Ant Colony Optimization for Real-world Vehicle Routing Problems

Dario Coltorti, dario.coltorti@antoptima.ch, AntOptima SA, Lugano, Switzerland
Andrea E. Rizzoli, andrea.rizzoli@antoptima.ch, AntOptima SA, Lugano, Switzerland

Metaheuristics like ant colony optimization (ACO) can be used to solve combinatorial optimization problems. In this paper we refer to its successful application to the vehicle routing problem (VRP). At the beginning, we introduce the VRP and some of its variants. The variants of VRP were designed to reproduce the kind of situations faced in the real-world. Further, we introduce the fundamentals of ant colony optimization, and we present in few words its application to the solution of the VRP. At the end, we discuss the applications of ACO to a number of real-world problems: a VRP with time windows for a major supermarket chain in Switzerland; a VRP with pickup and delivery for a leading distribution company in Italy and an on-line VRP in the city of Lugano, Switzerland, where clients' orders arrive during the delivery process.

Introduction

Most logistics problems are particularly challenging as their search space grows exponentially with the problem dimensions and no efficient algorithms to explore such space are known. For these problems, which are technically known as NP-hard, the time required to find an optimal solution might be simply too high for practical purposes.

Heuristics methods have been devised to explore parts of the search space, concentrating in those parts that appear to be most promising, thus reducing the time required to obtain a sub-optimal, but still good enough, solution. A heuristic makes use of peculiar characteristics of a problem and exploits them to find a solution. Therefore a heuristic has to be especially devised for each new problem.

A metaheuristic is a set of concepts that can be used to define heuristic methods that can be applied to a wide set of different problems [16]. Well known examples of metaheuristics include simulated annealing (SA), tabu search (TS), iterated local search (ILS), evolutionary algorithms (EC), and ant colony optimization (ACO), the subject of this paper.

Ant Colony Optimization (ACO) is based on the observation that ants can find the optimal path between a food source and their nest exploiting a mix of probabilistic behavior and pheromone depositing. In fact, in ACO a set of artificial ants somehow simulate the behavior of real ants; the artificial ants move on the graph representation of a combinatorial optimization problem and build solutions probabilistically. The probabilities are biased by artificial pheromones that ants deposit while building solution (for a recent overview of ACO see [3]; for a detailed description [8]). In this paper we discuss how ACO can be successfully applied to the solution of real-world vehicle routing problems.

The Vehicle Routing Problem

The vehicle routing problem can be designed as a combinatorial optimization problem: Finding optimal routes for a fleet of vehicles performing assigned tasks on a number of geographically sectored clients. An answer to this problem is the best route serving all clients using a fleet of vehicles, respecting all operational constraints, such as vehicle capacity and the driver's maximum working time, and minimizing the total transportation cost.

There are 3 main factors that define and constrain each model of the VRP: the road network, specifying the relatedness among clients and depots, the vehicles, transporting goods between clients and depots on the road network; the clients, which place orders and receive goods.

Joining the various factors of the problem, we can define a whole set of different VRPs (for a detailed overview of the various VRPs see [20]). All these variants have been created in order to bring the VRP closer to the kind of situations faced in the real-world. Table 1 shows some important VRP starting from basic version, continuing by static case (VRP with time windows, VRP with time windows and pick-up and delivery constraints) and finishing by dynamic case (time dependent VRP like on-line VRP).

Optimization Framework Inspired By Ants

Ant colony optimization [5] is a metaheuristic inspired by the observation, made by ethologists, that ants are able to find the shortest path to a food source by laying and following chemical trails. The chemical substance which ants use to communicate information regarding the shortest path to food is called pheromone. Communicate means that a moving ant lays some pheromone on the ground, thus marking a path with a trail of this substance. In the majority of cases an isolated ant moves randomly and when it discovers a previously laid pheromone trail it can decide, with high probability, to follow it, thus reinforcing the trail with its own pheromone. The group behavior that results is a form of self-organisational process where the more ants follow a trail, the more attractive for other ants it becomes. The process running by basic rules is characterised 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. Other positive characteristics of the above process are the flexibility (adaptability) and the robustness (system doesn't depend on one ant). This group behavior of ants with its positive attributes inspired the ACO metaheuristic. The main factors are artificial ants (called from now on ants), simple computational agents that individually and iteratively construct solutions on a graph, which has been modeled depending on the specific problem. A problem solution is an ordered sequence of nodes connected by edges visited by exploring ants. Ants compute a solution in parallel, deploying the search process over several constructive computational threads. A dynamic memory structure, inspired by the pheromone laying process, guides the construction process of each thread.

