An Exact Algorithm for the Petrol Station Replenishment Problem

Fabien Cornillier

Fayez F. Boctor

Gilbert Laporte

and

Jacques Renaud

August 2005

Working Paper DT-2005-JR-2

Network Organization Technology Research Center (CENTOR),

Université Laval, Québec, Canada

An Exact Algorithm for the Petrol Station Replenishment Problem

Fabien Cornillier [a], Fayez F. Boctor [a],
Gilbert Laporte [b] and Jacques Renaud [a]

[a]*Network Organisation Technology Research Center (CENTOR)*
Faculté des sciences de l'administration
Université Laval
Québec, Canada G1K 7P4
Email addresses:
`fabien.cornillier@centor.ulaval.ca`
`fayez.boctor@fsa.ulaval.ca`
`jacques.renaud@fsa.ulaval.ca` *(corresponding author)*

[b]*Canada Research Chair in Distribution Management*
HEC Montréal
3000 chemin de la Côte-Sainte-Catherine
Montréal, Canada H3T 2A7
Email address: `gilbert@crt.umontreal.ca`

Abstract

In the *Petrol Station Replenishment Problem* (PSRP) the aim is to jointly determine an allocation of petroleum products to tank truck compartments and to design delivery routes to stations. This article describes an exact algorithm for the PSRP. This algorithm was extensively tested on randomly generated data and on a real-life case arising in Eastern Quebec.

Keywords: replenishment, loading, vehicle routing, assignment, matching, column generation

Introduction

The purpose of this article is to develop an exact algorithm for the *Petrol Station Replenishment Problem* (PSRP). A truncated version of this algorithm can also be used as a heuristic. The problem is motivated by the situation prevailing in the Province of Quebec (for which statistics were made available to us), but it also applies to several other contexts. In Quebec, more than seven billion litres of fuel (petrol and diesel) are distributed yearly to approximatively 5 000 stations. Distribution costs account for a percentage of the sales price varying between 0.7% and 9.3%, depending on the region, for an average of 2.8%, representing 50 million dollars a year[1].

In North America, most petroleum companies subcontract their distribution operations to private regional transporters who receive an amount varying between a few tenths of cent to slightly more than a cent for each litre delivered. Transporters use tank trucks made up of a tractor and of one or two trailers divided into compartments. Each truck contains from three to six compartments whose capacities vary between 5 000 and 16 000 litres. The total capacity of a truck varies between 43 000 and 59 000 litres depending on the number of axles. Each petrol station uses between three and five underground tanks of standard capacities (22 700, 25 000, 31 800, 35 000, 45 400 or 50 000 litres).

The replenishment of petrol stations is carried out from refineries or, more rarely, from intermediate depots. The PSRP consists of determining least cost delivery routes to a set of stations which must be supplied once by a heterogeneous fleet of vehicle, subject to a number of constraints. Delivery costs are made up of a term proportional to mileage and of a vehicle-dependent fixed portion. Constraints specify that the quantity of each product should be sufficient to fulfill the entire demand (possibly including a safety stock), but no more than 95% of the petrol station tank capacities may be filled[2]. Also, no more than 75% to 80% of the vehicle capacity may be used during the thawing season. Compartments are not equipped with a flow meter, which implies that they must be entirely emptied once replenishment has started. Because it is sometimes necessary to use the content of two compartments to fill a tank, and stations generally require two or three products, the number of stations visited by a truck on any given trip will very rarely exceed two. In addition, the front part of each trailer must be emptied last to ensure more stability when driving. Finally, a limit is imposed on the duration of any trip.

The PSRP differs from most vehicle routing problems (see e.g. Toth and Vigo[3]) because of the presence of compartments which can only hold one product, and the absence of flow meters which means that the content of a compartment cannot be split between stations. In a sense the PSRP is more complicated than standard routing problems but, at the same time, the limit of two visits per trip leads to an interesting simplification which we will exploit. As far as we are aware, no previous article has addressed this particular problem, but related studies exist. Brown and Graves[4] have considered the planning of single-customer trips in the presence of time windows, while Brown et al.[5] have developed a computerized assisted dispatch system for a problem similar to ours. The system combines human solutions with heuristics to assist real-time decision making. Finally, Malépart et al.[6] have proposed a number of simple heuristics to handle a general petrol distribution problem with multiple delivery trips and workforce management constraints.

The remainder of this article is organized as follows. A first section provides a mathematical model for the PSRP. Then we develop an exact algorithm for the case when at most two stations are visited on any trip. This is followed by computational results and conclusions.

Mathematical model

Let $V = \{1, \ldots, n\}$ be the set of stations to be visited and define a symmetric travel cost matrix on V^2. The minimal and maximal demand of each product at each station are known. All stations require a visit but the minimum demand for some products may be zero. An unlimited heterogeneous fleet of vehicles is available and there always exists a vehicle in which the minimum demand of any station can fit. Thus only one visit is necessary for each station.

Let $S \subseteq V$ be a subset of stations, and $K(S)$ be the set of vehicles capable of delivering the minimal demand of the stations of S. Then the cost $d_{S,k}$ of using a vehicle $k \in K(S)$ to visit all stations of S can be computed as a *Traveling Salesman Problem with Precedence Constraints* (TSPPC). These constraints are dictated by the necessity to empty the front of each trailer last. Let $k^* \in K(S)$ be the vehicle yielding the least TSPPC cost. Given a set S of stations and a vehicle $k \in K(S)$, it is generally preferable to fill the vehicle as much as possible without exceeding the compartment capacities and the station maximal demands.

