Multi-Period Vehicle Routing

Bruce L. Golden
Robert H. Smith School of Business
University of Maryland

Presented at TRANSLOG Workshop
Reñaca, Chile
December 2009

Colleagues and Co-Authors

- Damon Gulczynski, Math Department, University of Maryland
- Edward Wasil, Kogod School of Business, American University
- Benjamin Dussault, Math Department, University of Maryland

Introduction

- In the literature, the single-period vehicle routing problem still receives most of the attention
- But, multi-period vehicle routing is also important
  - Decisions span multiple time periods
  - Decisions made for one period impact outcomes in other periods
- Our goal here is to introduce several well-known problems in multi-period vehicle routing and then focus on one of these

Multi-Period Vehicle Routing Problems

- Traditional period vehicle routing problem (PVRP)
- PVRP variants
- Inventory routing problem (IRP)
- IRP variants
- Other problems

Traditional Period Vehicle Routing Problem

- We work with a planning horizon of T days
- Each customer has frequency of visit requirements (e.g., k out of T days)
- Visits to customers must occur on allowed k-day combinations
  - If T = 5 for a 5-day week, possible 2-day combinations are M & W and T & Th
- Decision variables
  - Assign visitation schedule to each customer
  - Solve a vehicle routing problem for each day

Evolution of the PVRP as of 2007

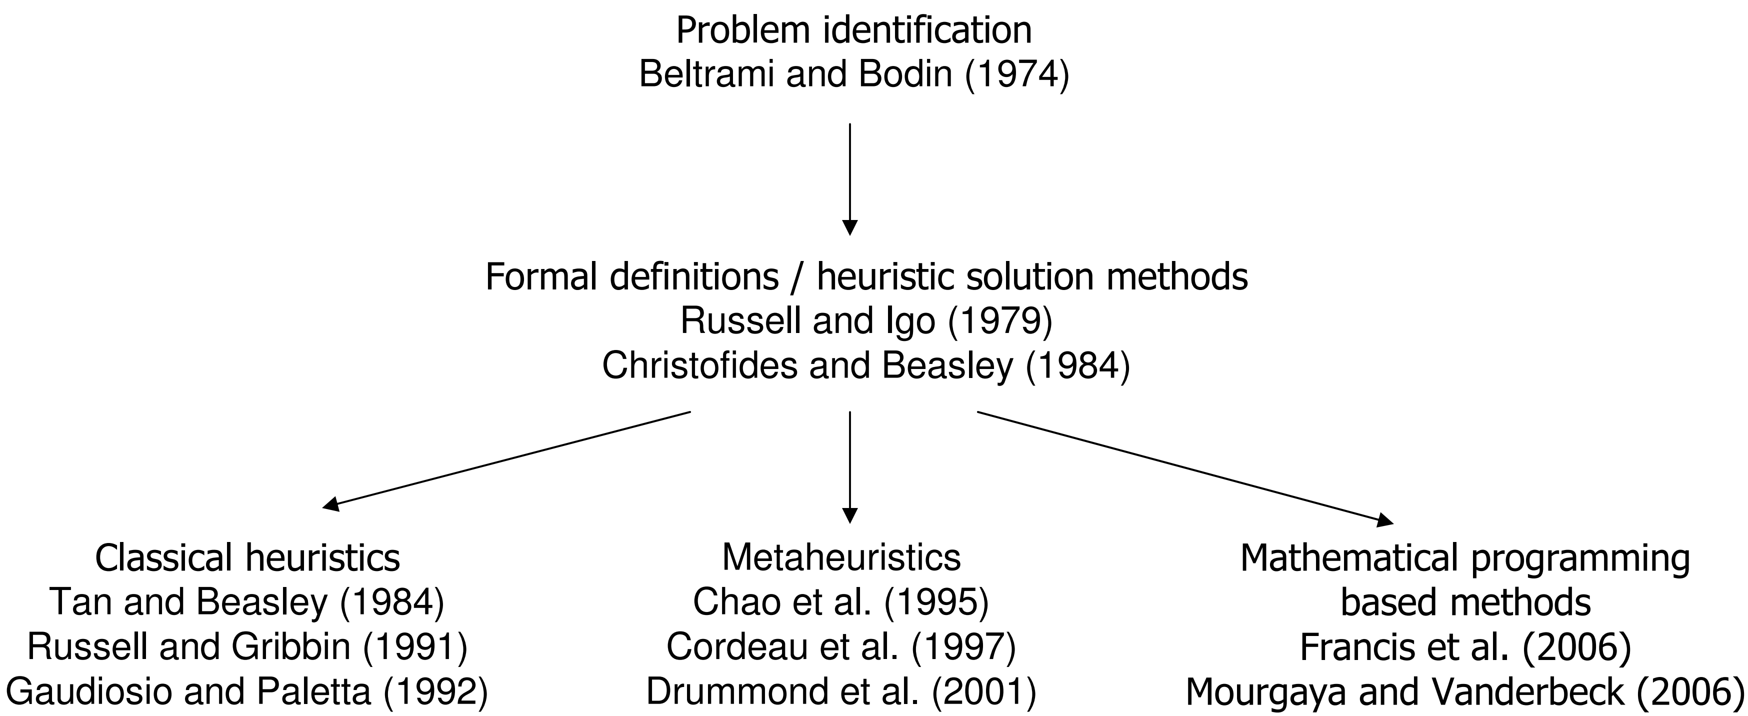

Above figure borrowed from Francis, Smilowitz, and Tzur (2008)

PVRP Variants

- Multiple depots
- Intermediate facilities for capacity replenishment
- Time windows
- Service choice (i.e., service frequency becomes a variable)
- Multiple routes (i.e., a vehicle can do more than one route per day)

The Inventory Routing Problem

- The IRP involves the repeated distribution of a single product from a single facility to n customers over T days
- Customers consume the product on a daily basis and maintain a small, local inventory