The memory structure incorporates information on the effectiveness of previously obtained results. Intermediate partial problem solutions are seen as states; at each iteration k of the algorithm each ant moves from state x k (i) to x k+1 (j), enlarging the partial solution from node i adding node j.

Based on these elements the first ACO algorithm to be proposed was Ant System (AS) [7]. It is organized in two main stages: construction of a solution, and update of the pheromone trail. Since its publication different variants have been proposed to improve the solutions of combinatorial optimization problems: elitist ant system [4], rank-based ant system [1], and Max – Min ant system [19] are variants, where the algorithm differs from the original mainly in the pheromone update rule. On the other hand, extensions of AS display more substantial changes in the algorithm structure. Ant Colony System (ACS, [6]) is one of them. ACS differs from AS for a revised rule used in the tour construction algorithm, and for the use of both local and global updates of the pheromone trails.

ACS has been shown to be very efficient in solving problems of the vehicle routing class, ranging from the static case (VRP with time windows, and VRP with time windows and pick-up and delivery constraints) to the dynamic case (on-line VRP). In the next section we describe how ACO has been applied in a number of cases to solve real world logistic problems.

Major supermarket chains: Distribution of goods from inventory stores to shops

In this business case one of the major supermarket chains in Switzerland has the following challenge: Palletized goods must be distributed to more than 600 stores, all over Switzerland. To replenish their local stocks each store orders daily quantities of goods, which have to be delivered within time windows. So each store can plan and allocate efficiently according to the daily availability of its personnel and the time requested for inventory management tasks. Further there are three types of vehicles: trucks (capacity: 17 pallets), trucks with trailers (35 pallets), and tractor units with semi-trailers (33 pallets). One practical restriction is the access of vehicles to the store, which depends on the store location. In some cases the truck with trailer can leave the trailer at a previous store and then continue to other less accessible locations. Moreover the number of vehicles is assumed to be infinite, since transport services can be purchased on the market according to the needs.

Problem Type	Constraints	Objective	NP-hard Problem (yes/no)
Capacitated vehicle routing problem (CVRP, basic version of the VRP)	- Having vehicles with limited capacity - Client demands are deterministic and known in advance - Deliveries cannot be split - Vehicle fleet is homogeneous	Minimise the total travel cost	Yes [13]
Vehicle routing Problem with time window (VRPTW) [15, 14, 12]	Each Client is associated with a time window and a service time	Minimise the total travel cost	Yes [18]
VRP with pick-up and delivery (VRPPD) [2]	The transport items are not originally concentrated in the depots, but they are distributed over the nodes of the road network. A transportation request consists in transferring the demand from the pick-up point to the delivery point. These problems always include time windows for pick-up and/or delivery.	Minimise the total travel cost	Yes
Probabilistic, dynamic and stochastic vehicle routing (assumed generic term: Dynamic VRP) like online VRP [10]	The assumption of time invariancy must be relaxed and data become time-dependent. Moreover, using data on current traffic conditions to estimate travel times requires the relaxation of the assumption of determinism, introducing uncertainty and adding another level of complexity to the problem.	Minimise the total travel cost.	Yes

Tab. 1: Important vehicle routing problems.

	Human Planner	AR-RegTW	AR-Free
Total number of tours	2056	1807	1614
Total km	147271	143983	126258
Average truck loading	76.91%	87.35%	97.81%

Tab. 2: Comparison of the man-made vs. the computer-generated tours in the VRPTW application.

The road network graph could be computed due to digital road maps. On the other hand the distance matrix between pairs of stores has been rescaled using a company speed model, based on many years of experiences and collecting data. For example, if the distance is less than 5 km, the average speed is 20 km/h; if the distance is more than 90 km, the speed is 60 km/h; in between there is a range of other speed values. Constant parameters are the time to set-up a vehicle for unloading and the time required to hook/unhook a trailer. A variable parameter is the service time, which depends on the number of pallets to unload. The main restriction is that all the routes must be performed in one day, and the company imposes an extra constraint stating that a vehicle must perform its latest delivery as far as possible from the inventory, since it could be used to perform extra services on its way back. These extra services were not included in the planning by explicit request of the company.

