij}^0 + x_{ij}^d \geq 2 c_{ij}^d \qquad\qquad i, j \in V^d, \ j \neq i, \ d \in N_D \setminus \{0\} \tag{1.11}$$

$$e_i^d \geq 0 \qquad\qquad i \in C^d, \ d \in N_D \tag{1.12}$$

$$l_i^d \geq 0 \qquad\qquad i \in C^d, \ d \in N_D \tag{1.13}$$

$$x_{ij}^d \in \{0, 1\} \qquad\qquad i, j \in V^d, \ i \neq j, \ d \in N_D \tag{1.14}$$

$$y_i^d \geq 0 \qquad\qquad i \in V^d, \ d \in N_D \tag{1.15}$$

$$z_i^d \in \{0, 1\} \qquad\qquad i \in C^d, \ d \in N_D \tag{1.16}$$

$$c_{ij}^d \in \{0, 1\} \qquad\qquad i, j \in V^d, \ i \neq j, \ d \in N_D \setminus \{0\} \tag{1.17}$$

The objective function (1.1) maximizes the number of customers served and minimizes the total time spent by the vehicles as well as the total earliness and lateness penalty, for each scenario including the master plan. In addition, the similarity of routes in each scenario with

the master plan is also maximized. The positive constants $\alpha_1 \ldots \alpha_5$ should be set according to the priority of these goals. In particular, α_1 should be adjusted with respect to α_3 and α_4 so that not visiting a customer and hence avoiding earliness or lateness penalty would not be desirable. Note that in many cases $\alpha_4 = 0$ since there is no explicit penalty for waiting time but it indirectly increases the total time spent. Constraints 1.2-1.5 are the routing constraints. Constraint 1.6 forces an artificial node to be served at the end of the route to keep track of the time spent by the vehicles. Constraint 1.7 is the subtour elimination constraint. Additionally, it enforces the arrival time constraint when a customer j is served right after customer i in scenario d. Otherwise, it is redundant. Constraint 1.8 imposes the earliness penalty, which is redundant if customer i is not served in the solution of scenario d. Constraint 1.9 imposes the lateness penalty and constraint 1.10 the common due date of the vehicles. Constraint 1.11 ties adjusted variables to a-priori variables to reward common arcs which are selected both in scenario d and the master plan (scenario 0). Lastly, constraints 1.12-1.17 are bounds on the variables.

In Problem 1, for each scenario except the master plan the uncertain service times and probabilistic customers are formulated through scenario-based stochastic optimization based on historical data. For the master plan (scenario 0), we need to provide the set of customers C^0 and the value of uncertain service time s_i^0 for each customer i. For the former, one idea is to select some of the customers with high probability of occurrence p_i. For the latter, we use robust optimization to construct worst-case service time values for the master plan. The methodology and the proofs directly follow from Sungur (2007) and Sungur et al. (2006), and they are not presented here for space considerations. Note that for each customer i, there is a total of $\sum_{d=1}^{D} O_i^d$ realizations of service time. Accordingly, let N_i be the set of these scenario indices with positive occurrence data for customer i: $d \in N_i$ if $O_i^d = 1$. We define s_i^0 of customer i as the mean value of its service time realizations as in $s_i^0 = \dfrac{\sum_{d \in N_i} s_i^d}{|N_i|}$ and we consider these realizations, s_i^d, as possible deviations from this nominal service time value s_i^0. We use these deviations to construct the following three scenario-based uncertainty sets to formulate the robust counterpart of constraint 1.7 for scenario 0: convex hull, ellipsoidal,

and box. Similar to what is shown in Sungur (2007) and Sungur et al. (2006) to derive robust counterparts of VRP with demand uncertainty, our derivations result only in modifications of the problem parameters (i.e. service times) and therefore the resulting robust counterpart is just an instance of the deterministic counterpart with modified service time data.

Proposition 1. *The robust counterpart is obtained by replacing constraint 1.7 in Problem 1 for scenario 0 with constraint 1.18 for convex hull, with constraint 1.19 for ellipsoidal, and with constraint 1.20 for box.*

$$y_i^0 + t_{ij} + s_i^0 + \max\{\max_{d \in N_i}(s_i^d - s_i^0), 0\} \leq y_j^0 + M(1 - x_{ij}^0) \quad i \in V^d,\ j \in V^d \setminus \{0\},\ i \neq j \quad (1.18)$$

$$y_i^0 + t_{ij} + s_i^0 + \sqrt{\sum_{d \in N_i}(s_i^d - s_i^0)^2} \leq y_j^0 + M(1 - x_{ij}^0) \qquad i \in V^d,\ j \in V^d \setminus \{0\},\ i \neq j \quad (1.19)$$

$$y_i^0 + t_{ij} + s_i^0 + \sum_{d \in N_i}|s_i^d - s_i^0| \leq y_j^0 + M(1 - x_{ij}^0) \qquad i \in V^d,\ j \in V^d \setminus \{0\},\ i \neq j \quad (1.20)$$

