

# Kent Academic Repository

## Full text document (pdf)

### Citation for published version

Salhi, Said (2008) An Application of Real Time Vehicle Scheduling: A Case Study. Working paper. University of Kent Canterbury, Canterbury

### DOI

### Link to record in KAR

<https://kar.kent.ac.uk/25480/>

### Document Version

UNSPECIFIED

#### Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

#### Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

#### Enquiries

For any further enquiries regarding the licence status of this document, please contact:

[researchsupport@kent.ac.uk](mailto:researchsupport@kent.ac.uk)

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

## ***Working Paper Series***

---

### **An Application of Real Time Vehicle Scheduling: A Case Study**

**Said Salhi**  
**Kent Business School**

**R H Currie**  
**Her Majesty's Revenue and Customs,**  
**Telford**

# An Application of Real-Time Vehicle Scheduling: A Case Study

<sup>1</sup>S Salhi and <sup>2</sup>RH Currie

<sup>1</sup> Centre for Heuristic Optimisation,

Kent Business School, University of Kent, Canterbury, UK.

<sup>2</sup> Her Majesty's Revenue and Customers, OR Section, Telford, UK.

## Abstract

An application of real-time changes in scheduling deliveries of road-making materials is conducted based on an implementation of a tabu search heuristic. This distribution problem deals with heterogeneous products and vehicles where the assignment of pickup points to requests needs also to be made. The problem is investigated as a full-load pickup and delivery problem with time windows. On-line as well as off-line experiments based on real data from a construction company in the UK are reported and discussed. Various practical issues that arise in this real time logistical problem are also discussed and analysed. Interesting and encouraging results are reported.

*Keywords:* real time routing, logistics, tabu search, case study.

## Introduction

This paper develops a systematic way of modifying a schedule in response to changes in customer's demands and resources that occur while the schedule is in operation. The problem arises from the delivery of road-making materials of a large construction company in the UK. In this application goods have to be taken from a number of depots (called "works") to a large number of customers where the orders have to be delivered in full (i.e., partial fulfillment of orders is not allowed). The objective

is to minimise the cost of making deliveries while satisfying all the necessary constraints. This problem fits into the class of real-time vehicle routing and scheduling. In recent years, this type of ‘on line optimisation’ appears to have gained a considerable amount of attention from researchers as well as practitioners. The proposed implementation which uses an existing tabu search heuristic already developed by the authors, as part of its research engine, is tested at two of this company dispatch centres. We would like to emphasize that the aim of this paper is not to put forward a new heuristic for the case of real time scheduling, but to focus on a real life implementation while highlighting some of the issues that are usually encountered by dispatchers when faced with changes in real-time routing and scheduling.

The paper is organized as follows. The remainder of this section describes the problem, highlights the types of change encountered in this particular logistical application and gives a brief review of the literature. The next section provides an illustrative example, describes briefly the heuristic adopted and outlines the proposed approach in handling changes in real time. The main steps of this approach are explained in section 3. Section 4 presents results of tests made off-line and on-line respectively, at Oldbury and Widnes, two of the company’s dispatch offices in the UK. Our findings and some suggestions for further research are summarised in the last section. Two appendices are also added for completeness, one for the tabu search algorithm used in this work, and the other for recording the detailed results.

## **A Brief Description of the Problem**

Works are locations from which a product may be collected for delivery. Journeys are normally made between works and customers. Customers making multiple requests are regarded as multiple customers. If an order is too large for one vehicle it is divided into known vehicle-loads as used by the company.

Loading times for a given works, and unloading times for a given customer, are assumed constant. The set of products is partitioned into ‘product groups’. A vehicle that has carried a member of one group may have to be cleaned, before a member of another can be loaded onto it. If so, the delay thus incurred is also assumed constant. If a vehicle arrives early at a works or customer site, it waits

until the start of the time window to begin loading or unloading. In some instances, time windows at customer sites are soft, so that if a vehicle arrives late, a penalty is incurred, up to a certain limit after which no allowance is regarded as practically feasible by the dispatcher.

