Sixth Sense Logistics

Sixth Sense Logistics – Challenges in supporting more flexible, 'human-centric' scheduling in the service sector

Fraser McLeod[1], Tom Cherrett[1], Duncan Shingleton[2], Tolga Bektas[1], Chris Speed[2], Nigel Davies[3], Janet Dickinson[5], Sarah Norgate[4]

[1]University of Southampton, University of Edinburgh[2], Lancaster University[3], Salford University[4], Bournemouth University[5]

Introduction

The need to conform to tight delivery and collection schedules can lead to increased pressures for logistics providers and inefficiencies in the services provided from an asset utilisation perspective. Transport behaviours, habits and practices in contemporary Westernized clock-time cultures are often situated around the notion of time being viewed as a scarce resource, with a high value placed on carving up this commodity into activities running punctually back-to-back (Norgate, 2006; Southerton et al., 2001).

Using an example from the service sector, this paper explores how service engineer jobs can be dynamically scheduled during the round in response to new job requests received to minimise client waiting time and transport costs (a multi-period dynamic vehicle scheduling problem (Angelelli et al., 2009)). The potential benefits of such an approach are presented along with the challenges of delivering it, with respect to the technology needed to help visualise the engineer's current and future trajectories in relation to the incoming client calls. The research is being undertaken as part of the RCUK 6th Sense Transport project which is investigating how mobile technologies coupled to social networking principals can be leveraged to provide individuals with different ways to relate to time (present and future) and new understanding of the relationships between their own future transport plans and those of others around them (Davies et al., 2012).

Background: Field service logistics and the scope for '6th Sense' transport

Service engineer or 'field logistics' companies typically receive daily calls from clients (both private and public sector businesses) requesting specific parts delivery/collection or related item servicing for a device (e.g. photocopier, cash register) being used under a contract service agreement. Site visits are categorised but usually involve either: i) *'Scripted work'* where the engineer has to do some form of software configuration on the customer's premises (system upgrades); ii*) 'Delivery drops'* where parts are delivered for a device e.g. toner cartridges; iii) *'Collection'* where a component is collected for subsequent re-conditioning; iv) *'Exchange'* where a component is swapped for a new one; v) *'Install'* which involves the physical installation of a part; vi) *'Survey'* which involves investigating a problem with a device at a customer premises); vii) *'Maintenance'* which involves a routine service of a device.

Engineers can often be self-employed and operate on either a flat daily fee basis, or on a piecemeal rate per job completed. As all of the seven service activity types with the exception of *'Install'* have fairly rigid book times, good prediction of dwell time once the engineer is on the client's premises is possible. Engineers will work for multiple customers and may also have different skill sets depending on the client portfolio. In terms of operating strategy, client calls are received via a central help desk at the service company and typically scheduled for the subsequent day's activity. In the case of a service company with national coverage, engineers will be allocated to a series of depots from where their parts are sourced and their daily rounds emanate. The client calls will be batched into regions and then assigned to the relevant depot manager who will issue them to their respective engineers, based on a manual scheduling approach, or using an optimisation package to obtain a set of workable rounds. Engineers can visit between 6 and 15 clients in a day (depending on the types of service required) and will take out the parts they need for those jobs at the start of the round. Service companies will often offer 'emergency call outs' where a logged client call must be resolved within a set time interval (often 4 or 8 hours from receipt). This entails engineers carrying certain stock parts which can be utilised in such circumstances and the call will have to be added into the round by the depot manager.

Service data were obtained from a major UK field logistics company for a 3-month period (7 Sep – 7 Dec 2011), relating to activity emanating from one depot servicing clients in the midlands (Table 1). The average and standard deviation of the recorded amount of time on site were calculated for each service type along with a frequency plot of time on site (**Error! Reference source not found.**). 'Exchanges' were the most common activity (83%) followed by collections (10%) with the mean time on site for the former being 10 minutes.