8   MADS Heuristic

To address the challenge of solving the large scale real-life problem, we develop a heuristic solution procedure based on insertion. Insertion has been proved to be a very efficient, simple, and therefore widely used heuristic for large scale VRP instances with time windows. In this section, we propose a problem specific solution procedure for our CDP. The problem formulation is given in Problem 1 together with related adjustments in Proposition 1. In our two-phase heuristic Master And Daily Scheduler (MADS), we develop concurrently a robust master plan and daily schedules for historical scenarios which are constructed based on the master plan according to a hybrid recourse action of partial rescheduling of routes. This recourse action combines the two recourse actions of omitting customers and rescheduling routes. Thus, MADS uses historical data modeled as scenarios and generates a robust master plan which is in return to be used for future outcomes during the planning horizon to generating daily schedules.

The first phase is the construction of initial solutions. First of all, a robust master plan is constructed by making insertions of customers to initially empty routes by minimizing the cost of insertion greedily. That is, among all the feasible insertions with respect to the common due date the cheapest one is done at each step. The cost of an insertion of a customer to a location in a vehicle is determined by the increase in the time spent and the increase in the total penalty of all the customers served by that vehicle. Note that the cost of insertion is dependent on the location of the inserted customer on the route. Second of all, the routes of the master plan are updated by each scenario following the two recourse actions to construct the daily schedules. First, a scenario modifies the initial robust routes of the master plan by following the recourse action of omitting customers, which is to follow the same sequence of customers in the robust routes by omitting (not visiting) the customers that do not occur in that particular scenario. Then the usual greedy insertion is used in rescheduling of these modified routes to insert brand new customers in that particular day data that do not occur in the robust routes.

The second phase is iterative. At each iteration, first, scenarios give feedback to the master plan about the customers that could not be scheduled in their daily routes; second, based on this feedback the master plan prioritorizes these unscheduled customers by the number of scenarios that they appear but could not be scheduled. Then the master plan updates its routes by performing feasible maximum-priority insertions in the cheapest way. Note that the selection of customers is based on the priority not on the cost of insertion. However, once a customer is selected then the cheapest possible insertion is done for this particular customer. Then the new master plan is re-dispatched to the scenarios which construct their daily schedules with respect to the hybrid recourse action as before. At the end of each iteration, the objective function is evaluated over all the scenarios based on the measures in Equation 1.1. The iterations of the second phase and thus the overall two-phase algorithm stop when there is no improvement in the objective, reaching a local optimum.

Note that for the master plan, the first phase is cost driven whereas the second phase is priority driven. For the scenarios, both phases are cost driven based on the current master

13

plan. The maximization of the number of customers served is mainly due to the recourse action of rescheduling of the routes; the minimization of the time spent and total penalty is mainly due to the greedy insertions; and maximization of the similarity of routes is due to the recourse action of omitting customers and to the feedback procedure between the master plan and daily schedules to prioritize the customers in the iterative second phase.

Between the two phases of the algorithm, we also implement the idea of buffer capacity. For the first phase, we reserve a *buffer capacity* by decreasing the common due date of vehicles. This slack time is later used in the second stage to schedule additional customers. This parameter of the algorithm is actually a tool to balance the cost driven stage and the priority driven stage in constructing the master plan which indirectly effects the daily schedules as well. We later explore experimentally the effect of this parameter on the aspects of cost and similarity of routes.

Another parameter of the algorithm is the set of customers to be considered in the master plan during the first phase. Given that in our real-life instance the total number of potential customers is several thousands, attempting to schedule all the customers is neither realistic nor feasible. Instead, a subset of customers should be selected. One intuitive idea is to select some of the customers which appear the most in the scenarios (customers with a high probability of occurrence). This should increase the similarity of routes since the master plan will be composed of mainly the customers which will not be omitted in most of the daily schedules. A reasonable and realistic number to pick for the cardinality of the set of customers with high probability of occurrence could be a value around the average number of customers per day. We also explore later the effect of this parameter on the quality of the solution.

The last parameters of the algorithm are the weights used to balance the cost of increase in time spent and the cost of increase in total penalty during the greedy insertion. Changing these weights in the selection of the cheapest insertion will alter the outcome by emphasizing more on minimizing either the time spent or the total penalty.