- The objective is to minimize the sum of transportation and inventory-related costs (stockouts can be costly)
- Distribution is often vendor managed in the petrochemical, industrial gas, and a growing number of other industries
- This is a very rich multi-period problem

The Inventory Routing Problem

- Decisions to be made by vendor
  - When to serve a customer (proximity vs. urgency)?
  - How much to deliver?
  - What are the delivery routes?
- Usage rates can be modeled in many ways
  - Deterministic usage
  - Stochastic usage
- IRP variants involve time windows and intermediate facilities for temporary storage of product

Other Multi-Period Problems

- Master-route design for small-package delivery

  Trans. Science (2009)
- Template-route design for small-package delivery

  MSOM (2009)
- The balanced billing cycle VRP for utility companies

  Networks (2009)

VRP and PVRP

- In contrast to the classical VRP, the PVRP is a multi-period, multi-level vehicle routing problem
- In the period vehicle routing problem (PVRP), customers might require service several times during a time period

PVRP Description

- We must first assign customers to patterns (certain days of the time period) and then find routes on each day servicing the customers scheduled on that day
- We seek to minimize total distance traveled throughout the time period
- For example, a waste management company has to assign customers to certain days of the week and then create daily routes
- Some customers might only need service once a week, some might need service multiple times

PVRP Example

- The time period is two days, $T = \{1, 2\}$
- Customer 1 must be serviced twice
- Customers 2 and 3 must be serviced once
- Node labels in parentheses are demands per delivery
- Edge labels are distances
- Vehicle capacity is 30

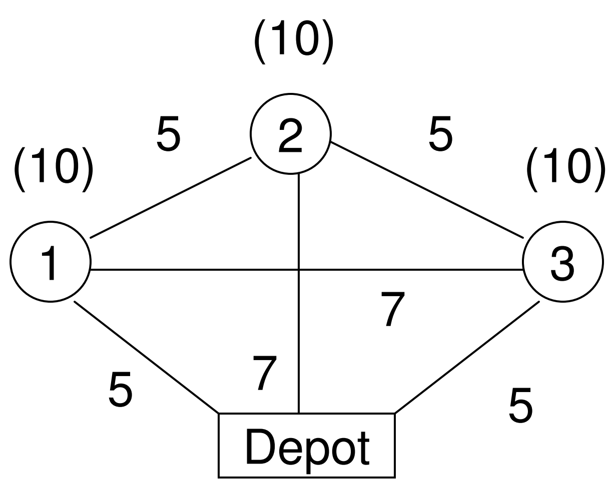

PVRP Example

- Customer 1 is assigned to day 1 and day 2
- Customer 2 is assigned to day 1
- Customer 3 is assigned to day 2

