

Planning and optimisation of vehicle routes for fuel oil distribution

Andrea E. Rizzoli^{a,c}, Norman Casagrande^a, Alberto V. Donati^a, Luca Maria Gambardella^{a,c}, Daniele Lepori^b, Roberto Montemanni^a, Piero Pina^b, Marco Zaffalon^{a,c}

^aIDSIA, Galleria 2, 6928 Manno, Switzerland

^bPina Petroli, 6916 Grancia, Switzerland

^cAntOptima, via Fusoni 4, 6900 Lugano, Switzerland

Abstract: Fuel oil distribution companies service their customers with fleets of tanker lorries. The problem is to find a set of routes minimising the number of travelled kilometres and the number of used vehicles, while satisfying customer demand. There are three major problems why traditional Operations Research techniques are not enough to deal with this problem, which is known as the Vehicle Routing Problem. First of all, it is inherently combinatorial, and exact algorithms fail when the dimension of the problem (number of customers and orders) reaches a reasonable size. Secondly, the problem can be extended and made more complex in many ways, for instance, adding more than one depot, considering more than one vehicle type, accounting for stochastic customer demand (the exact requested quantity is known only at delivery time), considering time windows during which the customers must be served, taking into account vehicle accessibility restrictions (some customers cannot be served by some vehicles). Finally, the problem can become very different when we consider on-line distribution, that is, we accept delivery orders for lorries which are en route. There, geolocation of customers and vehicles, online data transfer among lorries and the base station, have an impact as great as the solution strategy.

In this paper we present DyvOil, a software tool which assists the tour planner during the different stages of fuel distribution, from pre-planning on the basis of reorder forecasts, to online planning, through offline planning based on an advanced metaheuristic such as Ant Colony System. We also describe the case of Pina Petroli, a fuel oil distribution company located in Canton Ticino, Switzerland, which operates a fleet of 12 vehicles and serves thousands of customers using DyvOil.

Keywords: *Supply chain optimisation, vehicle routing problem, sales forecasting, ant colony optimisation*

1. INTRODUCTION

A traditional business model is articulated in three stages: production, distribution, and sales. Each one of these activities is usually managed by a different company, or by a different branch of the same company. Research has been trying to integrate these activities since the 60s when multi-echelon inventory systems were first investigated (Clark and Scarf, 1960), but, in the late 70s, the discipline which is now widely known as Supply Chain Management was not delivering what was expected, since the integration of data and management procedures was too hard to achieve, given the lack of real integration between the Enterprise Resource Planning (ERP) and the Enterprise Data Processing (EDP) systems (Sodhi, 2001). Only in the early 1990s did ERP vendors start to deploy products able to exploit the pervasive expansion of EDP systems at all levels of the supply chain.

The moment was ripe for a new breed of companies, such as i2, Manugistics and others, to put data to work and start to implement and commercialise Advanced Logistics Systems (ALS), whose aim is to optimise the supply chain seen as a unique process from the start to the end.

The first ALSs were the preserve of big companies, who could afford the investment in research and development required to study their case and to customise the application to interact with the existing EDP systems. Moreover, the available optimisation algorithms required massive computational resources, especially for combinatorial problems such as Vehicle Routing.

While ALSs were first deployed, researchers in the field of Operational Research were first investigating new “meta-heuristics”, heuristic methods that can be applied to a wide class of problems, such as Ant Colony Optimisation –ACO (Dorigo *et al.* 1996, Bonabeau *et al.* 2000).

Algorithms based on ACO draw their inspiration from the behaviour of real ants, which always find the shortest path between their nest and a food source, thanks to local message exchange via the deposition of pheromone trails. The remarkable advantage of ACO based algorithm over traditional optimisation algorithms is the ability to produce a good suboptimal solution in a very quick time. Moreover, for some problem instances, ACO algorithms have been proven to be the best overall (Gambardella and Dorigo, 2000)

The integration of optimisation algorithms based on innovative meta-heuristics, such as Ant Colony Optimisation, Tabu Search (Glover and Laguna, 1997), Iterated Local Search (Stützle and Hoos, 1999), Simulated Annealing (Kirkpatrick *et al.*, 1984), with ALSs for Supply Chain Management opens new perspectives of OR applications in industry. Not only big companies can afford ALSs, but also small and medium enterprises can use state-of-the-art algorithms, which run quickly enough to be adopted for online decision making.

improved using on-line information on traffic and vehicle locations. DyvOil has been designed, implemented, and tested in collaboration with Pina Petroli, a leading Swiss fuel oil distribution company.

