A VEHICLE ROUTING PROBLEM WITH STOCHASTIC TRAVEL TIMES

T. Van Woensel, L. Kerbache, H. Peremans and N. Vandaele

University of Antwerp, +32 3 220 40 69, tom.vanwoensel@ua.ac.be
HEC School of Management, +33 1 39 67 72 12, kerbache@hec.fr
University of Antwerp, +32 3 220 42 12, herbert.peremans @ua.ac.be
University of Antwerp, +32 3 220 41 59, nico.vandaele@ua.ac.be

Abstract:
Transportation is a main component of supply chain competitiveness since it plays a major role in the inbound, inter-facility, and outbound logistics. In this context, assigning and scheduling vehicle routing is a crucial management problem. Despite numerous publications dealing with efficient scheduling methods for vehicle routing, very few addressed the inherent stochastic nature of this problem. In this paper, we consider a vehicle routing problem and stochastic travel time due to potential traffic congestion. The approach developed introduces mainly the traffic congestion component that is combined with an Ants Colony Optimization heuristic. Standard test problems are used for illustrative purposes as well as for a discussion around the feasibility of proposed solutions if travel times were not constant.

Keywords:
Vehicle Routing Problems, stochastic travel times, queueing, Ants Colony Optimization

Introduction

The vehicle routing problem (VRP) can be described as a more general version of the well-known travelling salesman problem (TSP). In the TSP, a travelling salesman has to visit a set of customers located in different cities. The objective of the TSP is to find the shortest possible way to visit all customers. There are two constraints that the final route must satisfy: first, the route must start and end in the same town; secondly, he has to visit each customer, that is each city, exactly once.

The VRP aims to construct a set of shortest routes for a fleet of vehicles of fixed capacity. Each customer is visited exactly once by one vehicle which delivers the demanded amount of goods to the customer. Each route has to start and end at a depot, and the sum of the demands of the visited customers on a route must not exceed the capacity of the vehicle. Another constraint occuring in the real-world is that the customer may specify time intervals in which he will be able to receive the deliveries. This additional restriction leads to the vehicle routing problem with time window constraints (VRPTW). If time-window constraints are added to the TSP without a vehicle capacity constraint, the resulting problem is known as the travelling salesman problem with time windows (TSPTW). These time windows can be either soft or hard.

In this paper, it is assumed that there is only one depot from where the routes start and end for each vehicle, a homogeneous fleet consisting of several vehicles with fixed capacity, while the customers' demands are pre-determined. Formally, the vehicle routing problem can be represented by a complete weighted graph $G=(V,A,c)$ where $V=\{0,1,...,n\}$ is a set of vertices and $A=\{(i,j):i<>j\}$ is a set of arcs. The vertex 0 denotes the depot, the other vertices of V represent cities or customers. The non-negative weights c which are associated with each arc (i,j) represent the cost (distance, travel time or travel cost) between i and j. For each customer, a non-negative demand q_i and a non-negative service time d_i is given ($d_0 = 0$ and $q_0=0$). The aim is then to find the minimum cost vehicle routes where the following conditions hold: every customer is visited exactly once by exactly one vehicle; all vehicle routes start and end at the depot; every vehicle route has a total demand not exceeding the vehicle capacity Q; every vehicle route has a total route length not exceeding the maximum length L (Laporte 1992).

This paper is organized as follows: first, a brief literature overview of the stochastic travel time VRP is presented. Secondly, literature on the ants colony approach as an analysis and optimization methodology to cope with the VRP problem is also reviewed. The next section presents the ants colony approach to the deterministic VRP model. Then, the methodology to incorporate stochastic travel times in the deterministic VRP is described. The stochastic travel times are modelled by using a queueing approach to traffic flows. Computational results of the stochastic travel times VRP are presented. Finally, the paper ends with some conclusions and directions for future research.

Literature review

This literature review is split up into two parts: first the literature on vehicle routing problems with stochastic travel times is discussed and secondly, the literature related to ants colony optimization algorithms is overviewed.

