Tactical planning vehicle routing with call-in customers

Anirudh Subramanyam
Carnegie Mellon University, Pittsburgh, PA, USA

Frank Mufalli
Praxair Inc.

Jose M. Pinto
Praxair Inc.

Chrysanthos E. Gounaris
Carnegie Mellon University, Pittsburgh, PA, USA (gounaris@cmu.edu)

Abstract

During the tactical planning of multi-period vehicle routing operations, dispatchers must incorporate enough flexibility in their routing plans to accommodate potential customers who have not yet called in to place their order. We develop an Adaptive Robust Optimization framework to address this challenging problem and solve instances to guaranteed optimality.

Keywords: Vehicle Routing, Uncertainty, Robust Optimization

Motivation and Literature Review

Several transportation problems involve the tactical planning and routing of vehicles over a short-term horizon (typically several days). This is the case when customer requests are received dynamically over the planning horizon. Each request specifies a demand quantity and a set of days during which service can take place. The distributor must assign a visit day to each customer and plan the routes of each day so as to minimize the total routing cost. This is in contrast to typical vehicle routing problems (VRPs), where the planning horizon consists of one period (or one day) and all customer requests must be serviced within this period. The tactical routing plan is typically implemented in a rolling horizon fashion: routes of the first time period of the horizon are executed while new customer requests are recorded; unfulfilled requests at the end of the first day and new orders accumulated during the day constitute the new portfolio of requests to be considered for scheduling the following day. This problem has been previously referred to as the Tactical Planning VRP (Baldacci et al., 2011) or the Dynamic Multi-period VRP (Wen et al., 2010); it is a special case of the well-known Periodic VRP, where all customers must be visited exactly once throughout the planning horizon.

Such distribution problems are typical in systems in which logistics and other services are provided by appointment. Relevant examples include companies that provide on-site maintenance services and distribution companies that deliver high-end or custom-built products. In such applications, contracts might be established with costs determined by time-to-delivery or penalties associated with not visiting a customer on the first day of its requested window. Applications also arise in the food distribution sector (Wen et al., 2010), urban logistics (Archetti et al., 2014), automotive transportation (Cordeau et al., 2015) and in vendor-managed inventory systems (Coelho et al., 2012) .

In all of the existing applications, no knowledge is available about customers requesting service in future time periods or, in other words, complete uncertainty is assumed about future requests. Decisions are made using only information on service requests that have already been placed. While this hypothesis might make sense in some applications, this is usually not the case as companies often have large amounts of historical data which can be used to obtain forecasts that provide information regarding calls for service in future time periods. It is possible to take advantage of this information to generate better risk-averse solutions. Indeed, not accounting for this information while planning can create situations which can either be (i) infeasible because it could require too many vehicles on some day of the period or (ii) too expensive in terms of routing costs.

In this regard, the work by Albareda-Sambola et al. (2014) considers uncertainty in future customer requests. They assume a probabilistic description where at any time period they know the probability of demand in subsequent time periods of every possible customer that does not have a pending request. They then consider an adaptive policy by solving at each time period a VRP subproblem to identify the set of pending customers that will be serviced in the current time period along with the actual routes.

Problem Statement

In this work, we consider a novel robust optimization approach in which the uncertainty in future customer requests (used interchangeably with ‘orders’) is captured in a discrete uncertainty set. Our solution approach consists of designing routes for each day of the planning horizon such that all customers that have already placed a request are visited exactly once. Furthermore, any realization of future customer requests from the uncertainty set can be assigned to the designed routes in a way that all vehicle capacities are respected. That is, we guarantee that every possible realization of customer requests from the uncertainty set can be assigned to the routes in a feasible way. In this regard, the problem can be viewed as a two-stage robust optimization problem in which the routing plan for the known customer requests are first-stage decisions and the assignment of future customer requests to the routing plan constitute second-stage decisions.

Let h denote the length of the planning horizon and let $P = \{1,\ldots,h\}$ denote the set of days/time periods. We assume there is a single depot (denoted by 0), which consists of a homogeneous fleet of K vehicles, each of capacity $Q > 0$ available on each day of the horizon. Let $V_c = \{1,\ldots,n\}$ denote the set of customers who have a pending request. Each customer node $i \in V_c$ is associated with a demand $q_i > 0$ and a day window $[e_i, l_i]$ ($1 \leq e_i \leq l_i \leq h$), where e_i and l_i are the earliest and latest days by which customer i must receive service. Denote by $V_p \subseteq V_c$, the subset of customers that can be visited on day $p \in P$, i.e., $V_p = \{i \in V_C : e_i \leq p \leq l_i\}$. Let node 0 denote the depot. In the following, we describe how we model uncertainty in our context.