Table 1 – Time on site by type of service

Type of service	#	%	Average time (minutes)	Standard deviation (minutes)
Additional install same site	3	0.1%	123	203
Collection	354	10%	8	9
Collection with box	27	0.8%	6	4
Delivery	65	1.8%	7	10
Exchange	2991	83%	10	15
Install	147	4.1%	23	29
Install+survey	1	0.0%	50	n/a
Install+train	2	0.1%	34	11
Total	**3590**	**100%**	**11**	**16**

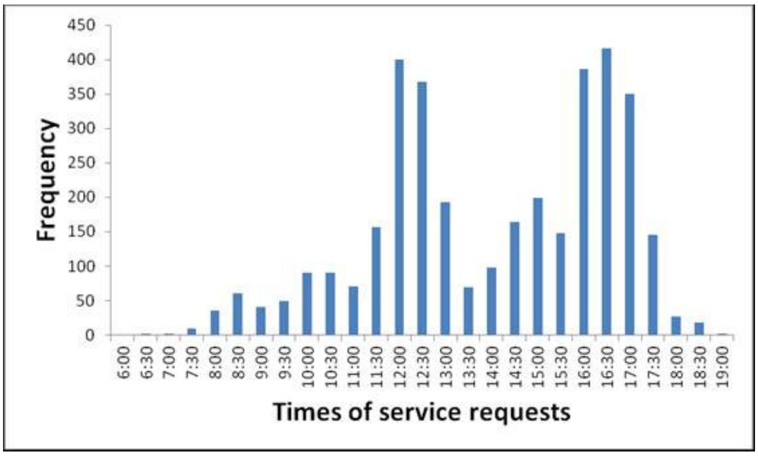

Figure 1 – Frequency plot of times of service requests (each bar is for time period [t-30,t])

Of interest are the ways in which mobile technologies are starting to be used in the sector, and how they could be used to provide greater visibility between engineer, depot and client to allow a more flexible attitude towards time and travel. Field logistics companies are beginning to experiment with tablet personal computers, primarily to allow the scripts (the product service manuals) to be carried in electronic format by the engineer. The functionality on the tablet allows not only the transfer of data between the engineer and the depot to log completed calls, but also the potential to track the vehicle and inventory on it with respect to the work schedule. Having the base schedule of the engineers planned activity for the day, their current locations at any time, real-time updates of client calls and a knowledge of likely service activity creates the potential for a '6th Sense' approach to scheduling activity to answer questions like:

1. Client: "When will an engineer be passing with a part I urgently need?"
2. Engineer: "I don't have the right part to fix this. I wonder which other engineers might be around in the next hour with the right part I could use, or whom I could swap the job with?"
3. Engineer: "I will be in the Birmingham area in 3 hours, which clients might request a service if they knew I was going to be there"

As new client calls come in continually during the working day, of particular interest is whether some could be scheduled into the current round (devised using the previous day's calls) where convenient for the engineer's vehicle (given its position relative to the time the call was received). This could allow engineers to react more dynamically to client calls, maximise their reward if working on a piecemeal rate and better utilise their working time in a more 'human-centric' approach to field logistics. This has been investigated in this paper using a database of daily client calls and linked engineer service visits

(2011) for a major technical field services company serving businesses across the midlands. Routing and scheduling software has been used to retrospectively quantify the benefits of reacting to new calls as they were received on the recorded rounds. Reactions to such an approach, using locative media to provide visibility between engineers and clients, has been gauged through a focus group of service engineers.

Literature Review

From an optimisation perspective, the problem being considered here is a multi-period dynamic vehicle scheduling problem, with new customer service requests being received over time. It has been assumed that if the client request is made before a specified cut-off time (12pm) then same-day servicing will be undertaken, otherwise next-day servicing will be undertaken (Angelelli et al., 2009). This approach is more complicated, as the scheduler has to decide on which day to include a new job, which requires a two-day '*look ahead period'* (today and tomorrow) to make this decision most effectively (Angelelli et al, 2009; Wen et al, 2010). The optimisations undertaken here only considered jobs for the current day.