Vehicle Routing Problems

Laporte (1992) surveys the main results of research on the general VRP. To solve the general VRP both exact and heuristic methods are available. Laporte classifies exact algorithms into three categories: direct tree search methods, dynamic programming and integer programming. Heuristic algorithms are more successful at solving VRP because of its complexity. Many heuristic methods for the VRP are derivations of methods developed for the TSP. We mention the nearest neighbour algorithm, insertion algorithms, and tour improvement procedures. These can be used for the VRP with only minor modifications. However, due to stronger constraints, the feasibility of the solution must be checked (Laporte 1992). Nevertheless, there are some algorithms that are specifically developed for the VRP. Among these, one can mention the savings algorithm (1964), the sweep algorithm (1974), the Christofides-Mingozzi-Toth two-phase algorithm (1979), and a tabu search algorithm by Genreau, Hertz and Laporte (1991). The heuristics are basically classified (see Osman 1993) as follows: constructive heuristics (e.g. the savings algorithm), two-step heuristics (e.g. cluster-first-route-second methods, route-first-cluster-second methods), incomplete optimisation, local search heuristics, metaheuristics and space filling curves.

Most of the above models, assume that all characteristics are deterministic and therefore are far from real-life applications. In this paper, the VRP problem considered deals with stochastic travel times. The motivation for using dynamic models is that the vehicles in the VRP operate in a traffic network which will be congested depending upon the time of the day. The literature related to vehicle routing with stochastic travel time is very scarce (Ichoua et al. 2003, Malandraki and Daskin 1992, Hill and Benton 1992, Malandraki and Dial 1996). The reason is that the time-dependent VRP is much harder to model and to solve.

In the stochastic VRP, the non-negative weights c wich are associated with each arc (i,j) represent the time-dependent travel time between i and j. Whereas in the deterministic VRP, these costs are associated with distance. The major shortcoming of the available literature is the modelling of the travel time function. It is often discretized into a limited fixed number of time intervals (e.g. morning, midday and afternoon) with a distinct associated fixed speed. However, these speeds are modelled in an arbitrary way. For instance, Brown et al. (1987) and Shen and Potvin (1995) used a rough approximation of travel time by manually resequencing the route taking into account congestion. Ichoua et al. used a model based on discrete travel time by adding correction factors to model the congestion. Models based on continuous travel times are very complex to solve (see Hickman and Bernstein 1997) and thus very simplistic assumptions had to be introduced to reach a solution.

Here the proposed methodology is to incorporate a queueing model to deal with the stochastic nature of travel times. It is shown that this queueing approach is a sound conceptual framework to model the stochastic, time-dependent speeds. Moreover, the VRP with stochastic travel times will be solved using the ants colony optimization heuristic. A brief literature review on both the ants colony optimization approach and its application to the VRP is discussed in the next section.

Ants Colony Optimization

The ACO algorithm is a stochastic optimization algorithm specifically intended to solve discrete optimization problems, like the above described VRP.

The inspiration of the ACO algorithm comes from the observation of the trail laying and the trail following behavior of a real ant species (Linepithaeme humile). As the ants move in search for food, they deposit an aromatic essence called pheromone on the ground. The amount deposited generally depends upon the quality of food sources found. Other ants, observing the pheromone are more likely to follow the pheromone trail, with a bias towards stronger trails. As such, the pheromone trails reflect the memory of the ant population and over time, trails leading to good food sources will be reinforced while paths leading to remote sources will be abandoned (Corne, Dorigo and Glover, 1999).

The above behavior of the real ants is then translated in the ACO algorithm. Artificial ants construct solutions for a given combinatorial optimization problem by taking a number of decisions probabilistically. In a first stage, these decisions are based on local information only (e.g. an heuristic rule). Gradually a trail emerges as the ants lay aritificial pheromone on the paths they followed. This pheromone dropping is dependent upon the quality of the solution they found: poor solutions get less pheromone than better ones. The other ants are then guided in their decision making not only by the local information but also by the collective memory based on the pheromone. Over time, paths with a high pheromone concentration will be reinforced and the artificial ants are guided to promising regions of the search space.

