Multi-Period Vehicle Routing with Call-In Customers

Anirudh Subramanyam, Chrysanthos E. Gounaris
Carnegie Mellon University

Frank Mufalli, Jose M. Pinto
Praxair Inc.

EWO Meeting
September 30th – October 1st, 2015

Vehicle Routing

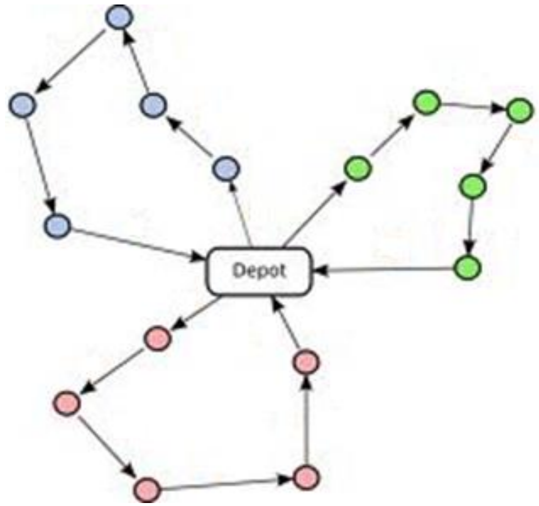

- Given a set of customer requests, determine minimum cost vehicle routes such that all requests are satisfied
  - Single vehicle: Traveling Salesman Problem (TSP)
  - Multiple (capacity-constrained) vehicles: Vehicle Routing Problem (VRP)

- Vehicle routing problems are central to transportation logistics and distribution planning

- Applications also arise in:
  - Production Planning and Scheduling
  - Network Design
  - Semiconductor Manufacturing

Tactical Planning in Vehicle Routing

- Given a set of customer requests
  - Assign a visit day to each customer over a short-term horizon (e.g., one week)
  - Design routes for each day in order to minimize sum of routing costs

- Scheduling and routing decisions are made simultaneously in a multi-period setting while respecting constraints on all days:
  - Customer availability (e.g., visit day windows)
  - Fleet capacity and other routing-specific constraints (e.g., time windows)

- Plan is executed in a rolling horizon fashion
  - Routes of the first day are executed
  - New requests are recorded and problem is re-solved using the updated portfolio of customer requests

Motivation

- A time horizon of one day reduces the problem to a standard VRP
  - Can be infeasible because it could require too many vehicles
  - Can be too expensive in terms of routing costs and overtime pay

- Actual setting may preclude single-day planning horizon

  …because of the need to set customer appointments

  …because of internal human resource constraints

- Applications arise in many fields:
  - Food and beverage distribution/chemical distribution [vehicle capacities are limiting]
  - Scheduling crews for planned maintenance of service equipment [hours-of-work constraints are limiting]

MILP Formulation

- Given
  - Planning horizon $p \in \{1, \dots, h\}$
  - Fleet $\mathcal{K}$ of m vehicles with capacity Q each, available every day
  - Customer set V_C with demands q_i and "day windows" $[e_i, l_i]$ (each customer must receive service within its day window)
  - Routing costs c_{ij} between every pair of nodes $(i, j) \in A$

- Determine
  - Visit day for each customer $y_i^p \in \{0,1\}$
  - Routes for each day $x_{ij}^p \in \{0,1\}$
  - Assignment of customers to fleet $z_{ik}^p \in [0,1]$

Carnegie Mellon ENGINEERING

CAPD CENTER

MILP Formulation

$$\underset{x,y,z}{\text{minimize}} \quad \sum_{p \in \mathcal{P}} \sum_{(i,j) \in A} c_{ij} x_{ij}^p$$

$$\text{subject to} \quad \sum_{\substack{j \in V: \\ j \neq i}} x_{ij}^p = \sum_{\substack{j \in V: \\ j \neq i}} x_{ji}^p = y_i^p \qquad \forall\, i \in V_C,\ \forall\, p$$

Fleet availability →

$$\sum_{j \in V_p} x_{0j}^p \leq m \qquad \forall\, p$$

Assign a visit day within day window →

$$\sum_{p=e_i}^{l_i} y_i^p = 1 \qquad \forall\, i \in V_C$$

Eliminate subtours →

$$\sum_{i \in V \setminus S} \sum_{j \in S} x_{ij}^p \geq y_v^p \qquad \forall\, v \in S,\ \forall\, S \subseteq V_C,\ \forall\, p$$

Assign vehicles to visited customers →

$$\sum_{k \in \mathcal{K}} z_{ik}^p = y_i^p \qquad \forall\, i \in V_p,\ \forall\, p$$

$$1 - x_{ij}^p - x_{ji}^p \geq \max\{z_{ik}^p - z_{jk}^p,\, z_{jk}^p - z_{ik}^p\} \qquad \forall\, i,j \in V_C : i < j,\ \forall\, k,\ \forall\, p$$