Usually each vehicle starts and finishes each day at its 'home' location (a works or a garage). Travelling cost is a function of the distance travelled and the load carried, and varies from vehicle to vehicle, and from works to works. Each vehicle used also incurs a fixed cost, which may vary with the size and the type of the vehicle.

Each works has an associated list of products that are available from it. A customer may insist on being supplied from a particular works. There is an unlimited supply of each product at each works that supplies that product. All the works can accommodate a vehicle of any capacity, but a customer may specify the maximum size of a vehicle that can unload at their site. Not all vehicles can carry all products.

### **The Types of Changes for this Logistical Problem**

In this particular scheduling problem, real-time changes occur to vehicles, works, customers or the road network, and the impact on the schedule may be different in each case. Usually, the scheduler is notified of the changes by telephone, and uses the same means to convey information to works managers and customers. For instance, at Oldbury, near Manchester in the UK, drivers are not given their daily schedules in advance, but are told, when they reach a works, which is the next customer, and the next works, on their schedule. The dispatcher telephones drivers in their vehicles to inform them of any possible changes.

#### *Customers-*

A customer may make a new request at the last minute, or defer, cancel, or change quantity or material. The frequency of changes of this kind varies with the weather. The scheduler usually tries to accommodate the change. An emergency at the site (a fire, for instance) may delay the vehicle at the site, or enforce cancellation or relocation. Even a short delay can have a 'knock-on' effect, if several loads are

needed in a short period. Cancellation may release the vehicle to service another customer more cheaply than the scheduled vehicle could, whereas other changes may prevent the vehicle from servicing customers assigned to it. In either case, the change could make it necessary or desirable to recast the rest of the schedule.

#### *Vehicles-*

Common causes of change include drivers phoning in sick, delays due to traffic, and breakdowns. Furthermore, a driver may decide at the last minute not to accept a journey, or to finish work for the day. On the other hand, independent hauliers sometimes enquire whether there is work for them (perhaps in a specified area) during the day. Any of these factors could necessitate rescheduling. In the event that transport unexpectedly becomes available, it is often acceptable to advance deliveries of dry goods that are due the following day, provided (for accounting reasons) that a month-end is not spanned.

#### *Road Network-*

Accidents, road works or disruption to public transport (as well as factors known in advance, such as school holidays) can affect the times of many journeys in the vicinity. Unexpected events of this kind may necessitate rescheduling.

#### *Works-*

Naturally, stocks can run low at a works, but the scheduler is usually forewarned of this. Less predictable events at the works can have far-reaching consequences. Natural disasters such as a flood could put a works out of commission for days. If a loading shovel fails, it can be replaced or repaired, but not generally until the following day. The Department of the Environment gives about a week's notice of its intention to check weighbridges. Should defects be found, they must be remedied before the weighbridge can be used again. Some but not all works have more than one weighbridge. In the worst case, a works might be out of action for a fortnight. Other calamities occur, such as stock becoming contaminated and therefore unusable.

Although the obvious and most commonly used response is to divert all routes from the affected works to the nearest alternative, a strategy that moves stock in

bulk from one works to another could be investigated in future. In this case, the receiving works is regarded as a customer site. Some works regularly need to import stock for production. For example, the works at Berkswell imports gravel, to mix with locally extracted sand to produce concrete.

### **Some Related Work**

We present some of the work that describes techniques related to changes in real-time. Due to the advances of IT and computing power, real time optimisation especially within freight management has become a challenging and a practical research avenue that undoubtedly will be one of the hot topics in the next decade for academics as well as practitioners. For instance in 2004 a special issue of the journal *Transportation Science* edited by Gendreau & Potvin (2004) was entirely devoted to this area. Real time rescheduling also occurs frequently in production scheduling where the customers change their order, decide to cancel their order, among others due to unplanned circumstances. An MPhil dissertation by Lee (2004) was devoted to responding in real time to frequent rescheduling for the case of a construction company in the UK. A hybrid of dynamic programming and simple insertion rules were used to deal with such a practical problem. Du, Li & Chou (2005) put forward some algorithms suitable for dynamic routing in the case of business to consumer electronic commerce. Very recently Li, Mirchandani & Borenstein (2007) present the so-called the vehicle rescheduling problem and develop fast algorithms including ideas on how to construct a common feasible network which defines initial prices to speed up the auction algorithm. Related references in rescheduling for flexible manufacturing can also be found in this paper.

