

Inventory Routing Problem description for ROADEF/EURO 2016 Challenge



Abstract: This document describes the model scope of the Air Liquide Inventory Routing Problem related to the distribution of bulk gases for the ROADEF/EURO 2016 Challenge. It provides a formulation of the optimization model and an explicit formulation of various business and technical constraints that should be respected and how the objective function is defined and computed for a given solution.

Keywords: Inventory Routing Problem, Industrial Gas Distribution

TABLE OF CONTENTS

1.	INTRODUCTION	3
1.1.	CONTEXT	3
1.2.	POSITION OF THE PROBLEM IN THE OR WORLD	3
2.	PROBLEM SCOPE	4
3.	PROBLEM SETTING	5
3.1.	KEY DEFINITIONS	5
3.2.	UNITS OF MEASURE FOR QUANTITY	5
3.3.	TIME REPRESENTATION AND HORIZON	5
3.4.	DRIVERS MASTER DATA.....	5
3.5.	TRAILERS MASTER DATA	6
3.6.	LOCATIONS MASTER DATA	6
4.	FORMULATION.....	7
4.1.	VARIABLES	7
4.2.	CONSTRAINTS	7
4.2.1.	BOUNDING CONSTRAINTS.....	7
4.2.2.	CONSTRAINTS RELATED TO DRIVERS	7
4.2.3.	CONSTRAINTS RELATED TO TRAILERS.....	8
4.2.4.	CONSTRAINTS RELATED TO SITES.....	8
4.2.5.	CONSTRAINTS RELATED TO SHIFTS.....	9
4.2.6.	CONSTRAINTS RELATED TO QUALITY OF SERVICE	10
4.3.	OPTIMIZATION GOAL	10
4.3.1.	OBJECTIVE FUNCTION	10
4.3.2.	TIME INTEGRATIONS OVER SCHEDULING OPTIMIZATION HORIZONS	11
4.4.	APPENDIX : SPECIFIC FUNCTIONS DEFINITION USED IN THE CONSTRAINT MODELING	12

1. INTRODUCTION

1.1. Context

Worldwide, Air Liquide operates Air Separation Units (ASUs) to produce pure Oxygen, Nitrogen and Argon products through separation from air. Air Liquide serves approximately 1 million customers worldwide in various industries, including healthcare, pharmaceuticals, food & beverage, electronics, welding & cutting, steelmaking, petrochemical & refining, and oil & gas production.

In this ROADEF/EURO Challenge, we will focus on example contexts from our Healthcare business (<https://www.airliquide.com/healthcare>) which deliveries large (bulk) volumes of liquid oxygen to 7500 hospitals worldwide. On-site vaporization units convert this cold, liquid oxygen into gaseous form (at around room temperature) which travels through the hospital in dedicated piping to supply patients with the oxygen they require.

For bulk delivery customers, Air Liquide installs on-site storage vessels (“tanks”), which are regularly refilled by trucks (“vehicles”) transporting the gases as cryogenic liquid (“liquid gas”) from our ASUs to our customers.

Through its remote telemetry system, Air Liquide monitors customer tank levels and consumption rates and forecasts each customer’s future consumption over the coming hours and days. Through this process, Air Liquide takes the responsibility to guarantee that sufficient inventory of product will be maintained on-site at the customer to meet their demand (in logistics, we call this policy Vendor Managed Inventory, VMI).

While VMI customers are in the majority, a smaller but still significant set of customers (called “call-in” customers) are supplied on an “on demand” policy, as they directly place orders with Air Liquide when they need product.



In this VMI context, Air Liquide must efficiently organize the safe and reliable roundtrips of its truck fleet.

Therefore, on a daily basis, Air Liquide’s dispatchers review the forecast of each customers’ consumption and then adapt the schedule of the transportation accordingly. Their goal is to reduce the cost per delivered unit (e.g., €/kg) over the long term, while avoiding product shortage (run-outs), satisfying the orders of the call-in customers and respecting safety and regulatory constraints (e.g., limits on continuous driving time).