This approach has been succesfully applied to a number of combinatorial optimization problems: the Graph Coloring Problem, the Travelling Salesman Problem, the Quadratic Assignement Problem, the Vehicle Routing problem and the Vehicle Routing Problem with Time Windows (see e.g. Corne, Dorigo and Glover, 1999 for references).

The VRP with deterministic travel times

In the following outline, the ACO meta-heuristic is described in pseudo-code.

Initialization
For I^{max} iterations do
    For each ant $k=1,...,m$ generate a new solution
    Daemon actions to improve the solution
    Update the pheromone trails

After initializing the pheromone level to a small value next to zero, the heuristic starts building routes for a certain prefixed number of iterations I^{max}. Each step in the iterations part of the ACO meta-heuristic, is discussed in more detail in the following sections. Bullnheimer, Hartl and Strauss (1999) showed that the ants colony optimization approach is competitive when compared with other meta-heuristics such as Tabu Search, Simulated Annealing and Neural networks. The ACO approach gives results that are within an average deviation of less than 1.5% over the best known solutions.

Generating new solutions

Ants start out in a randomly chosen first city. Next, they successively choose new cities to visit until all cities are visited. The depot is selected whenever the vehicle capacity or the total route length is exceeded. In this case, a new route starting at the depot is started. The selection of a new city to visit is based upon a probabilistic decision rule taking into account two aspects: how good was the choice of that city in the past and how promising is the choice of that city at this moment. The first aspect is information stored in the phermone trails associated with each arc (i,j), i.e. memory from previous iterations. The second aspect is called visibility which stores the local heuristic information (Bullnheimer, Hartl and Strauss, 1999).

An ant chooses its next city to visit only if there are potential savings to be gained. If no positive savings can be made, the ant returns to the depot and starts a new route. The savings are evaluated using the standard Clarke and Wright savings algorithm : $s_{ij} = d_{i0} + d_{0j} - d_{ij}$. With possible positive savings, ants probabilistically choose their next customer based on the following probability distribution:

$$pij = \frac{[t_{ij}]^a [h_{ij}]^b [k_{ij}]^l}{\sum [t_{ij}]^a [h_{ij}]^b [k_{ij}]^l}$$

$$pij = 0$$

If j is part of Λ= {j in V: j is feasible to be visited} then the probability is larger than zero. The above probability distribution combines both memory and visibility aspects. It is biased by the parameters a, b, l that determine the relative influence of the pheromone at the trails and the different parts of the visibility (h_{ij}, k_{ij}) respectively. The various variables t_{ij}, h_{ij}, k_{ij} are now discussed in more detail.

The first variable t_{ij} refers to the pheromone level at arc *(i,j)*. This variable gets updated for each iteration and will be a representation of the quality of the solutions found in the past. The pheromone updating procedure is discussed in subsection 3.3. The above function is then extended with VRP specific information. In the case of the deterministic VRP, h_{ij} is set equal to the reciprocal of the distance between i and j: $1/d_{ij}$. This formulation was the original approach used by Dorigo and Gambardella (1997) for the Travelling Salesman Problem.

Secondly, a parameter k_{ij} being the degree of capacity utilization is introduced. The idea is that selecting a city that leads to a higher degree of utilization of the vehicle, is to be prefered. The degree of capacity utilization k_{ij} (the vehicle is in city i and has used up sofar a capacity of Q_i and wants to go to city j), is defined as: $(Q_i + q_j)/Q$ (Bullnheimer, Hartl and Strauss, 1999c).

Daemon actions to improve the solutions

After an artifical ant k has constructed a feasible solution, there are two possible actions to take: either the pheromone trails get updated immediatly using the solutions found or the solutions are first improved using a daemon action. In the general ACO meta-heuristic, the daemon actions are optional, but experiments have shown (Bullnheimer, Hartl and Strauss 1999b; Bullnheimer, Hartl and Strauss 1999c) that in the case of VRP, the daemon actions greatly improve the solution quality and the speed of convergence.