The PSRP can now be formulated as a *Set Partitioning Problem* (SPP). Let x_S be a binary variable equal to 1 if and only if all stations of S are served by the same vehicle. The formulation is then:

$$\text{(SPP)} \qquad \text{Minimize} \sum_{S \subseteq V, S \neq \emptyset} d_{S,k^*} x_S \tag{1}$$

$$\text{subject to:} \qquad \sum_{S: i \in S} x_S = 1 \qquad (i \in V) \tag{2}$$

$$x_S = 0 \text{ or } 1 \qquad (S \subseteq V, S \neq \emptyset). \tag{3}$$

Solving the SPP optimally is impossible for all but trivial cases because of the large number of subsets S, and of the difficulty of determining k^* and d_{S,k^*}. However, an exact solution methodology can be developed for the special case where $|S| \leq 2$, which corresponds to current practice, as discussed in the introduction.

The Tank Truck Loading Problem

We now address the more difficult *Tank Truck Loading Problem* (TTLP) which consists of optimally assigning the demand of a set S of stations to a given vehicle $k \in K(S)$. The TTLP can be shown to be NP-hard by using the same argument as Smith[7] for the multiple inventory loading problem. More precisely, the TTLP is defined as follows. Let the tanks of all stations of S be indexed by t ($t \in \{1, \dots, T\}$). This index does not contain any information on the stations, so that the difficulty of the TTLP does not depend on $|S|$, but on the total number of underground tanks associated with S. Let the compartments of vehicle k be indexed by c ($c \in \{1, \dots, C\}$). Also define the constants:

s_t	the initial inventory level of tank t;
P_t	the usable capacity of tank t;
m_t	the minimum inventory level of tank t required to fulfill the demand for the planning horizon;
a_t	the minimum delivery for tank t: $a_t = \max\{0, m_t - s_t\}$;
b_t	the maximum delivery for tank t: $b_t = P_t - s_t$;
Q_c	the capacity of compartment c;

and the variables:

x_t	the amount delivered to tank t;
y_{tc}	a binary variable equal 1 if and only if compartment c is used to deliver the demand of tank t.

The TTLP is then formulated as follows:

$$\text{(TTLP)} \qquad \text{Maximize } \sum_{t=1}^{T} x_t \tag{4}$$

$$\text{subject to:} \qquad a_t \le x_t \le b_t \qquad (t \in \{1, \ldots, T\}) \tag{5}$$

$$x_t \le \sum_{c=1}^{C} Q_c y_{tc} \qquad (t \in \{1, \ldots, T\}) \tag{6}$$

$$\sum_{t=1}^{T} y_{tc} \le 1 \qquad (c \in \{1, \ldots, C\}) \tag{7}$$

$$y_{tc} = 0 \text{ or } 1 \qquad (t \in \{1, \ldots, T\}; c \in \{1, \ldots, C\}). \tag{8}$$

In this formulation the objective function (4) maximizes the total delivered quantity. Constraints (5) impose bounds on the amounts delivered. Constraints (6) specify that the delivery amount associated with tank t does not exceed the alloted compartment capacity. By constraints (7) at most one demand can be assigned to any compartment.

Note that this problem differs from related problems studied by Christofides et al.[8] and Smith[7]. In the loading problem described by Christofides et al.[8], there is only one liquid product and the objective function is to minimize the number of used compartments. These authors have also studied the unloading problem where a demand quantity may be unloaded from several tanks and a tank may be only partially unloaded. A value is associated with each compartment and the objective is to minimize the value of all used compartments. Smith[7] deals with the multiple inventory loading problem where a holding cost per unit of volume is associated to each product and a fixed cost is incurred for each delivery. Each demand correspond to a fixed number of units. In this problem, the objective is to minimize an aggregate objective function containing both delivery and storage costs, subject to the restrictions of determining a feasible loading arrangement within the vehicle.

Exact algorithm for the Tank Truck Loading Problem

We have devised the following exact algorithm for the TTLP. A first test is conducted in order to quickly identify some classes of infeasible instances (Step 1), and an attempt is then made to identify a feasible solution by means of a sequential allocation process (Steps 2-5). If this process fails, the TTLP is solved by means of a standard *Integer Linear Programming* (ILP) algorithm (Step 6). If a feasible allocation is known to exist an attempt is made to identify an even better solution by solving an *Assignment Problem* (AP) (Step 3), and by then applying an improvement step (Step 4). A test is then applied to check whether the solution is optimal (Step 5). If this is the case the algorithm terminates with a feasible and optimal solution ; otherwise the ILP solver is applied (Step 6).

Step 1 (Feasibility test)

Let T^+ be the number of tanks for which $a_t > 0$ and T^s the number of tanks that must be split between several compartments, i.e., those tanks for which $a_t > \max\{Q_c\}$. If $T^+ + T^s > C$ or $\sum_{t=1}^{T} a_t > \sum_{c=1}^{C} Q_c$, then no feasible solution exists : stop.