Break symmetry →

$$z_{ik}^p + \sum_{l=1}^{k} z_{jl}^p \leq 3 - x_{0i}^p - x_{0j}^p \qquad \forall\, i,j \in V_C : i < j,\ \forall\, k,\ \forall\, p$$

Respect vehicle capacity →

$$\sum_{i \in V_C} q_i z_{ik}^p \leq Q \qquad \forall\, k,\ \forall\, p$$

Call-In Customers

- Customers who are not scheduled to be visited now but might potentially "call-in", requesting service in a future time period

- The above framework does not explicitly account for the uncertainty of future orders within the planning horizon
  - Myopically optimizes based on current information
  - Can generate routing plans that are infeasible and/or too expensive

- Challenges
  - … Characterize discrete (yes/no) nature of uncertainty
  - … Build tractable models that insure against such uncertainty
  - … No existing methods in Robust Optimization can systematically and tractably treat general discrete uncertainty

Characterizing Discrete Uncertainty

- All call-in orders are assumed to come from a (possibly huge) database of "potential orders" V_O

- Each call-in customer c_o places an order o, associated with
  - Call-in date d_o
  - Demand q_o
  - Service Day Window $[e_o, l_o]$

- Can account for multiple orders placed by the same customer
  - Duplicate $o_1, o_2, \ldots$ with same demand but different day windows

- All call-in orders satisfy $e_o \geq 2$, since all customers of day 1 ("today") are known at time of optimization

Characterizing Discrete Uncertainty

- "Uncertainty set"

$$\Xi = \{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : A\boldsymbol{\xi} \le b\}$$

  - Finite collection of relevant "realizations" of customer orders
  - Each element is a 0 – 1 vector indicating which orders can be realized together throughout the planning horizon

- The above representation can capture practically-meaningful scenarios

…budget of orders throughout the week

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

…budget of calls on any day

$$\Xi = \left\{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : \sum_{o \in V_O : d_o = p} \xi_o \le \Gamma_p \quad \forall\, p = \{1, \ldots, h\}\right\}$$

…geographical budgets

$$\Xi = \left\{\boldsymbol{\xi} \in \{0,1\}^{|V_O|} : \sum_{o \in V_O \cap B_l} \xi_o \le b_l \quad \forall\, l = \{1, \ldots, L\}\right\}$$

…budget of orders from same customer

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

Robust Counterpart

- Routing plan must remain feasible for any realization of orders

  …must have enough fleet capacity to accommodate call-in demand in future time periods

- $\theta_{ok}^{p} \in \{0,1\}$ indicates which vehicle we "virtually" assign to order o

$$\sum_{p=e_o}^{l_o} \sum_{k \in \mathcal{K}} \theta_{ok}^{p} = 1 \qquad \forall\, o \in V_O$$

- Robust capacity constraint:

$$\sum_{o \in V_O} q_o \theta_{ok}^{p} \xi_o + \sum_{i \in V_C} q_i z_{ik}^{p} \leq Q \qquad \forall\, k,\ \forall\, p,\ \forall\, \boldsymbol{\xi} \in \Xi$$

  – Resulting model is a semi-infinite MILP

  – Can be reformulated into an MILP only for certain uncertainty sets: disjoint budgets

Algorithmic Framework

- Solve the robust problem with the deterministic capacity constraint, i.e., with $\Xi = \Xi' \leftarrow \{\mathbf{0}\}$, using a branch-and-bound solver

- At each node, apply the following cutting plane algorithm:

1. Obtain the current primal solution $(\boldsymbol{x}^*, \boldsymbol{y}^*, \boldsymbol{z}^*, \boldsymbol{\theta}^*)$
2. For each k and p, solve the following MILP and obtain the solution $\hat{\boldsymbol{\xi}}$

$$\hat{Z} = \max_{\boldsymbol{\xi} \in \Xi} \sum_{o \in V_O} q_o \theta_{ok}^{*p} \xi_o$$

3. (a) If $\hat{Z} > Q - \sum_{i \in V_C} q_i z_{ik}^{*p}$, then $\Xi' \leftarrow \Xi' \cup \left\{ \hat{\boldsymbol{\xi}} \right\}$ and add the robust capacity constraint corresponding to $\hat{\boldsymbol{\xi}}$ to the master problem and re-solve the node. Go to 1.
   (b) Else, the current node is feasible with respect to all robust capacity constraints. Stop.

