Vehicle routing with stochastic time-dependent travel times

4OR-Q J Oper Res (2009) 7:363–377
DOI 10.1007/s10288-009-0097-9

RESEARCH PAPER

Vehicle routing with stochastic time-dependent travel times

C. Lecluyse · T. Van Woensel · H. Peremans

Received: 15 January 2008 / Revised: 16 July 2008 / Published online: 25 February 2009

Abstract Assigning and scheduling vehicle routes in a stochastic time-dependent environment is a crucial management problem. The assumption that in a real-life environment everything goes according to an a priori determined static schedule is unrealistic. Our methodology builds on earlier work in which the traffic congestion is captured in an analytical way using queueing theory. The congestion is then applied to the *VRP* problem. In this paper, we introduce the variability in traffic flows into the model. This allows for an evaluation of the routes based on the uncertainty involved. Different experiments show that the risk taking behavior of the planner can be taken into account during optimization. As more weight is given to the variability component, the resulting optimal route will take a slightly longer travel time, but will be more reliable. We propose a powerful objective function that is easily implemented and that captures the trade-off between the average travel time and its variance. The evaluation of the solution is done in terms of the 95th-percentile of the travel time distribution (assumed to be lognormal), which reflects well the quality of the solution in this stochastic time-dependent environment.

Keywords Vehicle routing · Time-dependent travel times · Travel time reliability

MSC classification (2000) 90B06 · 90B15

C. Lecluyse (✉) · H. Peremans
Department of Applied Economics, University of Antwerp,
Prinsstraat 13, 2000 Antwerp, Belgium
e-mail: christophe.lecluyse@ua.ac.be

T. Van Woensel
Eindhoven University of Technology, Eindhoven, The Netherlands

1 Introduction

Most traffic networks in Europe face high utilization levels, and consequently, congestion occurs. When having a sufficiently high utilization, the smallest stochastic events (both in arrival processes or in service processes) cause waiting, which in the case of traffic systems, materializes in lower speeds. As speed changes, travel time will vary for a given distance. As such, all transportation problems which intend to minimize total time used, are subject to these physical considerations of congestion. Consequently, the road and traffic conditions and their resulting variability cannot be ignored in order to carry out a good quality route optimization. Uncertainty about the traffic conditions represented in travel times is a pervasive aspect of routing and scheduling, especially in a just-in-time environment or in highly congested regions like Europe. As the cost impact due to this uncertainty can be substantial, risk sensitive planners may want to evaluate to what extent their routes and schedules are risky in terms of travel times. Indeed, a slightly longer route in terms of expected travel time might be more interesting for a planner if the associated variance is considerably smaller. In this paper we deal with the *VRP* problem with stochastic time-dependent travel times. In a real-life environment the travel times on an individual link are stochastic in nature. Due to weather conditions, car accidents and congestion, among others, time and spatial fluctuations of traffic flows can be observed. The key issue considered in this paper is the variability of the travel times which we consider to be a good approximation of travel time reliability.

In Van Woensel et al. (2008), a dynamic vehicle routing problem with time-dependent travel times due to traffic congestion, was presented. The developed approach introduced the modeled traffic congestion component using a queueing approach to traffic flows. Explicitly making use of the time-dependent congestion results in routes that are (considerably) shorter in terms of travel time. Moreover, a first rough expression for the variance of the travel time was obtained (Van Woensel et al. 2008).

The main contributions of this paper are as follows:

1. We extend the objective function of a *VRP* with time-dependent travel times with the standard deviation of the travel time. As such, we control the degree of travel time variability during the optimization. The method presented is built on classical methods and as a consequence can be implemented quite easily.
2. We show that extending the objective function with this extra information about the stochastic travel time distribution provides interesting results when considering the reliability of travel times. Depending on environmental and road conditions as well as the risk-taking behavior of the planner, these improvements can be substantial.
3. The reduction of the travel time variability may come at the cost of an increase of the expected travel time. To evaluate a solution, we use the 95th-percentile of the travel time distribution (assumed to be lognormal) as a quality measure. Using this measure, the quality of the solution improves if the increase in expected travel time results in a lower travel time associated with the 95th-percentile of its distribution. Results show that the reliability improves as more weight is given to the variance

component during optimization. Different environmental and road conditions are compared and evaluated using this quality measure.
The trade-off between expected travel time and standard deviation of the travel time and the resulting reliability is demonstrated using an Arena simulation. The simulation validates the underlying assumption that the travel time over an entire tour is well approximated by a lognormal distribution.

This paper is organized as follows: in Sect. 2, the literature background on our *VRP* variant is presented, followed by a formal description of the *VRP*, the objective function and the time-dependency implementation in Sect. 3. The experimental design with the input data and a tabu search implementation are presented in Sect. 4. Results on the solution quality are presented in Sect. 5. Finally, conclusions and future research are presented in Sect. 6.