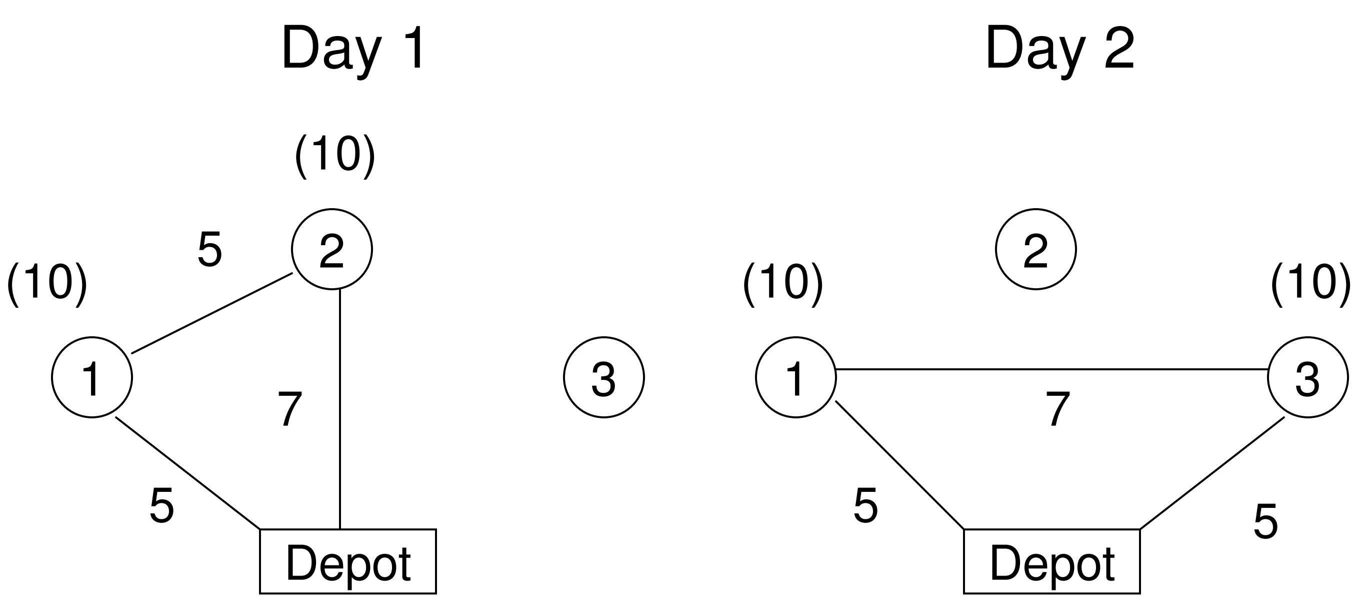

Total distance traveled is 34 units

PVRP Applications

- Commercial sanitation
- Grocery and soft drink distribution
- Fuel oil and industrial gas delivery
- Internal transport installation and maintenance
- Utility services
- Automobile parts distribution
- Oil collection from onshore wells

PVRP Literature

- PVRP literature dates back to the 1970s

- Recent papers on solving the PVRP
  - Cordeau, Gendreau, and Laporte (1997)
  - Alegre, Laguna, and Pacheco (2007)
  - Hemmelmayr, Doerner, and Hartl (2009)

IP Based PVRP Heuristic

- We develop an IP-based Heuristic algorithm for solving the PVRP, denoted IPH
- We easily adapt IPH to solve two variants later on
- We present computational results which demonstrate the effectiveness of our algorithm

IPH: Overview

1. Generate an initial PVRP solution *S*
2. Improve *S* by solving an IP that reassigns and reroutes customers in a way that maximizes total savings (main step)
3. Improve daily routes using a VRP heuristic
4. Remove and reinsert customers using an IP
5. Repeat steps 2-4 until a stopping condition is reached

IPH: Initial Solution

- We first assign customers to patterns in a way that balances the amount of demand serviced on each day (standard assignment IP)
- Next we find a VRP solution on each day using a quick heuristic (e.g., CW savings)
- The result is an initial PVRP solution

IPH: Improvement IP

- Given a solution *S*, we formulate an improvement IP (IMP) that maximizes the savings from reassigning customers to new service patterns, removing customers from their current routes, and inserting them into new routes
- We solve IMP repeatedly until no more improvement is achieved

IPH: IMP Formulation

Objective Function

- Maximize savings from reassignments and rerouting

Constraints

- We remove a customer from its current route if and only if we reinsert it elsewhere
- We move a customer to a new day if and only if we assign it to a pattern containing that day

IPH: IMP Formulation

- The total demand of the customers we move to a route minus the total demand of the customers we move from the route cannot exceed the residual capacity of the route
- If a customer *i* or its predecessor is removed from a route, we do not move any other customers immediately prior to *i* and we remove at most one of *i* and its predecessor (this ensures the objective function gives the savings accurately)
- We assign a customer to at most one feasible pattern