The degree of dynamism of a problem is usually measured by the number of changes occurring during a time period, the tightness of the time windows and the time of service of customers as found by the schedule. A more formal expression of this measure is reviewed in the thesis of Larsen (2001).

There are two main solution approaches that deal with dynamic routing problems. The first one, known as ‘reactive or event-based approach’, attempts to incorporate changes on the fly depending on their urgency, whereas the second one,

known as ‘time slicing’, transforms the problem into subproblems where each one refers to one time slot. Each subproblem is then solved and information is passed from one subproblem to the next. Changes are appended usually to the next subproblem. Authors working in the latter aspect consider the new arrival of customers as the only possible change and these include Montemanni, Gambardella, Rizzoli & Donati (2005), Hanshar & Ombuki-Berman (2007) and Wallace (2007). In this study, we concentrate on the first category as a variety of changes do occur in this particular practical problem. An interesting review based on the first category is given by Ichoua, Gendreau & Potvin (2000). The authors mention the following techniques.

*Stochastic Methods:* Models can be based on Markov decision processes, or these can be combined with classical network formulations.

*Expert Systems:* These involve recording and mimicking decisions taken by an experienced human dispatcher in a similar context.

*Diversion Strategies:* The paper contrasts strategies that retain the current destination with those that allow diversion from the arc being traversed when the new information arrives. It points out that a finite time has to be allocated for the schedule to be modified, before any diversion can begin. It suggests three different ways of setting this time, and presents a diversion procedure within a tabu search framework for a class of routing problems with time windows.

*Adapting Static Algorithms:* This implies a rolling horizon inside which future events are considered. When new information is input, either the optimization procedure can be re-run, or insertions and other local operations can be performed, with or without post-optimization procedures.

The method used in the current paper can be regarded as a modification of the last two approaches. It is thought suitable to this problem, because it deals with the variations caused by discrete events, and provides schedules independently of the dispatcher, who can accept or modify them.

In a later paper Ichoua, Gendreau & Potvin (2003), discuss different ways of modelling travelling times when the speed varies during the day. They stress the importance of the FIFO (first-in, first-out) property, and show that models that



do not guarantee this property may generate unnecessary waiting time at nodes. A model is presented that associates a speed with each pair of nodes and each time period. The travelling speed between each pair of nodes changes when the boundary between time periods is crossed. A similar idea is used in this research to measure the different times during the day and between zones (rural, urban, city centres etc). An approach that takes into account the dependency of the travel time using queueing concepts is developed by vanWoensel, Kerbache, Peremans & Vadaele (2003). Fu (2001) adapts a Dynamic Programming technique to determine the optimal route to a known destination via a series of nodes, where the times over each arc are stochastic variables. These times are predicted with reduced uncertainty as a vehicle approaches the starting node of the arc concerned. The development of a nonmyopic heuristic that considers advance information about future tasks is proposed by Spiney & Powell (2004), and a dynamic solution approach where different scenarios for future travel times are used is devised by Huisman, Freling & Wagelmans (2004). Thangiah & Awan (2004) adapt their heuristic, originally designed for split delivery pickup and delivery, to respond to real time events such as addition, deletion or change in information of a customer or driver. Potvin, Xu & Benyahia (2006) introduce the concept of tolerance to unforeseen delays within various dispatching strategies when studying vehicle routing and scheduling with dynamic travel times, and analyse the impact of different levels on the solution quality. Eglese, Maden & Slater (2006) demonstrate empirically the usefulness of using time-dependent travel time in vehicle scheduling and constructed a data base for road time that estimates more accurately the travel time. This interesting idea was tested successfully on a road network in the North of England.

## **Methodology**

In this section we first give an illustrative example and a brief description of the tabu search heuristic used. Our approach on how to respond to the changes is then briefly outlined. Some explanations of the main steps of this approach are provided in the next section.