Air Liquide’s dispatchers have strong experience with this process. However, while it is relatively easy to find feasible schedules avoiding run outs, obtaining the minimum cost per delivered unit is a challenge. The specific complexity of the problem is mainly linked to two factors:

- ➔ There is an exponential number of feasible solutions, making the problem hard to solve even for a few customers;
- ➔ Each short term decision impacts the future cost; indeed, a decision that can reduce the cost today may be not a good decision over the long term.

1.2. Position of the problem in the OR world

In the OR world, this problem can be classified as an “Inventory-routing problem” (IRP), first proposed by Bell et al in 1983. While many mathematical formulations already exist, along with a growing scientific literature, the Air Liquide IRP has several features that make it unique:

- the objective is rational: the goal is to minimize the logistic ratio (e.g., cost per unit delivered: €/kg) which represents the best use of each km driven and each hour spent and not the total number of km or total hours only.
- Secondly, the solution must satisfy specific business-related constraints that are generally not taken into consideration in the literature.
 - Consideration that drivers' Bases and sources of product are not always co-located: classical models assume that the drivers' Bases and sources are co-located.
 - No prior assignment of drivers to trailers.
 - Accurate modeling of time: Air Liquide IRP model considers an effectively continuous time (accurate to the minute) for the timing of operations (and discrete time for inventory control)
 - Multi-trip problem: a shift may be composed by several trips, i.e. it can alternate loading and deliveries
- Finally, the average size of the problem is bigger than those generally studied in literature.

2. PROBLEM SCOPE

The context described in the challenge is adapted from the Healthcare business environment for the supply of hospitals with bulk liquid oxygen.

The main features of this **Version 1 core model** are:

Product and sourcing: A Single Product with Unlimited Supply

- One product: Liquid Oxygen at cryogenic temperature (approximately -220°C)
- One single production site (source) with unlimited product available 24/7 (no capacity constraints on production and storage)
- Safety first: fixed loading times at sources to be able to safely handle cryogenic products

Customers: Do not leave our customers breathless!

- Vendor Managed Inventory (VMI) customers only, available 24/7 for delivery (no orders / call-in customers).
- Customer Consumption forecast is known in advance for the whole horizon at an hourly time step.
- Safety levels inside the cryogenic tanks are determined for the whole horizon to always guarantee high enough level of oxygen. No run-out are allowed below safety level.
- Safety first: fixed unloading operation time per customer to be able to safely access site and perform all tasks.

Transportation resources: Pairing Drivers and Tractor-Trailers

- All the transportation resources are located at one single Base (which is located at a different location to the production source)
- Several drivers with different availability time windows (Day and night drivers).
- A driver can only drive one trailer (typically each trailer are used by both day and night drivers)
- Heterogeneous Fleet: Each trailer has a specific capacity.

Shift definition: Happy Drivers, Safely Home Each Day.

- Each shift has only one driver and only one Trailer.
- A day or night driver has to start from the Base and come back to the Base
- Driver will be paid from the beginning of the shift at the Base to the end of the shift (within the drivers availability hours)
- Safety first: abide by legal regulations, particularly limitations on cumulative driving time

Objective function: Meet customer demands using time and km at best

- The objective is to minimize the Logistic Ratio, which is equal to the time and distance cost divided by the total quantity delivered over the whole horizon (€/km). The time cost is proportional to the duration of the shifts (mainly related to the salary of the driver), and includes the driving time, the idle time, and the loading/unloading times. The distance cost is proportional to the distance traveled by a vehicle (mainly related to the fuel consumption).

Instances sizes

- Instances are within a small/medium scale range: 1-4 trailers/drivers, 50 to 200 Customers, Time Horizons from 1 week up to 1 month

3. PROBLEM SETTING

The problem consists of planning bulk distribution shifts in order to minimize total distribution cost over the long term.

The goal is to build delivery shifts to match the demand requirements subject to given resources and technical constraints in order to **minimize the ratio between the total distribution cost dividing the total quantity delivered over a fixed horizon (*logistic ratio*)**.

3.1. Key definitions

Base	A base is both the starting and the ending location of shifts
Vehicle	A vehicle is the combination of a driver and a trailer.
Delivery	A delivery corresponds to a stop of a vehicle at a customer site: the delivery hose is hooked up to the vessel and product is delivered.
Load	A load corresponds to a stop of a vehicle at a source site, where the Vehicle is filled with products.
Shift	A shift is a chronological list of activities made by a driver during his working period. It starts and ends at the driver's Base. For example, activities would be structured such as "arrive at Base, source, or customer."