IPH: IMP Example

- Initial Solution

  Customer 1 is assigned to day 1 and day 2

  Customer 2 is assigned to day 1

  Customer 3 is assigned to day 2

  Vehicle Capacity is 30 units

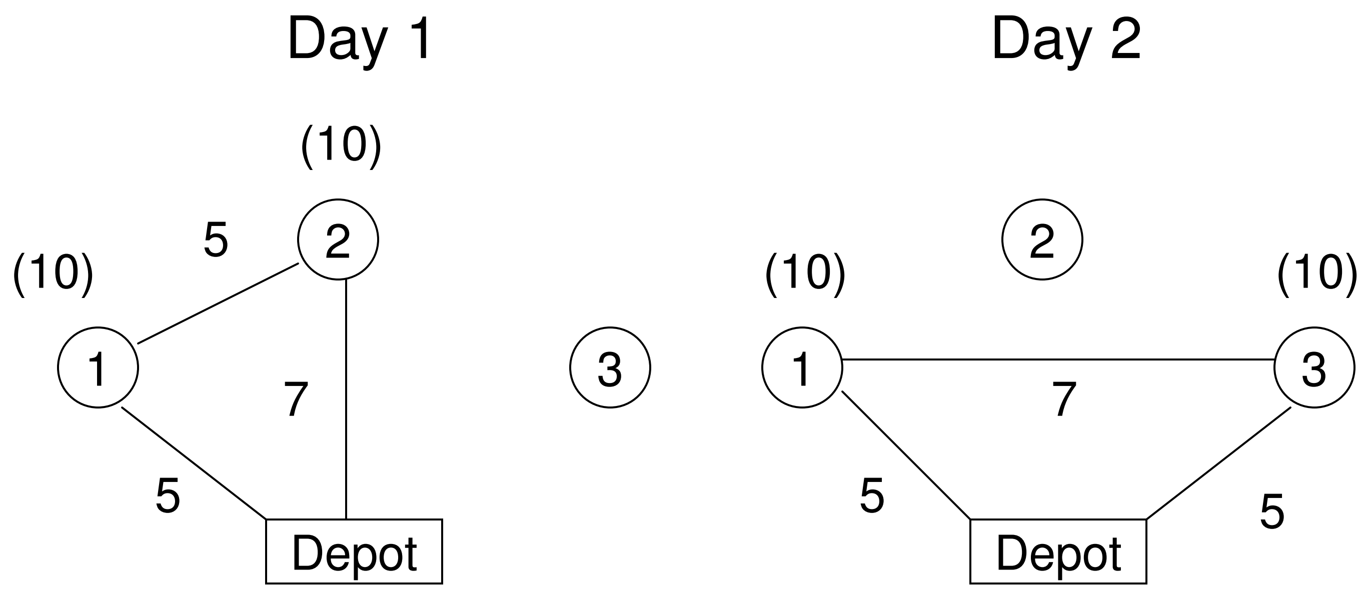

Total distance traveled is 34 units

IPH: IMP Example

- Improved Solution

  Customer 2 is removed from its route on day 1, reassigned to day 2 and inserted immediately prior to customer 3, for a savings of 4 units