Step 2 (Sequential assignment)

Sort the tanks in non-increasing order of the a_t and break ties by non-increasing order of the b_t; sort the compartments in non-increasing order of the Q_c. Iteratively assign the minimal demand a_t of each tanks to the next unused compartment. If a_t exceeds the capacity Q_c of the compartment being considered, split this demand into two demands t' and t'' with $a_{t'} = b_{t'} = Q_c$ and $a_{t''} = a_t - Q_c$, $b_{t''} = b_t - Q_c$. Then assign t' to compartment c and insert t'' in its appropriate position in the list and set $T := T + 1$ (increase the number of demands by one). If some demands cannot be assigned to a compartment through this process, go to Step 6.

Step 3 (Assignment algorithm)

Demands for which $a_t > 0$ (including split demands) must then be assigned to compartments in order to minimize the total unused capacity. The assignment costs e_{tc} are defined as $e_{tc} = \infty$ if $a_t > Q_c$, and $e_{tc} = \max\{0, Q_c - b_t\}$ otherwise. If $T < C$, create $C - T$ dummy demands t with $e_{tc} = Q_c$ for all c. If in the solution of the assignment problem t is assigned to c and $b_t > Q_c$, then define a new demand $\tilde{t}$ with $a_{\tilde{t}} = 0$ and $b_{\tilde{t}} = \max\{0, b_t - Q_c\}$. For each demand for which $a_t = 0$, define a new demand with $a_{\tilde{t}} = 0$ and $b_{\tilde{t}} = b_t$.

Step 4 (Assignment of remaining demands)

If all compartments have been used or all demands have been assigned, go to Step 5. Consider all non-assigned demands $\tilde{t}$ with $b_{\tilde{t}} > 0$. Iteratively assign the largest $b_{\tilde{t}}$ to the largest unused compartment c available, and set $b_{\tilde{t}} := \max\{0, b_{\tilde{t}} - Q_c\}$. Repeat this operation as long there remain positive $b_{\tilde{t}}$ and unused compartments.

Step 5 (Optimality test)

The solution is optimal and the algorithm terminates whenever any of the following conditions is satisfied: 1) all compartments are full, 2) all maximal demands have been assigned to a compartment, 3) $T^+ = C$, or 4) there exists a unique tank t with a demand $a_t > 0$ completely filling the largest $C - 1$ compartments and part of the smallest compartment, i.e., $T^+ = 1$ and $\sum_{c=1}^{C-1} Q_c < a_t \leq \sum_{c=1}^{C} Q_c$. Condition 4 simply states that if there is only one positive demand and if that demand have to use all compartments, its assignment is optimal.

Step 6 (ILP solution)

Solve the TTLP by means of an ILP solver. In this algorithm Step 6 always terminates with an optimal solution if it is entered from Step 5. However, it may terminate with an infeasible solution if it entered from Step 2. A simple TTLP heuristic consists of eliminating Step 6, which increases the risk of ending with an infeasible or suboptimal solution.

Solving the routing problem for $|S| \leq 2$

Two distinct strategies can be applied to solve the routing problem when $|S| \leq 2$. First observe that the number of non-empty feasible subsets is at most $(n^2 + n)/2$. As a result, all cases can readily be enumerated. *Strategy 1* consists of solving the TTLP for each S, and values of k in non-decreasing order of fixed costs until a feasible solution is obtained for vehicle k^*. The value of d_{S,k^*} is then readily determined. Because $|S| \leq 2$, the SPP (1)-(3) reduces to a *Matching Problem* (MP) over V (with possible self-matchings), with matching costs d_{S,k^*}, as shown by Christofides[9]. Under this strategy, the TTLP is solved $(n^2 + n)/2$ times, once for each set S.

Strategy 2 is based on a column generation scheme. Initially the least fixed cost vehicle is assigned to each set S and the MP is solved. A test is then performed to check TTLP feasibility on each of the selected routes. If all routes are feasible, the algorithm ends. Otherwise the cheapest feasible vehicle k^* is determined for each set S for which the TTLP was infeasible, and the MP is solved again with the new matching costs. This procedure is iterated until a feasible solution has been reached. Under the second strategy, more MPs may have to be solved but the number of calls to the TTLP is likely to be much less than under the first strategy.

A preprocessing step applicable to both strategies is to eliminate sets S which are *a priori* infeasible. These are sets for which the trip duration exceeds the prescribed limit and those for which Step 1 of the TTLP algorithm concludes that no feasible solution exists.

Numerical example

This section describes an example with nine stations and two products. Table 1 gives the coordinates of each station i (0 is the depot) as well as their minimum and maximum demands a_{ip} and b_{ip} for each of the two products p. Table 2 shows the travel costs c_{ij} between stations i and j (equal to the Euclidian distances), and Table 3 describes the routing cost of all possible routes visiting one or two stations. We use four compartments with capacities $Q_1 = 7$, $Q_2 = 3$, $Q_3 = 2$ and $Q_4 = 1$. This example shows that the algorithm can easily handle fixed demands as only stations 5 and 9 have different minimum and maximum demands.