Algorithm 1 Insertion Routine
___
Require: Initial routes

  Calculate insertion cost of possible insertion locations in the routes for non-scheduled customers

  repeat

    Pick the cheapest insertion (a feasible insertion of a non-scheduled customer which will result in the least increase of the total time spent by the vehicles and the total penalty) and insert it into the vehicle at the proper location

    Update insertion cost of possible insertion locations of that vehicle

  until All the occurring customers are inserted or no feasible insertion is possible

  return The resulting routes
___

Algorithm 2 Phase One
___
Require: Distance matrix of the network, scenario data for each day (customers appearing in that day together with earliest start, latest start, and service time data), master data (customers to be scheduled in the master plan during Phase One), percent buffer capacity

  Reserve a buffer capacity in the length of the routes

  Setup the master plan data (by calculating worst-case values of service times)

  Call Insertion Routine for the master plan with initially empty routes

  for Each scenario d **do**

    Take the master routes as initial solution and drop the non-occurring customers in the data of scenario d

    Call Insertion Routine for scenario d with these initial routes

  end for

  Calculate the objective value and save the current solution

  return The current solution
___

Algorithm 3 Phase Two
___
Require: Solution of Phase One as the initial solution

Remove reserved buffer capacity in the length of the routes from consideration

repeat

    Calculate priorities of customers which could not be inserted in the scenarios and create a sorted list with non-increasing priorities

    if The priority list is empty **then**

        return The current solution

    end if

    for Each customer i in the priority list **do**

        if Customer i can be feasibly inserted **then**

            Make the least cost insertion of customer i in the master schedule

        end if

    end for

    if No feasible insertion was possible **then**

        return The current solution

    end if

    for Each scenario d **do**

        Take the master routes as initial solution and drop the non-occurring customers in the data of scenario d

        Call Insertion Routine for scenario d with these initial routes

    end for

    Calculate the objective value

    if The objective value is improved **then**

        Save the current solution

    end if

until The objective value is not improved

return The current solution
___

The description of the MADS is given in Algorithm 1-3. Note that as a result of the second phase (and of MADS) the master plan is generated based on historical scenarios. Then the master plan can be used during the planning horizon for each future scenario in the following way: first non-occurring customers in the data of a particular scenario are dropped and then Algorithm 1 is executed with these initial routes in order to generate the daily schedule. This is repeated in each day during the planning horizon to generate daily schedules with respect to the master plan. At the end of each planning horizon, based on new observations of that particular planning horizon the MADS algorithm can be used in generating a new master plan for the next planning horizon.

9   Experimental Analysis

In this section, we first present the real-life problem data of the CDP in industry which is faced by a worldwide delivery company. We discuss both the characteristics of the data and investigate the nature of the uncertainty. We also analyze the sensitivity of the two-phase MADS algorithm to the problem data, effect of changing the service time distribution and effect of changing the probability distribution of the occurrences of customers, as well as to its parameters, effect of changing the set of customers to be scheduled in the master plan during the first phase and effect of changing the buffer capacity.

Throughout the experimental analysis, we assume that a past data for two planning horizons are available. We use the past data for the first planning horizon to *train* the master plan and we use the remaining past data to *evaluate* the performance by treating it as future outcomes. Also, all the experiments are performed on a Dell Precision 670 computer with a 3.2 GHz Intel Xeon Processor and 2 GB RAM running Red Hat Linux 9.0 and all the solutions could be obtained within one hour of CPU time.

We compare the quality of the solution of the MADS algorithm with a solution obtained by independently scheduling scenarios without a master plan which we refer as the independent daily insertion (IDI) algorithm. This solution is obtained by treating each scenario as

an independent CDP with no scenarios. This means that each scenario in the original problem is considered as a "master plan" and only Algorithm 1 is executed with initially empty routes and without reserving a buffer capacity. Therefore the IDI algorithm does not make any considerations in increasing the similarity of the resulting independent daily schedules but maximizes the number of customers served while minimizing total time spent and total penalty for each day. Lastly, we also evaluate the solution of the MADS algorithm on the real problem data and show that our solution can be better than the current practice.

9.1 Problem Data and Uncertainty

The CDP application that we study has 3715 potential customers in an urban area of northwest California. The locations of the customers are known and most of them can be enclosed in a box of 25 miles^2. Therefore, travel distances are not significant. At the beginning of each day, any of these potential customers can put a delivery request with an uncertain service time. The planning horizon of the service provider company is two weeks (14 days). The data of delivery requests is collected for a horizon of 28 days which is equivalent to two planning horizons. The average number of occurring customers per day is 472. There are a total of 4 couriers to serve the customers each day. The couriers are ready to leave the depot at 9am and they must return to the depot by 8pm. The travel time is considered deterministic and to convert the distance measures to time units, it is assumed that the couriers travel with an average speed of 35 mph in the city.

We adapt the CDP formulation given by Problem 1 and we consider each day as a scenario in formulating the uncertainty. For service time uncertainty in the robust counterpart of constraint 1.7, we use the convex hull uncertainty set by replacing constraint 1.7 when $d = 0$ with constraint 1.18. We also set $\alpha_4 = 0$ since our application does not consider an explicit earliness penalty. For the other four goals, we prioritorize α_1 as the highest, we weight α_2 and α_3 equally but with less priority than α_1, and we give α_5 the smallest priority.