## An Illustrative Example

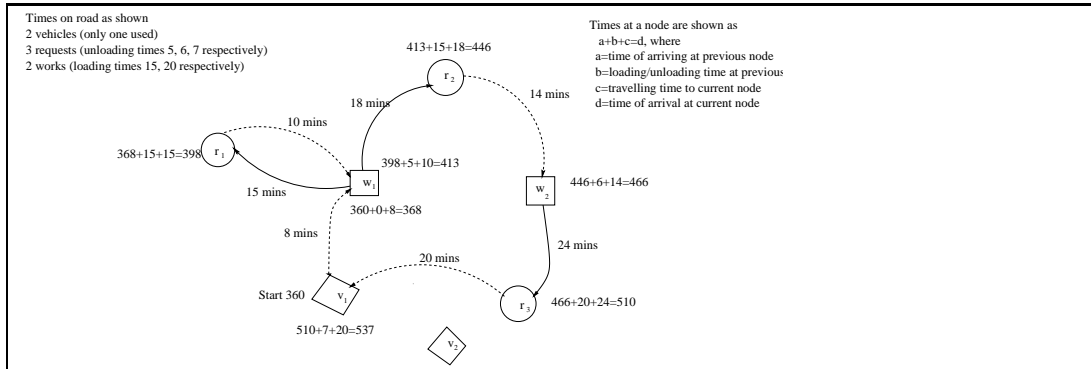


Figure 1: A Simple Example of a Route

Figure 1 represents an instance involving 3 requests, 2 works and 2 vehicles, of which only 1 is needed. The dashed lines indicate empty travelling, the solid lines full travelling. Times are shown as minutes past midnight. For example, vehicle 1 is available to start from  $v_1$  at 6 o'clock (ie 360). It takes 8 minutes to travel to works  $w_1$  (hence arrival time is 368), 15 minutes to load, and 15 minutes to travel to the site of request  $r_1$  (hence arrival time is 398). The optimal route is the one shown in figure 1, finishing at  $v_1$  at 8:57. This solution was obtained by an LP-based method developed by the authors and which uses the software MP-XPRESS, available from Dash Associates, for its execution. More details on the problem, including the mathematical formulation, can be found in Currie & Salhi (2003). Also in the same paper, the authors proposed a 3-level heuristic which produced relatively good solutions when compared to the LP-based method on small instances. In level 1 an initial solution is first created using an adaptive heuristic that uses greedy methods and regret costs. Levels 2 and 3 are refinement-based procedures.

### The method adopted

The tabu search (TS) method which was previously put forward by Currie & Salhi (2004) is the method that is used in this implementation. For convenience, the main steps of the method are briefly given here and the algorithm is reproduced in Appendix 1 for completeness.

The initial solution is first constructed as in Currie & Salhi (2003) though any constructive heuristic could be used for this purpose. Such a solution is then improved using descent-based improvements such as the Or-Opt move which is a quicker version of a 3-opt move, the SHIFT move which considers a chain of customers from the start or the end of a route to be inserted into another route, and finally the CROSS move which swaps chains of customers from any part of two different routes (each chain may consist of 1, 2 or 3 customers). Full details of the CROSS, the Or-Opt and the SHIFT moves may be found, respectively in Taillard, Badeau, Gendreau, Guertin & Potvin (1997), Potvin, Kervahut, Garcia & Rousseau (1995) and Currie & Salhi (2004). These three moves are the ones used within the TS heuristic of Currie & Salhi (2004). In that paper diversification strategies and a scheme that deals with controlling infeasibility are also embedded in this search. In addition, a soft aspiration criterion within tabu search, as described in Salhi (2002), is also used. Here, a move which is tabu with a relatively smaller tabu tenure but which yields a cost close to, but not better than, the overall best is also considered as a potential move and not discarded from the search as ‘usually’ implemented in tabu search. This tabu search heuristic was made even more efficient due to suitable data structures and neighbourhood reductions that were embedded into the search. These computational schemes, see Currie (2003), cut down on the computing time significantly without a detriment on the solution quality.