Table 1. Station coordinates, minimum and maximum demands

i	x_i	y_i	a_{i1}	b_{i1}	a_{i2}	b_{i2}
0	3	2	-	-	-	-
1	3	5	1	1	3	3
2	0	4	2	2	7	7
3	2	4	1	1	3	3
4	3	3	2	2	4	4
5	5	3	0	1	6	8
6	5	2	1	1	8	8
7	1	1	0	0	2	2
8	3	0	3	3	7	7
9	4	0	0	2	4	4

We present in Table 4 TTLP data associated with $S = \{5, 9\}$, where (a_1, b_1) corresponds to $(a_{5,1}, b_{5,1})$, (a_2, b_2) to $(a_{5,2}, b_{5,2})$, (a_3, b_3) to $(a_{9,1}, b_{9,1})$, and (a_4, b_4) to $(a_{9,2}, b_{9,2})$. In Step 1 of the TTLP algorithm, the feasibility test is successful because $T^+ + T^s < C$ and $\sum_{t=1}^{T} a_t < \sum_{c=1}^{C} Q_c$.

Table 2. Travel cost matrix between stations

	j = 0	1	2	3	4	5	6	7	8	9
i = 0	0.00	3.00	3.61	2.00	1.00	2.24	2.00	2.24	2.00	2.24
1		0.00	3.16	1.00	2.00	2.83	3.61	4.47	5.00	5.10
2			0.00	3.00	3.16	5.10	5.39	3.16	5.00	5.66
3				0.00	1.00	2.24	2.83	3.61	4.00	4.12
4					0.00	2.00	2.24	3.83	3.00	3.16
5						0.00	1.00	4.47	3.61	3.16
6							0.00	4.12	2.83	2.24
7								0.00	2.24	3.16
8									0.00	1.00
9										0.00

Table 3. Cost of routes containing stations i and j

	j = 1	2	3	4	5	6	7	8	9
i = 1	6.00	9.77	6.00	6.00	8.06	8.61	9.71	10.00	10.30
2		7.21	8.61	7.77	10.90	11.00	9.00	10.60	11.50
3			4.00	4.00	6.47	6.83	7.84	8.00	8.36
4				2.00	5.24	5.24	6.06	6.00	6.40
5					4.47	5.24	8.94	7.84	7.63
6						4.00	8.36	6.83	6.47
7							4.47	6.47	7.63
8								4.00	5.24
9									4.47

Table 4. Tanks in non-decreasing order of the a_t and compartments in non-decreasing order of the Q_c

Tanks	Compartments
$(a_2, b_2) = (6, 8)$	$Q_1 = 7$
$(a_4, b_4) = (4, 4)$	$Q_2 = 3$
$(a_1, b_1) = (0, 1)$	$Q_3 = 2$
$(a_3, b_3) = (0, 2)$	$Q_4 = 1$

Table 5. Sequential assignment

Tanks	Compartments
$(a_2, b_2) = (6, 8)$	$Q_1 = 7$
$(a_{4'}, b_{4'}) = (3, 3)$	$Q_2 = 3$
$(a_{4''}, b_{4''}) = (1, 1)$	$Q_3 = 2$
$(a_1, b_1) = (0, 1)$	$Q_4 = 1$
$(a_3, b_3) = (0, 2)$	

In Table 4 tanks are sorted in non-increasing order of the a_t and compartments in non-increasing order of the Q_c. We split the second demand as it exceeds the second compartment capacity, and we now have five demands (Table 5). We then assign a compartment to each demand in order to minimize the total unused capacity (Step 3). Figure 1 depicts the associated bipartite graph with costs e_{tc} on which the assignment algorithm is applied. The bold edges are those which are selected.

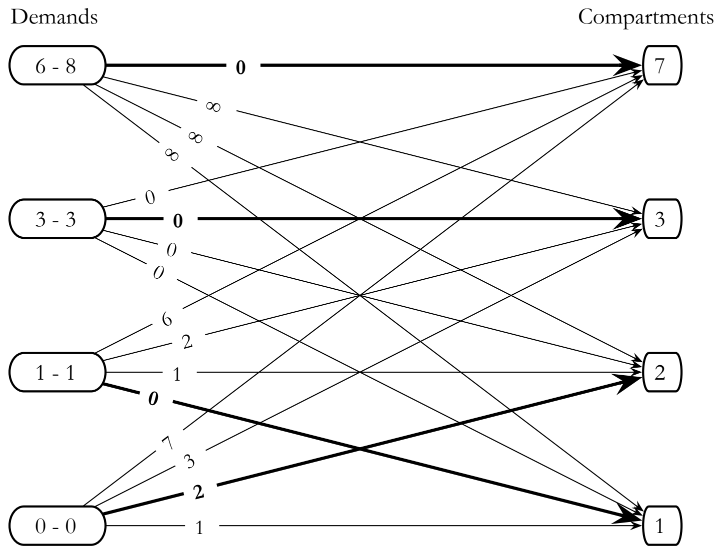

Figure 1. Graph of the assignment problem

Table 6. Assignment of residual demands

Demands	Compartments
$(a_{\tilde{3}}, b_{\tilde{3}}) = (0, 2)$	$Q_3 = 2$
$(a_{\tilde{2}}, b_{\tilde{2}}) = (0, 1)$	
$(a_{\tilde{1}}, b_{\tilde{1}}) = (0, 1)$	
$(a_{\tilde{4}}, b_{\tilde{4}}) = (0, 0)$	

As there are residual demands (one unit of non-satisfied demand for tank 2 and three units for tank 3), and one unused compartment (Table 6), we assign demand 3 to compartment 3 (Step 4). All compartments are now full and this solution passes the optimality test (Step 5). There is therefore no need to solve the TTLP by means of an ILP solver.