Solution method and results

This planning challenge was modeled as a VRPTW, and solved by an implementation of the MACS-VRPTW algorithm [9], named ANTROUTE. MACS-VRPTW is the most efficient ACO algorithm for the VRPTW and one of the most efficient metaheuristics overall for this problem. ANTROUTE adds to MACS-VRPTW the ability to handle the choice of the vehicle type: at the start of each tour the ant chooses a vehicle. To prevent vehicles arriving too early at the stores a waiting cost was also introduced. The central idea of the MACS-VRPTW algorithm is to use two ant colonies (MACS stands for multi ant colony system) to optimize two objectives: One colony, named ACS-VEI, minimizes the vehicles while the other one, named ACS-TIME, minimizes time.

Human tour planners evaluated the first tours computed by ANTROUTE and the tours were not accepted as feasible, even if the performance was considerably higher than theirs and no explicit constraints were violated. That's the reason why a further modeling step was required, to let "invisible" constraints emerge. One of them was a regional planning strategy, that led to petal shaped tours, as the human planners were currently doing. This way of doing tours was included in the reformulation of the problem, but at the same time the project team tried to loosen the constraint a bit. Stores would be attributed to distribution regions, allowing at the same time stores near the border of the distribution region to also belong to the neighbouring region. This new generation of tours were a bit worse than the unconstrained solution, but nonetheless better than the solutions found by the human planners. Table 2 presents the results obtained by ANTROUTE compared with those of the human planners.

ANTROUTE was run under two scenarios: AR-RegTW, with regional planning and 1-hour time windows; AR-Free, where the regional and the time windows constraints were detached. The challenge was to distribute 52000 pallets to 6800 clients over a period of 20 days. ANTROUTE was run on the available set of orders daily and it took about 5 minutes to find a solution. At the same time, the planners were at work and it took them at least 3 hours to find a solution. After the testing period, the performances of the algorithm and of the planners have been compared using the same objective function. A further advantage of an algorithm able to find the solution to a very hard problem in such a short time is the possibility of using it as a strategic planning tool beside of the operative role. Figure 1 indicates how running the algorithm with wider time-windows at the stores returns a smaller number of tours, which can be interpreted in a significant reduction of transportation costs. The logistic manager can therefore use the optimization algorithm as a tool to check how to re-design the time-windows in the stores.

Major logistics operator: Distribution from factory to inventory stores

In this business case the company is a major logistics operator in Italy. The distribution process comprises moving palletized goods from factories to inventory stores, before they are after distributed to shops. A customer in this vehicle routing problem is either a pick-up or a delivery point. A central depot doesn't exist, and approximately 1000 – 1500 vehicles per day are used.

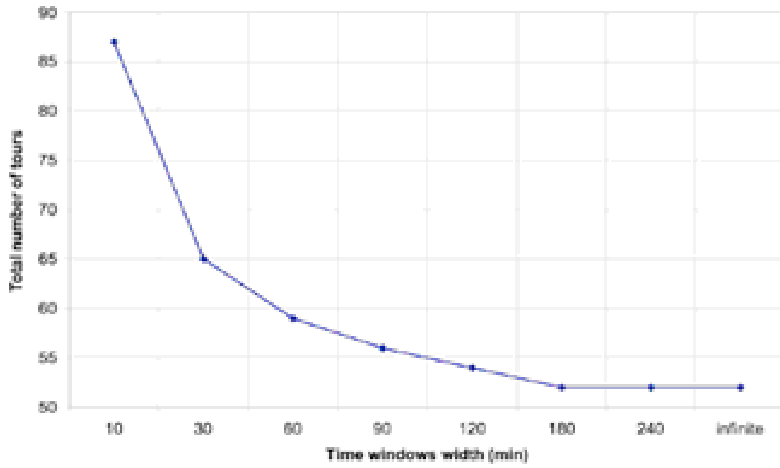

Fig. 1: The relationship between the number of vehicle routes and the time window width.