2 Literature review

Travel times between any two customers are a stochastic process related to traffic congestion. Depending on the time of the day the traffic network will face a different level of congestion. The number of vehicles, the road capacity, road conditions, weather conditions, …influence the speed of the vehicles. There has been limited research on solving the *VRP* problem with stochastic time-dependent travel times. In one of the first approaches Malandraki and Daskin (1992) treated the travel time between two customers as a function of distance and the time of the day (if this temporal component causes more travel time variation than travel time variation due to accidents, weather conditions,…), resulting in a piecewise constant distribution of the travel time. Although they only incorporate the temporal component of traffic density variability, they acknowledge the importance of the traffic density variability due to accidents, weather conditions and other random events. However, the FIFO principle is not necessarily satisfied with this approach (Ichoua et al. 2003).

Ichoua et al. (2003) introduced a model that guarantees that if two vehicles leave the same location for the same destination (along the same path), the one that leaves first will never arrive later than the other. This is satisfied by working with step-like speed distributions and adjusting the travel speed whenever the vehicle crosses the boundary between two consecutive time periods. To reduce computational run times, they limited the number of time slices to three. The speed differences are then modeled using correction factors on the weights of the links. Donati et al. (2003) extended this line of research by indicating the importance of optimizing the starting time in addition to optimizing the routes in a time-dependent environment. They show that the degree of feasibility (defined as not violating a time constraint) and optimality decreases for the best solutions for the constant speed model when they are used in a time-dependent context with increasing variability of the traffic conditions. Similar results were also observed by Haghani and Jung (2005). In contrast with Ichoua et al. (2003) they model travel time as a continuous function that can accept any kind of travel time variation.

As indicated by Ichoua et al. (2003) the literature on time-dependency in a *VRP* context is limited. Stochastic and time-dependent travel times are more extensively operated on in shortest path analysis (e.g., Hall 1986; Fu and Rilett 1998; Gao and Chabini 2002, 2006; He et al. 2005). He et al. (2005) indicate that although mean and variance contain the most important information about path travel time, finding the single route with expected shortest travel time is not appropriate for routing when planners are not risk neutral. The entire travel time distribution contributes to the routing choice. Chen et al. (2003) propose using the standard deviation and the 90th percentile travel time in addition to the mean to measure service quality.

Stochastic travel times are introduced in the vehicle routing problem by Laporte et al. (1992). Following Gendreau et al. (1996b) a stochastic *VRP* arises whenever some elements of the problem are random. A stochastic model usually involves two stages. In the first stage, a route is planned a priori, followed by a realization of the random variables. In the second stage, a recourse or corrective action is then applied to the solution of the first stage. The cost/saving generated through the recourse may have to be considered when designing the first stage solution.

3 Problem formulation

Formally, the routing problem considered can be represented by a complete directed graph $G = (V, A)$ where $V = \{0, 1, \ldots, n\}$ is a set of nodes representing the depot (0) and the customers $(1, \ldots, n)$, and $A = \{(i, j)|i, j \in V\}$ the set of directed links. For each customer, a fixed non-negative demand q_i is given ($q_0 = 0$). The aim is then to find routes with best travel time where the following conditions hold (Laporte 1992): every customer is visited exactly once by exactly one vehicle; all vehicle routes start and end at the single depot; every vehicle route has a total demand not exceeding the maximum vehicle capacity Q. Define a solution as a set S with m routes $R_1, \ldots, R_m$ where $R_r = (0, r_1, r_2, \ldots, 0)$ and each vertex $i \geq 1$ belongs to exactly one route. For ease of notation, write $i \in R_r$ if the node is part of the route R_r and write $(i, j) \in R_r$ if i and j are two consecutive nodes of R_r. Also define $E(\widetilde{TT}_{ij}^{t_0})$ and $Var(\widetilde{TT}_{ij}^{t_0})$ as the expected travel time and the variance of the travel time needed to cover the distance (i, j) leaving vertex i at time t_0.

Only taking the expected travel times into account ignores the risk profile of the planner. An extension of the objective function thus involves adding the standard deviation of the travel time. The weight of the latter is controlled by a (positive) parameter β, i.e.,

$$Min \sum_{r=1}^{m} \sum_{(i,j)\in R_r} E(\widetilde{TT}_{ij}^{t_0}) + \beta \sqrt{\sum_{r=1}^{m} \sum_{(i,j)\in R_r} Var(\widetilde{TT}_{ij}^{t_0})} \quad (1)$$

Note that the proposed approach is similar to the mean-variance analysis used in financial planning of portfolios (Best and Grauer 1991; Grauer and Hakansson 1993). In this literature it is argued that risk aversion can be modeled through the inclusion of a variance term in the objective function (Mulvey et al. 1995).

As the travel times are time-dependent, we need to be careful not to violate the FIFO principle (Ichoua et al. 2003). Any vehicle starting after time t, cannot arrive earlier than any vehicle starting before time t. To adhere this principle, the day is discretized into P time zones of equal length $\triangle p$ with a different travel speed distribution associated with each time zone p $(1 \leq p \leq P)$. A new speed is adopted as the vehicle crosses the boundary of a given time zone (Ichoua et al. 2003; Van Woensel et al. 2008).