Dynamic routing is still considered to be an under-researched area (Angelelli et al., 2007) particularly when compared with research into static problems (Wen et al., 2010). It is a multi-objective problem with the typical goals being to: (i) provide customers with fast response times (here, including same-day response); (ii) minimise transportation costs; and (iii) balance workloads between available drivers. Wen et al (2010) considered the trade-off between these objectives using a weighted linear combination; they also used a quadratic penalty function, applied to the customer waiting time so that, one customer having to wait three days would incur a higher penalty than three customers waiting one day. An alternative approach is to minimise costs subject to meeting the minimum service level agreements (Angelelli et al, 2009).

Variants of the problem relate to whether vehicle capacity is a constraint or not (not an issue here) and whether allocated tasks may be redistributed among the drivers at successive re-optimisations of the schedules. Angelelli et al. (2009) considered a combined problem where delivery tasks could not be exchanged between drivers (as the items to be delivered were specific to the customers) but other tasks could be. In our exploratory study, it is assumed that all tasks can be exchanged, including deliveries, as drivers carry some common parts, however, in reality, there may be some jobs which could not be readily swapped once allocated. Customer demands are most often specified with some form of time window constraint, be it relating to time of day or day of the week, depending on the service level agreement. Sometimes these constraints are 'soft', in the sense that they can be broken but subject to a penalty cost, which may relate to sub-contracting the work or the customer paying a reduced price due to the lower level of service. Drivers are usually subject to working and driving time restrictions, in accordance with national and EU regulations.

A measure of the 'degree of dynamism' (DoD) of any particular problem can be made as:

$$N_d / (N_d + N_s)$$

where N_d is the number of same-day service jobs (dynamic) and N_s is the number of next-day service jobs (static) that are undertaken (Khouadjia et al, 2012). Thus the DoD can range from 0, where all jobs for the day are known in advance ($N_d = 0$), to 1, where no jobs are known at the start of the day ($N_s = 0$). If the cut-off time is half-way through the working day and if job requests are uniformly distributed throughout the working day, then typically the DoD will be 0.5, with equal numbers of static and dynamic tasks ($N_d = N_s$). The more dynamic the problem, the more unknowns there are, which will naturally tend to result in increased costs. Results by Khouadjia et al (2012) suggested that their dynamic problem (with a DoD of 0.5) resulted in a total driving distance of around 20% higher than the equivalent static problem. The effect of lack of information on performance can be measured by the *competitive ratio* which is defined as the ratio of the achieved performance against the performance achieved by an optimal off-line algorithm (Angelelli et al, 2007).

The most common solution approach to the dynamic scheduling problem is to repeatedly solve static scheduling problems based on the new service requests received. If new job requests are frequent then it makes sense to wait until a certain time period has elapsed before rescheduling, rather than try to react to every new demand. Khouadjia et al (2012) adopted this approach and divided the working day into 25 intervals with rescheduling taking place approximately every 30 minutes, which, they reported, provided the optimal trade-off between solution quality and computational time. This approach is adopted here but with considerably less frequent updating of schedules (3 schedules per day).

When one or more new jobs are agreed for same-day servicing, they can either be added to the tasks already allocated to drivers, without reallocation of existing jobs between drivers, or a complete reorganisation of routes/schedules can be undertaken. The former approach is easier and quicker to achieve than the latter but will tend to produce worse solutions; the latter approach may become impractical, due to computation time, if the size of the problem is large (i.e. many customers and many drivers) and/or there are frequent new job requests (Ichoua et al, 2000). Another issue with a computation time of any significant length is that by the time the new schedules have been devised, the drivers may be in new locations. In this case it would be desirable to be able to estimate driver locations *X* minutes ahead, to be used in the scheduling, where *X* is the time needed to devise and communicate the new schedules to drivers. This is where mobile technology could help estimate likely trajectories' of drivers into the immediate future.

The task of revising schedules during the working day requires knowledge of the jobs already serviced (so that they can be removed from the problem specification) as well as current locations of vehicles or predicted locations *X* minutes ahead (to be used as the starting points for the revised schedules). Ichoua et al (2000) stated that, in most applications, each driver's next destination is usually fixed and used as the starting point for their onward schedule. They considered a potential improvement whereby an immediate diversion to a new customer is permitted. This option would require detailed knowledge of current vehicle positions, schedules and routes to enable prediction of near future vehicle positions. Such knowledge may be facilitated through the use of mobile computing technology.