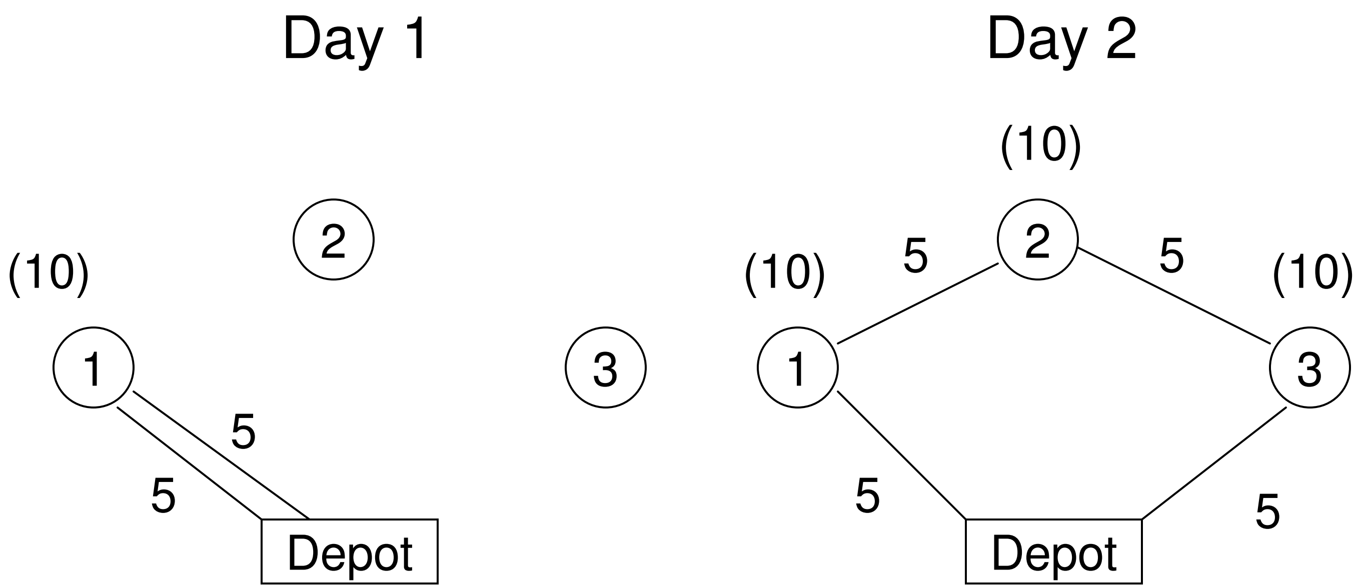

Total distance traveled is 30 units

IPH: ERTR

- We improve daily routes by solving VRPs using the enhanced record-to-record travel algorithm (ERTR)
- ERTR was developed by Groer, Golden, and Wasil (2008)
- One-point move, two-point exchange, two-opt move
- Three-point move, OR-opt move

IPH: Re-initialization

- We remove some customers randomly from their routes
- We create a fictitious day F and assign all removed customers to routes on day F
- We solve IMP adding a constraint forcing all customers visited on F to be reassigned to a feasible pattern
- Re-initialization allows us to further explore the solution space from a new starting solution

IPH: Best Solution

- We iterate IMP, ERTR, and the re-initialization procedure until a stopping condition is reached

- The best solution found throughout this process is returned at the end

IPH Results

- On 32 problems from the PVRP literature, IPH had an average deviation from the best-known solution of 1.14%
- For Cordeau, Gendreau, and Laporte (1997) it was 1.58%
- For Alegre, Laguna, and Pacheco (2007) it was 1.36%
- For Hemmelmayr, Doerner, and Hartl (2009) it was 1.39%

IPH Run-times

- On the 32 problems IPH had an average run-time of 686 seconds
- For Cordeau, Gendreau, and Laporte it was 139 (comparable machine)
- For Alegre, Laguna, and Pacheco it was 1449 (slower machine)
- For Hemmelmayr, Doerner, and Hartl it was 149 (comparable machine)
- Problem size ranges from 50 to 417 customers

IPH Virtues

- Comparable performance to the best algorithms on PVRPs
- Easily modified to handle real-world variants
- In the real world, we rarely start from scratch

PVRP Variants

- In practice, companies have pre-existing solutions that over time, due to the addition and deletion of customers and other small modifications, become inefficient
- Some companies have fleets of thousands of vehicles that service millions of customers annually
- Not practical economically, logistically, or perhaps contractually to reroute from scratch
- Instead, the major inefficiencies in the pre-existing routes should be eliminated in a way that does not cause widespread disruption

PVRP Variants

- From a pre-existing solution S' we find an improved solution S^* while constraining the total amount of disruption

- We consider three types of constraints

  1. Hard constraint: we set a limit W and only accept solutions S^* such that the number of customers assigned to different patterns from S' to S^* is at most W
  2. Soft constraint: in the objective function we penalize S^* for each customer assigned to a different pattern than the one in S'
  3. Restricted reassignment constraints: we fix all multi-day customers to their patterns in S' but allow one-day customers to be freely reassigned in S^*

- Collectively, we call these variants the PVRP with reassignment constraints (PVRP-RC)

PVRP-BC

- Because route balance is important in industry we also consider the period vehicle routing problem with balance constraints (PVRP-BC)
- In the PVRP-BC we start with a pre-existing solution S' that is well-balanced or has low cost routes (but not both)
- In the objective function, we penalize a solution for having imbalanced routes
- Starting from S', we wish to find a solution S^* that minimizes routing distance plus an imbalance penalty

PVRP-RCH and IPH-RCH

- We modify IPH to solve the PVRP-RC with a hard constraint (PVRP-RCH)
- In the PVRP-RCH, we allow no more than W customer reassignments from the pre-existing initial solution S'
- In practice, it helps to do this in stages
- We denote the modified IPH algorithm by IPH-RCH