In addition, we also analyzed the service time characteristics of all the potential customers in general. As it is common in the routing literature and in industry, service times

follow a lognormal distribution, see Dessouky et al. (1999). In our particular case, we use a transformation by shifting and scaling a lognormal distribution with mean $\mu_s = 0.0953$ and standard deviation $\sigma_s = 0.25$. The mean and the standard deviation of the actual sample data are respectively about 4 and 7 minutes.

Lastly, we analyzed the probability distribution of the occurrence of all the potential customers in general. In Figure 1, P2 is the actual discrete distribution and P4 is its continuous approximation which is a shifted power function. Without loss of generality, customers are represented by non-increasing IDs with respect to their probability of occurrence. Note that there are customers with probability 1 (i.e. occurring in all of the scenarios). In addition, we present two more discrete distributions, P1 and P3, which are generated by modifying P4. In P1, the probabilities of occurrence are decreased with respect to original P2; and in P3, they are increased. These three distributions, P1, P2 and P3, are used in our experiments to sample scenarios.
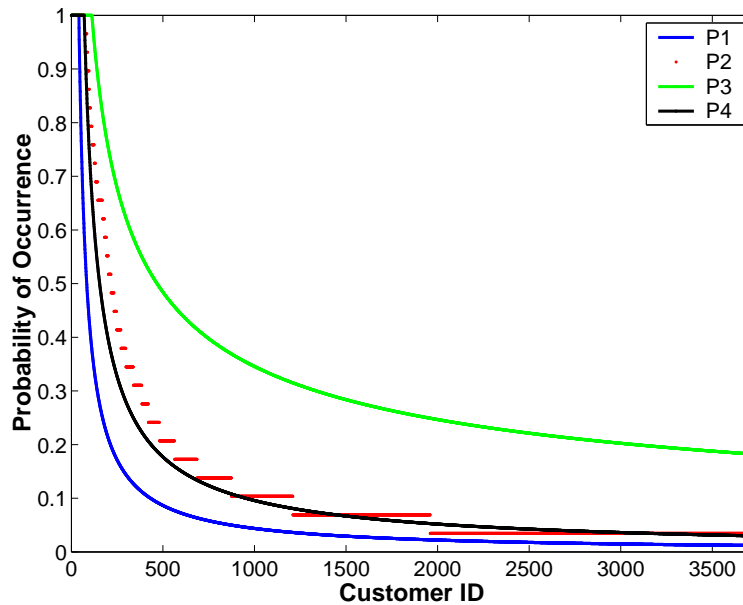

Figure 1: Actual and modified probability distributions of the occurrence of customers

9.2 MADS versus Independent Daily Insertions

In this first set of experiments, we explore the effect of changing problem data on the two-phase MADS algorithm and we compare the quality of the solutions with the independent daily insertion (IDI) algorithm, leaving the effect of changing the parameters of our heuristic to the next set of experiments. Therefore for the current experiments we need to fix the value of these parameters of the MADS algorithm, the buffer capacity and the set of customers to be scheduled in the master plan during phase one. For the latter, we refer to the distribution P2 in Figure 1 and we define a cut to select the customers with high probability of occurrence. In order to keep the resulting number of total customers around the average number of customer per day, 472, we fix the value of this cut to 0.25. That is, only the customers which have a higher probability of occurrence than 0.25 are considered in scheduling the master plan during phase one. This total number of customers turns out to be 422 for our application. For the former problem parameter, we set the buffer capacity to 50%, meaning that only half of the total allowed time for the couriers is considered during the first phase. This way, we weight the cost driven first phase and priority driven second phase equally.

For the uncertainty in service time and probabilistic customers, we consider a base case with respect to our fitted lognormal distribution and the probability distribution P2 in Figure 1. For each problem instance, we sample data of scenarios for the planning horizons with respect to these two distributions. First, the occurrence data of each scenario is generated by randomly selecting customers according to P2 until 472 customers are selected in each scenario. Then each customer is assigned a random service time following the lognormal distribution. Recall that the time windows and travel times are deterministic. Thus, all the required data for a problem instance is generated by this process. Also, recall that for the MADS algorithm, only the first half of the total data is used to generate the master plan for a planning horizon and the remaining half is used to evaluate its performance; whereas for the IDI algorithm, only the second half is used as *future* outcomes.

We generate additional cases by deviating from the base case in two ways. First, we change the standard deviation σ_s of the lognormal distribution to see the effect of increased

service times, with $\sigma_s = 0.500$, and decreased service times, with $\sigma_s = 0.125$. Second, instead of P2 we sample customers from P1 and P3 in Figure 1. The effect of sampling from P1 with respect to P2 is the following: in each scenario most of the customers are selected among the customers with high probability of occurrence since the others have a significantly less chance to occur. Note that since the customers considered in the master plan during the first phase of the MADS algorithm also have high probability of occurrence, the similarity in routes of this heuristic solution should improve on such instances. The effect of sampling from P3 is simply the reverse: this time the customers which have relatively smaller probability of occurrence have in fact a higher chance to occur in each scenario with respect to P2. Thus, in the solution of MADS, the similarity measure is expected to be low for such instances.

