Vehicle routing problem with transfers

Mgr. Václav Sobotka, supervisor doc. Hana Rudová PhD.
Masaryk University Brno

Motivation

The goal in vehicle routing problems is to find an optimal set of routes for a given fleet of vehicles so that all customer requests are served and potential constraints are respected. The solution quality is usually given by the runtime costs or the overall distance covered by the vehicle fleet. Since computed solutions tend to be of considerably higher quality than those of humans, employing algorithms allows for achieving substantial savings.

Vehicle routing problems exist in numerous variants with wide variety of constraints. One of the more general variants allows for exchange of load between vehicles at designated transfer points. This practice is referred to as transfers and its presence allows for finding more flexible routes with the potential of significant impacts on solution quality. Unfortunately, taking transfers into account greatly increases the complexity of the problem at hand, especially for problem instances with larger number of customer requests and vehicles.

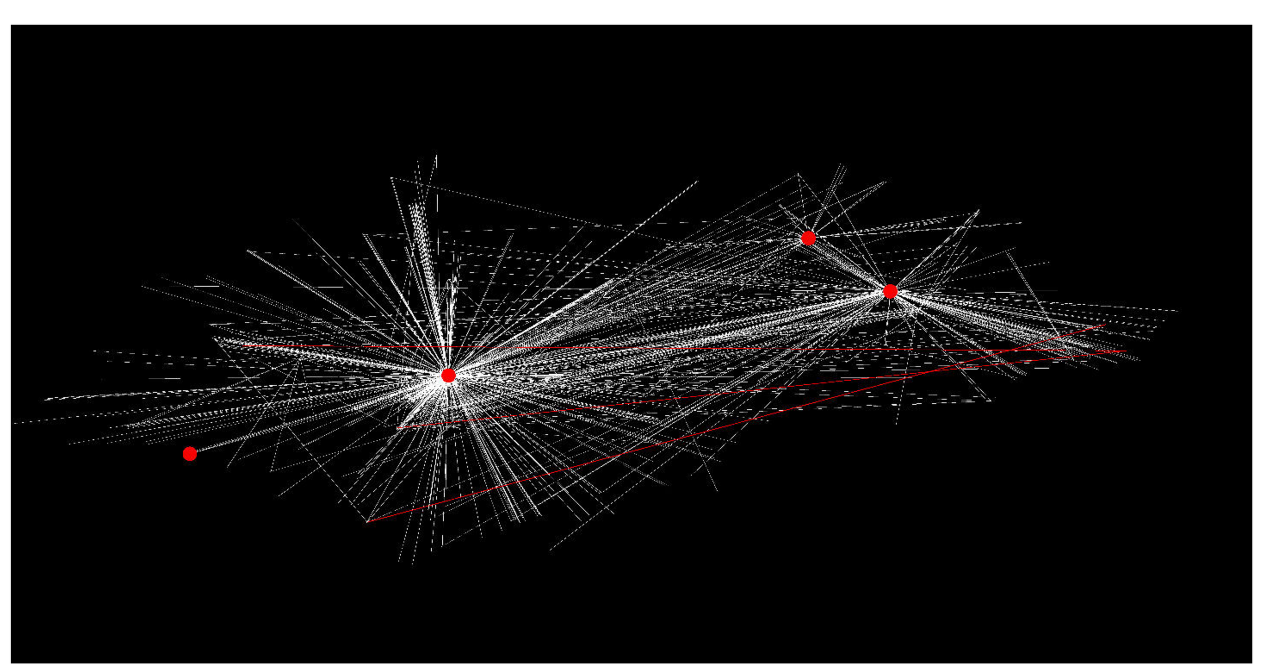

Thesis goals

The aim of the work was to study and ultimately address the phenomena of transfers in very large problem instances:

- Study existing approaches to addressing transfers
- Design a method capable of tackling very large problems
  - More than 1,200 requests and 300 vehicles
- Evaluate the proposed solver on sets of both real-world and synthetic instances