IPH-RCH Results

- For comparison, we implement a greedy algorithm that randomly selects a reassignment from the current three best-savings reassignments and repeats until W reassignments are made or there is no improvement
- On 26 problems, with W = 10% of the number of customers, we run IPH-RCH once and the greedy algorithm 151 times, recording the best solution and the median solution, respectively
- We started with an initial solution S' that was on average 16.64% above the baseline (solution in which no restriction was put on the number of reassignments)
- For IPH-RCH the solution was 9.91% above, for Greedy Best it was 11.51% above, and for Greedy Median it was 12.62% above

Tradeoff Between Distance and Disruption

- We modify IPH for the PVRP with soft reassignment constraints (IPH-RCS)

- For an example problem, we run IPH-RCS for different reassignment penalties and see how the solution is impacted

IPH-RCS Results

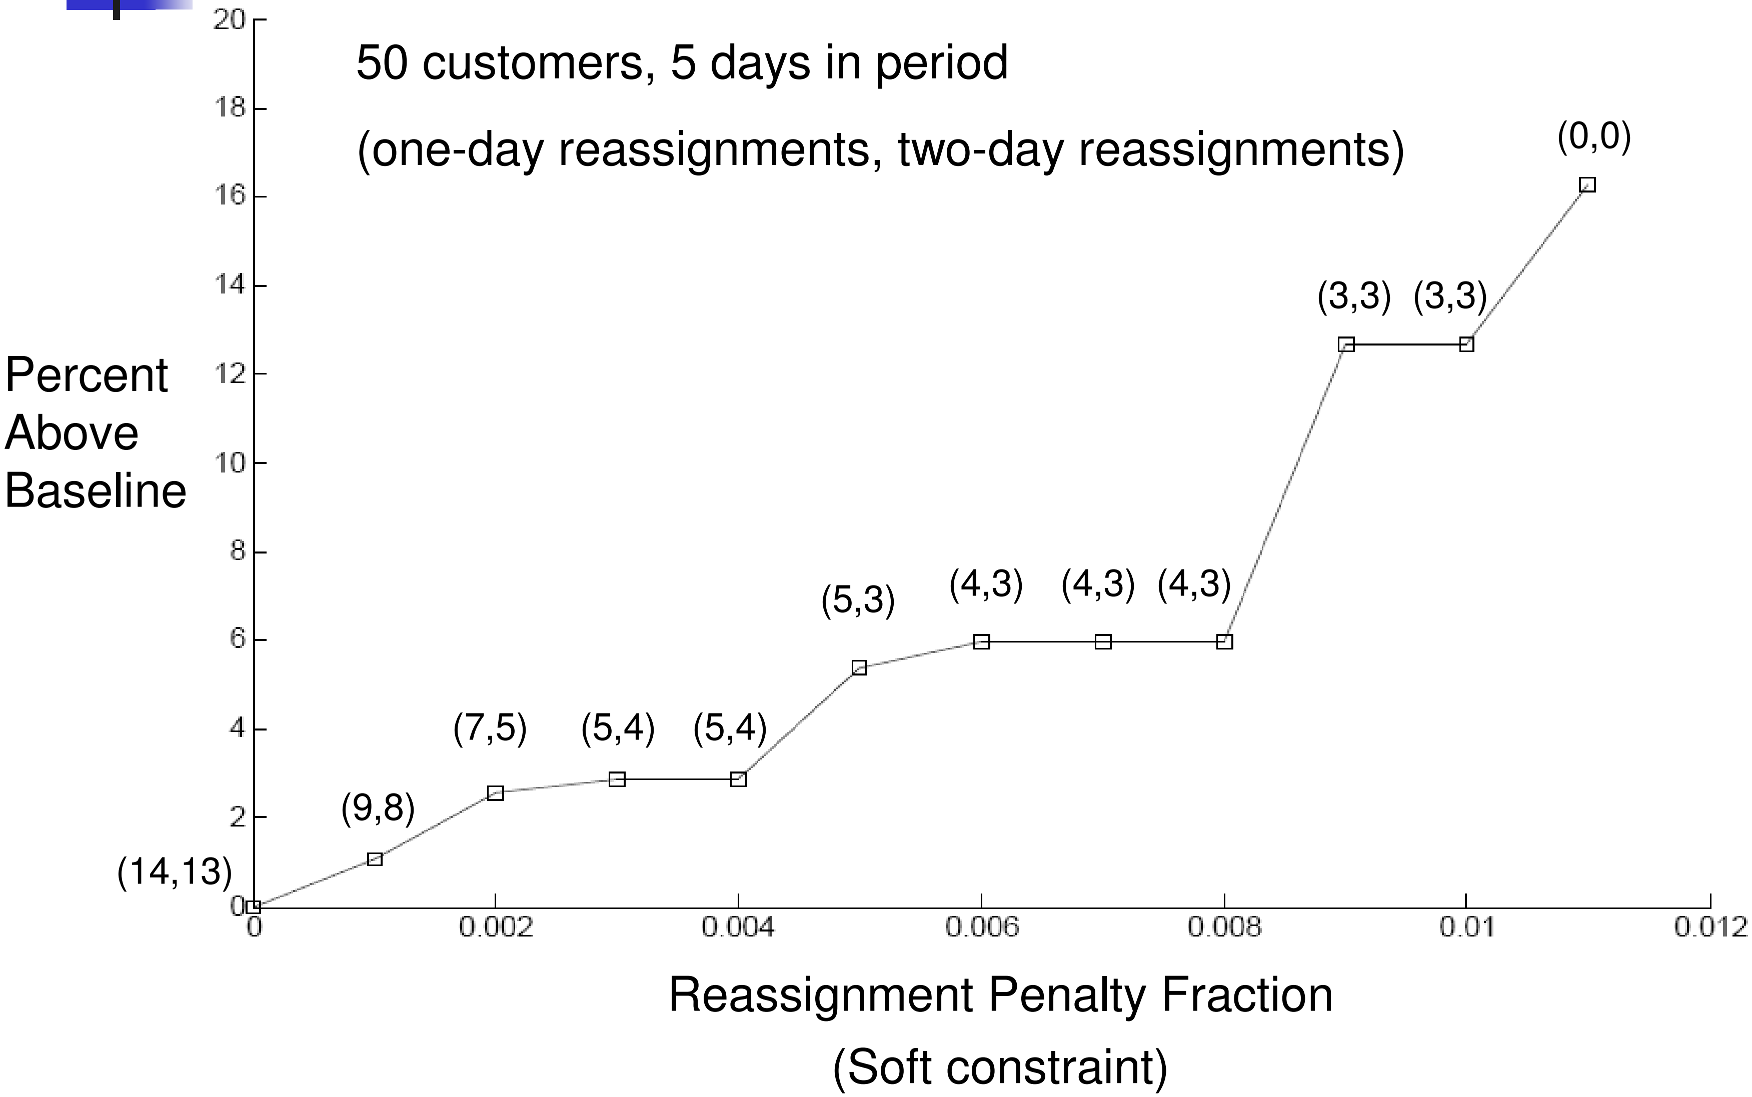

Reassigning one-day customers

- One-day customers are the easiest and most convenient to reassign
- We consider the PVRP in which we fix all multi-day customers to their initial patterns, but allow one-day customers to be reassigned freely
- In our computational experiments, the initial solution can be substantially improved using this practical compromise

PVRP-BC

- Route balance is of key importance in industry
- We consider the problem of improving a maximally balanced initial solution or a low cost initial solution without inducing too much imbalance
- Minimize: routing distance + $\rho C(U - L)$

  ρ = imbalance penalty fraction

  C = distance of the initial solution

  U = the most customers on a route in a solution

  L = the fewest customers on a route in a solution
- We denote this problem PVRP-BC
- We modify IPH for the PVRP-BC (IPH-BC)