2. CLOSING THE LOOP BETWEEN SALES AND DISTRIBUTION

Sales and distribution processes require the ability to *forecast* customer demand and to *optimally plan* the fine distribution of the products to the consumers. These two strategic activities, forecast and optimisation, must be tightly interconnected in order to improve the performance of the system as a whole (Gambardella *et al.*, 2001).

In Figure 1 the workflow process of a distribution-centred company is sketched.

The *sales department* generates new orders by contacting the customers (old and new ones) to check whether they need a new delivery. The effectiveness of this operation can be increased thanks to the DyvOil FORECAST module, which estimates the consumption of every customer, indicating the best re-order time for each of them.

New orders are then processed by the *planning department*, which, according to the quantity requested, the location of the customers, the time windows for the delivery, decides how many vehicles to employ and computes the best routes for the delivery, in order to minimise the total travel time and space. This task is assisted by the ACO algorithm, represented by the OPTIMISE block.

The vehicle tours are then assigned to the fleet, which is monitored by the *fleet operational control station*, which monitors the evolution of deliveries in real time. This process is assisted by the SIMULATE/MONITOR/RE-PLAN module, which allows re-planning online in face of new urgent orders, which were not yet available during the previous off line planning phase.

Finally, after vehicles have returned to the depot, delivery data are off-loaded and transferred back to the company database.

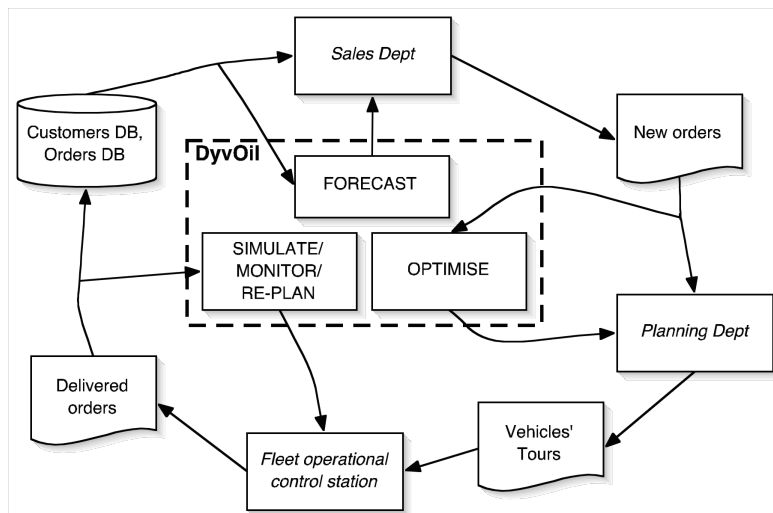


Figure 1. The workflow loop in a distribution-centred company.

In this paper, we introduce *DyvOil*, an ALS for planning the sales and distribution process of fuel oil, which includes an advanced ACO-based algorithm for the optimisation of the Vehicle Routing Problem. *DyvOil* does not take into account the production stage of the supply chain, since it is aimed at the level of fuel distribution companies, where the production and distribution stages are intrinsically decoupled.

In the next sections, we detail the workflow in a distribution centred company, then we describe the sales forecasting and the vehicle routing optimisation components of *DyvOil*. We also describe how vehicle routing can be sensibly

3. THE FORECASTING MODULE

Distribution of a product can be improved if warehouse stock-outs are foreseen, avoiding urgent orders which perturb the standard planning of delivery operations. Moreover, predicting when a customer is likely to order is fundamental to anticipate the competitors, if the customer has the option to ask for service from more than one company.

The approach adopted is based on the extraction of data from the past order history of each customer. This allows estimation of the dynamics of the warehouse management of each customer, thus allowing prediction of the time a re-order is most likely to be issued.

The past order history data hide valuable information about the customers' habits that can be used to predict the period of the next order. Customer profiling aims at extracting the skeleton of the customer's behaviours to the extent of predicting the next action. In the specific case, we can roughly devise two major types of behaviours. In the first case, a customer asks for refuelling when the reservoir is low, in a fairly independent way with respect to the period of the year (*a-periodic behaviour*). The remaining type of customers refuels in specific periods of the year, thus exhibiting a *periodic behaviour*.

The periodic behaviour can be identified by a straight analysis of the past orders. We can devise three types of periodic behaviours.

- *Regular customers in a certain period.* By analysing the data, we observed that some customers ask for refuelling in precise periods in ways quite independent of the consumption. There are customers that preventively buy fuel during summer, others that always fill the reservoir at the end of winter and others that wait for winter to refuel.
- *Regular customers in two periods.* As in the point above, except that we allow a customer to periodically refuel in two distinct periods of the year.
- *Customers that never buy fuel during winter.* In this case the customers prefer to buy fuel in the months that are not cold. These people usually buy it either before winter or after the cold period.