Computational Results

- Database of 10 benchmark instances from VRP literature
  - 20 – 50 fixed/deterministic customers, 2 – 4 vehicles
  - Extended into 5-period problems

- For each instance, we constructed a large database V_O of orders
  - Day windows randomly assigned with maximum window of 3 days
  - Demands vary between $\left[\underline{q}, \overline{q}\right]$ with respect to deterministic customers

- Uncertainty is described by a budget of orders throughout the week

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

  - Γ varies between 0% and 25% of the no. of fixed customers

Computational Results

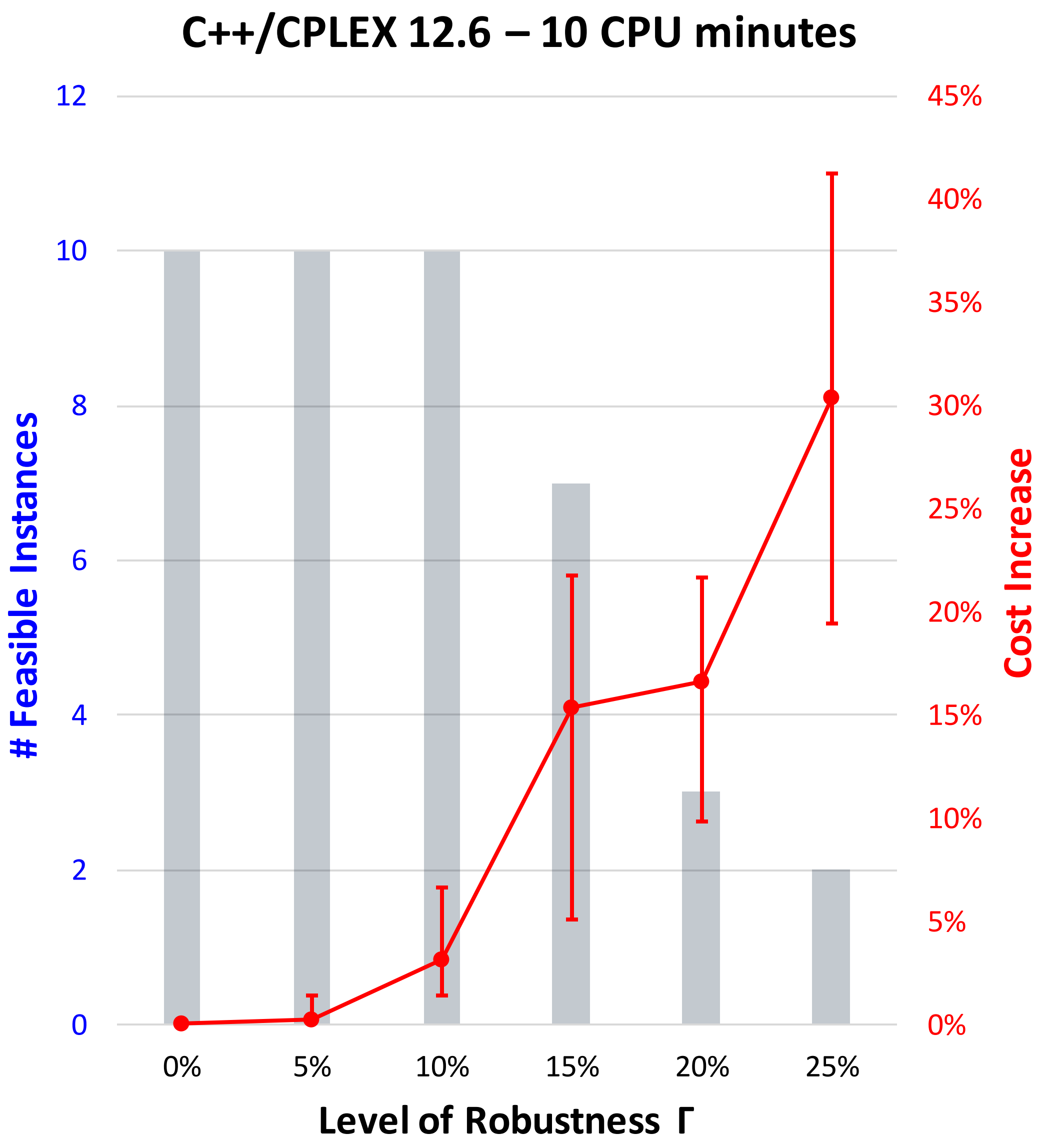

Conclusions and Future Work

- Ignoring uncertainty in future customer requests can lead to infeasible or highly expensive routing plans in multi-period tactical VRPs
- We developed a robust optimization framework to account for the uncertainty of future call-in customers
  - This is the first approach to systematically address discrete uncertainty in robust optimization
- A robust plan can be obtained with a small increase in routing costs
- Future work:
  - Improve the conservatism of the proposed framework
  - Procedurally reduce the number of "relevant" realizations Ξ'
  - Better formulations/decomposition techniques to solve larger instances