In our model the speed in each time zone is obtained by applying queueing theory to traffic flows. Based on traffic counts, analytical queueing models model the behavior of traffic flows as a function of the most relevant physical and geographical determinants (i.e., free flow speed, maximum road capacity, variability due to the weather, ...). The travel times can then be modeled much more realistically using these speeds (i.e., expressed in kilometer per hour) and are directly related to the physical characteristics and the geographical location on the arc. For a discussion on the queueing model, the interested reader is referred to Vandaele et al. (2000), Heidemann (1996) and Van Woensel and Vandaele (2006).

One of the earliest studies explicitly dealing with the travel speed distribution is that of Berry and Belmont (1951) who looked into the distribution of the measured speed of a vehicle as it crosses a particular point on the highway. Speed was found to be normally distributed. Travel times, taken as the reciprocal of speed, are shown to be also roughly normal, although slightly skewed indicating that a lognormal distribution might be interesting as an alternative (Kharoufeh and Gautam 2004). Other empirical results (Taniguchi et al. 2001; Kwon et al. 2000) show that there is always a certain minimum time needed to cover the distance (i.e., it is impossible to traverse the distance in a time shorter than this minimum time). After this minimum time, the probability increases rapidly to a maximum after which the probability slowly decreases with a long tail (i.e., skewed to the right). Due to these characteristics, Taniguchi et al. (2001) proposed to use a lognormal distribution rather than a normal distribution.

It can be shown that the convolution of k lognormal distributions is again (approximative) lognormal (Beaulieu and Xie 2004). We assumed a lognormal distribution of the travel time, but the analysis could also be applied if another distribution was chosen. Simulation results in Sect. 5 show that the travel time distribution over an entire tour is again well approximated by a lognormal one.

4 Experimental design

In this section, we first describe the input data used. Secondly, we explain the impact of the problem as defined supra on the classical Tabu Search implementation.

4.1 Input data

4.1.1 Selection of speed profiles

For the subsequent analysis, we will create four test situations, based on two dimensions of interest. The first dimension is the level of congestion, the second dimension deals with the road and/or weather conditions.

For the first dimension, we will use two different speed profiles. The first speed profile is the result of a congested flow during the entire day, whereas the second one has heavy congestion only during the morning and evening peak hours. The reasons for this choice of speed profiles are twofold.

First, today there exists roads that follow those patterns. Some roads have a distinct morning and evening rush-hour, whereas other roads are congested for the better part of the day.

Secondly, in the near future it is very likely that the traffic volume will increase. The 10 Year Plan in the UK confirms this (DfT 2000). On top of that, it states that despite all governmental policies, the congestion level on all roads in the UK will increase. In 2000 however, the 10 Year Plan indicated that by 2010 the current level of congestion (base 2000) would decline by 5–6%. New measurements and predictions in 2003, however, indicated that by 2010 the traffic will increase by 11–20%, even if all the measures in the 10 Year Plan work as predicted and will increase by 27–32% otherwise (DfT 2000, 2003). As a consequence, we could also look at the chosen speed profiles as being the situation today vs the (nearby) future.

The second dimension deals with the impact of weather (and road conditions) on the travel times and especially the reliability of the best tour. This will provide us with four test situations, that will be investigated in depth in Sect. 5.

4.1.2 Empirical data

Since the number of links in a fully connected directed graph is enormous, collecting data for each link and each time zone is a monumental challenge. Instead, we assume that the same flow profile applies to all links. Indeed it seems reasonable to assume that most motorways, on average, follow the same pattern of having a morning and evening congestion period (see also Ichoua et al. (2003) for a similar reasoning).

We assume that the flow on a road segment is given as well as the free flow speed. Given an empirical dataset, the remaining queueing parameters can be tuned to represent the relevant environmental conditions as close as possible (Van Woensel and Vandaele 2006). As this is not the objective here, we choose the queueing parameters such that the resulting speed profiles matches the four scenarios. As shown in Figs. 1 and 2, this results in a reasonable speed profile both when there is a congested flow during the entire day (Fig. 1, *Congested flow*) as when there is only heavy congestion during the morning and evening peak hours (Fig. 2, *Rush-hour flow*).

Figures 1 and 2 represent the two different speed profiles used for the analysis. Figure 1 (*Congested flow*) is the result of a congested flow during the entire day, Fig. 2 (*Rush-hour flow*) is the result of heavy congestion only during the morning and evening rush hours. To simulate the effect of bad weather, the queueing parameters are adjusted in such a way that during bad weather, a large variability of speeds becomes apparent (Van Woensel and Vandaele 2006).

We model two types of roads for taking specific constraints in real cases into account, representing highways and rural roads (Ichoua et al. 2003). Due to the lack of data, we reduced the maximum allowed free flow speed, as such artificially creating multiple road types. All arcs between even nodes represent highways. All other links represent rural roads.

