Heuristics for Rich VRP Models

Geir Hasle Oddvar Kloster Atle Riise

SINTEF Applied Mathematics
P.O. Box 124 Blindern, N-0314 Oslo, Norway
{Geir.Hasle,Oddvar.Kloster,Atle.Riise}@sintef.no

1 Introduction

In basic and applied research in transportation optimization, there is now a strong trend towards formulating and solving richer models. This trend partly originates from external requirements of tool vendors and end users, but also from within the scientific community itself. Many idealized models that have been studied extensively over the past 4-5 decades are, despite their computational complexity, regarded as being solved from a pragmatic point of view. This fact, together with recent methodological advances and the general increase in computing power, motivates a shift to novel, scientifically even more challenging and industrially more adequate problem formulations.

A case in point is the Vehicle Routing Problem (VRP, see Toth and Vigo 2002 for a survey). In real-life applications, there are extensions related with the type of operation and various idiosyncrasies on the supply, demand, and infrastructure sides. Common extensions are:

- multiple tours per vehicle
- multiple depots
- heterogeneous fleet
- split deliveries
- mixing of pickup and delivery orders
- periodic orders
- inventory constraints
- compatibility constraints
- constraints on excess travel time for passengers
- time dependent travel times

A recent survey on rich models of the VRP (Bräysy et al 2004) shows that numerous extensions of the well-known capacitated vehicle routing (CVRP) problem have been studied.

Variants of the VRP in the literature originate from the type of decision that is to be supported, and the intended frequency of decision-making. In location-allocation routing, the goal is normally a strategic, long-term decision involving the location of depots and transit points in addition to route design. Fleet size and mix problems address tactical, medium term decisions related with the composition of a non-homogeneous fleet. In dynamic (on-line) VRPs, the goal is to provide guidance for continuous, operational management of a fleet of vehicles, and the problem will possibly change during resolution. The VRP "world" is inherently uncertain. Factors such as travel time and customer demand will be more or less uncertain. Stochastic VRP models have been given increased attention in the scientific community lately.

The conventional approach in OR has been to focus on specific extensions to the VRP and study these in isolation. This basically reductionistic approach has been highly successful in terms of understanding the effect of such extensions and has lead to development of effective and efficient algorithms for several VRP variants. However, this approach may be hard to continue as the number of extensions increase. An alternative approach is to study a richer model, with a goal of developing robust algorithms that give good performance across all instances. The research reported here follows the latter approach.

The remainder of this paper is organized as follows. Section 2 describes the context of our work on Rich VRPs and presents the overall approach. In Section 3, we describe the generic VRP model. A uniform algorithmic approach for the generic VRP model is presented in Section 4. In Section 5, performance of our approach on several types of VRP is described. Conclusions and further work are found in Section 6.

2 Context and Overall Approach

In order to capture a large variety of VRP variants, we have developed a rich VRP model and a corresponding, generic VRP solver. Our work has partly been driven by concrete end user requirements. The VRP model enables modeling of a large variety of real-life problem aspects and extends the basic CVRP in several ways. For resolution, a uniform algorithmic approach has been selected, i.e., all types of problem instances that conform to our generic VRP model are basically solved in the same way. The resolution process consists of 3 phases:

1. Construction
2. Tour Depletion
3. Iterative Improvement

For each phase, novel heuristics inspired from existing heuristics for classical VRP variants have been extended to accommodate our problem model.

3 A Rich VRP Model

The basic CVRP may be informally defined as follows. A number of identical vehicles with a specific capacity are located at a central depot. A number of customers with specified locations and demands (all either pickup or delivery) are given. The goal is to design a set of least cost routes for the available vehicles in such a way that:

- all customers are visited exactly once
- vehicle capacities are adhered to
- a constraint on maximal length or duration of each route is satisfied (for the distance constrained CVRP)

The objective is normally a hierarchical one. The primary objective is to minimize the number of vehicles used. Given solutions with the same number of vehicles, the secondary objective is to minimize the total distance of the routes.

The CVRP captures essential characteristics of many real-life routing problems. Most often, however, there are important aspects that require a richer model for an adequate study. Below, we briefly describe the basic extensions of our generic VRP model relative to the CVRP.