A further difficulty with scheduling servicing work lies in estimating the amount of time required to undertake the work. While some tasks are simple and readily predictable (e.g. a delivery) other tasks, such as installation work, may be inherently unpredictable. In some cases, one or more jobs may have to be carried over to the next day, with possible associated penalty costs, if earlier jobs took longer than expected.

Description of the Problem and Research Challenge

The problem under consideration here is the dynamic scheduling of service engineer tasks to allow same-day response for some of the tasks, with vehicle schedules being revised on the day in response to new job requests that are received. The two main objectives considered here, which may conflict with one another, are: i) to minimise the amount of time that clients have to wait before being serviced ii) to minimise transport costs.

Although the literature review suggested updating schedules every 30 minutes, this was considered to be impractical for this application and only three optimisations each day were used: one calculated for the start of the day (some drivers could set off from as early as 4.30am if early access to premises was possible), one mid-morning (10.30am) and one at mid-day (12pm). It was assumed that all service requests made before these times would be included in the newly devised schedules. Any service request made after 12pm would not be serviced until the following day. An analysis of times of day when service requests were registered (Figure 1) indicated that 10.6% of all requests were made before 10.30am and 28% before 12pm, indicating a degree of dynamism of 0.28. It should be borne in mind that the shape of this profile would likely change once customers became aware of a same-day servicing option if registering before 12pm.

Time window data for customers were not available to this study, however, observation of actual arrival times at sites indicated the vast majority were visited between 6am and 5.30pm so this time window was assumed for all sites. It was also assumed that all tasks could be reallocated to a different driver when rescheduling, although, in reality, some jobs may not be swapped in this way, which will be considered in future research. The problem comprised a fleet of five identical vans, without any capacity considerations, starting and ending at a single depot (in Coventry, UK). The time required for each job was taken to be the actual time taken which was recorded in data collected between September and December 2011. A commercially available software package was used to devise routes and schedules. The procedure that was used involved:

(i) Creating the initial schedule for the day based on the orders placed after 12pm on the previous day
(ii) Creating the mid-morning schedule by including orders placed before 10.30am, specifying an opening time window of 10.45am for these jobs (assuming it will take at least 15 minutes to calculate the new routes and convey them to drivers). For convenience, this schedule (and the mid-day schedule) included the jobs already done. This was done by fixing their time windows to the servicing times found in the initial schedule. Similarly, jobs due to start within the next 15 minutes were fixed in this way on the basis that the driver was almost at the next destination and should not be diverted.
(iii) Creating the mid-day schedule (12pm) by including orders placed before 12pm. In this case, all jobs already done and the jobs due to start within the next 15 minutes had their time windows fixed to the servicing times found in the mid-morning schedule. Since this is the final schedule for the day it represents the routes undertaken from which the reported results were taken.

A consequence of the adopted approach is that while it is possible for one of the new jobs to be scheduled as a driver's next job (as long as the previously scheduled next job is not due to start within the next 15 minutes), the model does not correctly represent the possibility that the driver could have been en route to the previously scheduled next job and would have had to divert to the new job, resulting in the additional mileage, a+c-b, illustrated in Figure 2. Addressing this issue in a modelling environment is a challenging task, as estimating the driver's current location is difficult bearing in mind the various routes that may be possible with varying travel speeds. In practice, current locations of drivers would be available through mobile technology and could be used as the starting points for the new schedules.

