Route-Centroid Partitioning with AILS-II Global Search and Parallel HGS Intensification for Large-Scale CVRP

Sobihan Surendran

Sorbonne Université and Université Paris Cité,
CNRS, Laboratoire de Probabilités, Statistique et Modélisation, F-75005 Paris, France
LOPF, Califrais' Machine Learning Lab, Paris, France

Abstract

This report describes a CVRP solver developed for the CVRPLIB BKS challenge, targeting improvements over published best-known solutions on large-scale instances. The proposed approach combines Hybrid Genetic Search (HGS) with the AILS-II metaheuristic to exploit complementary strengths in diversification and intensification within a scalable computational pipeline.

1 Motivation

Hybrid Genetic Search (HGS) (Vidal, 2022; Vidal et al., 2012) is a state-of-the-art CVRP heuristic that combines population-based recombination with intensive local search. It performs particularly well when enough runtime is available to balance diversification and intensification. However, on very large-scale instances, achieving top performance can require a substantial computational budget. AILS-II (Máximo et al., 2024) is an adaptive iterated local search that alternates perturbation and local improvement to rapidly obtain high-quality feasible solutions, making it especially effective under tight time budgets and a natural candidate for producing warm-start solutions for HGS. Motivated by these complementary strengths, we use AILS-II to quickly generate strong incumbents and decompose them into smaller, structured subproblems on which HGS can intensify efficiently in parallel, while periodic global refinement allows improvements across subproblem boundaries.

2 Method

The proposed solver follows a multi-cycle hybrid scheme that uses AILS-II for global exploration and warm-starting, and parallel HGS runs for local intensification. In each cycle, AILS-II produces an incumbent solution S, which is decomposed into routes. We represent each route by the centroid of its served customers and cluster these centroids into P groups, yielding P route-based subproblems. For each group, we build a restricted CVRP instance containing only the customers covered by the routes in that cluster, and run HGS to improve the corresponding sub-solution; all P runs are executed in parallel.

To reduce artifacts caused by hard boundaries, we generate several alternative partitions in each cycle, so that routes/customers near partition borders can be optimized under different groupings and neighborhood contexts. After each parallel HGS round, the improved routes are merged into an updated incumbent solution, from which a new partition is computed for the next round. Finally, we apply a global AILS-II pass on the full instance to enable cross-group moves that are not visible within restricted subproblems. This loop is repeated for several cycles, and we return the best solution obtained.

References

V. R. Máximo, J.-F. Cordeau, and M. C. Nascimento. Ails-ii: An adaptive iterated local search heuristic for the large-scale capacitated vehicle routing problem. *INFORMS Journal on Computing*, 36(4):974–986, 2024.

T. Vidal. Hybrid genetic search for the cvrp: Open-source implementation and swap* neighborhood. *Computers & Operations Research*, 140:105643, 2022.

T. Vidal, T. G. Crainic, M. Gendreau, N. Lahrichi, and W. Rei. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60(3):611–624, 2012.