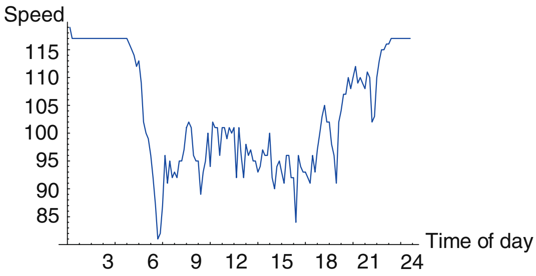

Fig. 1 Speed distribution with congested flows during the entire day

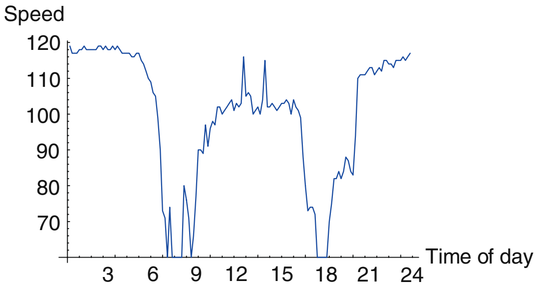

Fig. 2 Speed distribution with congested flows during the morning and evening rush hours

4.2 Tabu search implementation

In this paper Tabu search, first proposed by Glover (1986), is used to generate solutions as it has a number of advantages: general applicability of the approach, flexibility for taking into account specific constraints in real cases and ease of implementation (Pirlot 1996). For this Tabu Search implementation the following references where used as a basis: Gendreau et al. (1994, 1996a), Hertz et al. (2000) and Van Woensel et al. (2008).

The first change made to this basic algorithm consists of replacing distance by travel time. The main change however consists of extending the basic objective function with the standard deviation of the route travel time. The objective function thus becomes $E(\widetilde{TT}) + \beta\sqrt{\sum Var(\widetilde{TT})}$ (see also Eq. 1), with $E(\widetilde{TT})$ ($Var(\widetilde{TT})$) the expected travel time (variance of the travel time) and β a positive parameter to account for the risk aversion of the planner. This objective function deals with the distribution of the travel time of the solution and cannot be reduced to the classical objective function (i.e., $\sum_{ij\in A} x_{ij}c_{ij}$).

The neighborhood structure used is based on the λ-interchange. With λ limited to 2, up to 2 costumers are exchanged between two routes (Osman 1991, 1993). It must be noted that the links after the exchanged nodes in the modified routes must be re-evaluated completely, because of the stochastic time-dependent nature. Indeed, the travel time distribution of a given link from i to j depends on time of departure at node i (Donati et al. 2003). This also implies that the best solution of leaving the depot at 6h00 is not necessarily the best solution of leaving the depot earlier or later that day. Therefore, the neighboring solutions of the λ-interchange structure are also evaluated in a limited number z of adjacent starting times. In case of improvement, the starting time of the tours are updated. The rationale behind this is that a truck can decide to leave earlier or later to avoid periods of (anticipated) high congestion. Initially, the parameter "z" is set to 3 and the starting time at the depot is 6h00, as such all neighboring solutions are evaluated with starting times: 5h30, 5h40…6h20,6h30. The parameter "z" is further used for diversification and intensification. As this requires a full re-evaluation of the tours, we limit the number of non-improving solutions by using a soft constraint on the excess capacity during the search procedure. The practical objective function then becomes

$$Min \sum_{r=1}^{m} \sum_{(i,j)\in R_r} E(\widetilde{TT_{ij}^{t_0}}) + \beta \sqrt{\left(\sum_{r} \sum_{(i,j)\in R_r} Var(\widetilde{TT_{ij}^{t_0}})\right)} + \gamma \sum_{r} \left[\sum_{i\in R_r} q_i - Q\right]^{+} \tag{2}$$

where $[x]^{+} = max(0, x)$ and γ is a positive parameter. If the solution is feasible with respect to capacity the third part of Eq. 2 reduces to zero. If the solution is infeasible with respect to capacity a penalty proportional to the excess capacity (γ) is added. Decreasing the parameter γ will make it easier for infeasible solutions with respect to capacity to be picked as best solution for that iteration as the excess load will not have a large impact, thus enables diversification. Increasing γ punishes solutions with excess capacity in their tours, forcing the best solution for this iteration to become feasible again. The best solution evidently has no excess capacity in any of the tours.

5 Results

The impact of the road and environmental conditions are evaluated in this section. Substantial gain in terms of travel time reliability is found by giving more weight to this component during optimization. The decrease in variability can be offset by an increase in expected travel time. To measure the effect of this phenomenon, we use the 95th-percentile of the travel time distribution as a measure of the quality of a solution. Giving more weight to the standard deviation improves the 95th-percentile of the distribution. Finally, an Arena Rockwell Software Inc. (2005) based simulation is presented, which confirms the results.