For each case, we generate 30 random problem instances and report the averages of the solutions. When moving from one case to another we modify only one parameter at a time keeping the rest of the problem instance the same, which allows observing the sole effect of changing this particular parameter. Table 1 provides these experimental results. The left part is the input parameters and the right part is the output measures. The headings are as follows: "Alg" is the algorithm used in generating the solution; "Prob" is the probability distribution in Figure 1 used to sample customers; "Std" is the value of the standard deviation of lognormal distribution of service times, σ_s; "NS" is the total number of customers which could not be served in the daily schedules, "Time" is the total time spent by the couriers in daily schedules which is composed of travel, waiting, and service times; "Penalty" is the total lateness penalty in daily schedules; and "Arcs" is the total length of common arcs in daily schedules which is the similarity measure. Since we compare the quality of the solutions with the independent daily insertion algorithm which does not generate a master plan, we slightly modify the similarity measure "Arcs" to make the solutions comparable. Instead of calculating the similarity of a scenario with respect to the master plan, we re-calculate this number based on the average similarity of scenarios with each other and report to total of these averages.

The general observations based on Table 1 suggest that IDI covers customers with smaller

Table 1: Comparison of MADS with IDI

Alg	Prob	Std	NS	Time	Penalty	Arcs
MADS	P1	0.125	0.0	96817.1	65.7	461.8
MADS	P1	0.250	0.0	99676.4	99.5	363.4
MADS	P1	0.500	18.8	109592.6	646.2	267.4
MADS	P2	0.125	0.0	96827.8	86.7	226.2
MADS	P2	0.250	0.0	99798.9	109.7	215.4
MADS	P2	0.500	18.4	109593.5	638.7	160.4
MADS	P3	0.125	0.0	97168.8	147.1	46.1
MADS	P3	0.250	0.0	99913.2	228.4	43.3
MADS	P3	0.500	15.9	109288.0	571.1	29.5
IDI	P1	0.125	0.0	93392.3	5159.4	17.0
IDI	P1	0.250	0.0	98436.3	2919.4	17.5
IDI	P1	0.500	2.9	107644.5	154.5	16.4
IDI	P2	0.125	0.0	93365.5	4228.0	12.6
IDI	P2	0.250	0.0	98335.7	2365.4	12.9
IDI	P2	0.500	2.9	107621.9	189.0	11.5
IDI	P3	0.125	0.0	93381.5	2079.3	4.8
IDI	P3	0.250	0.0	97983.9	1247.0	5.7
IDI	P3	0.500	3.0	107425.9	133.5	4.2

time spent but with an increased lateness penalty and decreased similarity of routes compared to MADS. Therefore, generally speaking when MADS can cover all the customers, it improves the similarity of routes and lateness penalty at the expense of increased time spent; the only exceptions are the cases with $\sigma_s = 0.500$. The improvement on similarity is expected due to the feedback process in MADS between the master plan and the scenarios. However, the improvement on the lateness penalty is mainly due to the buffer capacity. During the greedy insertion routine, most of the lateness penalty is incurred at the latest insertions in the routes since the cheapest ones are performed in the beginning delaying the costly ones whose penalty increases even more due to previous cheaper insertions. Reserving a buffer capacity prevents the insertion routine from incurring these high penalties during phase one; and during the iterative phase two since some of the insertions are done based on priorities, again the greedy nature of the insertion routine is avoided up to some extend. Another general comment is that high values of σ_s result in customers which could not be served in

the solution. In such cases, IDI performs better in covering customers than MADS because there is no effort done in creating common arcs in scenarios, which provides flexible schedules to serve more customers.

When we analyze Table 1 in particular for MADS, we see that in general increasing the probability of occurrence for a given standard deviation increases time spent and lateness penalty, and decreases similarity, making all these measures worse. The result is intuitive and expected since sampling from a wider likely range of customers results in diverse scenarios, which increases the cost aspects and decreases the similarity. On the other hand, increasing the standard deviation for a given probability of occurrence has also the same effect on these measures as before with the addition that high values result in unserved customers. This is also expected because increased and dispersed service times make the problem worse with respect to all measures. As a general conclusion, for MADS it is desired to have a combination of minimum values of σ_s and "Prob" in the problem data.