### **The approach to responding to changes**

If desired, minor amendments can be made manually to the schedule, whether they arise from real-time changes or are simply the exercise of corporate preference. Our methodology, which recasts the schedule on notification of change, consists of 3 main steps namely the generation of the original schedule (step 1), the amendment of the data to respond to the changes (step 2) and finally the rerun of the heuristic on the remaining data (step 3). The pseudo-code of this approach is given in Figure 2 and the main steps are amplified where necessary in the next section.

```

STEP 1 Generate a schedule at start of the day.
STEP 2 Occurrence of a change

    2.1 Estimate time (T', say) from which the revised schedule will start.

    2.2 For each request
        If loading is due to have commenced before T', then
            - Erase the request from the input to the generation of the revised schedule.
        Otherwise
            - Include such a request in the input.

    2.3 For each vehicle
        If at time T', a vehicle is approaching a works or its depot, then
            - Calculate the arrival time (TW, say), and
            - Record vehicle as available and at its destination as from TW.
        Otherwise (vehicle is at or approaching a customer site)
            - Calculate time unloading should be completed (TC, say), and
            - Record vehicle as available and at the site as from TC.

STEP 3 Recast schedule

    3.1 If the only change is a cancellation, then
        - Use the uncompleted part of the existing schedule as the initial solution
    3.2 Otherwise
        - Run the heuristic from the beginning.

```

Figure 2: The Real Time Rescheduling Heuristic

## The Main Steps of the Rescheduling Heuristic

In this section we describe briefly the main three steps of this rescheduling heuristic.

**The original schedule (step 1)** This schedule can be generated the night before, using any suitable heuristic. In this work the tabu search heuristic briefly described in the previous subsection (see also Currie & Salhi (2004) for details) is used.

### *Generation of the inputs*

In practice the validity of the data constitute an integral part of the success of any heuristic. In other words, if an input is ignored or wrongly measured it will systematically yield unwelcome outcomes regardless of the power or the efficiency of the method used. The main items that are worth noting are briefly outlined here.

### *Computation of travel time*

The travelling speed may vary with the locality or the time of day. The calculation of costs and times from distances or Cartesian coordinates is therefore separated from the heuristic, to allow the user to alter them if necessary.

The day is divided into 3 time zones. Travelling speeds are assumed constant within each zone. Consider an arc that starts at time  $a$  in Zone 1. Its travelling

time ( $\theta_1$ , say) is calculated at the speed obtaining during Zone 1. Suppose that the boundary between Zones 1 and 2 is at time  $b$ , and that  $a + h + \theta_1 > b$ , where  $h$  is the loading or unloading time at the starting node of the arc. Then the travelling time is recalculated at the speed used during Zone 2 ( $\theta_2$ , say). The total time ( $t$ ) from arrival at starting node to arrival at finishing node is then calculated as  $t = h + (b - a) + \frac{\theta_2}{\theta_1} \times (a + \theta_1 - b)$

This model allows speeds to vary during the day and satisfies the FIFO property as mentioned by Ichoua et al. (2003). The process of transforming speeds into journey times forms a part of the interface between input and output, as represented in Figure 3.