In this paper, at the end of the solution generation, all routes are checked for 2-optimality and are improved if possible. A route is 2-optimal if it is impossible to improve the route by exchanging two arcs. This is the VRP implementation of the so-called 2-opt heuristic for the TSP. When all solutions are improved by the 2-opt heuristic, these solutions are then used to update the pheromone trails.

Update the pheromone trails

After the route construction and the improvement of these original routes, the pheromone trails deposited on the different links depending upon the solution quality. The solution quality is the objective value obtained which equals the total cost (e.g. either distance, travel time or cost) of the route for ant k defined as L_k. For each arc *(i,j)* part of the route used by ant k, the pheromone is increased by $\Delta t^k_{ij}=1/ L_k$. Using this updating rule, more costly routes will increase the pheromone levels on the arcs less than lower cost routes.

In addition to the pheromone updates using the routes of all the ants, all arcs belonging to the so far best solution are emphasized as if σ ants (the so-called elitist ants) had used them. One elitist ant increased the pheromone level by an amount $\Delta t^*_{ij} = 1/L^*$ if arc *(i,j)* is part of the best route so far (Bullnheimer, Hartl and Strauss, 1999b). In a last step, part of the existing pheromone trails evaporates with a factor $(1\text{-}r)$, r being the trail persistence and $0 £ r £ 1$. Pheromone evaporation is needed to avoid a too rapid convergence of the algorithm towards a sub-optimal region. It implements a form of forgetting, making the exploration of new areas of the search space possible.

Summarizing, the pheromone update after iteration t is done as follows:

$$t_{ij}^{t} = r\,t_{ij}^{t-1} + \sum_{k=1}^{m} \Delta t_{ij}^{k} + s\Delta t_{ij}^{*}$$

Where r being the trail persistence and 0 £ r £ 1, σ the number of elitist ants, Δt^{k}_{ij} ¹ 0 if arc (i,j) is used by ant k and Dt^{*}_{ij} ¹ 0 if arc (i,j) is used by the ant with the best solution so far. Using this pheromone update rule, arcs used by many ants and which are contained in shorter tours will receive more pheromone and will be more likely to be chosen in future iterations of the algorithm (Dorigo and Stutzle).

The VRP model with stochastic travel times

In this section, the methodology to incorporate stochastic travel times in the deterministic VRP is elaborated in detail. Both the deterministic and the stochastic VRP will be modelled using an Ants Colony Optimization (ACO) algorithm. This section is organized as follows: first the queueing approach to determine the stochastic travel times is explained in detail and secondly, the integration of the stochastic travel times in the ACO models for the VRP is elaborated.

In the stochastic travel time VRP, the key issue is the computation of the travel times on arcs dependent upon the time period. The travel time T^{p}_{ij} during time period p is determined as: d_{ij}/v_{ij}^{p}. Hence, to determine the travel time on arc (i,j), one needs information on the distance between (i,j) and on the travel speed on that arc at time p: v_{ij}^{p}. The distance is already available in the deterministic VRP models, but the speed still needs to be specified. In this paper, the time-dependent speeds are obtained using queueing models for traffic flows (Vandaele, Van Woensel and Verbruggen, 2000 and Heidemann, 1997).

Queueing approach

It is often observed that the speed for a certain time period tends to be reproduced whenever the same flow is observed. Based on this observation, it seems reasonable to postulate that, if traffic conditions on a given road are stationary, there should be a relationship between flow, speed, and density. This relationship results in the concept of speed-flow-density diagrams. These diagrams describe the interdependence of traffic flow (q), density (k) and speed (v). The seminal work on speed-flow diagrams was the paper by Greenshields in 1935.

