

Ant Colony Optimization for vehicle routing in advanced logistics systems

Luca Maria Gambardella^{a,b}, Andrea E. Rizzoli^{a,b}, Fabrizio Oliverio^b,
Norman Casagrande^a, Alberto V. Donati^a, Roberto Montemanni^a, Enzo Lucibello^b

^aIDSIA, Galleria 2, 6928 Manno, Switzerland

^bAntOptima, via Fusoni 4, 6900 Lugano, Switzerland

URL: <http://www.idsia.ch>, <http://www.antoptima.com>

Email: luca@idsia.ch

ABSTRACT

Many distribution companies service their customers with non homogeneous fleets of trucks. Their 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 and AntRoute, two software tools which assist the tour planner during the different stages of goods 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 customers using DyvOil and Migros, the largest Swiss supermarket chain, which uses AntRoute operates daily a fleet of hundreds of trucks distributing goods to its shops.

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

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 SAP, 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 are multi-agent systems that exploit artificial stigmergy for the solution of combinatorial optimization problems: they 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 short time. Moreover, for some problem instances, ACO algorithms enhanced with local optimisation capabilities, 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.

In this paper we present a modular approach to ALS design and implementation, driven by the user needs. We show how different algorithms and modules can be implemented in an ALS and how tailor-made solutions can be integrated into traditional supply chain management software.

In the next sections, first we detail the workflow in a distribution centred company, then we introduce and discuss the off-line and on-line vehicle routing problems, which are common to most distribution companies. The former is solved when the orders to be delivered are known in advance, the latter when new orders arrive while the distribution process is on. Finally, we briefly report on how two pieces of software, DyvOil and AntRoute, have been designed, implemented, and tested in collaboration with Pina Petroli, a leading Swiss fuel oil distribution company, and Migros, the largest supermarket chain in Switzerland.

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

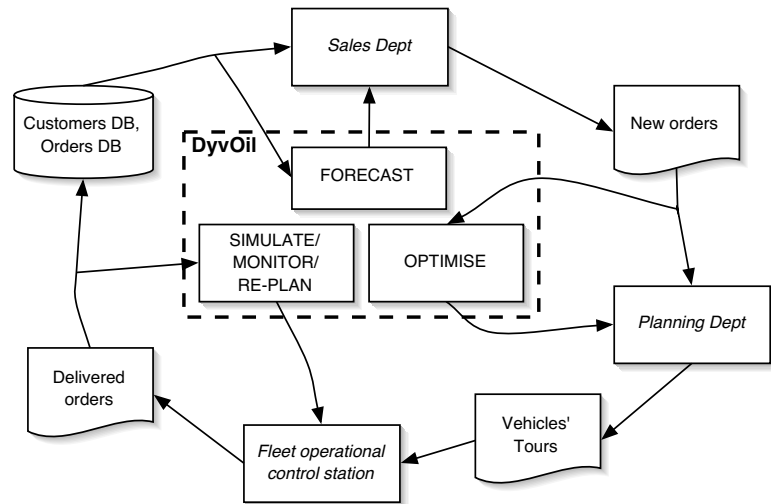


Figure 1. The workflow loop in a distribution-centred company.

need a new delivery. The effectiveness of this operation can be increased thanks to the 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 ACO algorithms solving the Static Vehicle Routing Problem (SVRP), embedded in the OPTIMISE block in our schema.

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. This module uses ACO algorithms designed to solve the Dynamic Vehicle Routing Problem (DVRP)

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

THE STATIC VEHICLE ROUTING PROBLEM

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.

ACS for SVRP

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.

In MACS-VRPTW, a very efficient algorithm developed by Gambardella et al. (1999), two ant colonies work in parallel. The first colony builds tours trying to minimize the total traveled distance, given a number of vehicles imposed from the outside. The second colony tries to reduce the number of used vehicles, in order to minimize the cost of the used resources. If a solution which uses a smaller number of vehicles is found, this is passed to the other colony, which tries again to minimize the distance. The process is repeated iteratively until no further improvements can be obtained.

THE DYNAMIC VEHICLE ROUTING PROBLEM

In Dynamic Vehicle Routing Problems (DVRP) new orders dynamically arrive when the vehicles have already started executing their tours and they may have already left the depot. The vehicle tours have to be re-planned at run time in order to include these new orders, since we do not want vehicles to travel back to the depot when they are assigned new orders. A communication system must exist between vehicles and the depot.