IPH-BC Results

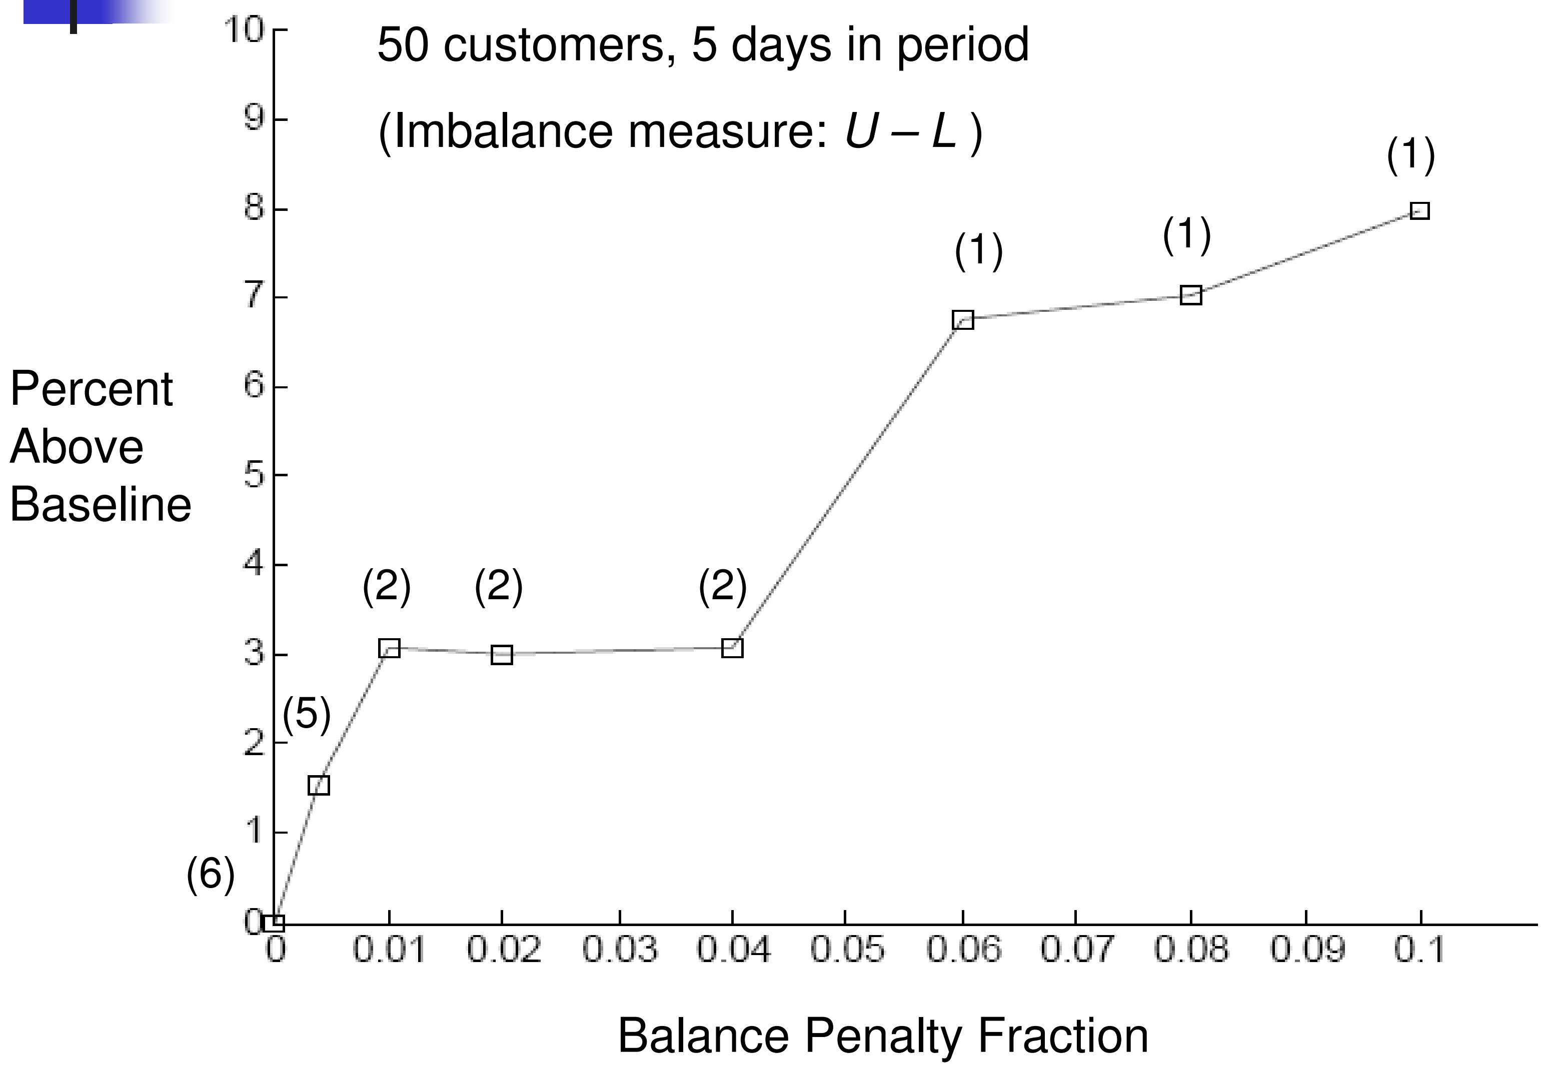

Conclusion

- We develop a new IP-based algorithm for solving the PVRP
- We adapt our algorithm to solve several variants – the PVRP-RC, the PVRP-BC, and both
- These PVRP variants are important in modeling real-world problems