Using well-known formulas of queueing models, these speed-flow-density diagrams can be constructed. Traditionally, these speed-flow-density diagrams are modeled empirically: speed and flow data are collected for a specific road and econometrically fitted into curves (Daganzo 1996). This traditional approach is limited in terms of predictive power and sensitivity analysis. Vandaele, Van Woensel and Verbruggen (2000) and Heidemann (1997), showed that queueing models can also be used to explain uninterrupted traffic flows and thus offering a more practical approach, useful for sensitivity analysis, forecasts, etc. Jain and Smith (1999) describe in their paper a state-dependent M/G/C/C queueing model for traffic flows. Part of their logic is used here to extend our queueing models to state-dependent ones.

In a queueing approach to traffic flow analysis, roads are subdivided into segments, with length equal to the minimal space needed by one vehicle on that road.

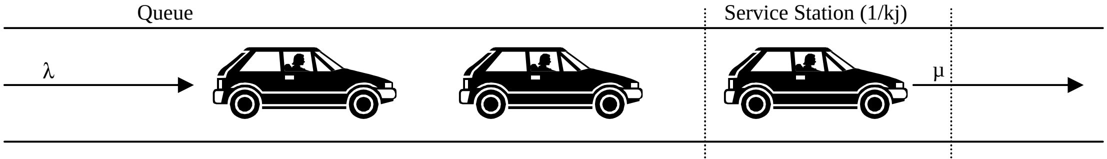

Define k_j as the maximum traffic density (i.e. average maximum number of cars on a road segment). This length is then equal to $1/ k_j$ and matches the minimal space needed by one vehicle on that road. Each road segment is then considered as a service station, in which vehicles arrive at a certain rate lambda and get served at another rate μ (Vandaele, Van Woensel and Verbruggen, 2000; Van Woensel, Creten and Vandaele, 2001; Heidemann, 1997).

Vandaele, Van Woensel and Verbruggen (2000) developed different queueing models. The M/M/1 queueing model (exponential arrival and service rates) is considered as a base case, but due to its specific assumptions regarding the arrival and service processes, it is not useful to describe real-life situations. Relaxing the specifications for the service process of the M/M/1 queueing model, leads to the M/G/1 queueing model (generally distributed service rates). Relaxing both assumptions for the arrival and service processes results in the GI/G/m queueing model. Moreover, following Jain and Smith (1997) a special case of the GI/G/m queueing model is derived: a state dependent GI/G/m queueing model. This model assumes that the service rate is a (linear, exponential,etc.) function of the traffic flow. In this case vehicles are served at a certain rate, which depends upon the number of vehicles already on the road. It can be shown that the speed v is calculated by dividing the length of the road segment $1/ k_j$ by the total time in the system (W). The total time in the system W is then different depending upon the queueing model used (Vandaele, Van Woensel and Verbruggen 2000).

In general, the speed can be expressed in the following basic form: $v = v_f /(1+W)$. This formula shows that the speed is only equal to the maximum speed v_f if the factor W is zero. For positive values of W, v_f is divided by a number strictly larger than 1 and speed is reduced. The factor W is thus the influence of congestion on speed. High congestion (reflected in a high W) leads to lower speeds than the maximum. The factor W is a function of a number of parameters depending upon the queueing model chosen: the traffic intensity, the coefficient of variation of service times and coefficient of variation of interarrival times. High coefficients of variation or a high traffic intensity will lead to a value of W strictly larger than zero. Actions to increase speed (or decrease travel time) should then be focussed on decreasing the variability or on influencing the traffic intensity, for example by manipulating the arrivals (arrival management and ramp metering).

Results show that the developed queueing models can be adequately used to model traffic flows (Van Woensel, 2003). Moreover due to the analytical character of these models, they are very suitable to be incorporated in other models, e.g. the VRP. For a more detailed discussion of the queueing models and their results, the interested reader is referred to Vandaele, Van Woensel and Verbruggen (2000) and Van Woensel, Creten and Vandaele (2001).