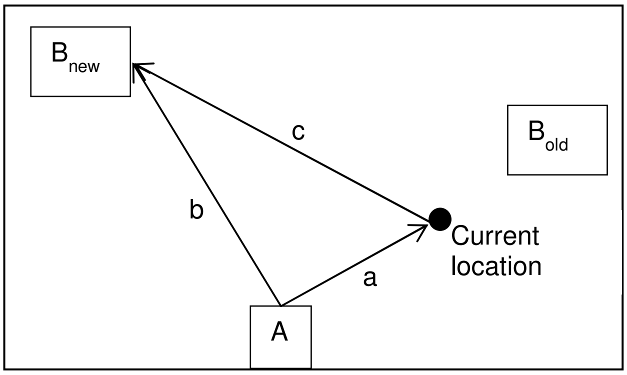

Figure 2 – Diversion en route

<u>Results</u>

Results are presented for one day's operation (8 September 2011) to give a flavour of the potential performance of dynamic scheduling. The actual rounds that were undertaken on that day, all of which were registered on the previous day, are shown in Figure 1 **alongside other service requests that were registered on the 8th September, some of which could, potentially, have been suitable for same-day servicing, depending on the times they were made. A total of 44 service visits were made by four drivers on 8/9/2011; however, if dynamic scheduling was in operation, the numbers and locations of sites visited would be different since jobs registered before 12pm the day before (9 of them) would have been undertaken the previous day, and jobs registered before 12pm on 8 September (24 of them) are included here, giving rise to a total of 59 jobs (44+24-9) in the dynamic scenario. This means that a direct before-and-after comparison of round times and distances is difficult to make and is certainly not meaningful for the one day being considered here, however, it was possible to make a direct comparison of the response times for the jobs that were in common between the two, namely, the jobs registered after 12pm on 7/9/2011 and before 12pm on 8/9/2011 (**
Table). Modelling runs over an extended period of several weeks are needed to provide more robust results and are planned as follow-up work.

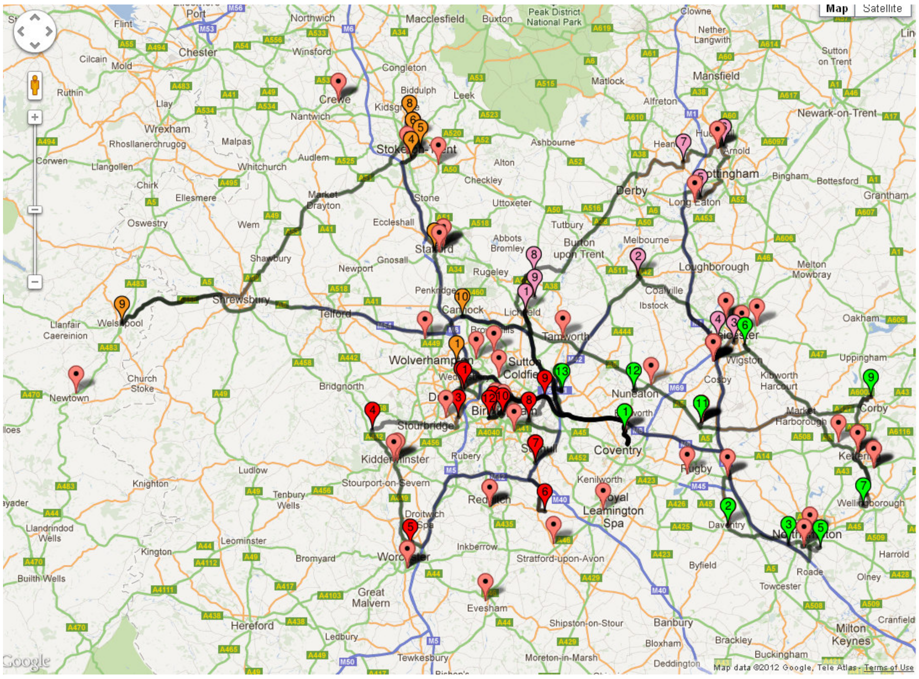

Key: Different colours refer to different drivers

Figure 1 – Driver rounds (numbered) and jobs logged on 8/9/2011 (unnumbered)