The a-periodic behaviour may be more difficult to predict, since consumption also depends on some exogenous variables (such as external temperature) and this, in turn, is related to the geographical zone of the customer. Nevertheless, the availability of historical data on temperature,

jointly with a substantial amount of other data, allows a customer's next order to be predicted.

In conclusion, we observed that about $\frac{1}{3}$ of the customers in the records of Pina Petroli's database can be classified and predicted according to one of the patterns reported in Table 1. This positive result shows that useful predictions can be carried out by an approach based on data analysis, which can then support the activity of the sales representatives.

Table 1. The results of the empirical analyses on the Pina Petroli database

	Num.	Percentage
Fixed period	134	9.31%
Never in winter	116	8.06%
Regularity in several periods	47	3.27%
Unclassified	378	26.27%
Classified	1061	73.73%
Examined	1439	100%

4. THE OPTIMISATION MODULES

DyvOil has two optimisation modules; the first one is used every evening to *off-line* plan vehicle tours for the next day. The second one is used *on-line*, during distributions, to face urgent orders, which are quite frequent in winter for fuel distribution, since it often happens that some customers run out of fuel.

4.1. Off-line optimisation

The most elementary version of the vehicle routing problem is the capacitated vehicle routing problem (CVRP) where n customers must be served from a unique depot, each customer asks for a quantity q , while the vehicles have a capacity Q . Since the vehicles' capacities are limited, they must periodically return to the depot for refilling. Therefore a CVRP solution is a collection of *tours*, where each customer is visited only once and the total quantity delivered in a tour is at most Q .

This problem can be made more complex by adding accessibility constraints (not all vehicles can serve all customers), and time windows (customers can be served only during fixed time intervals during a day). A VRP problem with time windows is denoted by the acronym VRPTW.

A tour is an ordered collection of tour deliveries. Each tour starts from and ends at the depot. Lorries may return more than once to the depot for reloading, but we consider these reloading

stops as a particular type of delivery and we define a tour as the ordered list of deliveries served by a vehicle during a day. In Figure 2 we represent a vehicle serving a set of customers.




Figure 2. A vehicle completes a set of tours, composed of the three tours a, b, and c.

The CVRP can be formulated in terms of graph theory, and it can be shown to be an extension of the Travelling Salesman Problem, thus the CVRP is also NP-hard. We are interested in solving an extension of the CVRP, the vehicle routing problem with time windows, where each customer has a time window during which s/he can be served. Every tour has a maximum duration and the vehicle fleet is non-homogeneous.

The solution of the VRPTW problem is a set of tours over a given time horizon, referred to as the planning horizon.

The Ant Colony System

The Ant Colony System (ACS) algorithm is based on a computational paradigm inspired by the way real ant colonies function. The medium used by ants to communicate information regarding shortest paths to food, consists of *pheromone trails*. A moving ant lays some pheromone on the ground, thus making a path by a trail of this substance. While an isolated ant moves practically at random, an ant encountering a previously laid trail can detect it and decide, with high probability, to follow it, thus reinforcing the trail with its own pheromone. The collective behavior that emerges is a form of autocatalytic process where the more the ants follow a trail, the more attractive that trail becomes to be followed. The process is thus characterized by a positive feedback loop, where the probability with which an ant chooses a path increases with the number of ants that previously chose the same path. The ACS paradigm is inspired by this process.

The algorithm

The DyvOil off-line optimisation algorithm is similar to that described in Gambardella et al.

(1999), which solves very efficiently the VRPTW.

The main elements of the algorithm are ants, simple computational agents that individually and iteratively construct solutions for the problem. At each step, every ant k computes a set of feasible expansions to its current partial solution and selects one of these probabilistically, according to the following probability distribution. For ant k the probability p_{ij}^k of visiting customer j after customer i , the last one of the current partial solution, depends on the combination of two values: the attractiveness τ_{ij} , computed by some heuristic indicating the a priori desirability of that move and the trail level η_{ij} , indicating how proficient it has been in the past to visit j right after i (it represents therefore an a posteriori indication of the desirability of that move). Trails are updated at each iteration, increasing the level of those associated with arcs contained in “good” solutions, while decreasing all the others.