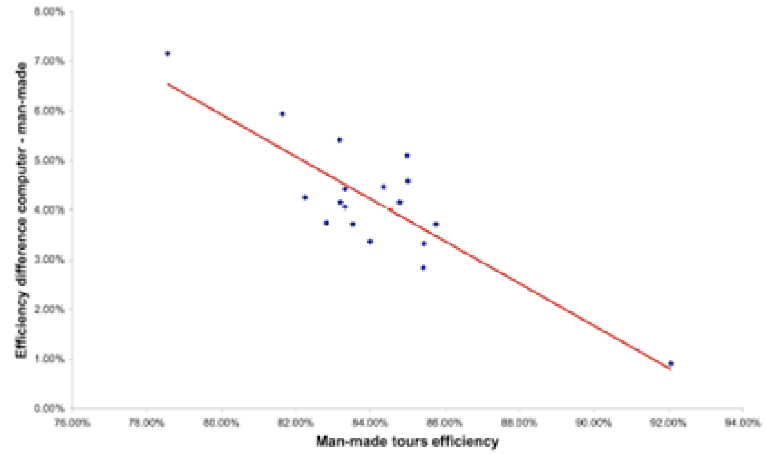

Fig. 2: Comparing man-made and computer-generated tours. Higher efficiency improvements are observed when the human planner performance is lower. The dots are experimental values, and the solid line is a regression on those values.

Because of the geographical constitution of Italy and the strict legal constraint on the maximum number of hours per day that a driver can travel, routes can be performed within the same day, over two days, or over three days. All pick-ups of a tour must happen before deliveries, and orders cannot be split among tours. Further are time windows associated with each store and there is only one type of vehicle: tractor with semi trailer. The load is measured in three units: Pallets, kilograms and cubic meters. Each one of these units has a capacity constraint and the first one that is passed causes the violation of the constraint. Since they are provided by flexible sub-contractors, the availability of trucks is assumed to be infinite. These sub-contractors are distributed all over Italy, and therefore trucks can start their routes from the first assigned customer, and for the company doesn't result any traveling cost to the first client in the route. The road network graph could be found out due to digital road maps, computing the shortest path between each couple of stores. The travel times are calculated according to the travelled distance, given the average speed that can be obtained on each road segment according to its type (highway, extraurban road, urban road).

Further the loading and the unloading times are assumed to be constant parameters, since the company has been unable so far to provide better estimates. Consequently this assumption is a rough approximation imposed by the company, which also imposed another constraint, related to the same problem, setting a maximum number of cities to visit per tour (usually less than six). Note that more than one customer can reside in a city. Furthermore, the company requested that the distance between successive deliveries should be limited by a parameter.

Solution method and results

This planning challenge was modeled as a VRP with pickup and delivery and time windows (VRPPDTW). The objective function quantifies the average tour efficiency. The ANTROUTE algorithm for using in this context has been modified: Since for this problem there is a single objective instead of two (business case before) — to maximise average efficiency — the ant colony minimising the number of vehicles was removed.

	Human Planner	ANTROUTE	Absolute difference	Relative difference
Total nr of tours	471.5	460.8	-10.7	-2.63%
Total km	175441	173623	-1818.2	-1.32%
Efficiency	84.08%	88.27%	+4.19%	-

Tab. 3: Comparison of the man-made vs computer-generated tours in the VRPPDTW application.

Table 3 resumes the comparison between man-made and computer-generated tours over a testing period of two weeks. A significant enhancement in the efficiency of computer-generated tours can be noticed. Another interesting point is to observe how the algorithm performance is correlated with the difficulty of the problem, which is related to the number of orders to satisfy. Figure 2 shows on the x-axis the efficiency of the man-made tours, and on the y-axis the efficiency improvement obtained using the computer-generated tours. When the problem is easy, because it containes a limited number of orders, and the human planner schedules well, the computer is not able to deliver a significative enhancement, but when the planner starts to fail coping with the problem complexity, and the performance falls, the gain in using the algorithm sensibly rises.

Fuel oil distributor: On-line VRP for fuel distribution

This case study treats a fuel oil distribution company in Switzerland, which serves its customers from its main depot located near Lugano with a fleet of 10 trucks. The fuel oil distributor noted that during every Winter season there was always a subset of their customers that ran out of fuel and had to place urgent orders. These unanticipated orders have an impact on the planned delivery routes of the trucks, and the vehicle routing problem becomes very “dynamic”. This means that a considerable percentage of orders must be fulfilled after the trucks have already left the depot. The goal of this case study was to evaluate the impact of a reactive strategy for vehicle routing, starting from data analysis collected in periods when urgent deliveries were in high request. A sample of 50 customers from the company data base was randomly selected and travel times among them were computed. In the company records, customers randomly appeared during the working day with random requests for a quantity of fuel to be supplied.

An 8 hours working day was considered and a service time of 10 minutes for each customer was supposed. The cut-off time was set to 4 hours. Thenceforward the new orders received were deferred to the following working day.

