A hybrid algorithm for general Vehicle Routing Problems

Anand Subramanian, Luiz Satoru Uchi, Eduardo Uchoa

This work proposes a hybrid heuristic algorithm for a large class of Vehicle Routing Problems. The developed approach consists of a combination of the Iterated Local Search (ILS) metaheuristic and a Set Partitioning (SP) formulation. A sequence of SP models, with columns corresponding to routes found by the ILS, are solved, not necessarily to optimality, by means of a Mixed Integer Programming (MIP) solver, that may interact with the ILS heuristic during its execution.

The algorithm was extensively tested on benchmark instances of the following VRP variants: (i) Capacitated VRP (CVRP); (ii) Asymmetric VRP (AVRP); (iii) Open VRP (OVRP); (iv) VRP with Simultaneous Pickup and Delivery (VRPSPD); (v) VRP with Mixed Pickup and Delivery (VRPMPD); (vi) Multi-depot VRP (MDVRP); (vii) Multi-depot VRP with Mixed Pickup and Delivery (MDVRPMPD); (vii) Heterogeneous Fleet VRP (HFVRP), both limited and unlimited fleet versions. Although we do not consider other classical variants such as the VRP with Time Windows (VRPTW) or the VRP with Backhauls (VRPB), our solution approach can be easily adapted to solve these problems.

The ILS based heuristic, called ILS-RVND, is a multi-start procedure that uses insertion heuristics in the constructive phase and a Variable Neighborhood Descent with Random neighborhood ordering (RVND) in the local search phase. The RVND is composed by well-known VRP inter-route neighborhood structures, namely Shift(1,0), Shift(2,0), Swap(1,1), Swap(2,1), Swap(2,2), Cross; and also by specific ones, namely *K*-Shift , ShiftDepot and SwapDepot. The *K*-Shift is employed in the HFVRP, while ShiftDepot and SwapDepot are employed in the MDVRP and MDVRPMPD. The best improvement strategy was adopted and the neighborhood structures are explored exhaustively. Every time a route is modified due to an inter-route move we perform an intra-route local search using classical TSP neighborhood structures, more precisely, Or-opt, 2-opt and exchange. Finally, the perturbation mechanisms consist of performing multiple Swap(1,1) or Shift(1,1) moves. The Shift (1,1) consists of moving a costumer from a route r1 to a route r2 and vice-versa.

The hybrid algorithm, called ILS-RVND-SP, extends the ILS-RVND heuristic by incorporating an improvement procedure based on a SP formulation into it. This strategy is quite similar to the classical two-phase petal algorithm. The idea is to store a pool of routes generated during the heuristic execution and then solve a SP problem in order to extract the best combination of routes. However, unlike traditional petal algorithms and other SP based approaches to VRPs, there is a cooperation between a MIP solver and the ILS-RVND heuristic while solving the SP problem. More precisely, whenever a new incumbent solution is found by the MIP solver, we call the ILS-RVND heuristic providing this incumbent solution as an initial solution for the metaheuristic, which, in turn, may be able to find improved solutions using routes that are not in the SP model. A time limit is imposed for the MIP solver. While the ILS-RVND heuristic seeks an equilibrium among solution quality, computational time, simplicity and flexibility, the ILS-RVND-SP gives preference to the solution quality at the expense of the computational time and simplicity, but still holding all the flexibility.

One of the challenges of designing a unified hybrid solution approach is to ensure that the MIP model is computationally tractable, regardless of the instance. For example, a SP model that exceeds the time limit only to solve its linear relaxation (e.g. due to an excessive number of routes) is not a suitable improving mechanism. On the other hand, a SP model that contains relatively few routes, is easily solved, but seldom finds improved solutions. Hence, it is necessary that the SP models generated throughout the algorithm find a balance between computational tractability and improvement potential. Experiments carried out in many instances with distinct characteristics indicated that the following simple pieces of data are crucial for estimating the dimension (number of routes) of a properly balanced SP model: (a) number of customers; (b) the average number of customers per route; and (c) the vehicle fleet (homogeneous or heterogeneous).

With respect to (a) and (c), we developed the following strategies: (i) when the number of customers is less or equal to 150 we build and solve just one SP model at the end of the algorithm; (ii) otherwise, we build and solve a SP model after each iteration of the ILS-RVND heuristic; and (iii) if the fleet is heterogeneous and unlimited we first solve the SP model without specifying the fleet composition, but if the gap of the linear relaxation of the root node with respect to the incumbent solution is greater than a given value or a time limit is exceeded we fix the fleet using the current best known solution and solve the SP again. In the second strategy, we make use of short-term memory by storing a pool of routes generated at each iteration of the algorithm and a long-term memory by keeping a subset of routes permanently in the pool.

With respect to (b), we do the following. When the ratio between the number of customers and the number of vehicles is smaller than 11, the SP models tend to become harder. Hence, in such cases, we only add the routes of a solution to the SP model if its gap when compared to the incumbent solution is smaller than a given threshold gap. However, this parameter is difficult to tune, specially in the second strategy. To overcome this issue we implemented a reactive approach that dynamically adjusts its value throughout the execution of the algorithm. We start with a given initial threshold gap and we update this value as follows: (i) if the SP model is solved at the root node, it means that the problem is easy and we increase the threshold gap by one tenth of its initial value; (ii) if the time limit is exceeded we decrease the threshold gap by one tenth if its initial value; and (iii) if the SP model is solved within the time limit we do not perform any modification.

The ILS-RVND-SP algorithm was evaluated in hundreds of well-known instances of the variants considered in this work, with up to 480 customers. The same parameter tuning was adopted for all runs and the CPLEX 12 was used as a MIP solver. The results obtained were quite competitive with those found by heuristics devoted to specific variants. A number of new best solutions were obtained.