Using this algorithm, we are able to determine whether a route is feasible or not. Some sets are eliminated by means of the preliminary test ; this is the case for the route visiting stations 1 and 8 for which $\sum_{t=1}^{T} a_t > \sum_{c=1}^{C} Q_c$, and for the route visiting stations 1 and 6 for which $T^+ + T_S > C$. One demand must be split as $a_{6,2} > \max\{Q_c\}$. Other routes (like the route visiting stations 8 and 9) are eliminated only after solving the ILP in Step 6.

Solving the TTLP on all possible combination of stations leads to the elimination of the following sets : {1,6}, {1,8}, {2, 4}, {2,5}, {2,6}, {2,8}, {3, 6}, {3, 8}, {4, 6}, {4, 8}, {5, 6}, {5, 8}, {6, 8}, {6, 9}, and {8, 9}. These infeasible sets are shaded in Table 3. The cost of each of these combinations is set to ∞ and the matching is found on the resulting cost matrix. The resulting distribution plan is {1,2}, {3,4}, {5,9}, {6.6}, and {7,8}, with a total cost of 31.9.

Computational results

The algorithms just described were coded in Objective-C and run on an Apple iBook G3 700Mhz computer. The MPs were solved with an implementation of Gabow's version[10] of Edmonds's algorithm[11]. The ILPs in Step 6 of the TTLP algorithm were solved by means of GLPK 4.2 (GNU Linear Programming Kit: http://www.gnu.org/software/glpk/glpk.html).

We present three sets of tests. We have first evaluated the performance of the TTLP algorithm. Second, we have tested the complete algorithm for the PSRP. Finally, we have solved a real case provided by a local distributor.

Results for the TTLP algorithm

We have first randomly generated instances with 5 demands and 5 compartments, under several values of the ratios T^+/C, $R_a = \sum_{t=1}^{T} a_t / \sum_{c=1}^{C} Q_c$ and $R_b = \sum_{t=1}^{T} b_t / \sum_{c=1}^{C} Q_c$ which significatively affect problem difficulty. To assess the behavior of the TTLP algorithm, we have solved a total of 8000 instances of the TTLP: 100 for each combination of T^+, R_a and R_b with $T^+ \in \{1, \ldots, 4\}$,

$R_a \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, and $R_b \in \{1.0, 1.5, 2.0, 2.5\}$. We used a tank truck with five compartments: $Q_1 = 15500$, $Q_2 = 5500$, $Q_3 = 5500$, $Q_4 = 9000$ and $Q_5 = 14500$. Each instance was solved by using the TTLP model (4)-(8) to ensure its feasibility and to determine its optimal solution value.

To create an instance, we first generate the a_t values such that $\sum_{t=1}^{T} a_t / \sum_{c=1}^{C} Q_c$ equals a given constant R_a:

(1) choose $T^+ - 1$ random numbers h_t from a discrete uniform distribution $U(0, R_a \sum_{c=1}^{C} Q_c)$;
(2) sort these numbers in non-decreasing order;
(3) set $a_1 = h_1$, $a_t = h_t - h_{t-1}$ for all $t \in \{2, \ldots, T^+ - 1\}$, and $a_{T^+} = \sum_{c=1}^{C} Q_c - h_{T^+-1}$;
(4) set $a_t = 0$ for all $t \in \{T^+ + 1, \ldots, T\}$.

We are now able to generate the b_t values in such a way that $\sum_{t=1}^{T} b_t / \sum_{c=1}^{C} Q_c$ equals a given constant R_b, and $b_t \geq a_t$ for all $t \in \{1, \ldots, T\}$:

(1) choose $T - 1$ random numbers g_t from a uniform distribution $U\left(0, (R_b - R_a) \sum_{c=1}^{C} Q_c\right)$;
(2) sort these numbers in non-decreasing order;
(3) set $b_1 = a_1 + g_1$, $b_t = a_t + g_t - g_{t-1}$ for all $t \in \{2, \ldots, T-1\}$, and $a_T = a_T + (R_b - R_a) \sum_{c=1}^{C} Q_c - g_{T-1}$.

The results are reported in Tables 7 and 8. The column headings are as follows:

T^+	number of demands for which $a_t > 0$;
R_a	ratio of the sum of the minimal demands to the total capacity of vehicle: $\sum_{t=1}^{T} a_t / \sum_{c=1}^{C} Q_c$;
R_b	ratio of the sum of the maximal demands to the total capacity of vehicle: $\sum_{t=1}^{T} b_t / \sum_{c=1}^{C} Q_c$;

Feasible Steps 1-5	number of instances solved by means of the heuristic part of the TTLP algorithm (Steps 1 to 5);
Optimal Steps 1-5	number of instances that were actually optimal after Step 5 of the TTLP algorithm (this is known because each instance was optimally solved during the generation process);
Optimal proven Steps 1-5	number of instances for which a provably optimal solution was determined by Step 5 of the TTLP algorithm;
Average optimality gap	average deviation of the heuristic solutions value (Steps 1 to 5) from the optimum;
%Capacity	average vehicle capacity used in the solution;
Seconds Heuristic	average time in seconds required for the resolution by the heuristic;
Seconds ILP	average time in seconds required for the resolution by the ILP solver.