Lastly, when we analyze Table 1 in particular for IDI, the results are more obscure and there is no overall general trend. However, increasing the probability of occurrence for a given standard deviation decreases lateness penalty; and increasing the standard deviation for a given probability of occurrence increases time spent and decreases lateness penalty with the addition that high values result in unserved customers. This decrease in lateness penalty with respect to increased standard deviation is mainly due to the load of customers in the vehicles. High standard deviation results in high service times yielding more balanced routes because of the tie-breaking and the nature of the cheapest insertion during the IDI algorithm. However, the effect of low standard deviation is just the opposite. This is intuitive since under high standard deviation, inserting several customers with high service times in one vehicle is not feasible and therefore these customers are evenly distributed to all the vehicles. On the other hand, in case of low standard deviation, several customers with small service times can be fit in a single vehicle leaving only few other customers with small service times to be inserted into the other vehicles. As a result, when the vehicles are balanced in case of high standard deviation, it is easier to meet the deadlines (and hence to incur smaller

lateness penalty) at the expense of increased time spent compared to the case of low standard deviation having an uneven balance of customers in the vehicles.

9.3   MADS versus Real-life Solution

In this second set of experiments, we both explore the effect of changing parameters of the MADS algorithm and compare the quality of its solution with the current practice. For the percent buffer capacity, we explore the following range: 0, 20, 40, 60, 80, 100. Note that 0% buffer capacity bypasses the first phase of the heuristic by making insertions in the master plan only based on priorities; and 100% is the reverse making all the insertions based on costs. For the cut of the probability distribution P2 in Figure 1, we use 0.25 as the base case resulting in 480 customers and consider the following deviations: 0.20 resulting in 563 customers and 0.30 resulting in 384 customers. As before, the first half of the real-life data is treated as past realizations for the MADS algorithm and the second half as future outcomes to evaluate and compare the solutions.

Table 2 shows the solutions on the real-life data instance for different parameter settings of our heuristic. The new headings are "Cut" for the cut of the probability distribution and "BF" for the percent buffer capacity. Also, "NA" stands for not applicable. In this particular real-life data instance all the customers could be feasibly served in each day of the planning horizon. In addition, we provide the solution obtained by the IDI algorithm and the real-life solution, as well as average results of each cut over the 6 buffer capacity values for the MADS algorithm indicated by "Avg" under the "BF" column. Also we use as before the slightly modified version of the similarity measure "Arcs" to make the solutions comparable.

When we look at the individual solutions, we see that both the real-life solution and MADS are able to improve similarity at the expense of increased time spent and lateness penalty with respect to IDI. When we compare the quality of the solution by MADS with the current practice, we see that in general MADS provides a better improvement in similarity and time spent with a possibly higher lateness penalty. However, we observe that in a total

Table 2: Comparison of MADS with Real-life solution

Alg	Cut	BF	NS	Time	Penalty	Arcs
Real-life	NA	NA	0	109009.8	5091.8	334.2
IDI	NA	NA	0	98749.0	905.2	37.6
MADS	0.20	0	0	101687.8	3279.1	352.3
MADS	0.20	20	0	101895.6	7333.8	361.0
MADS	0.20	40	0	101640.8	4927.4	371.3
MADS	0.20	60	0	102106.4	3887.1	911.5
MADS	0.20	80	0	102446.7	5134.6	965.4
MADS	0.20	100	0	102373.7	8218.8	1093.0
MADS	0.20	Avg	0	102025.2	5463.5	675.8
MADS	0.25	0	0	101894.9	4530.7	485.7
MADS	0.25	20	0	102022.5	13208.1	749.3
MADS	0.25	40	0	102045.0	5377.5	720.0
MADS	0.25	60	0	101872.6	2920.5	870.6
MADS	0.25	80	0	101828.7	4705.0	634.2
MADS	0.25	100	0	102373.7	8218.8	1093.0
MADS	0.25	Avg	0	102006.2	6493.4	758.8
MADS	0.30	0	0	101779.4	2414.9	675.1
MADS	0.30	20	0	102075.3	6197.4	600.7
MADS	0.30	40	0	102208.8	5039.8	694.8
MADS	0.30	60	0	101882.4	3357.0	872.3
MADS	0.30	80	0	102014.0	3604.5	905.5
MADS	0.30	100	0	102373.7	8218.8	1093.0
MADS	0.30	Avg	0	102055.6	4805.4	806.9

of 10 cases out of 18, MADS outperforms the real-life solution in all of the measures. This suggest that our heuristic can be tuned to provide improvements over the current practice.

When we look at the average results for MADS, we see that as the cut increases the similarity measure is improved whereas the effect on time spent and lateness penalty is not clear. The reason in the increase of similarity is due to the fact that the customers with relatively less probability of occurrence are eliminated by increasing the cut, resulting in more common arcs in the master plan with the scenarios. When it comes to the buffer capacity, there is no clear trend on time spent and lateness penalty. However for a given value of cut, increasing buffer capacity clearly increases similarity of routes since this would put more emphasize on the priority driven second phase of the heuristic improving the similarity measure. Lastly note that for a given value of cut, 0% buffer capacity behaves just like IDI resulting in the smallest time spent and lateness penalty with the worst similarity of routes in general; on the other hand, 100% is just the opposite, resulting in the highest similarity of routes with the worst time spent and lateness penalty in general. Also for the latter, the solution is independent from the value of cut since the first phase of the MADS algorithm is bypassed.