Types of Order
There are four types of order: Single Visit, Delivery, Pickup, and Direct. A Single Visit order does not have a size. It is used to model service orders. Direct orders are used to model pickup and delivery problems, where the vehicle will not go visit the depot between pickup and delivery, and there is a causal precedence constraint between the two. All order types are modeled as two tasks (pickup and delivery), except single visit orders.
Time Windows and Alternative Periods
Orders (through their tasks) and tours may have single or multiple hard time windows. Customer service must commence within (one of) the time window(s) for the order. Tours must start and

finish within its time window.

Service Times

Each order has a service time. Service times may depend on the location of the preceding order.

Capacity Dimensions

An unlimited number of capacity dimensions may be defined (e.g. volume and weight).

Vehicles

An arbitrary number of vehicles may be defined. The fleet may be homogeneous or heterogeneous.

Tours

Tours may have arbitrary start and end locations. There may be multiple tours for each vehicle. These are linked with temporal precedence relations.

Alternative locations

Orders (through their tasks) may have alternative locations. Tours may have alternative start and end locations.

Topologies

Travel times, distances and costs may be defined through alternative types of topology: Euclidean topology, Table topology, or Electronic Road Network (GIS) with various restrictions and speed models.

Cost Model

A variety of cost elements may be defined, including travel costs on distance and time, waiting time cost, unserviced costs on orders, initialization costs on tours, cost per order in a tour, and costs for breaking working time regulations.

Constraints

In addition to the standard capacity constraints in CVRP, constraints on total capacity over a set of tours may be defined. Compatibility constraints between tour / order and tours / location may be set up.

Locks

Parts of a solution may be locked, for instance to model history in dynamic routing.

In addition to the ability to model a large variety of real-life routing problems, the generic VRP model generalizes many VRP variants described in the literature, i.e., VRPTW, PDPTW, MDVRP, and FSMVRPTW. However, it is easy to find interesting VRP variants that are not covered by our model.

4 A Unified Algorithmic Approach to Rich VRPs

The CVRP is a NP-hard problem. For the VRPTW, finding a feasible solution is NP-complete. As our goal has been to be able to solve a large variety of VRP variants and large size instances, we have selected a heuristic approach. Moreover, we have chosen to follow a basically uniform algorithmic approach, i.e., we basically use the same algorithm for all types of instances.

A novel construction heuristic has been developed for generating one or more initial solutions in the Construction Phase. It builds upon ideas from classical construction heuristics (see for instance Bräysy and Gendreau 2002) but includes non-trivial extensions to address heterogeneous fleet, multiple depots, and tours with different start and end locations. The construction heuristic has a structure similar to the I1 heuristic (Solomon 1987), namely sequential best insertion of the most critical unserviced order, but the heuristics for determining best insertion and most critical order are different. An instance analysis is used to determine whether the problem is heterogeneous, and, in case, to determine the sequence in which new tours will be opened. Tour preferences are used to reduce the number of orders to be considered for insertion.

In the Tour Depletion Phase, a greedy tour removal heuristic is invoked. A single tour is depleted, and insertion of the unassigned orders in the remaining tours is attempted. The new solution is accepted if all unassigned orders are successfully inserted in the remaining tours. Depletion is attempted on all tours in sequence, and the process is repeated until quiescence.

The Iterative Improvement Phase is based on Variable Neighborhood Descent (VND, Hansen and Mladenović 2003), using a selection of 12 available intra-tour and inter-tour operators. The operators are mostly extensions of well-known operators for the VRPTW. Operators have been extended to accommodate the extended model. Moreover, exact and heuristic filters have been applied to increase speed. For some operators, additional opportunistic search focus has been added by analyzing the current solution and removing all moves except those that seem promising. For the Exchange operator, promising segment end points are identified by analysis of arc length, and only segments with these end points are included in the exchange neighborhood.