Solution method and results

The problem description above fits the on-line VRP variant, where new orders can be allocated to vehicles which have already left the depot (e.g., parcel collection, feeder systems, fuel distribution, etc.). Montemanni et al. [17] have developed an ACO-inspired algorithm, ACS-DVRP, adapted from the decomposition of the on-line VRP into a sequence of static VRPs. ACS-DVRP solves the on-line fuel oil distribution problem and its algorithm architecture consists of three main elements: the event manager, the ant colony algorithm and the pheromone conservation strategy. The event manager obtains new orders and maintains track of the already served orders and of the position and the remaining capacity of each truck. This information is used to build the sequence of static VRP-like instances. The working day is split into time slices and for each of them a static VRP is created. Every static VRP considers all the already received (but not yet executed) orders. New orders received during a time slice are deferred until its end. At the end of each time slice, customers whose service time starts in the next time slice are assigned to the trucks. They will not be considered in the following static VRPs.

The ant colony algorithm applied based on the MACS-VRPTW implementation, named ANTROUTE, is described in former sections. Instead of two ant colonies there is only one, which is in charge of minimizing the total travel time. Furthermore the pheromone conservation strategy is characterised as follows. Once a time slice is over and the relative static problem has been solved, the pheromone matrix comprised information about good solutions. Since each static problem is potentially very similar to the next one, this information is transferred to the next problem [11]:

if a couple of customers is in both the previous and the current time slice, the pheromone on the arcs connecting two nodes is brought forward as a fragment of its value in the previous problem.

Several test problems were created, where the algorithm ACS-DVRP was applied, due to varying the number n_{ts} of time slices into which the working day was divided. As the size of each problem in a time slice increases as the length of the time slice decreases, the time t_{acs} assigned to executing the ant colony system and the time t_{ls} allocated to local search improving the solution were adapted accordingly. Especially the ratio between t_{acs} and t_{ls} was kept around equal to 10. Table 4 and its first three rows with the values parameters n_{ts}, t_{acs} and t_{ls} define the settings of the experiments. The final row shows the total travel time of the solutions calculated by the ACS-DVRP algorithm. The results show that, for this specific case study, good values for n_{ts} are between 10 and 50. Especially, 25 appears to be the best choice. Large values of n_{ts} did not lead to satisfying results because optimization was restarted too often, before a good local minimum could be obtained. Otherwise, when n_{ts} was too small, the system was not able to take advantage of information on new incoming orders.

Conclusions

This contribution describes the metaheuristic ant colony optimization and how it can be successfully used to solve a number of variants of the basic vehicle routing problem. The main part presents two industrial-scale applications of ACO for the solution of static VRP problems: a VRP with time windows and a VRP with pickup and delivery. Then the contribution focuses its attention on one important dynamic variant of the VRP: the on-line VRP. The problem is receiving increasing attention based on its relevance to real world problems, in particular for distribution in urban environments. The applications of ACO on real-world VRP shows that this metaheuristic inspired by ants has become an important tool in applied operations research.

n_{ts}	200	100	50	25	10	5
t_{acs}	144	288	576	1152	2880	5760
t_{ls}	15	30	60	120	240	480
Travel time	12702	12422	10399	9744	10733	11201

Tab. 4: Experimental results on the case study of Lugano. n_{ts}: number of time slices into which the working day was divided; t_{acs}: the time allocated to executing the ant colony system; t_{ls}: the time dedicated to local search improving the solution.

Bibliography

[1] B. Bullnheimer, R.F. Hartl, and C. Strauss. A new rank-based version of the ant system: a computational study. Central European Journal of Operations Research, 7(1):25–38, 1999.

[2] G. Desaulniers, J. Desrosiers, A. Erdmann, M.M. Solomon, and F. Soumis. VRP with pickup and delivery. In Paolo Toth and Daniele Vigo, editors, The Vehicle Routing Problem, pages 225–242. SIAM, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 2000.

[3] M. Dorigo, M. Birattari, and T. Stützle. Ant Colony Optimization: Artificial Ants as a Computational Intelligence Technique. IEEE Computational Intelligence Magazine, 1(4):28-39, 2006.

[4] M. Dorigo. Optimization, learning and natural algorithms. PhD thesis, Dipartimento di Elettronica, Politecnico di Milano, Italy, 1992.

[5] M. Dorigo, G. Di Caro, and L.M. Gambardella. Ant algorithms for discrete optimization. Artificial Life, 5:137–172, 1999.

[6] M. Dorigo and L.M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. IEEE Transactions on Evolutionary Computation, 1(1):53–66, 1997.