3.2. Units of measure for quantity

For quantities, we use the weight (Kg).

3.3. Time representation and horizon

Let us consider the scheduling horizon T . We use two discrete time breakdowns of the interval $[0, T]$: H (Hours) and M (Minutes). Hourly timesteps give the inventory management granularity while minute timesteps give the time granularity for drivers and customers.

3.4. Drivers master data

The set of drivers is referred to as **DRIVERS**. By convention, indices referring to drivers will be denoted by d . It represents a driver with his/her characteristics. Each driver $d \in \mathbf{DRIVERS}$ is defined by:

- $\text{TIMEWINDOWS}(d)$: the set of availability intervals of driver d , each included in $[0, T]$ (in minutes)
- $\text{TIMECOST}(d)$: the cost per working time unit of driver d (in €/minute)
- $\text{MAXDRIVINGDURATION}(d)$: the maximum driving duration for driver d , before ending the shift at the base (in minutes)
- $\text{MININTERSHIFTDURATION}(d)$: the minimum time interval for driver d between 2 shifts (in minutes)
- $\text{TRAILER}(d)$: the only **trailer** which can be driven by the driver d .

3.5.Trailers master data

The set of trailers is denoted as **TRAILERS**. By convention, indices referring to a trailer will be denoted by t . Each trailer $t \in \mathbf{TRAILERS}$ is defined by:

- **DISTANCECOST** (t): the cost per distance unit for trailer t (in €/km).
- **CAPACITY** (t): the capacity of trailer t in mass (Kg). This capacity corresponds to the usable capacity which is the quantity of product that can be loaded in the trailer and delivered to customers. It does not include the minimum quantity of product that must remain in the trailer (to maintain a head of liquid in order to operate the pump) and the volume (i.e., vapor space) that must remain empty for pressure considerations.
- **INITIALQUANTITY** (t): the mass (in Kg) of usable product in trailer t at time 0. Must belong to the interval $[0, \text{CAPACITY}(t)]$.

3.6.Locations master data

A Location denoted p (or alternatively q when the distance is calculated between two locations) may be a base, a source, or a customer:

- **Bases** are the starting and ending locations of any shifts.
- **Sources** and **Customers** are loading and delivery locations respectively. For sources, quantities loaded to the trailers have a negative sign, while for customers, delivered quantities have a positive sign, as do their consumption forecasts.

Since all customers considered in this problem are VMI customers (forecastable customers), Air Liquide forecasts their tank level and decides to deliver product and whenever necessary and cost efficient.

By convention, the numbering of the locations p starts with the bases (e.g. 0,1,2...) then the sources (e.g. 3, 4,5...) and eventually the customers.

For this version 1, we will have only one base (location 0) and one source (location 1) and then the customers (from location 2).

All Locations (Base, Source and customers) common characteristics

- **DISTMATRIX** (p, q): distance between locations p and q (in km);
- **TIMEMATRIX** (p, q): travel time from location p to location q (in km).

Source and customers characteristics

- **SETUPTIME** (p): the fixed part of loading/delivery time for point p (delivery for a customer or loading for a source) (in minutes)

CUSTOMER-only characteristics:

- **SAFETY LEVEL** (p): The level at the customer tank must always be above the Safety Level to avoid product shortage;
- **FORECAST**(p, h): the amount of product in mass (Kg) that is used by the customer at location p during the timestep h ;
- **INITIALTANKQUANTITY**(p): the amount of product in mass (Kg) available in the customers tank at the beginning of the horizon;
- **CAPACITY**(p): the maximal amount of product in mass (kg) that can be delivered to a customer at location p ;
- **ALLOWEDTRAILERS** (p): set of trailers that allowed to supply the customer p .

4. FORMULATION

4.1. Variables

A solution of the problem is a set of shifts (denoted by **SHIFTS**).