10   Conclusions and Recommendations

In this study, we considered a real-life Courier Delivery Problem (CDP), a variant of VRPTW, for which we proposed a recourse model for the robust counterpart of the problem formulation and we developed an efficient two-phase heuristic based on insertion. We addressed the uncertainty in service times by using robust optimization and the probabilistic nature of the customers by using scenario-based stochastic optimization with recourse. Thus, we benefited from the simplicity of the robust model and the flexibility of recourse actions.

We first adapted a nominal VRPTW model for the CDP. We then defined a problem specific recourse action of partial rescheduling of routes which is a hybrid of two known recourse actions in the literature: omitting non-occurring customers and complete rescheduling of

routes. We showed how to develop a recourse model to provide a master plan for the planning horizon as well as daily schedules for scenarios which are modified versions of the master plan based on the recourse action. We also showed how to incorporate robust optimization for the service time uncertainty in the master plan. For each scenario, our overall CDP formulation maximizes the number of customer served while minimizing time spent by the couriers and total penalty by increasing similarity of routes with the robust routes of the master plan. For this model with hybrid recourse action of partial rescheduling of routes, we eventually developed a two-phase heuristic, MADS, using insertion as a subroutine to solve efficiently the large scale real-life problem.

We explored experimentally the sensitivity of our heuristic to uncertain problem parameters as well as to some control parameters. We also compared the quality of the solution with an independent daily insertion algorithm which does not make efforts in providing a master plan and thus in increasing similarity of routes. We observed that the MADS heuristic improves in general the similarity measure at the expense of increased time spent and the lateness penalty. We also compared the solution of our heuristic with a real-life problem solution. We showed that by tuning the parameters of our heuristic, it is possible to outperform the current practice in all of the measures.

In this study, we showed how to generate a robust master plan which would increase the similarity of daily schedules. However using different approaches in generating the master plan is still an open research direction. For instance, using an expected value approach for the uncertainty in master plan may not be a viable option in terms of similarity but other probabilistic approaches can be used such as deviating from the expected value by a fraction of the standard deviation of the uncertain parameter instead of considering the worst case value of robust optimization. Similarly, the selection of customers to be served by the master plan is subject to different methods than picking the most likely ones.

11   Implementation

The model and algorithm proposed here work on real-world data obtained from a courier industry company. Therefore the findings of this research can be applied immediately to study different routing strategies for problems in the courier industry. This will require collecting data and extracting input parameters from this data for the model and MADS algorithm. Then, the solution of the algorithm will provide a routing strategy which takes into account customer coverage, total time spent, lateness penalty, and similarity among daily routes.

As part of this project we presented this report to the courier company that supplied the real world data to help the process of implementing this research. Although impressed with the results, the company personel suggested additional issues that have not yet been taken into account. For instance, the routing solutions should also consider the fuel consumption in deciding the routes; and although familiarity with the routes is key, the order in which these are traversed is also important as this will influence service time and provide the basis of how the truck should be loaded. This research should take into account these additional features before it can be implemented by this courier delivery company.

References

Baldacci, R., E. Hadjiconstantinou, and A. Mingozzi (2004). An exact algorithm for the capacitated vehicle routing problem based on a two-commodity network flow formulation. *Operations Research 52*(5), 723–738.

Beasley, J. E. and N. Christofides (1997). Vehicle routing with a sparse feasibility graph. *European Journal of Operational Research* (98), 499–511.

Ben-Tal, A., A. Goryashko, E. Guslitzer, and A. Nemirovski (2003). Adjustable robust solutions of uncertain linear programs. Technical Report, Minerva Optimization Center, Technion. http://iew3.technion.ac.il/Labs/Opt/index.php?4.

Ben-Tal, A. and A. Nemirovski (1998). Robust convex optimization. *Mathematics of Operations Research 23*(4), 769–805.

Ben-Tal, A. and A. Nemirovski (1999). Robust solutions to uncertain programs. *Operations Research Letters 25*, 1–13.

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., P. Jaillet, and A. R. Odoni (1990). A priori optimization. *Operations Research 38*(6), 1019–1033.

Bertsimas, D. and M. Sim (2003). Robust discrete optimization and network flows. *Mathematical Programming 98*(1-3), 49–71.

Bertsimas, D. and M. Sim (2004). The price of robustness. *Operations Research 52*(1), 35–53.

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

Cordeau, J. F., M. Gendreau, G. Laporte, J. Y. Potvin, and F. Semet (2002). A guide to vehicle routing heuristics. *Journal of the Operational Research Society 53*(5), 512–522.