[7] M. Dorigo, V. Maniezzo, and A. Colorni. Ant system: optimization by a colony of cooperating agents. IEEE Transactions on Systems, Man and Cybernetics - Part B: Cybernetics, 26(1):29–41, 1996.

[8] M. Dorigo and T. Stützle. Ant Colony Optimization. MIT Press, Cambridge, Massachusetts, 2004.

[9] L.M. Gambardella, E. Taillard, and G. Agazzi. MACS-VRPTW: a multiple ant colony system for vehicle routing problems with time windows. In D. Corne, M.Dorigo, and F. Glover, editors, New ideas in optimization, pages 63–76. McGraw-Hill, London, United Kingdom, 1999.

[10] M. Gendreau and J.-Y. Potvin. Dynamic vehicle routing and dispatching. In T.G. Crainic and G. Laporte, editors, Fleet management and logistic, pages 115–226. Springer, Berlin, Germany, 1998.

[11] M. Guntsch and M. Middendorf. Pheromone modification strategies for ant algorithms applied to dynamic TSP. In E.J.W. Boers et al., editor, Applications of Evolutionary Computing: EvoWorkshops 2001: EvoCOP, EvoFlight, EvoIASP, EvoLearn, and EvoSTIM, Como, Italy, April 18-20, 2001, Proceedings, volume 2037 of Lecture Notes in Computer Science, pages 213–222. Springer, Berlin, Germany, 2001.

[12] J. Kytöjoki, T. Nuortio, O. Bräysy, and M. Gendreau. An efficient variable neighborhood search heuristic for very large scale vehicle routing problems. Computers and Operations Research, 34(9):2743–2757, 2007.

[13] M. Labbe, G. Laporte, and H. Mercure. Capacitated vehicle routing on trees. Operations Research, 39(4):616–622, 1991.

[14] F. Li, B. Golden, and E. Wasil. Very large-scale vehicle routing: New test problems, algorithms, and results. Computers and Operations Research 32(5):1165-1179, 2005.

[15] D. Mester and O. Bräysy. Active guided evolution strategies for the large scale vehicle routing problem with time windows. Computers & Operations Research, 32(6):1593–1614, 2005.

[16] Metaheuristics Network web site. Coordinated by Prof. M. Dorigo: http://www.metaheuristics.org. Accomplished in 2004.

[17] R. Montemanni, L.M. Gambardella, A.E. Rizzoli, and A.V. Donati. Ant colony system for a dynamic vehicle routing problem. Journal of Combinatorial Optimization, 10:327–343, 2005.

[18] M.W.P. Savelsbergh. Local search in routing problems with time windows. Annals of Operations Research, 4:285–305, 1985.

[19] T. Stützle and H.H. Hoos. MAX-MIN ant-system. Future Generation Computer Systems, 16(8):889–914, 2000.

[20] P. Toth and D. Vigo. An overview of vehicle routing problems. In P. Toth and D. Vigo, editors, The Vehicle Routing Problem, pages 1–26. SIAM, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 2001.

About the authors

Dario Coltorti received the Master degree in Business Administration from the University of Berne (lic.rer.pol.), Switzerland in 2005. Since 2005 he is one of the responsible for Business Development of AntOptima, the spin-off company of Swiss research Institute Dalle Molle for Artificial Intelligence (IDSIA). He accelerated the market introduction of innovative optimization products of AntOptima. With AntRoute AntOptima became laureate of Swiss Technology Award 2007 and was invited at the Hannover Fair "get new technology first" 2007 to show the latest and most advanced products for the logistics, transportation and production sector. Meantime AntOptima continued to acquire further European and Swiss customers. One of the latest clients of AntOptima is Kamps, Germany's leading bakery specialist.

Company webpage: www.antoptima.ch
Email: dario.coltorti@antoptima.ch

Dr Andrea E. Rizzoli is a co-founder of AntOptima. Dr Rizzoli is also senior researcher at IDSIA and professor at SUPSI (University of Applied Sciences of Southern Switzerland). He holds a M.Sc. in Control Engineering (1989) and a Ph.D. in Control Engineering and Informatics (1993). His research interests are in modeling and simulation of dynamic systems, in the development of decision support systems for resources management, and in the application of artificial intelligence and operations research techniques to transportation and logistics problems. Dr Rizzoli has authored more than 80 scientific publications. He is editor for informatics of the international journal “Environmental Modelling & Software”.

Company webpage: www.antoptima.ch
Email: andrea.rizzoli@antoptima.ch