5.1 Impact of the variance component

Minimizing the expected total travel time assumes that the planner is risk neutral in his planning behavior, i.e., the planner does not care about the risk involved. Ignoring the variance of travel time can be costly since this variance might be unacceptably large from a managerial or planning point of view. Indeed, one might prefer having a route that is on average slightly worse, but has a reduced variance, as such increasing the reliability of the predicted arrival times at all destinations. Depending on the risk attitude a different route will be chosen (He et al. 2005). By adjusting the parameter

β in the objective function, the planner can easily insert his risk attitude in the model. Higher values of β will result in routes that have more reliable travel times. Table 1 shows that the probability that the travel time is smaller than the travel time at TT_β (defined as $E(\widetilde{TT})+\beta\sqrt{\sum Var(\widetilde{TT})}$) increases as β increases. In addition, the tail of the distribution to the right of TT_β contributes to the total mass of the distribution. The higher β, the less mass there is left that contributes to the total mass of the distribution (Finkel 1990). For instance, for dataset 32k5 from Augerat (Augerat et al. 1998), the optimal route has a travel time distribution with σ (scale parameter of the lognormal travel time distribution) equal to 0.376 ($E(\widetilde{TT})$: 1308.87 minutes, $SD(\widetilde{TT})$: 509.62 min). When $\beta = 2.0$, 95.73% of the population of travel times is below TT_β. The remaining 4.27% however still contributes 8.93% of the total mass. Therefore, we will examine β-values up to 3.0, where the remaining mass is about 3% for this set.

The values in Table 2 indicate the relative decline of the standard deviation of the total travel time of the best solution found (with associated β) compared to the standard deviation of the travel time found by a minimization with $\beta = 0$ (i.e., not taking variability into account). The values are an average over 27 Augerat datasets. If $\beta = 0$, then the planner has a risk neutral behavior and treats the variance of the route as something residual, not worth optimizing. However, as the weight of the standard deviation of the travel time adopts higher values, the standard deviation of the associated best route continues to decrease, regardless of the environmental and road conditions, meaning that the planner can control the degree of variance of the travel time of the eventual solution in an easy way. The best improvement is obtained by increasing the value of β from 0 to 0.5, whereas the additional improvement of further steps reduces in magnitude. It is thus important to include the variability of the travel times in the objective function. Better improvements will be expected when the road and/or weather conditions are bad. If road conditions are bad, the speed will fluctuate

Table 1 Probability for $TT < TT_\beta$ and associated remaining mass in the tail of the travel time distribution for different β-values when σ (scale parameter of the lognormal travel time distribution) = 0.376

β	$\beta = 0.0$	$\beta = 0.5$	$\beta = 1.0$	$\beta = 1.5$	$\beta = 2.0$	$\beta = 2.5$	$\beta = 3.0$
$p(TT < TT_\beta)(\%)$	57.4	74.58	85.61	92.11	95.73	97.71	98.77
$mass(TT > TT_\beta)(\%)$	57.45	38.76	24.59	15.00	8.93	5.25	3.06

Table 2 Impact of β on the standard deviation of the travel time, compared to the standard deviation of the travel time with $\beta = 0$

Flow	Road conditions	$\beta = 0.5$ (%)	$\beta = 1.0$ (%)	$\beta = 1.5$ (%)	$\beta = 2.0$ (%)	$\beta = 2.5$ (%)	$\beta = 3.0$ (%)
Congested	Bad	−1.61	−2.18	−2.52	−2.75	−3.14	−3.47
Rush-hours	Bad	−1.45	−2.05	−2.50	−2.95	−3.32	−3.46
Congested	Good	−3.30	−4.18	−4.72	−5.17	−5.40	−5.56
Rush-hours	Good	−3.17	−3.83	−4.20	−4.55	−4.71	−5.04

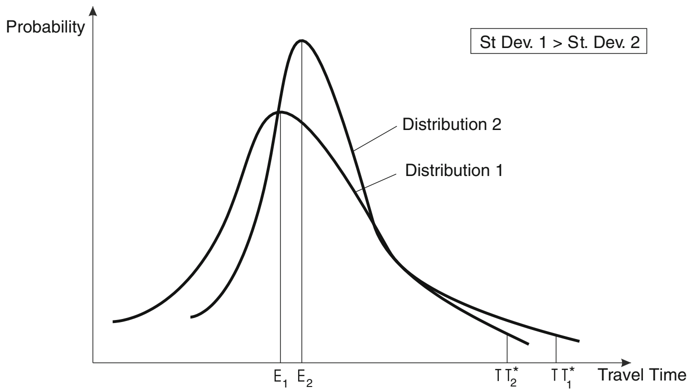

Fig. 3 Impact of 95th-percentile on solution quality (Lognormal distribution)

more, which makes it more difficult to predict when a tour is over, as opposed to better road conditions. If the flow is congested during the entire day, the improvement is also more substantial compared to a flow which is characterized by two rush-hours. This is due to the fact that between the two congestion periods, drivers are able to uphold free flow speed, leading to less variability.