Model of Uncertainty

We assume that there is a set of potential customer requests, denoted by V_O, each of which may or may not realize throughout the planning horizon. Each request/order $o \in V_O$ is associated with a certain customer c_o, a "call-in" day $d_o \geq 1$, a day window $[e_o, l_o]$ such that $d_o < e_o \leq l_o \leq h$,[1] and demand $q_o > 0$, if the order does materialize. Note that we do not require $c_o \in V_C$; indeed, c_o could be any customer registered in some sort of "database" of potential customers.

Orders from V_O are realized from a (bounded) uncertainty set $\Xi = \{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : A\boldsymbol{\xi} \leq b\}$, which consists of a finite (but possibly very large) number of discrete points. The random variable $\boldsymbol{\xi}$ in this uncertainty set is a $0-1$ vector and represents the materialization of orders from V_O; more specifically, ξ_o denotes whether order o materializes or not. Note that uncertainty in the actual call-in day d_o can be accommodated in the above modeling framework by duplicating orders from the same customer into different orders with the same demand but different day windows.

This model of uncertainty is very flexible as it allows us to capture various practically relevant scenarios and linear correlations in the uncertainty. Moreover, these correlations can be combined in any way possible. For example, we can easily account for the following scenarios:

- a budget of $\Gamma > 0$ requests throughout the horizon,

$$\Xi = \{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : \sum_{o \in V_O} \xi_o \leq \Gamma\}$$

- a budget of $\Gamma_p > 0$ requests on day p,

$$\Xi = \{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : \sum_{o \in V_O : d_o = p} \xi_o \leq \Gamma_p \quad \forall\, p \in P\}$$

- budgets of orders from the same customer (e.g., at most one placed call),

$$\Xi = \{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : \sum_{o \in V_O : c_o = c} \xi_o \leq 1 \quad \forall\, c\}$$

[1]Note that, if an order o for which $l_o > h$ was considered, the optimizer would push it to be served outside the horizon because it is cheaper to simply visit the customer after time period h, while still being feasible for all scenarios for which this order indeed materializes.

Robust MILP Formulation

Let $G_p = (V_p \cup \{0\}, E_p)$ denote the complete undirected graph in period p, where $E_p = \{(i,j) \in \{V_p \cup \{0\} \times V_p \cup \{0\} : i < j\}$. Each edge $(i,j) \in E_p$ is associated with a travel cost $c_{ij} > 0$. We use variables $x_{ijp} \in \{0,1\}$ to denote if edge $(i,j) \in E_p$ is traversed on day $p \in P$. Furthermore, let $y_{ip} \in \{0,1\}$ denote if customer request $i \in V_c$ is serviced on day $p \in \{e_i, e_i+1, \ldots, l_i\}$. We use auxiliary pseudo-binary variables $z_{ikp} \in [0,1]$ to indicate if customer request $i \in V_c$ is serviced by vehicle $k \in K := \{1, \ldots, K\}$ on day p. The following is a complete formulation for the deterministic problem:

$$\min_{x,y,z} \sum_{p \in P} \sum_{(i,j) \in E_p} c_{ij} x_{ijp} \tag{1}$$

$$\text{s.t.} \quad \sum_{j:(i,j) \in E_p} x_{ijp} = 2 \sum_{k \in K} z_{ikp} = 2 y_{ip} \qquad \forall\, i \in V_p,\ \forall\, p \in P \tag{2}$$

$$\sum_{j \in V_p} x_{0jp} = 2K \qquad \forall\, p \in P \tag{3}$$

$$\sum_{p=e_i}^{l_i} y_{ip} = 1 \qquad \forall\, i \in V_C \tag{4}$$

$$\sum_{(i,j) \in \delta_p(S)} x_{ijp} + \sum_{i \in S} (1 - y_{ip}) \geq 2 \left\lceil \frac{1}{Q} \sum_{i \in S} q_i \right\rceil \qquad \forall\, S \subseteq V_p,\ \forall\, p \in P \tag{5}$$

$$1 - x_{ijp} \geq \max\{z_{ikp} - z_{jkp}, z_{jkp} - z_{ikp}\} \qquad \forall\, (i,j) \in E_p,\ \forall\, k \in K,\ \forall\, p \in P \tag{6}$$

$$z_{ikp} + \sum_{l=1}^{k} z_{jlp} \leq 3 - x_{0ip} - x_{0jp} \qquad \forall\, (i,j) \in E_p,\ \forall\, k \in K,\ \forall\, p \in P \tag{7}$$