Computation of the stochastic travel times

To compute the travel times, one should note that in the stochastic case, the travel speeds are no longer constant over the entire length of the arc. More specifically, one has to take into account the change of the travel speed when the vehicle crosses the boundary between two consecutive time periods. The time horizon is discretized into P time periods of equal length Dp with a different travel speed associated to each time period p (1 £ p £ P). The travel speeds are obtained using the above discussed queueing models for traffic flows. In general the travel time T_{ij} from customer i to customer j, starting at time p_0 can be obtained as:

$$T_{ij} = j\,\Delta p_{first} + (k-2)\Delta p + f\Delta p_{last}$$

With Dp_{first} the first time zone which contains p_0 and Dp_{last} the last time zone used to cover the distance d_{ij}. The travel time is thus the sum of the following components: the fraction of travel time still available in the first time zone, given by $j\,\Delta p_{first}$, the travel times of the $(k-2)$ intermediate time zones passed: $(k-2)D$p and the fraction of the travel time in the last time zone, given by $f\Delta p_{last}$.

Stochastic travel times in ACO

In this section, the integration of the stochastic travel times in the ACO approach for the VRP is discussed. This section is split up into three parts based on the original ACO heuristic: first, the solution generation is discussed, then the daemon actions are described and finally, the pheromone updating is elaborated in detail.

Generating new solutions

Again as before, the choice of the next city is based on the probability function defined above. However, we now have to take into account the time-dependent travel times. Consequently, the objective function shifts now from minimizing the total distance to minimizing the total travel time. An ant chooses its next city to visit only if there are potential savings to be gained, otherwise the depot is selected. The savings are however now based on travel times. The new savings formula is then changed to:

$$s_{ij} = T_{i0} + T_{0j} - T_{ij}.$$

With possible positive savings, ants probabilistically again choose their next customer based on a similar probability distribution defined above. The following table compares the deterministic formulae in terms of distance and the stochastic formulae in terms of travel times for the different parameters in the ACO model. Note that in the stochastic case, both the pheromone level t_{ij} and the visibility h_{ij} have an extra time dimension based on the different time zones and are now denoted as t^p_{ij} and h^p_{ij}.

	Deterministic ACO	Stochastic ACO
Objective function	Min(Total Distance)	Min(Total Travel Time)
t_{ij}	t_{ij}	t^p_{ij}
h_{ij}	$1/d_{ij}$	$1/T_{ij}$
k_{ij}	$(Q_i + q_j)/Q$	$(Q_i + q_j)/Q$

Daemon actions to improve the solutions

At the end of the solution generation, all routes are checked as before for 2-optimality and are improved if possible. A route is 2-optimal if it is not possible anymore to improve the route by exchanging two arcs. Unlike in the deterministic VRP where the gain is calculated based on distances, in the stochastic VRP, the gain is calculated in terms of travel time.

Moreover, extra improvement heuristics are performed taking into account explicitly the time-dependent nature of the problem: all starting times of the different subtours that make up a complete VRP solution, are shifted in time to evaluate the effect of the start time on the total travel time. In case of improvement, the starting time of the associated subtour is updated. The rationale behind this optimization is that in a dynamic reality, a truck can decide to leave earlier or later to avoid periods of high congestion.

Update the pheromone trails

After the route construction and the improvement of the original routes, the pheromone trails are laid depending upon the solution quality. The pheromone t^p_{ij} is now updated on each link when the ant crosses that link *(i,j)* during the time interval *p*. Moreover, the neighbouring time intervals for the same link *(i,j)* are also incremented using the following pheromone allocation rule:

$$t^p_{ij} = \frac{1}{L_k}$$

$$t^p_{ij} = g^z \frac{1}{L_k}, z \in [p - Z, p + z]$$

With Z being the range of the pheromone propagation over the different neighbouring time zones. The above pheromone allocation rule results in a symmetrical, bell-shaped distribution over all relevant time zones. This mechanism is utilized to incorporate the variation around the expected travel time. As such, it can be used as a rough approximation of a confidence interval around the mean travel time.