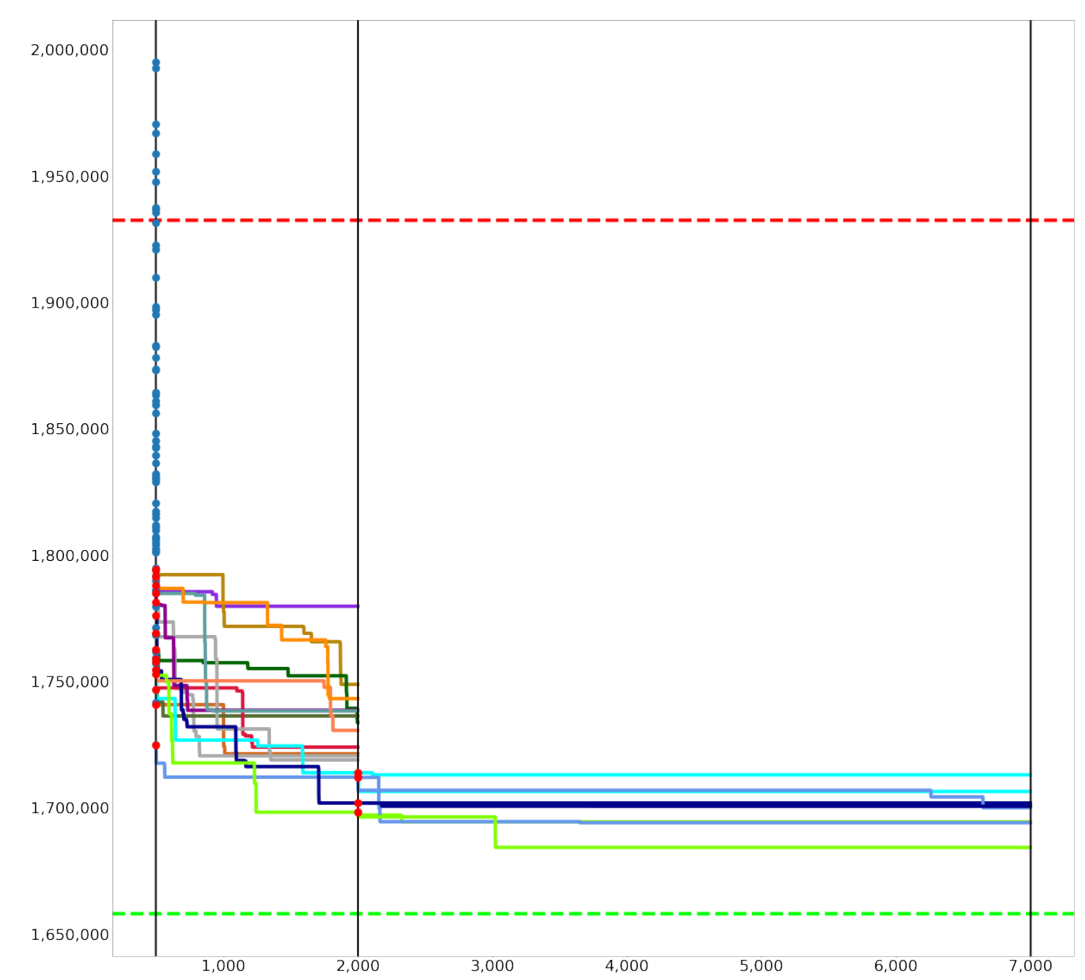

Proposed approach

The reviewed literature indicated that widespread methods based on meta-heuristics are not tractable for problem instances larger than 300 requests. The choice was to elaborate on the concept of apriori decided transfers present in the work of Petersen and Ropke:

- The transfer-or-not decision is made apriori before any routes are constructed
- Each request to be transferred is replaced by two new requests via the transfer point
- The newly derived problem can be solved by standard means without transfers

The thesis largely extends the basic concept:

- The concept is generalized to multiple transfer points
- More elaborate strategy for making the apriori decision
- Diverse set of problems is derived instead of just one

Considering several tens of derived problems allows for more robust apriori approach to transfers, but requires additional computational costs:

- The calculations may be efficiently paralellized
- Incremental pruning of non-perspective derived problems

Data

The proposed approach was evaluated on both real-world and synthetic problem instances. The real-world instances were provided by an industrial partner. The synthetic instances were obtained by using a custom implementation of instance generator utilizing OpenStreetMaps geographical data and characteristics extracted from the real-world instances.

Results

The aim of the experiments was to compare the quality of solutions with and without transfers:

- Total of 17 problem instances solved multiple times in both settings
- Substantial savings achieved for all of the 17 instances
  - Average savings of 15 % in costs
  - Best savings of 24.5 %
- Most prominent benefits on very large and real-world instances
- Favourable run times of the transfer-aware solver
  - Comparable or shorter calculation with 4+ CPUs
  - Convenient scaling by providing more CPUs

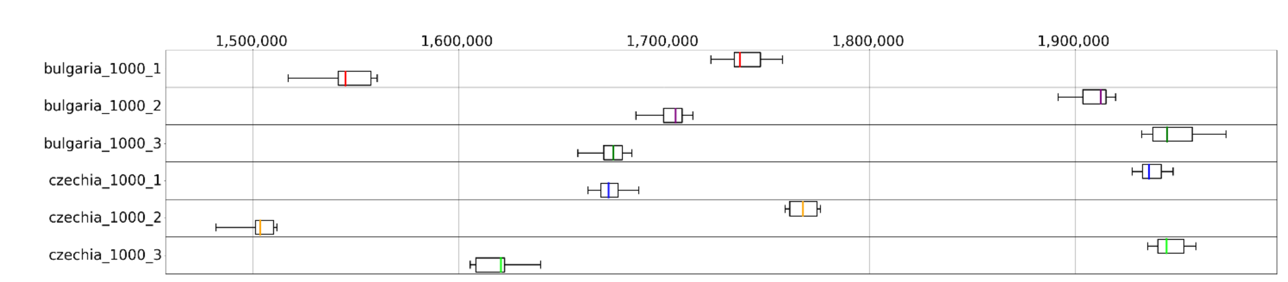

Contribution highlights

- Novel approach to transfers for very large problems
- Unique evaluation of apriori-decided transfer benefits
  - The work of Petersen and Ropke does not quantify benefits of the apriori-based approach
- Savings of up to 24.5 % in costs achieved on both real-world and synthetic instances
- Instance generator based on OpenStreetMaps for further academic use