Tests were performed for various combinations of T^+, R_a and R_b. For the sake of conciseness, we only report in Table 7 extensive results for the case $R_b = 1.5$ which appears to be the most realistic value. Average statistics computed over the 2000 instances are reported for all values of R_b in Table 8. Results reported in Table 7 indicate that the TTLP heuristic (Step 1 to 5) identifies a feasible solution in 95.9% of all cases, and a proven optimum 43.0% of the time. We know that the average percentage of optimal solutions after Step 5 is in fact 61.45%. This means that the TTLP heuristic identifies 69.97% (43.0/61.45) of the optimal solutions. The average optimality gap after Step 5 is only 2.09% and computation times per instance are insignificant. The high values in the column %Capacity indicate that our instances are tightly constrained and our solutions make good use of compartment capacity. Average results reported in Table 8 show that similar conclusions extend to other values of R_b (except for the case $R_b = 1$ where optimality can rarely be proven after Step 5).

Results for the PSRP

In the second series of tests, we have solved the PSRP under the two strategies described for the solution of the routing problem. We have also solved instances of the PSRP with an homogeneous and an heterogeneous fleet. The tank and truck characteristics, station demands and distances were randomly generated in a way that reflects real-life situation which served as a basis for the study. For each station, we chose a total consumption of three products equals to $M = \sum_{t=1}^{3} m_t$ (litres per day) from a discrete uniform distribution $U(10\,000, 50\,000)$. More specifically, $m_1 = 0.7M$, $m_2 = 0.1M$ and $m_3 = 0.2M$. The initial inventory levels s_t were chosen from a discrete uniform distribution $U(0, P_t)$, with $P_1 = 35\,000$, $P_2 = 25\,000$ and $P_3 = 25\,000$. The depot and stations coordinates were uniformly generated in a 100×300 Euclidian space (distances are symmetric and Euclidian). We assume that all stations must be replenished, so we only retained those for which there was at least one strictly positive minimal demand a_t. We generated 30 problems with $n = 50, 100$ and 200. No limit was imposed on the length of vehicle routes.

We also considered three types of tank trucks with four or five compartments:

- Type 1: five compartments with capacities 16 000, 16 000, 10 000, 6 000, and 6 000 litres (total: 54 000 litres);
- Type 2: five compartments with capacities 15 500, 5 500, 5 500, 9 000, and 14 500 litres (total: 50 000 litres);

Table 7. Computational results for the TTLP with $R_b = 1.5$

T^+	R_a	Feasible Steps 1-5	Optimal Step 1-5	Optimal proven Step 5	Average optimality gap	%Capacity	Seconds Heuristic	Seconds ILP
1	0.1	100	49	49	5.51	94.47	0.000	0.010
1	0.3	100	100	96	0.00	99.98	0.000	0.007
1	0.5	100	66	66	1.77	98.20	0.000	0.008
1	0.7	100	66	66	1.41	98.56	0.000	0.004
1	0.9	100	100	100	0.00	97.10	0.000	0.003
2	0.1	100	34	16	6.49	92.52	0.000	0.007
2	0.3	100	42	36	3.81	95.52	0.000	0.008
2	0.5	100	51	43	3.09	96.09	0.000	0.010
2	0.7	100	62	50	1.16	97.87	0.000	0.009
2	0.9	88	59	42	0.96	97.81	0.000	0.005
3	0.1	100	31	14	5.26	92.54	0.000	0.013
3	0.3	100	46	24	3.31	94.94	0.000	0.013
3	0.5	100	53	39	2.06	96.51	0.000	0.015
3	0.7	100	72	42	0.91	97.42	0.000	0.017
3	0.9	65	57	32	0.40	98.32	0.000	0.010
4	0.1	100	56	25	2.54	92.79	0.000	0.023
4	0.3	100	65	27	1.71	95.13	0.000	0.019
4	0.5	100	78	34	0.81	96.29	0.000	0.020
4	0.7	97	80	26	0.52	96.05	0.000	0.025
4	0.9	67	62	32	0.12	98.89	0.000	0.011
Average:		95.85	61.45	43.0	2.09	96.35	0.000	0.012

Table 8. Aggregate computational results for the TTLP for differents values of R_b

R_b	Feasible Steps 1-5	Optimal Step 1-5	Optimal proven Step 5	Average optimality gap	%Capacity	Seconds Heuristic	Seconds ILP
1	95.85	53.60	8.05	2.27	90.00	0.000	0.037
1.5	95.85	61.45	42.95	2.09	96.35	0.000	0.012
2	95.85	72.55	57.35	1.31	97.65	0.000	0.011
2.5	95.85	77.40	65.25	0.98	98.24	0.000	0.010

– Type 3: four compartments with capacities 15 000, 15 000, 10 000, and 10 000 litres (total: 50 000 litres).

For the homogeneous fleet case, we chose the type 1 tank truck. Results are reported in Table 9 under the following column headings:

Nb. vehicle types	number of vehicle types (1 corresponds to an homogeneous instance; 3 corresponds to an heterogeneous instance);
Nb. Stations	number of stations to replenish;

The following headings are averages over 30 instances:

TTLPs solved	number of TTLPs solved;
Infeasible TTLPs	number of infeasible TTLPs;
Seconds	time in seconds required for the resolution of the PSRP;
Sets eliminated	number of infeasible TTLPs eliminated in the preprocessing step under strategy 2;
MPs solved	number of matching problems solved under strategy 2.