When VND reaches a local optimum, several diversification mechanisms are employed. First, Very Large Neighborhood Search (Ahuja et al 2002) is used to generate alternative local optima from the incumbent solution. After a number of local optima have been found without improvement, a new initial solution is attempted in a new search thread. In both types of diversification, a solution space distance metric based on arc similarity is used to determine whether the current solution is too close to earlier visited solutions. In case, a new diversification will be enforced. The overall strategy may be regarded as a hybrid of VND and Iterated Local Search (Lourenço et al 2003).

The search engine uses a flexible framework for operator sequencing and composition of macro operators. Empirical investigation, both on VRP benchmarks from the OR literature and industrial cases, has revealed that the performance of VND may depend heavily on the detailed sequencing of operators. Moreover, the best sequence varies with size and type of problem, as well as search phase. These observations have suggested a probabilistic selection of neighborhood operator, based on statistics on the recent performance of each operator during search. The approach may be regarded as a simple form of learning, and constitutes a hyper-heuristic (see Burke et al 2003) for rich VRPs.

5 Experimental Results

A comprehensive empirical investigation has been performed on a variety of test instances taken from the literature, as well as real-life instances from industry. The investigation shows that our uniform approach produces solutions of high quality in reasonable time for the much studied CVRP and VRPTW instances from the literature. For the somewhat richer VRP variants, such as the PDPTW and the FSMVRP(TW), our approach yields very good results, and in many cases, the best solutions known. At the time of writing, our approach yields the best known solutions for 80 out of the 354 PDPTW test cases by Li and Lim (2001).

6 Summary, Conclusions and Further Work

We have developed a generic, rich VRP model that accommodates many of the idiosyncrasies found in real-world routing problems, and that generalizes many of the extended VRP variants found in the literature. A uniform, heuristic approach consisting of Construction, Tour Depletion and Iterative Improvement phases is used for resolution. The Iterative Improvement phase uses a hybrid meta-heuristic search strategy that combines Variable Neighborhood Descent and Iterated Local Search. Diverse solutions are generated through a combination of Very Large Neighborhood Search and restart from alternative initial solutions. A solution space distance metric is used to cut off search threads that will probably end up in basins of attraction that have been visited earlier. Hyper-heuristics based on the recent merit of the available neighborhood operators are utilized to guide dynamic selection of neighborhood operator.

Empirical investigation on test problems from literature and industry has shown that the overall approach is robust and efficient over a large variety of VRPs. However, we believe there is still room for improvement. More specifically, future work will include development of more efficient filters on neighborhoods, more informed selection of promising moves, novel diversification mechanisms, compound neighborhood operators, more sophisticated hyper-heuristics, and new opportunistic strategies for operator selection.

References

R. Ahuja, O. Ergun, J. Orlin, A. Punnen, "A Survey of very large-scale neighborhood search techniques" *Discrete Applied Mathematics* Vol 1 23, 75-102 (2002).

E. Burke, G. Kendall, J. Newall, E. Hart, P. Ross, S. Schulenburg, "Hyper-heuristics: an Emerging Direction in Modern Search Technology". In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds), 145-184, Kluwer, 2003.

O. Bräysy, M. Gendreau, G. Hasle, A. Løkketangen, "A Survey of Heuristics for the Vehicle Routing Problem, Part I: Basic Problems and Supply Side Extensions" Preliminary SINTEF Report, Oslo, Norway (2004).

O. Bräysy, M. Gendreau, G. Hasle, A. Løkketangen, "A Survey of Heuristics for the Vehicle Routing Problem, Part II: Demand Side Extensions" Preliminary SINTEF Report, Oslo, Norway (2004).

P. Hansen and N. Mladenović, "Variable Neighborhood Search". In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds), 145-184, Kluwer, 2003.

H. Li, A. Lim, "A Metaheuristic for the Pickup and Delivery Problem with Time Windows". Presented at ICTAI'01, Dallas, Texas (2001).

H. R. Lourenço, O. C. Martin, T. Stützle, "Iterated Local Search". In *Handbook of Metaheuristics*, F. Glover and G. Kochenberger (eds), 321-354, Kluwer, 2003.

M. M. Solomon, “Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints”, *Operations Research* 35, 254–265. (1987)

P. Toth and D. Vigo (eds), *The Vehicle Routing Problem*, SIAM Monographs on Discrete Mathematics and Applications, Philadelphia (2002).