Preliminary computational results

The VRP with stochastic travel times was tested on one benchmark problems described in Christofides et al. (1989). These problems contain between 50 and 199 customers in addition to the depot. The customers in problems 1-5 are randomly distributed in the plane, while they are clusters in problem instances 6 and 7. For this paper, the 50 customer dataset is used to gain some initial insights in the specific behavior of the VRP with stochastic travel times. For this instance, the truck capacity is 160 and there is no restriction on the tour length.

The problem instances are first used to obtain the best solution for the deterministic VRP, i.e. minimizing total distance. Then, using the speeds obtained from the queueing models, the obtained deterministic route can be recalculate in terms of total travel time. In a last step, the VRP which immediately takes into account stochastic travel times is solved and compared with the latter.

Following Bullnheimer et al. (1999c), the different parameters are set for the deterministic case to: $a=b=l=5$ and σ = the number of customers in the problem instance. The number of iterations I^{max} is also set equal to the number of cities in the problem instance and the evaporaton rate r is set to 0.75. For the stochastic VRP, the different parameters are set $a=b=l=5$ and σ = the number of customers in the problem instance. Based on the experiments, the number of iterations however is now equal to two times the number of customers, the parameter γ is set equal to 0.3 and the evaporation rate r is set to 0.45. The number of artificial ants is set equal to the number of cities in both the deterministic as the stochastic VRP. The length of the time period is set equal to ½ hour. The maximum number of time periods considered is 196.

The following table, shows the results for the deterministic case in terms of travel time, compared with the best, average and worst solution obtained for the stochastic VRP over 10 runs. The first experiment shows the results without the daemon action of shifting the subtours in time. The second experiment shows the results with this daemon action activated. The starting time was shifted over 10 time periods, which is equivalent with a decision starting in the morning or in the afternoon.

	Deterministic	Stochastic		
		Best	Average	Worst
No daemon	49.42	49.20	50.49	52.83
With daemon	49.42	48.02	48.69	49.24

Results show that the results improve when adding the daemon actions that explicitly take into account the dynamic character of the problem. Moreover, the spread between the best and the worst solution reduces considerably when improving the solutions with the daemon actions. Concluding, starting your subtour earlier or later results in an improvement of almost 2 hours or almost 5%.

Conclusions and Future research

In this paper, we considered a vehicle routing problem with stochastic travel time due to potential traffic congestion. The approach developed introduces mainly the traffic congestion component in the standard VRP models. The traffic congestion component was modelled using a queueing approach to traffic flows. Both the deterministic as the stochastic VRP are then modelled using the Ants Colony Optimization heuristic. Results show that the total travel times can be improved significantly when explicitly taking into account congestion in the Ants Colony Optimization heuristic. Using the information of congestion, will result in routes that are shorter in terms of travel time. Shortly, results for more and different datasets (e.g. more customers, etc.) will be made public.

In future, more extensions will be made to this basic framework to solve the stochastic VRP. First, the daemon actions that explicitly take into account congestion seem to contribute substantially to the solution quality. Consequently, new *dynamic* improvement heuristics should be developed. For instance, once the total route is known, there is the possibility that it can be improved by adding the depot in each subtour. This will result in a shorter route in terms of travel time but the route will be covering more distance with an additional truck. The developed approach will also be applied to the VRP with the addition of time windows.