The results of the three schedules for 8/9/2011 (start of day, mid-morning and mid-day) are shown in Table 2, alongside the optimal schedule that could have been achieved with perfect predictive knowledge of the jobs to be undertaken. It can be seen that, at the start of the day, there were 35 jobs that had been registered after 12pm on the previous day and that these jobs could be undertaken by 3 vehicles (assuming a 10-hour maximum shift). By 10.30am, four new jobs had come in and these could be readily accommodated by the same three vehicles. However, by, 12pm, there were a 'rush' of 20 more jobs, which resulted in two additional vehicles being dispatched (at 12.15pm). However, only 58 of the 59 jobs could be scheduled with strict adherence to the various operating constraints that were specified. In practice, the service company may not wish to offer same-day servicing in some situations if, for example, the transport manager has the visibility of driver rounds and logistical expertise to see that some jobs would be better left to the next day. Alternatively, the experienced transport manager may have been able to anticipate this type of problem and may have decided to utilise a fourth/fifth vehicle earlier in the day to try to prevent its occurrence; however, this strategy would be inefficient if the number of incoming service requests was relatively low. Comparing with the optimal off-line solution, it can be seen that the lack of knowledge of the incoming calls resulted in a *competitive ratio* of 1.14, in terms of distance, and of 1.04, in terms of time taken, although these are underestimates, as dynamic solution failed to schedule one task.

Table 2 – Results of the 3 schedules for 8/9/2011

	# Jobs	Km	Time taken	# Vehicles
Start of day	35	929.9	22:04:39	3
Mid-morning	39	954.1	25:10:49	3
Mid-day	58	1382.2	35:25:00	5
Optimal off-line solution	59	1212.9	33:58:37	5

In terms of responding to client calls, the results suggested that for the test day, the mean 'response time' (the time between the service call being lodged by the client to an engineer being on site) was 11 hours 39 minutes, a 42% reduction on the actual case for the clients who logged calls during the period (Table 3).

Table 3 – Comparison of before and after response times

	Response time (hh:mm)	
	Before	**After**
Average	19:55	11:39
Standard deviation	03:12	07:45
Minimum	13:45	0:35
Maximum	28:32	23:21

<u>The role for mobile technology to enable dynamic visibility</u>

Identifying the engineer's current location can be achieved using the GPS capability in mobile technology. Estimating where the engineer might be in a future time frame is more challenging and will depend on likely routes, journey speeds and congestion. There are many existing approaches to the problem of finding an optimal path based on distance. Other methods also exist for handling moving obstacles (Reif and Sharir 1994), which can be considered as having time as another dimension in the planning space. The challenge in this case is in providing sufficiently accurate predictions of engineer mobility patterns that the logistics scheduler, on behalf of the individual engineers, or even the individual engineers themselves, can adjust their travel plans to take advantage of opportunistic encounters (e.g. client calls that come in during the round relevant to their current or immediate future location). In so doing, we are introducing increased uncertainty into their travel schedules which may cause anxiety (confirmed through an engineer focus group), not only for the engineers themselves, but also for clients who now expect a more timely response to a call request.

Smart phones and tablets provide a rich visual platform to offer individual engineers and the logistics scheduler a number of insights to the dynamic data available. One key challenge is that the predicted future traffic patterns which dictate potential engineer route choice are based on past behaviour and are therefore uncertain. While one recent previous work on this topic (Sanyal et al., 2009) uses size, or colour, or traditional error bars to visualise the uncertainty of values, there seems to be no consensus of what method is best, as different methods are needed depending on the data being visualised. To handle such high-dimensional data, one would need to develop simple low-dimensional interfaces for the user to perceive the right amount of information in an efficient manner. How to show such four-dimensional data with lots of details on a screen with a limited size is a challenge. We envisage an app where engineers' current locations are visible and using a slide bar, allow their projected locations to be viewed into the future so that possible match-ups with new client calls can be gauged (Figure 4). This could be an interface into the dynamic vehicle scheduling algorithm and allow proposed new schedules to be viewed by the logistics scheduler and engineers at each optimisation point.