The following decision variables are defined on each shift $s \in \mathbf{SHIFTS}$:

- *driver* (s): the driver for shift s
- *trailer* (s): the trailer (and associated tractor) for shift s
- *start* (s): the starting time for the shift s (within $[0, T[$ in minutes)
- *Operations*(s) : it is a list of operations (loadings, deliveries) performed during the shift

Each operation $o \in \mathbf{OPERATIONS}(s)$ is defined by:

- *arrival* (o): the arrival time of operation o (within $[0, T[$ in minutes)
- *point* (o): the location at which operation o takes place either sources or customers
- *quantity* (o): the quantity to be delivered or loaded in operation o . It is negative for loading operations at sources, positive for delivery operations at customers. (in kg)

4.2. Constraints

Notes:

- *all the operators prev, final, start, end... used in this section are defined in the appendix.*
- *The labelling of the constraints corresponds to the Version 1 scope and will be completed once the Version 2 will be released. This labelling is consistent with the labelling available in the checker.*

4.2.1. Bounding constraints

All unary constraints on variables (e.g., non-negativity, inclusion in $[0, T[$, lower and upper bounds specified in previous section) must be satisfied.

4.2.2. Constraints related to drivers

[DRI01 | Inter-Shifts duration]

For each driver d , two consecutive shifts assigned to d must be separated by a duration of $\text{MININTERSHIFTDURATION}(d)$.

For all $d \in \mathbf{DRIVERS}$ For all $s_1, s_2 \in \text{shifts}(d)$,

$\text{Start}(s_2) > \text{end}(s_1) + \text{MININTERSHIFTDURATION}(d)$ OR $\text{start}(s_1) > \text{end}(s_2) + \text{MININTERSHIFTDURATION}(d)$

[DRI03 | Respect of maximal driving time]

For each operation associated with a shift s (including the final operation performed at the base), the cumulated driving time is the total travel time on the shift up to the previous operation plus the travel time from the location of the previous operation to the location of the current operation:

For a given $s \in \mathbf{SHIFTS}$, For all operations $o \in \{\text{Operations}(s)\}$

If operations(s) is not final(s)

$\text{cumulatedDrivingTime}(o) = \text{cumulatedDrivingTime}(\text{prev}(o)) + \text{timeMatrix}(\text{prev}(o), o)$

```

else
    cumulatedDrivingTime(o) = cumulatedDrivingTime(prev(o))+ timeMatrix (o,final(s))
endif
Endfor

```

Then, the constraint formulation is: for each operation, the cumulated driving time of the shift cannot exceed (by law) the maximum allowed driving time. While there are similar regulatory constraints on the driver's continuous working time from the beginning of the shift to the end, the cumulated driving time is generally the limiting constraint – and is therefore the only constraint considered in this simplified problem.

```

For all s ∈ SHIFTS, For all o ∈ Operations(s) and {final(s)},
    cumulatedDrivingTime(o) ≤ MAXDRIVINGDURATION(drivers(s))
EndFor

```

[DRI08] Time windows of the drivers

For each shift s , the interval $[start(s), end(s)]$ must fit in one of the time-windows of the selected driver.

```

For all s ∈ SHIFTS
    It exists at least a tw ∈ TIMEWINDOWS(Drivers(s)), start(s) ≥ start(tw) and end(tw) ≥ end(s)

```

4.2.3. Constraints related to trailers

[TL01 | Different shifts of the same trailer cannot overlap in time]

Consider any two shifts s_1 and s_2 that use the same trailer. Then, either s_1 ends before the start of s_2 or s_2 ends before the start of s_1 .

```

For all tl ∈ TRAILERS For all s1, s2 ∈ shifts(tl),
    start(s2) > end(s1) or start(s1) > end(s2)

```

[TL03] The trailer attached to a driver in a shift must be compatible

For each shift s , the assigned trailer must be the trailer that can be driven by the driver.

```

For all s ∈ SHIFTS, trailer(s) = TRAILER(driver(s))

```

4.2.4. Constraints related to sites

[DYN01 | Respect of tank capacity for each site]

For each site p , the tank quantity at each time step h must be contained into the interval $[0, CAPACITY(p)]$.

```

For all p ∈ CUSTOMERS, For all h ∈ [0, H[
    0 ≤ tankQuantity(p, h) ≤ CAPACITY(p)

```

For each customer p , the tank quantity at each time step h is equal to the tank quantity at time step $h - 1$, minus the forecasted consumption at h , plus all the deliveries performed at h .