Results presented in Table 9 indicate that the second algorithmic strategy for the routing problem is far superior to the first. The preprocessing step leads to the elimination of over 21.09% of all station sets, representing 98.51% of the infeasible sets. As a result very few matching problems have to be solved. The computing time of the second strategy is about 0.28% of the first. This is mostly due to the fact that very few TTLPs are solved under the second strategy.

The comparison between the homogeneous and the heterogeneous fleet cases reveals that the second type of problem is more difficult for the first strategy but

Table 9. Behavior of the routing algorithms

		Strategy 1			Strategy 2			
Nb. vehicle types	Nb. Stations	TTLPs solved	Infeasible TTLP	Seconds	Sets eliminated	MPs solved	TTLPs solved	Seconds
1	50	1 275	293	44.8	287	1.37	27	0.0
1	100	5 050	1 108	212.3	1 088	1.37	53	0.5
1	200	20 100	4 399	723.7	4 304	2.07	220	4.8
3	50	1 844	285	137.4	281	1.10	26	0.0
3	100	7 173	1 061	483.1	1 048	1.30	52	0.5
3	200	28 443	4 171	1 964.1	4 140	1.37	103	4.1

easier for the second. This can be explained as follows. When the fleet is heterogeneous and the first strategy is used, a larger number of TTLPs must be solved, each using a different vehicle type. In contrast, with the second strategy, the presence of many vehicle types means that fewer matching problems need to be solved because the likehood of being able to serve a given set S of stations is higher in the heterogeneous case.

Results for the real life instance

We have also solved a real-life instance arising in Eastern Quebec, with a depot located in Quebec City. The area covered by this region is about 130 000 km^2 (see Figure 2). We used the data relative to deliveries made to 42 Esso stations on a single day. We used the drivers' worksheets to determine the ordered and delivered quantities of three products for that day. On that day 26 routes using eight vehicle types were used to make deliveries, resulting in a total distance of 7827.5 km. Currently all routes are determined by the dispatchers and vehicle loads are determined by the drivers.

To generate an equivalent instance (called scenario A), we reconstructed the distance matrix from postal codes using the commonly used package PC*MILER[12]. We set the a_t values equal to the delivered quantities and the b_t values equal to the ordered quantities. We also generated three other instances with the same distance matrix but using different a_t and b_t values, as shown in Table 10.

Computational results are reported in Table 11 and compared with the actual solution. In scenario A, which corresponds to the actual case, the solution uses 24 vehicles instead of the current 26, and these vehicles travel 17.2% fewer km. The total quantity delivered is 1.16% higher and the number of litres per kilometer is 22.12% higher. These statistics clearly confirm the efficiency of our solution methodology on this example. The three columns B, C and D shows the sensitivity of the solution to variations in the minimal and maximal demands. Increasing the b_t values leads to larger delivered quantities but vehicle routes remain the same. In scenario D, the results show that decreasing the a_t values leads a higher quantity delivered with fewer vehicles and fewer kilometers. This can be explained by the fact that lowering the minimal demands gives a higher likehood of being able to identify a better solution of the corresponding routing problem since more sets S of stations are feasible.

Table 10. a_t and b_t values for the four real-life instances

Scenario	a_t	b_t
A	$a_t =$ delivered quantity	$b_t =$ ordered quantity
B	a_t	$b_t + 5\,000$
C	a_t	$b_t + 10\,000$
D	$\max\{1\,000, a_t - 5\,000\}$	$b_t + 5\,000$

Table 11. Results for the real-life instances

	Actual solution	A	B	C	D
#routes	26	24	24	24	22
min	1 167 500	1 167 500	1 167 500	1 167 500	863 500
max	1 187 000	1 187 000	1 817 000	2 447 000	1 502 000
qty	1 167 500	1 181 000	1 401 000	1 440 000	1 230 000
km	7827.5	6481.9	6481.9	6481.9	5972.3
qty/km	149.2	182.2	216.1	222.2	206.0

Conclusions

We have developed an exact algorithm for the Petrol Station Replenishment Problem which decomposes into two subproblems: the Tank Truck Loading Problem