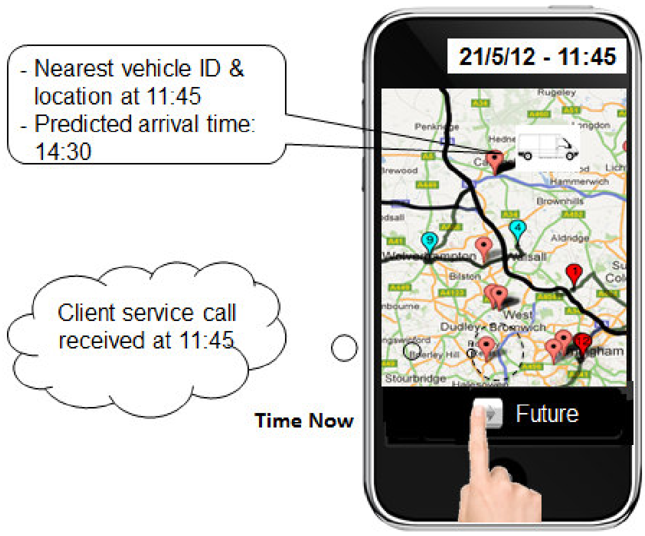

Figure 4. Prototype 6th Sense app visual showing current and future engineer locations

There is also scope here for the client to have a view of engineer location in a sense that notifications could be raised when an engineer is within a certain proximity of a client's base. The focus group highlighted that such a concept would be of benefit to engineers and that issues of privacy and 'being tracked' were not a concern in this specific case.

Conclusions

Dynamic scheduling, aided by mobile technology to help visualise new transport options appears to have considerable potential, although it is recognised that the issues are complex and drivers/ managers are often tightly constrained in what they are able to do. The case study presented here has indicated the levels of improvement in response time, in offering same day servicing to some customers, at the expense of increased vehicle mileage and time taken. Further work is planned in better representing the possibility of diversions mid-route and in obtaining more robust results over a longer time period. Further research is also needed in determining optimal vehicle dispatching strategies in a dynamic environment.

References

- Angelelli, E., Bianchessi, N., Mansini, R. & Speranza, M. G. (2009), 'Short term strategies for a dynamic multi-period routing problem', *Transportation Research Part C: Emerging Technologies,* 17(2)**,** 106-119.
- Angelelli, E., Speranza, M. G. & Savelsbergh, M. W. P. (2007), 'Competitive analysis for dynamic multiperiod uncapacitated routing problems', *Networks,* 49(4)**,** 308-317.
- Davies, N., Lau, M., Speed, C., Cherrett, T., Dickinson, J. and Norgate, D. (2012) Sixth Sense Transport : Challenges in Supporting Flexible Time Travel, at HotMobile'12 February 28–29, 2012, San Diego, CA, USA. ACM 978-1-4503-1207-3
- Ichoua, S., Gendreau, M. & Potvin, J. (2000), 'Diversion Issues in Real-Time Vehicle Dispatching', *Transportation Science,* 34(4)**,** 426-438.
- Khouadjia, M. R., Sarasola, B., Alba, E., Jourdan, L. & Talbi, E. (2012), 'A comparative study between dynamic adapted PSO and VNS for the vehicle routing problem with dynamic requests', *Applied Soft Computing,* 12, 1426-1439.
- Norgate, S. (2006), 'Beyond 9 to 5: Your Life in Time', Weidenfeld and Nicolson.
- Reif, J., and Sharir, M (1994) Motion planning in the presence of moving obstacles. Journal of the ACM 41 (July 1994), 764-790.
- Sanyal, J., Zhang, S., Bhattacharya, G., Amburn, P., and Moorhead, R (2009) A user study to compare four uncertainty visualization methods for 1d and 2d datasets. EEE Transactions on Visualization and Computer Graphics 15 (November-December 2009), 1209-1218.
- Southerton, D., Shove, E. & Warde, A. (2001), 'Harried and hurried: time shortage and the co-ordination of everyday life', CRIC Discussion Paper No. 47, The University of Manchester and UMIST.
- Wen, M., Cordeau, J. F., Laporte, G. & Larsen, J. (2010), 'The dynamic multi-period vehicle routing problem', *Computers and Operations Research,* 37, 1615-1623.