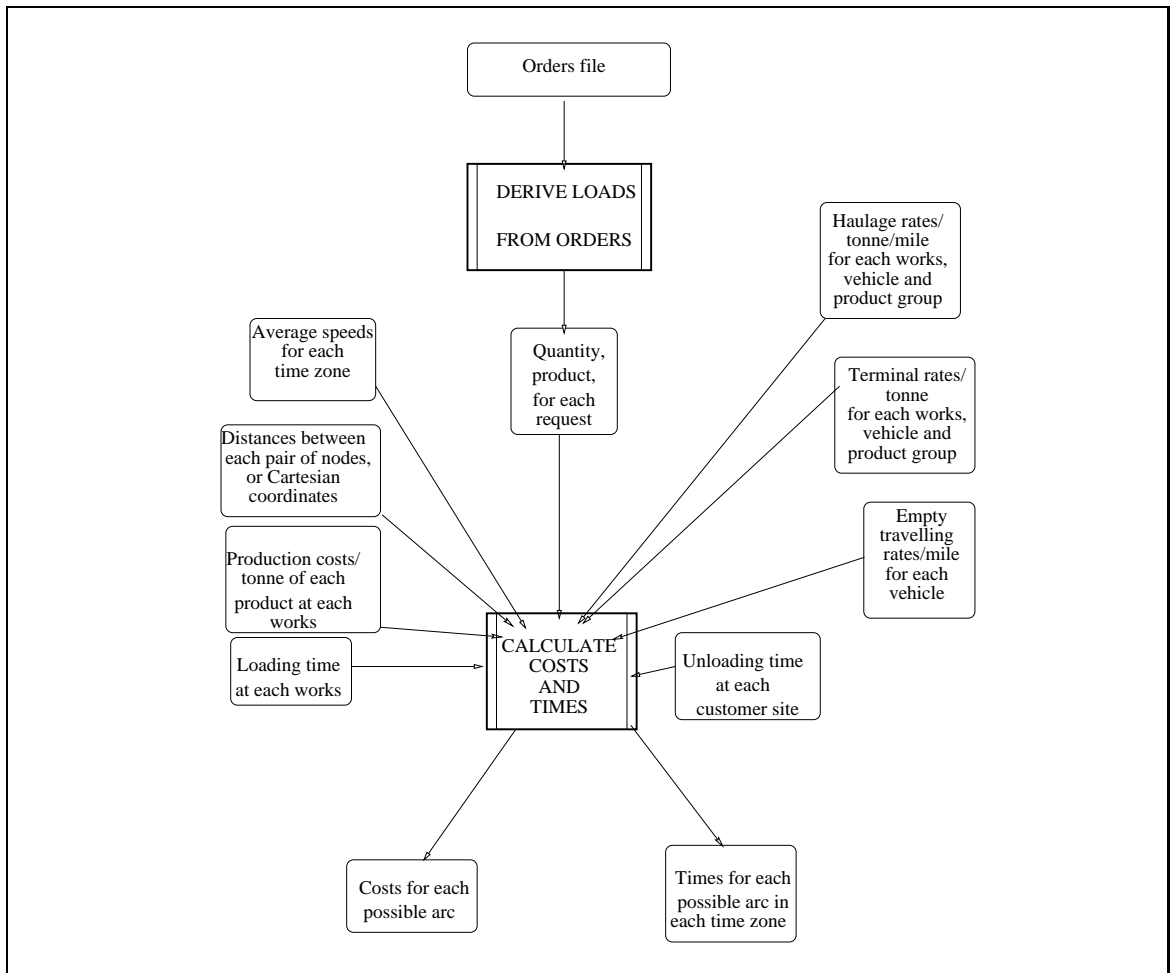


Figure 3: Calculation of Times and Costs

Figure 4 shows the required inputs to and outputs from the program that uses

the heuristic to create a schedule, and depicts the process of making changes while the schedule is in operation.

