

Multi-Product Vehicle Routing Problem with Split Deliveries and Changeover Costs

1. Overview

The **Multi-Product Vehicle Routing Problem with Split Deliveries and Changeover Costs (MPVRP-CC)** studies how a shared fleet can distribute several products from a set of depots to geographically dispersed customers.

The aim is to decide which vehicles to use, where they should load, which customers they should visit, how much they should deliver, and in which order the products should be transported. A good solution must satisfy every demand while balancing travel distance with the operational costs associated with preparing and loading vehicles.

The problem applies to many sectors, including petroleum distribution, chemicals, food, agriculture, pharmaceuticals, and waste collection. These examples differ in practice, but they share the same planning challenge: a vehicle may perform several trips and carry different products over the course of its schedule.

2. Logistics network

The network contains:

- **vehicles**, each with a capacity, a home garage, and an initial product configuration;
- **products** to be distributed;
- **garages**, where vehicles begin and end their schedules;
- **depots**, where products are stored and loaded;
- **service stations**, or customers, with demand for one or more products.

Every depot, garage, and station has a geographical position. Travel cost is measured using the distance between these locations. Depot stocks are limited, and no vehicle may carry more than its capacity.

3. How a vehicle operates

A vehicle leaves its home garage, performs one or more delivery trips, and returns to the same garage:

Garage → Depot → Customers → Depot → ... → Garage

Each trip begins with a loading operation at a depot. The vehicle then visits one or more stations and delivers a single product before returning to a depot or ending its schedule at the garage. A vehicle carries only one product during a trip, but it may carry another product on a later trip.

4. Changeover and loading costs

The **changeover cost** is an operational transition cost associated with preparing a vehicle for the product loaded on its next trip. It should not be understood only as a penalty for switching from one product to another. The loading operation itself may require preparation, handling, inspection, or equipment setup, including for the first trip of the day.

Depending on the application, this cost may include:

- cleaning, purging, washing, drying, or decontamination;
- preparation and handling during loading;
- reconfiguration of tanks, compartments, hoses, pumps, or temperature settings;
- quality, safety, inspection, or certification procedures;
- labor, consumables, and equipment use;
- waiting time, vehicle downtime, and loss of availability;
- administrative and coordination work before departure.

The cost is described by a directed matrix. Its value depends on the vehicle's current product configuration and on the product that will be loaded. A transition from product **p** to product **q** may therefore cost more or less than the reverse transition.

The initial configuration of each vehicle is part of the instance. The first loading is evaluated from that initial configuration, so an initial preparation cost is incurred before any delivery takes place. Later costs are evaluated at each new loading. The diagonal is deliberately positive: loading the same product again still incurs a low preparation cost that reflects the loading operation.

The cost-bearing benchmark uses three off-diagonal regimes: **normal** values range from 1001 to 3500, **high** values range from 4501 to 15000, and a **mixed** matrix samples both ranges arc by arc. Every diagonal value is sampled from the **low** range, from 25 to 150. All bounds are inclusive and the matrix entries are integers.

5. Split deliveries

A station's demand for one product may be larger than a vehicle's capacity or may be more efficiently shared among several vehicles. The problem therefore allows **split deliveries**: several vehicles may contribute to the same station-product demand.

The full requested quantity must still be delivered. A vehicle may visit the same station several times only when those visits concern different products. For any given product, that vehicle may serve the station at most once during its complete schedule.

For example, a vehicle may visit station 4 once with product 1 and return later with product 2. It may not return to station 4 a second time with product 1. If the demand for product 1 must be split, another vehicle has to deliver the remaining quantity.

6. Objective

The objective is to minimize the total of:

- the distance traveled by the fleet;
- the operational transition costs incurred during initial and subsequent loading operations.

Mathematically, the objective can be written as:

$$\min Z = \underbrace{\sum_{k \in K} \sum_{(i,j) \in A} d_{ij} n_{ijk}}_{\text{total travel distance}} + \underbrace{\sum_{k \in K^+} \left(C_{p_k^0, p_{k1}} + \sum_{t=2}^{|T_k|} C_{p_{k,t-1}, p_{kt}} \right)}_{\text{initial and subsequent loading-transition costs}}$$

where:

- K is the set of vehicles and K^+ is the set of vehicles that perform at least one trip;
- A is the set of possible travel arcs;
- d_{ij} is the distance between locations i and j ;
- n_{ijk} is the number of times vehicle k travels from i to j ;
- T_k is the ordered set of trips performed by vehicle k ;
- p_k^0 is the initial product configuration of vehicle k ;
- p_{kt} is the product loaded by vehicle k for trip t ;
- C_{pq} is the preparation and loading-transition cost from configuration p to loaded product q .

This creates the central trade-off of the problem. The shortest routes are not always the least expensive: a slightly longer plan may reduce costly preparations, while a compact route may require more product transitions or loading setups.

7. Conditions for a feasible solution

A solution is feasible when:

- every station receives exactly the quantity requested for each product;
- vehicle capacities are respected;
- quantities loaded at each depot remain within available stocks;
- every trip carries exactly one product;
- a station is visited only for a product it requires;
- each vehicle visits a station at most once for any given product, although it may return with a different product;
- each used vehicle starts and ends at its home garage;
- consecutive trips form a complete, connected schedule.

The benchmark does not include time windows, explicit service times, or depot replenishment. Distances are Euclidean and rounded to the nearest integer, and all locations are considered accessible.

8. Benchmark scenarios

The benchmark contains 100 pairs of instances. Every instance has between two and six products, and the cost-bearing set comprises 34 **normal**, 33 **high**, and 33 **mixed** matrices. Each pair has two versions:

- **with changeover costs**, using the generated cost-bearing transition matrix;

- **without changeover costs**, using the same data with every transition cost set to zero.

The paired instances have the same fleet, locations, stocks, demands, capacities, vehicle configurations, and identifier. Comparing their solutions shows how loading and transition costs influence vehicle use, product sequences, depot choices, routes, and total distance.

Only the instances with changeover costs are included in the official ranking. The zero-cost versions are provided for comparative experiments.

9. Official score

For each feasible official instance, the score is the sum of the total travel distance and the total transition cost. Participants may submit any subset of the 100 solutions in a ZIP archive; submitting the complete set at once is not required.

The platform identifies the solutions present in the archive and evaluates them independently. A solution that is absent, unresolved, unreadable, or infeasible receives the same penalty of 100000 for its instance.

The final score is the sum obtained across all 100 instances. Lower scores are better.