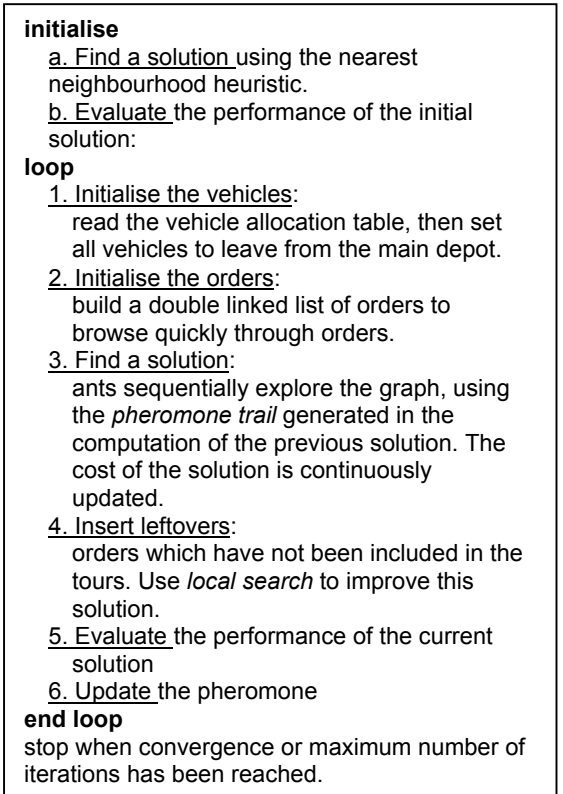


Figure 3. The optimisation algorithm

The principle behind this algorithm is that each vehicle is represented by an ant which incrementally builds a tour. Before adding a customer to a tour, the ant checks the constraints (vehicle type, accessibility). Then, if the constraints are satisfied, the insertion costs are evaluated in terms of distance, service time, and

its attractiveness measured by the pheromone intensity.

In Figure 3 we report a schematic outline of the ACO algorithm implemented in DyvOil. In the initialisation stage it must be noted that the evaluation of a solution is the travelled distance. In the main algorithm loop, at point 3, the ant colony algorithm works by *intensification* and *diversification* to build a solution. Intensification means that the pheromone trail is used to reinforce the solution; diversification means that random deviations are allowed to avoid local minima.

Calibration and validation

The optimisation algorithm depends on a speed model, which estimates the time required to travel from customer to customer, and from a service model, which computes the time required to deliver a given amount of fuel, including set-up time. A good model produces tours which are feasible in reality.

The calibration and validation of model parameters has been made using two sets of data collected from the 17 November to 21 November 1997 and from the 15 July to the 29 July 2002 by Pina Petroli. Delivery data are automatically recorded on a cassette¹, using the TDL standard format (DIN, 2000).

The calibration results show a percent error on the travel time estimate equal to 13.7% and a percent error on the service time equal to 14.6%.

The validation returned a percent error of -18.8% on the aggregate of service and travel times. Simulation plays an important role in assessing the effective feasibility of the optimized tours.

Performance evaluation

During the week from the 15 to the 19 July 2002 and the following week from the 22 to the 27 of July we compared the plans made by the human operator with plans generated by the algorithm. Summer is a low demand period, and Pina Petroli used only two vehicles to serve 127 customers during the two weeks.

First, we entered the tours, as ordered sequences of deliveries, as planned by the human operator. We then used the algorithm to compute the time required to perform all deliveries, assuming that the planner used the same speed model as the algorithm. This assumption is quite reasonable, in light of the fact that the speed model data has been elicited from the fleet planning department.

¹ The cassette is actually an EEPROM memory card,

We ran the algorithm with the same vehicles and used the same time windows, which were available to the Pina Petroli planning department. The algorithm is implemented in C++ and a good solution is obtained in two-three minutes on a Pentium PC at 1Ghz. The results are reported in Table 2:

Table 2. Comparison of performances 15-27 July 2002

	Distance (km)	Time (min)
Human	1578.4	7099.9
Algorithm	1302	6273.5

The algorithm allows a reduction of 21.3% in distance and a reduction of 13.2% in travel time. We expect to obtain even better results in high demand periods, when the human operator is under pressure because of urgent deliveries and a much greater number of orders has to be delivered.

5. SIMULATION, MONITORING, AND RE-PLAN

Once the tours have been assigned to the vehicles, the distribution operations can start; vehicles leave the depot and the computer-optimised tours have to face the real world, where the deterministic assumptions of the speed and the service models do not hold anymore. To assist the operational control of fuel distribution, DyvOil provides a *simulation* algorithm to simulate the distribution process in the face of variations in traffic and meteorological conditions (the former impacts the speed model, the latter the service one). The simulation module allows operators to verify the feasibility of distribution routes under different working conditions. It also allows estimation of the vehicles' position in absence of geolocation devices.