5.2 The 95th-percentile as a quality measure

The reduction of the standard deviation comes at a certain cost, i.e., a likely increase of the average travel time. To check whether this cost is acceptable, we propose the use of the 95th-percentile as a quality measure assuming a lognormal distribution for the travel time. We use the 95th-percentile as a single measure of the quality of the solution. Figure 3 represents two lognormally distributed solutions. The Figure illustrates that if the 95th-percentile of the travel time of the solution with worse average travel time, but better standard deviation (Distribution 2) is lower than the one with best average travel time (Distribution 1), we have nevertheless managed to improve solution quality.

This can also be derived from our test cases. The impact of β on the improvement in the 95th-percentile can be observed in Table 3. The travel time associated with the 95th-percentile decreases when more weight (higher β) is given to the standard deviation in the objective function. The best improvement is observed in the first step, regardless of the test situation. The additional improvement of higher β values reduces in magnitude. This means that although the average travel time will become larger with increasing β, the total travel time will be better in 95% of all cases but with decreasing importance.

If the road conditions are good, the relative improvement of the travel time of the 95th-percentile is more substantial for the congested flow throughout the day compared to a flow with two rush-hours for equal β values. From Table 4, we see that if weather conditions are good, the squared coefficient of the travel times of the two flow types are of the same magnitude. Therefore, since the standard deviation of the travel

Table 3 Impact of β on the 95th-percentile of the travel time, compared to the 95th-percentile of the travel time with $\beta = 0$ (lognormal distribution)

Flow	Road conditions	$\beta = 0.5$ (%)	$\beta = 1.0$ (%)	$\beta = 1.5$ (%)	$\beta = 2.0$ (%)	$\beta = 2.5$ (%)	$\beta = 3.0$ (%)
Congested	Bad	−0.62	−0.82	−0.94	−0.98	−1.12	−1.16
Rush-hours	Bad	−0.52	−0.72	−0.84	−0.91	−1.06	−1.07
Congested	Good	−1.32	−1.62	−1.75	−1.94	−1.97	−1.99
Rush-hours	Good	−1.38	−1.66	−1.79	−1.90	−1.99	−2.07

Table 4 Squared coefficient of variation of the travel times for given test situation and β values

Flow	Road conditions	$\beta = 0.0$	$\beta = 0.5$	$\beta = 1.0$	$\beta = 1.5$	$\beta = 2.0$	$\beta = 2.5$	$\beta = 3.0$
Congested	Bad	0.038	0.037	0.037	0.036	0.036	0.036	0.036
Rush-hours	Bad	0.038	0.037	0.037	0.036	0.036	0.036	0.036
Congested	Good	0.110	0.103	0.101	0.100	0.099	0.098	0.098
Rush-hours	Good	0.113	0.106	0.104	0.103	0.102	0.102	0.101

Table 5 Improvement (in min) of the 95th-percentile of the tour travel time when comparing the optimal routes with $\beta = 3$ and $\beta = 0$

Flow	Road conditions	Average	Minimum	Maximum
Congested	Bad	35.11	6.90	93.57
Rush-hours	Bad	32.26	0.33	68.10
Congested	Good	88.32	15.58	166.32
Rush-hours	Good	81.61	15.85	259.98

times is higher for the congested flow, better improvements can be expected for this flow type with increasing β.

On the other hand, if road conditions are bad, the best relative improvement is observed for the two rush-hour flow. In bad weather, the squared coefficient of variation of the travel times for the flow with two rush-hours is larger than the congested flow (Table 4). This means that for the flow with two rush-hours the standard deviation is relatively large compared to the mean. Adding some weight to it will thus result in better relative results.

Table 5 represents the gain in travel time (in min) of the 95th-percentile for the test cases when comparing $\beta = 3$ with $\beta = 0$. For instance, the gain over all Augerat sets for the congested flow in bad road conditions is on average 88.32 min. The minimum improvement for that test situation is 15.58 min and the maximum improvement is almost 3 h (166.32 min). It is clear that the reduction of the standard deviation of the travel time is substantial enough to overcome the increase in average travel time. Extending the objective function to account for the travel time variability provides results with better overall reliability, especially when road conditions are bad.

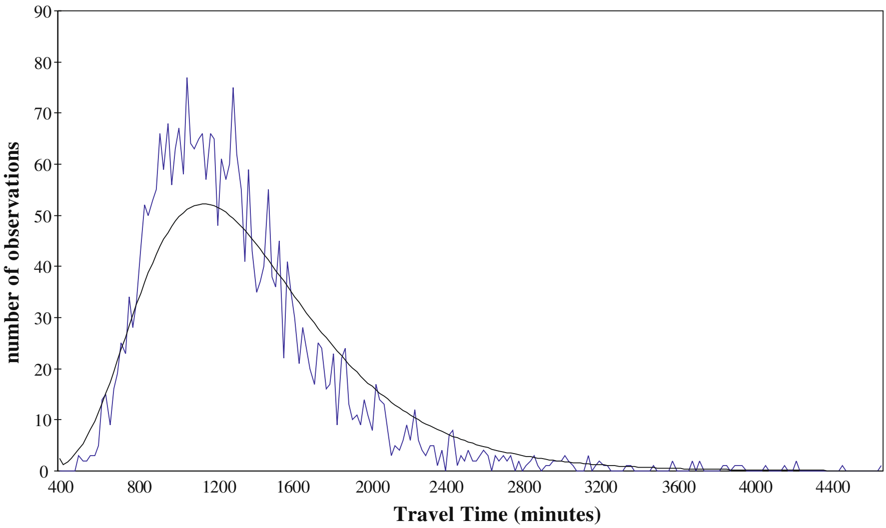

Fig. 4 Travel time distributions of set $32k5$ with a congested flow in good road conditions. Results following a simulation with Arena are plotted together with the expected lognormal distribution (derived from Tabu Search solution)