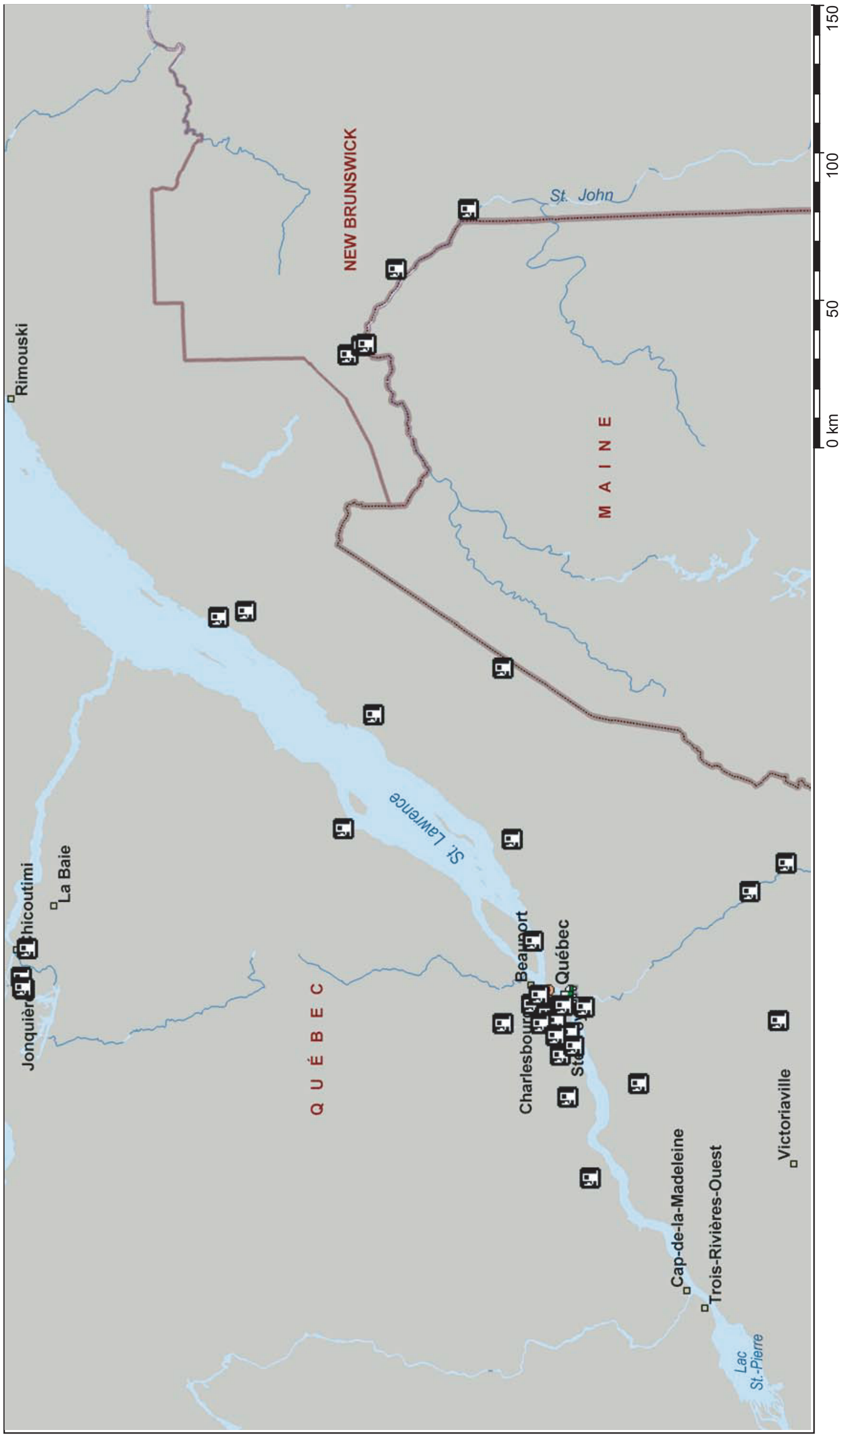

Figure 2. Station locations

and the Routing Problem. The TTLP is NP-hard but can often be solved to optimality by a heuristic. Otherwise an optimal solution is easily obtained by solving an integer linear program. This approach is appropriate for any instance size arising in practice. The Routing Problem is also NP-hard but reduces to a polynomial matching problem when at most two stations are visited on each route, as in most real-life instances we have encountered. Both algorithms were extensively tested on randomly generated data and on a real-life example arising in Eastern Quebec. Results show that the proposed solution algorithms perform remarkably well.

Acknowledgement

This work was partially supported by the Canadian Natural Sciences and Engineering Research Council (NSERC) under grants OGP0036509, OGP0039682 and OGP0172633. This support is gratefully acknowledged.

References

1 Ministère des ressources naturelles du Québec. Prix de l'essence ordinaire par région administrative. Technical report, 2004.

2 V. Malépart, J. Renaud, and F. F. Boctor. La distribution des produits pétroliers au Québec : État de la situation. Technical report, Université du Québec, 1998.

3 P. Toth and D. Vigo, editors. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, Philadelphia, 2002.

4 G. G. Brown and G. W. Graves. Real-time dispatch of petroleum tank trucks. *Management Science*, 27:19–32, 1981.

5 G. G. Brown, C. J. Ellis, G. W. Graves, and D. Ronen. Real-time, wide area dispatch of Mobil tank trucks. *Interfaces*, 17(1):107–120, 1987.

6 V. Malépart, F. F. Boctor, J. Renaud, and S. Labilois. Nouvelles approches pour l'approvisionnement des stations d'essence. *Revue Francaise de Gestion*

Industrielle, 22:15–31, 2003.

7 J. C. Smith. A genetic algorithm approach to solving a multiple inventory loading problem. *International Journal of Industrial Engineering*, 10:7–16, 2003.

8 N. Christofides, A. Mingozzi, and P. Toth. Loading problems. In P. Toth and N. Christofides, editors, *Combinatorial Optimization*, pages 339–369. Wiley, Chichester, 1979.

9 N. Christofides. Vehicle routing. In E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D.B. Shmoys, editors, *The Traveling Salesman Problem. A Guided Tour of Combinatorial Optimization*, pages 431–448. Wiley, Chichester, 1985.

10 H. N. Gabow. An efficient implementation of Edmonds' algorithm for maximum matching on graphs. *Journal of the ACM*, 23:221–234, 1976.

11 J. Edmonds. Paths, trees and flowers. *Canadian Journal of Mathematics*, 17: 449–467, 1965.

12 *PC*MILER User's Guide*. ALK Technologies, Inc., 2001.