Dantzig, G. B. and J. H. Ramser (1959). The truck dispatching problem. *Management Science 6*, 80–91.

Dessouky, M., R. Hall, A. Nowroozi, and K. Mourikas (1999). Bus dispatching at timed transfer transit stations using bus tracking technology. *Transportation Research Part C 7*, 187–208.

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).

El-Ghaoui, L., F. Oustry, and H. Lebret (1998). Robust solutions to uncertain semidefinite programs. *SIAM J. Optimization 9*(1).

Erera, A. (2000). *Design of Large-scale Logistics Systems for Uncertain Environments*. Ph. D. thesis, University of California, Berkeley.

Fisher, M. (1995). Vehicle routing. In M. Ball, T. Magnanti, C. Monma, and G. Nemhauser (Eds.), *Network Routing*, Volume 8 of *Handbooks in Operations Research and Management Science*, pp. 1–33. Elsevier Science, Amsterdam.

Fukasawa, R., H. Longo, J. Lysgaard, M. P. de Aragão, M. Reis, E. Uchoa, and F. R. Werneck (2006). Robust branch-and-cut-and-price for the capacitated vehicle routing problem. *Mathematical Programming 106*, 491–51.

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

Gendreau, M., G. Laporte, and R. Seguin (1996). Stochastic vehicle routing. *European*

Journal of Operational Research 88(1), 3–12.

Goldfarb, D. and G. Iyengar (2003). Robust portfolio selection problems. *Mathematics of Operations Research 28*(1), 1–38.

Hall, R. W. (1996). Pickup and delivery systems for overnight carriers. *Transportation Research A* (30), 173–187.

Haughton, M. (2000). Quantifying the benefits of route reoptimization under stochastic customer demands. *Journal of Operational Research Society* (51), 320–322.

Haughton, M. and A. Stenger (1998). Modeling the customer service performance of fixed-rotes delivery systems under stochastic demands. *Journal of Business Logistics* (9), 155–172.

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

Jaillet, P. (1988). A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Operations Research 36*(6), 929–936.

Jaillet, P. and A. Odoni (1988). The probabilistic vehicle routing problem. In B. L. Golden and A. A. Assad (Eds.), *Vehicle Routing: Methods and Studies.* Elsevier, North-Holland, Amsterdam.

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

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 Ssystems 37*, 51–63.

Lambert, V., G. Laporte, and F. Louveaux (1993). Designing collection routes through bank branches. *Computers & Operations Research 20*(7), 783–791.

Laporte, G., M. Gendreau, J. Potvin, and F. Semet (2000). Classical and modern heuristics for the vehicle routing problem. *International Transactions in Operations Research 7*, 285–300.

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 I. Osman (1995). Routing problems: a bibliography. *Annals of Operations Research 61*, 227–262.

Lenstra, J. K. and A. H. G. Rinnooy Kan (1981). Complexity of vehicle routing and scheduling problems. *Networks 11*, 221–227.

Lu, Q. and M. M. Dessouky (2006). A new insertion-based construction heuristic for solving the pickup and delivery problem with hard time windows. *European Journal of Operational Research 175*, 672–687.

Lysgaard, J., N. A. Letchford, and W. R. Eglese (2004). A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming 100*, 423–445.

Morales, J. C. (2006). *Planning Robust Freight Transportation Operations*. Ph. D. thesis, Gerogia Institute of Technology.

Quadrifoglio, L. and M. M. Dessouky (2007). An insertion heuristic for scheduling mobility allowance shuttle transit (mast) services. *Journal of Scheduling 10*, 25–40.

Ralphs, T., L. Kopman, W. Pulleyblank, and L. Trotter (2003). On the capacitated vehicle routing problem. *Mathematical Programming 94*(2-3), 343–359.

Savelsbergh, M. W. P. and M. Goetschalckx (1995). A comparison of the efficiency of fixed versus variable vehicle routes. *Journal of Business Logistics* (19), 163–187.

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

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research 35*, 254–265.

Sungur, I. (2007). *The Robust Vehicle Routing Problem.* Ph. D. thesis, University of Southern California.

Sungur, I., F. Ordóñez, and M. M. Dessouky (2006). A robust optimization approach for the capacitated vehicle routing problem with demand uncertainty. Working paper 2006-06, Industrial and Systems Engineering, USC.

Toth, P. and D. Vigo (2002). Models, relaxations and exact approaches for the capacitated vehicle routing problem. *Discrete Applied Mathematics 123*, 487–512.

Waters, C. D. J. (1989). Vehicle-scheduling problems with uncertainty and omitted customers. *The Journal of the Operational Society 40*(12), 1099–1108.

Zhong, H., R. W. Hall, and M. M. Dessouky (2007). Territory planning and vehicle dispatching. *Transportation Science* (41), 74–89.