Figure 3. The architecture of the DVRP algorithm

ACS for DVRP

Montemanni et al. (2003) have developed the algorithm ACS-DVRP based on the decomposition of the DVRP into a sequence of static VRPs. There are three main elements in the algorithm architecture: *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 the section titled “ACS for SVRP”. *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 Guntsch and Middendorf (2001).

The performance of the tours depends on the strategy adopted by the Event Manager to split the day in time slices. If the algorithm runs as soon as a new order is received, the resulting tour could be myopically designed, since one gives up the possibility of including another (still unknown) order in the modified tour. On the other hand, waiting too long can lead to a sensible decrease in customer satisfaction, since the time between an order is placed and when it is executed can become too long.

TWO EXEMPLAR VEHICLE ROUTING APPLICATIONS

The first application, DyvOil, supports planning the sales and distribution processes of fuel oil. DyvOil covers all the phases of the workflow loop of Figure 1, since it aims at small and medium enterprises willing to integrate sales and goods distribution. DyvOil includes a Sales Forecasting module and a Vehicle Routing Planning module. Moreover, DyvOil is currently being enhanced with an on-line vehicle routing module, to provide the ability to respond to unexpected delivery requests. DyvOil is intended for use by small and medium enterprises, which deliver their goods in cities and urban centres, subject to high volatility of order batches.

The second application, AntRoute is a software component embedded in the existing ALS of a medium/big enterprise. It has been designed with the clear objective of providing high quality solutions of large scale off-line vehicle routing problems, where the order batches are known in advance with sufficient certainty.

DyvOil, forecasting demand

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.

DyvOil is an application for the distribution of fuel oil, the consumption of which is related to external factors, such as temperature, and customer re-order preferences, but the same principle can be applied for a great variety of goods. Data mining techniques can be used when large enough base of past orders data is available.

Thanks to the use of DyvOil Forecasting module, Pina Petroli claims that the number of successful calls (telephone calls to customers leading to an order) has sensibly increased from one success out of four to one out of two.

DyvOil: optimising distribution

DyvOil has two optimisation modules: off-line and on-line. Particular care has been given to the integration of these advanced algorithmic modules within the user interface of DyvOil (see Figure 4). DyvOil is a Windows

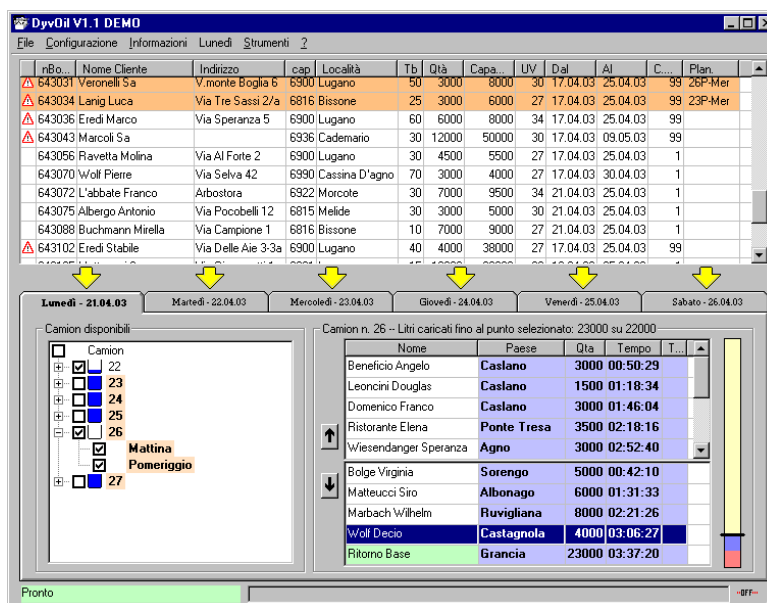


Figure 4. A sample of DyvOil user interface

based application which can be interfaced to most ERP and EDP systems thanks to its flexible and easy to configure ODBC interface to external databases.

DyvOil displays an interface that lets the human planner manually enter vehicle routes and, on-demand, these routes can be automatically optimised by the algorithm, which returns its results in a few minutes. The planner can then decide whether the computer generated results

can be accepted as they are, or they require further refinements and adaptations. This kind of interactive usage has proven very successful in getting the technology adopted by planners in small-medium enterprises, where the obstacles to technology adoption are usually higher.