Constraints (2)–(3) are standard degree equations, (4) are assignment constraints, while (5) are a novel generalized version of the rounded capacity inequalities, while constraints (6) and (7) are responsible for the encoding of the auxiliary z_{ikp} variables (Gounaris et al., 2013). We remark that inequalities (5) are exponential in number and, hence, are added dynamically in the context of a branch-and-cut framework. We separated these inequalities using a simple modification in the tabu search-based separation procedure for the standard rounded capacity inequalities (Augerat et al., 1998).

In order to derive the robust formulation of our problem, we need to encode the following two-stage robust constraint, which states that every realization of orders from Ξ should be assignable to the currently existing routes. Here, $O_p := \{o \in V_O : e_o \leq p \leq l_o\}$ for each $p \in P$.

$$\forall\, \boldsymbol{\xi} \in \Xi \quad \exists\, \boldsymbol{\vartheta}(\boldsymbol{\xi}) : \left[\begin{array}{l} \vartheta_{okp}(\boldsymbol{\xi}) \in \{0,1\} \quad \forall\, o \in O_p,\ \forall\, k \in K,\ \forall\, p \in P \\ \displaystyle\sum_{k \in K} \sum_{p=e_o}^{l_o} \vartheta_{okp}(\boldsymbol{\xi}) = \xi_o \quad \forall\, o \in V_O \\ \displaystyle\sum_{i \in V_p} q_i z_{ikp} + \sum_{o \in O_p} q_o \vartheta_{okp}(\boldsymbol{\xi}) \leq Q \quad \forall\, k \in K,\ \forall\, p \in P \end{array} \right] \tag{8}$$

Note that the decision-maker does not really care about the decisions $\boldsymbol{\vartheta}$ in the current time period. Ideally, they must be chosen after the realization of the uncertainty so as to obtain a routing plan that is least conservative. In that regard, the routing decisions can be viewed as first stage decisions that will be decided here-and-now and will actually be executed–at least for the first period–,while decisions $\boldsymbol{\vartheta}$ can be viewed as second-stage recourse decisions that can adapt depending on the realization of uncertainty. In fact, it is not difficult to show that a static $\boldsymbol{\vartheta}-$policy would be overly conservative in most cases.

As a first approach, we consider the following robust model: $\vartheta_{okp}(\boldsymbol{\xi}) = \theta_{okp}\xi_o$, which always assigns order o to the same vehicle k in the same time period p irrespective of the realization that o is a part of. The robust formulation is then:

$$
\begin{aligned}
(R_0) \qquad & \text{Eq. (1)} \\
\text{s.t. } & \text{Eq. (2)} - (7) \\
& \theta_{okp} \in \{0,1\} && \forall\, o \in O_p,\ \forall\, k \in K,\ \forall\, p \in P && (9) \\
& \sum_{k \in K} \sum_{p=e_o}^{l_o} \theta_{okp} = 1 && \forall\, o \in V_O && (10) \\
& \sum_{i \in V_p} q_i z_{ikp} + \sum_{o \in O_p} q_o \theta_{okp} \xi_o \leq Q && \forall\, k \in K,\ \forall\, p \in P,\ \forall\, \boldsymbol{\xi} \in \Xi && (11)
\end{aligned}
$$

Note that the number of constraints (11) is typically very large and therefore, they are treated in a cutting plane fashion within our branch-and-cut framework. For a fixed primal solution $(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},\boldsymbol{\theta})^*$, the separation problem to be solved is to compute $\zeta^*(k,p) = \max_{\boldsymbol{\xi}\in\Xi} \sum_{o\in O_p} q_o \theta^*_{okp}\xi_o$ for all $k \in K, p \in P$. If for some k, p, it happens that $\zeta^*(k,p) > Q - \sum_{i\in V_p} q_i z^*_{ikp}$, then the separation problem gives us a violating realization $\xi \in \Xi$ to be added back to problem (R_0). This separation algorithm is used at every node of our branch-and-bound tree and run after the separation of all VRP-specific inequalities.