5.3 Simulation

To validate the approximations, used when building the variance estimating model presented above, we constructed a simulation in Arena in which we reconstructed the best solution of a given dataset as a sequence of lognormal distributions (representing the link travel times) with the individual link-based mean and standard deviation. A number of trucks then have to complete the routes. On each link, a random travel time is generated according to the distribution for that link. For set $32k5$, 3001 trucks completed the tours and their travel times have been plotted in Fig. 4. The results indicate that the resulting total travel time distribution is indeed lognormally distributed. In addition, the plotted results are close to the total travel time used in the Tabu Search. For instance, the travel time associated with the 95th-percentile is 2205.5 min (Fig. 4, Table 6), which corresponds with what we expect from the travel time distribution of the best solution (2262.71 min (95th-percentile of lognormal travel time distribution with $E(\widetilde{TT})$: 1308.87 min and $SD(\widetilde{TT})$: 509.62 min)).

The positive impact in terms of travel time reliability when optimizing the *VRP* for a more heavily weighted standard deviation is validated by the simulation results provided in Table 6. The best solutions of a Tabu Search optimization with β values equal to 0.0 and 3.0 are reconstructed. For each set the average travel time increases and the standard deviation of the travel time decreases, as such increasing the travel time reliability. The decrease in the standard deviation is substantial enough to improve the overall quality of the solution (better travel time associated with the 95th-percentile). Values of $\beta = 0.0$ and 3.0 are two extreme situations. The planner can use any value in between depending on his own risk attitude. From a planning point of view, it is

Table 6 Comparing the best routes (through Tabu Search optimization with respective β-values) for three sets in congested flow during the entire day using Arena

		$E(\widetilde{TT})$	$SD(\widetilde{TT})$	95th-percentile TT
32k5	$\beta = 0.0$	1320.87	550.77	2205.5
	$\beta = 3.0$	1332.51	514.51	2169.5
38k5	$\beta = 0.0$	1286.60	480.72	2087.7
	$\beta = 3.0$	1296.14	471.02	2049.4
80k10	$\beta = 0.0$	2701.11	746.35	4007.0
	$\beta = 3.0$	2720.64	685.73	3972.9

Average travel time, standard deviation of the travel time and 95th-percentile are provided after 3001 trucks completed the best routes

better to have more predictability in the routing than a potentially faster route. The uncertainty about the actual arrival time will be avoided as the planner becomes more risk averse.

6 Conclusions and future research

We have argued that taking time dependent travel speeds into account can be of much interest in routing problems. Indeed, minimizing the expected travel time does not deal with the true stochastic nature of the travel times. Since the real speed is a realization of a stochastic process, it is important to account for the variability of the speed and thus the travel time uncertainty when planning a route. This paper has tackled the problem and has shown how more reliable routes can be obtained. These routes have the potential to reduce real operating costs for a broad range of industries which face daily routing problems.

Including variance as well as expected value of travel time in the objective function has many potential applications. We have shown that our proposal gives a manager a simple yet powerful tool to incorporate congestion uncertainty in the planning of routes and allows for various attitudes towards risk. The resulting routes were shown to be more reliable and predictable. Although the gain in terms of less travel time variability will be offset by a higher average travel time, the travel time associated with the 95th-percentile will improve. Depending on the road and environmental conditions, this improvement will be more or less substantial. These conclusions are confirmed by independent simulation studies.

It must be noted however, that in some cases the reduction in variability will not be substantial enough to compensate for the reduction of the expected travel time. If for instance the initial route has already a small travel time distribution (associated with a high speed), then it will be hard to find a new route/starting time with a better travel time distribution.

As there is not much information available on how to model the variance of the travel times in literature, most analyses are in terms of expected travel time, we have heuristically determined variances in the analysis presented in this paper. Hence, we are currently deriving general conditions on speed profiles that will guarantee the validity of the conclusions derived here.