Alternatively, with a limited investment in hardware, vehicles can be monitored with GPS location devices. Their position is communicated to the base station via GSM/GPRS modems and displayed on a map.

The information on the current vehicle position is fundamental to manage "urgent" orders, which imply a re-plan of the distribution routes, shifting the problem focus to *dynamic* vehicle route optimisation.

5.1. On-line optimisation

In Dynamic Vehicle Routing Problems (DVRP) new orders dynamically arrive when the vehicles have already started executing their tours, which consequently have to be re-planned at run time in

order to include these new orders. This problem is quite frequent in winter, when demand and consumption are high, and customers run out of their storage. Although the use of the FORECAST module helps in sensibly reducing the number of such events, it may still happen that a customer places an urgent order.

We have developed the algorithm ACS-DVRP based on the decomposition of the DVRP into a sequence of static VRPs (Montemanni et al. 2003). There are three main elements in the architecture we propose: *Event manager*, collects new orders and keeps trace of the orders already served, the position of each vehicle and their residual capacity. This information is used to construct the sequence of static VRP-like instances. The *Ant Colony System* algorithm used to solve the static instances is the one described in Section 4.1. *Pheromone conservation*, once a static problem has been solved, the pheromone matrix contains information about good solutions to this problems. As each static problem is potentially very similar to the next one, this information is passed to the next problems, which includes the new orders that arrive in the meantime. This passage is efficiently implemented, following a strategy inspired by Guntch and Middendorf (2001).

6. CONCLUSIONS

We have presented DyvOil, an Advanced Logistic System for the sales and distribution of fuel oil. It is compact, fast, easy to install and use, able to satisfy the requirements of both large and small vehicle fleets thanks to innovative algorithms based on the Ant Colony System metaheuristic. DyvOil helps forecast customer demand, thus allowing better pre-planning of distribution routes. It provides fast and efficient solutions to the off-line vehicle routing problem, but it also assists in dynamic vehicle routing, in the face of unexpected, urgent orders.

DyvOil, developed as a CTI-KTI (Swiss Commission for Technology and Innovation) sponsored project, is now a commercial product of AntOptima and Pina Petroli.

7. REFERENCES

Bonabeau, E., M. Dorigo, and G. Theraulaz, Inspiration for optimization from social insect behaviour, *Nature* 406, 39–42, 2000.

Clark, A.J., and H. Scarf, Optimal policies for a multi-echelon inventory problem, *Management Science*, 6, 475–490, 1960.

DIN - Deutsches Institut für Normung, *Datenschnittstelle für Tank- und Transportfahrzeuge – Teil 1*, DIN 2605-1, 2000-07.

Dorigo, M., V. Maniezzo, and A. Coloni, The ant system: optimization by a colony of cooperating agents, *IEEE Transactions on Systems, Man, and Cybernetics--Part B*, 26(2), 29–41, 1996.

Gambardella, L., É. Taillard, G. Agazzi. MACS-VRPTW: a multiple ant colony system for vehicle routing problems with time windows. In: D. Corne et al. (eds.), *New ideas in optimization*, McGraw-Hill, London, 1999.

Gambardella L.M. and M. Dorigo, An Ant Colony System Hybridized with a New Local Search for the Sequential Ordering Problem, *INFORMS Journal on Computing*, 12(3), 237–255, 2000.

Gambardella, L.M., M. Mastrolilli, A.E. Rizzoli M. Zaffalon, An integrated approach to the optimisation of an intermodal terminal based on efficient resource allocation and scheduling, *Journal of Intelligent Manufacturing*, 12 (5/6), 521–534, 2001.

Glover, F. and M. Laguna, *Tabu Search*. Kluwer Academic Publishers, Boston, MA, 1997.

Guntch, M., and M. Middendorf, Pheromone modification strategies for ant algorithms applied to dynamic TSP. In: *Lecture Notes in Computer Science*, 2037, 213–222, Springer-Verlag, 2001.

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science*, 220, 671–680, 1983.

Montemanni, R., L.M. Gambardella, A.E. Rizzoli, A.V. Donati. A new algorithm for a Dynamic Vehicle Routing Problem based on Ant Colony System (Extended abstract). Accepted at *ODYSSEUS 2003: Second International Workshop on Freight Transportation and Logistics*, Palermo, Italy, 27-30 May 2003.

Sodhi, M. S., Applications and Opportunities for Operations Research in Internet-Enabled Supply Chains and Electronic Marketplaces, *Interfaces*, 31(2), 56–69, 2001.

Stützle, T., and H. Hoos, Analyzing the Run-time Behaviour of Iterated Local Search for the TSP, In *Proceedings of MIC'99, III Metaheuristics International Conference*, 1999.