Computational Results

We derived 15 5-period instances from five benchmark CVRP instances, ranging from 20 to 100 customers, and 2 to 4 vehicles. We used specifications similar to those provided in Baldacci et al. (2011) to generate the deterministic data of our problem (i.e., the set of fixed/deterministic customers and their day windows). For each instance, we considered a classical budget uncertainty set, $\Xi = \left\{\xi \in \{0,1\}^{10n} : e^T\xi \leq \Gamma n\right\}$, where n is the number of deterministic customers appearing throughout the planning horizon, $\Gamma \in \{0, 0.05, 0.10, 0.15, 0.20, 0.25\}$ and e is the vector all ones. This uncertainty set models that no more than Γn customer requests will appear throughout the planning horizon. With increasing Γ, the level of robustness increases, albeit at an increased routing cost. The scenario $\Gamma = 0$ represents the deterministic equivalent of our robust problem, i.e., the scenario where decisions are made purely based on requests that have already been placed.

In Table 1, we report, for each level of Γ: the number of instances (out of 15) for which a robust feasible solution exists (denoted by *# Feasible*); the average cost increase over the $\Gamma = 0$ scenario, defined as $(\text{cost}_\Gamma - \text{cost}_0)/\text{cost}_0$, across all robust feasible solutions (denoted by *% Cost increase*); the number of (feasible) instances which could be solved to optimality (denoted by *# Proven optimal*); the average time to prove optimality of these instances in seconds (denoted by *t (sec)*). Finally, for those instances which were feasible but could not be solved in the time limit of

700 seconds, the residual gap (i.e., the optimality gap at the time limit) is reported as an average.

Our study allows the decision-maker to quantify the cost increase in implementing a robust plan, and to decide the level of robustness that is appropriate for the application. For the instances we considered, robust routing plans with $\Gamma = 10\%$ can be obtained with a modest cost increase of less than 5%.

Table 1: Summary across 90 problems.

Γ (%)	# Feasible (out of 15)	% Cost increase	# Proven optimal	t (sec)	# Reached time limit	Residual gap (%)
0	15	0.0	11	23.5	4	3.1
5	15	0.0	9	17.5	6	4.8
10	15	2.0	9	81.4	6	8.2
15	12	5.9	5	36.6	7	9.1
20	8	16.6	2	0.1	6	8.3
25	3	30.4	2	8.6	1	11.4

References

Albareda-Sambola, M., Fernández, E., and Laporte, G. (2014). The dynamic multiperiod vehicle routing problem with probabilistic information. *Computers & Operations Research*, 48:31–39.

Archetti, C., Archetti, C., Jabali, O., and Speranza, M. G. (2014). Multi-Period Vehicle Routing Problem with Due Dates. Technical Report CIRRELT-2014-41, CIRRELT.

Augerat, P., Belenguer, J., Benavent, E., Corberán, A., and Naddef, D. (1998). Separating capacity constraints in the CVRP using tabu search. *European Journal of Operational Research*, 106:546 – 557.

Baldacci, R., Bartolini, E., Mingozzi, A., and Valletta, A. (2011). An Exact Algorithm for the Period Routing Problem. *Operations Research*, 59:228–241.

Coelho, L., Cordeau, J.-F., and Laporte, G. (2012). Multi-Period Vehicle Routing Problem with Due Dates Multi-Period Vehicle Routing Problem with Due Dates. Technical Report CIRRELT-2012-37, CIRRELT.

Cordeau, J.-F., Dell'Amico, M., Falavigna, S., and Iori, M. (2015). A rolling horizon algorithm for auto-carrier transportation. *Transportation Research Part B: Methodological*, 76(June):68–80.

Gounaris, C. E., Wiesemann, W., and Floudas, C. A. (2013). The Robust Capacitated Vehicle Routing Problem Under Demand Uncertainty. *Operations Research*, 61(3):677–693.

Wen, M., Cordeau, J.-F., Laporte, G., and Larsen, J. (2010). The dynamic multi-period vehicle routing problem. *Computers & Operations Research*, 37:1615–1623.