Let us remind that the values of $Forecast(s, h)$ and $\sum_{o \in Operations(p, h)} quantity(o)$ are positive for customers. This allows writing the same basic inventory dynamic equation for sources and customers.

```

For all p ∈ { c ∈ CUSTOMERS }
    tankQuantity(p, -1) = INITIALTANKQUANTITY(p)
    For all h ∈ [0, H[

```


Given that $dyn = tankQuantity(p, h-1) - FORECAST(p, h) + \sum_{o \in Operations(p, h)} quantity(o)$
 $tankQuantity(p, h) = \max(dyn, 0)$

Note: We assume that the entire quantity delivered in an operation is available in the customer tank as soon as the truck arrives at the customer (even if the trailer should stay a fixed time to complete the delivery, as explained in constraint SHI03).

4.2.5. Constraints related to shifts

[SHI02 | Arrival at a point requires traveling time from previous point]

For $o \in Operations(s)$, for all shift s ,

$arrival(o) \geq departure(prev(o)) + TIMEMATRIX(prev(o), o)$

Since each shift ends at the base, we need to take in account the travel time from the last operation to the base:

For all shift s ,

$arrival(s) \geq departure(last(Operations(s)) + TIMEMATRIX[point(last(Operations(s)), point(final(s))])$

Note: the waiting time of the drivers (at the gate of a source or a customer) is not defined as such in this model. As a consequence, there is no maximum to this waiting time. Therefore, an arbitrary long "idle time" (where the driver "rests at the door of the customer or source doing nothing") can precede any operation belonging to a shift, as long as all of the constraints are respected.

[SHI03 | Loading and delivery operations take a constant time]

$departure(o) = arrival(o) + SetupTime(point(o))$

[SHI05 | delivery operations require the customer site to be accessible for the trailer]

For all $s \in SHIFTS$, $o \in Operations(s)$: if $point(o) \in CUSTOMERS$ then $trailer(s) \in ALLOWEDTRAILERS(point(o))$

[SHI06 | trailerQuantity cannot be negative or exceed capacity of the trailer]

For a given shift $s \in SHIFTS$, For all $o \in Operations(s)$ with $\{final(s)\}$,

$trailerQuantity(o) = trailerQuantity(prev(o)) - quantity(o)$

$trailerQuantity(o) \geq 0$

$trailerQuantity(o) \leq CAPACITY(trailer(s))$

[SHI07 | Initial quantity of a trailer for a shift is the end quantity of the trailer following the previous shift.]

$endTrailerQuantity(s) = trailerQuantity(last(Operations(s)))$

From these definitions we can derive this constraint:

if $(s = first(shifts(trailer(s))))$

$startTrailerQuantity(s) = INITIALQUANTITY(trailer(s))$

else

$startTrailerQuantity(s) = endTrailerQuantity(prev(s, shifts(trailer(s))))$

endif

[SHI11 | Some product must be loaded or delivered]

For each source, some product (a negative quantity) must be loaded and for each customer, the some product (a positive quantity) must be delivered.

For all $p \in \text{CUSTOMERS}$, For all $h \in [0, H-1]$, For all $o \in \text{Operations}(p, h)$
 $quantity(o) > 0$

For all $p \in \text{SOURCES}$, For all $h \in [0, H-1]$, For all $o \in \text{Operations}(p, h)$
 $quantity(o) < 0$

4.2.6. Constraints related to quality of service

[QS02 | Run-out avoidance]

For each VMI customer p , the tank level must be maintained at a level greater than or equal to the safety level $SafetyLevel(p)$, at all times.

For all $p \in \text{CUSTOMERS}$, For all $h \in [0, H-1]$
 $SafetyLevel(p) \leq tankQuantity(p, h)$

4.3. Optimization goal

4.3.1. Objective function

The goal is to minimize the distribution costs required to meet customer demands for product over the long term horizon. In order to tend to that final goal, we will minimize the **logistic ratio**.