References

1. Brown, G.G., C.J. Ellis, G. Lenn, W. Graves and D. Ronen (1987), Real time, wide area dispatch of mobile tank trucks, Interfaces 17, 107-120
2. Bullnheimer B., R.F. Hartl and C. Strauss (1999). An Improved Ant system Algorithm for the Vehicle Routing Problem. Paper presented at the Sixth Viennese workshop on Optimal Control, Dynamic Games, Nonlinear Dynamics and Adaptive Systems, Vienna (Austria), May 21-23, 1997, to appear in: Annals of Operations Research (Dawid, Feichtinger and Hartl (eds.): Nonlinear Economic Dynamics and Control, 1999.
3. Bullnheimer B., R.F. Hartl and C. Strauss (1999b). Applying the Ant System to the Vehicle Routing Problem. In: Voss S., Martello S., Osman I.H., Roucairol C. (eds.), Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization, Kluwer:Boston.
4. Bullnheimer, B., Hartl, R.F., Strauss, C.(1997c), An improved ant system algorithm for the vehicle routing problem, POM Working Paper 10/97
5. Clarke, G. and Wright, J.W. Scheduling Of Vehicles From A Central Depot To A Number Of Delivery Points, Operations Research, 12(4), 1964, pp. 568-581
6. Christofides, N. and Mingozzi, A., Vehicle Routing: Practical and Algorithmic Aspects in C.F.H. van Rijn (ed.) Logistics where the ends have to meet, Pergamon Press, Oxford, 1989
7. Corne D., M. Dorigo, and F. Glover, New ideas in optimization, 1999, McGraw-Hill
8. Daganzo, C.F., 1997, Fundamentals of Transportation and Traffic Operations, Elsevier Science Ltd., Oxford
9. Dorigo M. & L.M. Gambardella (1997). Ant Colonies for the Traveling Salesman Problem. BioSystems, 43:73-81.
10. Dorigo M. and T. Stützle. The ant colony optimization metaheuristic: Algorithms, applications and advances. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, volume 57 of *International Series in Operations Research & Management Science*, pages 251-285. Kluwer Academic Publishers, Norwell, MA, 2002.
11. Gillett, B.E. and Miller, L.R. A Heuristic Algorithm for the Vehicle-Dispatch Problem, Operations Research, 22(2), 1974, pp. 340-349
12. Heidemann, D., 1996, A queueing theory approach to speed-flow-density relationships, Transportation and Traffic Theory, Proceedings of the 13th International Symposium on Transportation and Traffic Theory, Lyon, France, 14-26 July 1996
13. Hickman, M.D. and N.H.M. Wilson (1995), Passenger travel time and path choice implications of real-time transit information, Transportation Research C 3 (4), 211-226
14. Hill, A.V. and W.C. Benton, Modeling intra-city time-dependent travel speeds for vehicle scheduling problems, European Journal of Operational Research, 43(4), 1992, pp. 343-351
15. Ichoua, S., M. Gendreau and J-Y. Potvin, Vehicle dispatching with time-dependent travel times, European Journal of Operational Research, 144, 2003, pp. 379-396
16. Jain, R. and Smith J.M., 1997, Modelling vehicular traffic flow using M/G/C/C state dependent queueing models, Transportation Science, 31, 324-336
17. Laporte, G. The Vehicle Routing Problem: An overview of exact and approximate algorithms, European Journal of Operational Research, 59(3), 1992, pp. 345-358
18. Malandraki, C. and M.S. Daskin (1992), Time dependent vehicle routing problems: formulations, properties and heuristic algorithms, Transportation science 26(3), 185-200
19. Malandraki, C. and R.B. Dial (1996), A restricted dynamic programming heuristic algorithm for the time dependent traveling salesman problem, European Journal of Operational Research 90, 45-55

20. Osman, I., Vehicle Routing and Scheduling: Applications, Algorithms and Developments, Proceeding of the International Conference on Industrial Logistics, Rennes, 1993
21. Shen Y. and J.-Y. Potvin, A computer assistant for vehicle dispatching with learning capabilities, Operations Research 61, 189-211
22. Vandaele N., T. Van Woensel and A. Verbruggen, 2000, A queueing based traffic flow model, Transportation Research-D: Transport and environment, January, vol. 5 nr. 2, pp. 121-135
23. Van Woensel T. (2003), Modelling uninterrupted traffic flows, a queueing approach, Ph.D. Dissertation, University of Antwerp, Belgium
24. Van Woensel T., Creten R. and N. Vandaele (2001), Managing the environmental externalities of traffic logistics: the issue of emissions, POM journal, Special Issue on Environmental Management and Operations, Vol. 10, nr 2, Summer 2001