References

Augerat P, Belenguer JM, Benavent E, Corber A, Naddef D (1998) Separating capacity constraints in the CVRP using tabu search. Eur J Oper Res 106:546–557

Beaulieu NC, Xie Q (2004) An optimal lognormal approximation to lognormal sum distributions. IEEE Trans Vehicular Technol 53(2):479–489

Berry DS, Belmont DM (1951) Distribution of vehicle speeds and travel times. In: Proceedings of 2nd Berkeley symposium on mathematical and statistical probabability, pp 589–602

Best MJ, Grauer RR (1991) Sensitivity analysis for mean-variance portfolio problems. Manage Sci 37(8):980–989

Chen C, van Zwet E, Varaiya P, Skabardonis A (2003) Travel time reliability as a measure of service. Technical report, Transporation Research Board

DfT (2000) Transport 2010: the 10 year plan. Technical report, DETR, July

DfT (2003) Delivering better transport—progress report. Technical report

Donati AV, Montemanni R, Casagrande N, Rizzoli AE, Gambardella LM (2003) Time dependent vehicle routing problem with a multi ant colony system. Technical Report IDSIA -02-03, International IDSIA, 2003

Finkel AM (1990) A simple formula for calculating the mass density of a lognormally-distributed characteristic: Applications to risk analysis. Risk Anal 10(2):291–301

Fu L, Rilett LR (1998) Expected shortest path in dynamic and stochastic traffic networks. Transp Res Record 32(7):499–516

Gao S, Chabini I (2002) The best routing policy problem in stochastic time-dependent networks. Transp Res Record 1783:188–196

Gao S, Chabini I (2006) Optimal routing policy problem in stochastic time-dependent networks. Transp Res B 40:93–122

Gendreau M, Hertz A, Laporte G (1994) A tabu search heuristic for the vehicle routing problem. Manage Sci 40(10):1276–1290

Gendreau M, Laporte G, Séguin R (1996) A tabu search heuristic for the vehicle routing problem with stochastic demands and customers. Oper Res 44(3): 469–477

Gendreau M, Laporte G, Séguin R (1996b) Stochastic vehicle routing. Eur J Oper Res 88(1):3–12

Glover F (1986) Future paths for integer programming and links to artificial intelligence. Comp Oper Res 13(5):533–549

Grauer RR, Hakansson NH (1993) On the use of mean–variance and quadratic approximations in implementing dynamic investment strategies: a comparison of returns and investment policies. Manage Sci 39(7):856–871

Haghani A, Jung S (2005) A dynamic vehicle routing problem with time-dependent travel times. Comp Oper Res 32:2959–2986

Hall RW (1986) The fastest path through a network with random time-dependent travel times. Transp Sci 20(3):182–188

He RR, Kornhauser AL, Ran B (2005) Essentially best routes in dynamic and stochastic transportation network. Int J Vehicle Inf Commun Syst 1(1–2):1–14

Heidemann D (1996) A queueing theory approach to speed-flow-density relationships. In: Proceedings of the 13th international symposium on transportation and traffic theory, Lyon, France, 1996. Transporation and traffic theory

Hertz A, Laporte G, Mittaz M (2000) A tabu search heuristic for the capacitated arc routing problem. Oper Res 48(1):129–135

Ichoua S, Gendreau M, Potvin J-Y (2003) Vehicle dispatching with time-dependent travel times. Eur J Oper Res 144:379–396

Kharoufeh JP, Gautam N (2004) Deriving link travel-time distributions via stochastic speed processes. Transp Sci 38(1):97–106

Kwon J, Coifman B, Bickel P (2000) Day-to-day travel time trends and travel-time prediction from loop-detector data. Transp Res Record 1717:120–129

Laporte G (1992) The vehicle routing problem: An overview of exact and approximate algorithms. Eur J Oper Res 59(3):345–358

Laporte G, Louveaux F, Mercure H (1992) The vehicle routing problem with stochastic travel times. Transp Sci 26(3):161–170

Malandraki C, Daskin MS (1992) Time dependent vehicle routing problems: Formulations, properties and heuristic algorithms. Transp Sci 26(3):185–200

Mulvey JM, Vanderbei RJ, Zenios SA (1995) Robust optimization of large-scale systems. Oper Res 43(2):264–281

Osman IH (1991) Metastrategy Simulated Annealing and Tabu Search Algorithms for Combinatorial Optimization Problems. PhD Thesis, Imperial College London, The Management School

Osman IH (1993) Vehicle routing and scheduling: Applications, algorithms and developments. In: Proceeding of the international conference on industrial logistics, Rennes

Pirlot M (1996) General local search methods. Eur J Oper Res 92:493–511

Rockwell Software Inc. (2005) Arena user's guide. Rockwell Software Inc., USA

Taniguchi E, Thompson RG, Yamada T, Van Duin R (2001) City logistics: network modelling and intelligent transport systems. Pergamon, New York

Van Woensel T, Vandaele N (2006) Empirical validation of a queueing approach to uninterrupted traffic flows. A Quart J Oper Res 4(1):59–72

Van Woensel T, Creten R, Vandaele N (2001) Managing the environmental externalities of traffic logistics: the issue of emissions. POMS J Spec Issue Environ Manage Oper 10(2)

Van Woensel T, Kerbache L, Peremans H, Vandaele N (2008) Vehicle routing with dynamic travel times: A queueing approach. EJOR 186(3): 990–1007

Vandaele N, Van Woensel T, Verbruggen A (2000) A queueing based traffic flow model. Transp Res D 5(2):121–135