The logistic ratio is defined as the total cost of the shifts divided by the total quantity delivered on those shifts:

$$LR = \frac{\sum_{s \in SHIFTS} Cost(s)}{TotalQuantity}$$

The cost of a shift represents the **distribution costs** related to the shift, including:

- the **distance cost**, applied to the total length of the shift, which is generally related to the trailer used (covering fuel consumption and maintenance)
- the **time cost** applies to the total duration of the shift which is generally related to the driver (covering the driver salary and charges)

$$Cost(s) = (DISTANCECOST(trailer(s)) \times travelDist(s) + TIMECOST(driver(s)) \times (end(s) - start(s)))$$

Where:

$$travelDist(s) = \sum_{o \in OPERATIONS(s)} DISTMATRIX(prev(o), o) + DISTMATRIX(last(OPERATIONS(s)), poin(final(s)))$$

The total quantity delivered over all shifts is computed as follows:

$$TotalQuantity = \sum_{s \in SHIFTS} \sum_{\substack{o \in Operation(s) \\ quantity(o) > 0}} quantity(o)$$

Note that if $TotalQuantity = 0$, we will consider also $LR = 0$.

4.3.2. Time integrations over scheduling optimization horizons

The objective of an IRP is to minimize distribution costs over an enough long period (horizon) covering one or more replenishment cycles for all the customers.

If the time horizon is too short, the optimization can be short-sighted because of the “end-of-period side effect”: customers that do not strictly require delivery within the horizon, which might increase the risk of shortage just beyond the horizon.

Considering a long period makes the relative impact of this effect negligible, but complexifies the problem and requires longer forecast on the customers, which could be far from the reality.

A “good” value for the time horizon cannot be generalized as it depends on many factors, particularly the “confidence” in the customer forecast at various points over that horizon. For this reason, in this challenge, we ask the contenders to try different horizons, even for the same dataset.

4.4.APPENDIX : Specific Functions definition used in the constraint modeling

For the sake of readability, we use the following notations or intermediate variables.

- $shifts(d) := \{s \in Shifts \text{ such that } driver(s) = d\}$, i.e., the list of shifts that use driver d .
- $shifts(tl) := \{s \in Shifts \text{ such that } trailer(s) = tl\}$, i.e., the list of shifts that use trailer tl .
- $operations(s)$ = the union of all operations on the tours of this shift, including initial and final operation of each tour.
- $operations(tl)$ = the union of all operations of each trailer tl sorted by time.

3 generic functions are defined and used to navigate within any non empty lists:

- $first(list)$ returns the first item in the list
- $last(list)$ returns the last item of the list.
- $prev(e, list)$ returning the previous item e in list **list** (or e if e is the first item of the list)

Several customized functions can be used:

- $prev(o)$: For each operation o we define $prev(o)$ as the predecessor of o in the list of operations of the $prev(o) := shift(prev(o, Operations(shift(s))))$. If o is the first operation of the shift then we extend this definition by defining $prev(o)$ as a “virtual” first operation f with the following characteristics
 - $shift(f) := shift(o)$
 - $point(f) := Base(shift(o))$
 - $trailerQuantity(f) = startTrailerQuantity(shift(o))$
 - $quantity(f) := 0$
 - $arrival(f) := start(shift(o))$
 - $cumulatedDrivingTime(f) := 0$
 - $departure(f) := departure(shift(o))$
- $final(s)$: a final operation is created at the end of the shift with the same attributes as $operations$ defined above but with null volume. This last stop is only defined in order to simplify the subsequent writing of constraints by making the number of operations equal to the number of arcs (travel from one point to another) in the shift. It hosts variables related to the last travel from the last operation point to the base. The following attributes of this operation have constrained values:
 - $point(final(s)) = base(s)$
 - $quantity(final(s)) = 0$
 - $arrival(final(s)) = arrival(s)$
 - $departure(final(s)) = end(s)$
- For any (delivery or load) point p and timestep h , the set of all Operations at point p taking place during timestep h is defined by:
 $Operations(p, h) := \{ o \in Operations \text{ such that } point(o) = p \text{ and } round(arrival(o) / U) = h \}$

Finally the start and end of a time window tw are referred to as $start(tw)$ and $end(tw)$ respectively.