The *off-line module* is used every evening to plan vehicle tours for the next day. This module solves the SVRP using ACO algorithms, as described in section “ACS for SVRP”. The particular implementation of ACO in DyvOil is aimed at catching the finer details constraining the delivery process, since DyvOil is often used in urban environments. This requirement imposes strong accessibility constraints to vehicles: not all lorries can access all customers, who are often located in the city centre where both architectural constraints and traffic rules limit the vehicle dimensions.

The off-line planning module has been tested in a real world setting at Pina Petroli, a leading Swiss fuel oil distribution company. It has been observed an increase in the vehicle routing performance of 20% up to 30% with respect to human generated plans (Rizzoli *et al.* 2003).

The volatility of order batches, the fact that some orders arrive during delivery and they must be urgently executed leads to the requirement of on-line route planning. This module solves the DVRP, as described in section “ACS for DVRP”. This module is still a prototype, but first results, obtained using data from Pina Petroli, show that the algorithm is very efficient, also when compared to heuristic techniques (Montemanni *et al.* 2003)

AntRoute

An always increasing number of large and medium-large distribution companies have already adopted ALS to manage their whole supply chain. The basic information processing infrastructure is in place and many supply chain management suites already provide optimisation modules for some components. The objective of the AntRoute software component is to integrate a state-of-the-art optimisation algorithm within an existing supply chain management structure. AntRoute has been implemented in C++ and it has been deployed as a windows DLL, but its code can be recompiled under most operating systems.

AntRoute has been designed for the inclusion in the supply chain of Migros, the largest supermarket chain in Switzerland, which has recently concentrated its inventories in the new logistic hub of Suhr, in central Switzerland. All the department stores of Central, Western and Eastern Switzerland send their orders in Suhr, via a SAP-based application. The orders are processed to prepare the pallets, ready to be distributed by lorry to the stores. Migros planners use CADIS, an application which includes the high-performance

AntRoute algorithms to compute the tours of the distribution vehicles.

Vehicles are non-homogeneous and AntRoute must also cope with accessibility restrictions, stores impose soft time windows on delivery time. The objectives are cost minimization (cost per km) and tour minimization (to limit the number of vehicles). The objective function evaluates each solution (a set of tours) according to the following scheme:

- first minimize the number of tours,
- then minimize the costs per kilometer and the cost of violating the soft time windows

The algorithm has been modeled after MACS-VRPTW (Gambardella *et al.* 1999).

In this application the orders are known with sufficient advance to be able to run an off-line optimisation.

AntRoute outperforms the human planners creating tours which are shorter, use less vehicles, and are less expensive.




Figure 5. A sample of the tours computed by AntRoute

CONCLUSIONS

Advanced Logistic Systems ask for high performance optimisation algorithms. In the case of vehicle routing, algorithm based Ant Colony Optimization prove to be among the best and most performing. We have discussed the classical VRP problem where orders are known in advance, before optimization starts, and its extension to on-line planning (dynamic VRP) where orders arrive during the distribution process. A new heuristic based on ACO for the solution of DVRP has been briefly described.

Finally, we have presented DyvOil, an Advanced Logistic System for the sales and distribution of fuel oil, and AntRoute, a software component embedded in a third-party ALS.

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

AntRoute is a highly specialised software module dedicated to the solution of medium to large scale static vehicle routing problems. It has been applied with remarkable success to the optimisation of the vehicle routes of Migros, the largest supermarket chain of Switzerland.

AntRoute has been developed by AntOptima for Migros Logistik Transport with the collaboration of Cantaluppi & Hug Software and Consulting, Zürich.

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.
- 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.
- Guntsch, 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.
- Rizzoli, A.E., Casagrande, N., Donati, A.V., Gambardella, L.M., Lepori, D., Montemanni, R., Pina, P., and Zaffalon, M. Planning and optimisation of vehicle routes for fuel oil distribution. In Post, D.A. (ed) *MODSIM 2003 International Congress on Modelling and Simulation*, Volume 4, pp. 2024--2029. Modelling and Simulation Society of Australia and New Zealand. July 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.

BIOGRAPHY

Luca Maria Gambardella is a Research Director at IDSIA and Professor of Artificial Intelligence at SUPSI. His major research interests are in the area of optimisation, simulation, multi-agent learning and adaptation, applied to both academic and real-world problems. In particular, he has studied and developed several Ant Colony Optimization algorithms to solve travelling salesman problems, quadratic assignment problems, sequential ordering problems and vehicle routing problems. In these domains, the best-known solutions for many benchmark instances have been computed. He leads IDSIA research on routing and optimization. He has led several research and industrial projects both at Swiss and European level. He is member of the board of the Swiss Association of Operational Research.