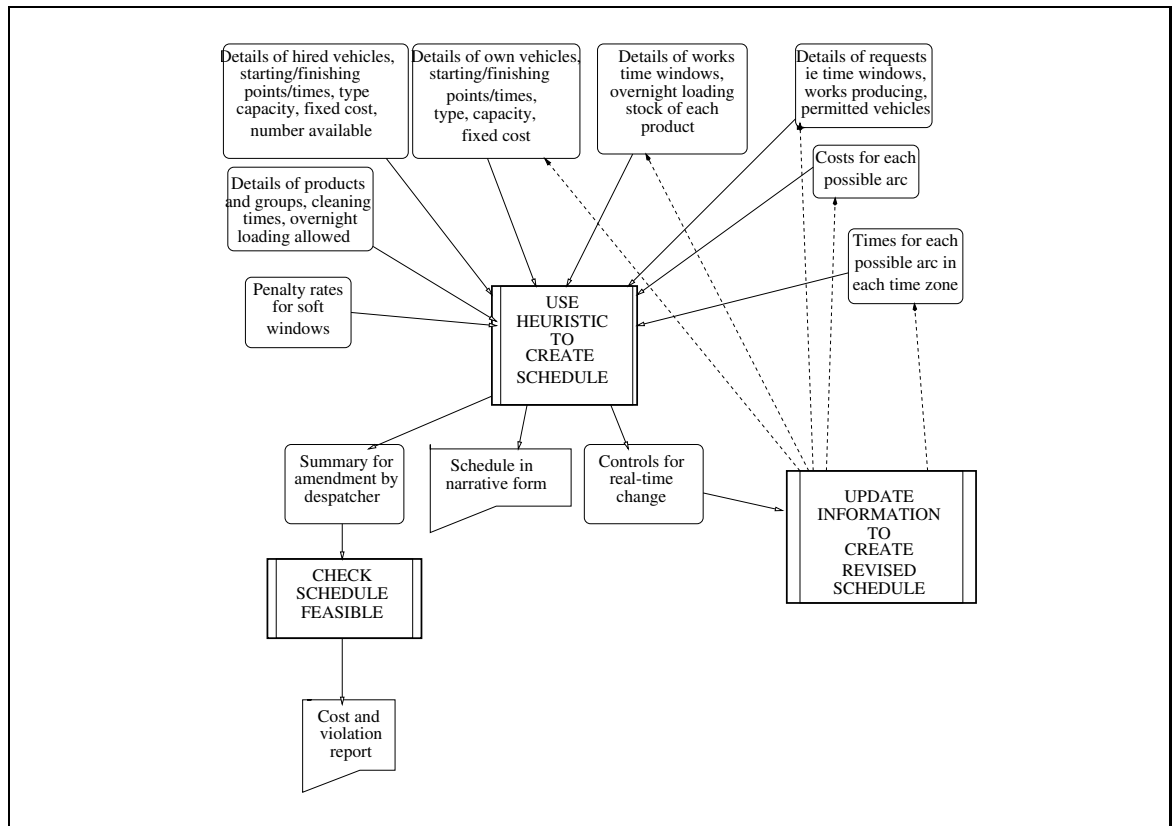


Figure 4: Inputs to and Outputs from the Heuristic

**Amend data (step 2)** In order to include those customers that have not yet been served, in the light of real-time changes, it is necessary to update the data files to reflect the current position, and re-run the program. To facilitate this, the creation of the original schedule includes routines that file details of vehicles used including hired vehicles, large demands that need to be divided where necessary into loads, and the schedule in a form that can be input to a program for revision. The data files have to be amended to reflect any significant changes of the type discussed above, before the program is rerun to revise the schedule.

*Time allowance*

Let

$t_1$  : the anticipated time to record changes to the data,

$t_2$  : the allowed running time of the heuristic,

$t_3$  : the time required to inform drivers of changes to their next destination.

Note that  $t_3$  is estimated based on the advice from the schedulers of this particular company. The value of  $t_1$  is based on how quickly the new information is recorded and the way the files are organised whereas  $t_2$  relates to the computation time the heuristic usually takes. A cut-off time for  $t_2$  is then set accordingly and the best solution recorded up to that time is returned as the new solution.

Let  $\Delta = t_1 + t_2 + t_3$  be the total time spent to allow for the new schedule to be practically implementable. In our tests, we set a 15 minute total allowance for  $\Delta$  by using approximately 5 minutes each for  $t_1$ ,  $t_2$  and  $t_3$  though  $t_2$  could be easily reduced without affecting the solutions. If a change is notified at time  $T$ , then the input to the revision is time  $T' = T + \Delta$ . Changes that affect customers for which loading is due to have commenced at the works before  $T'$  are deemed to be unalterable, and are omitted from the data input to the revision. All other customers are rescheduled. Such a restriction, though it is somewhat limiting, can be considered practical from the viewpoint of implementation. Such an assumption is also used by other researchers such as Ichoua et al. (2000). Vehicles approaching a works or their finishing point at  $T'$  are available from the time of arrival. Vehicles at or approaching a customer site at  $T'$  are available from the site from the time unloading is complete.

If it is decided not to make further use of a vehicle that is away from its base at the time of rescheduling, the cost of returning the 'stranded' vehicle to its depot must be included.

On the other hand, the product group of the last delivery made by each vehicle before the interruption must be recorded, in case the first delivery required by the amended schedule is of a different group. In this case a cleaning delay may be incurred.

**Rerun the heuristic (step 3)** If the only change is the cancellation of one or more customers, then the existing schedule is a feasible solution to the amended instance though it may not be very good. It can therefore be input as the initial

solution to our tabu search heuristic or to any improvement heuristic available to the user. If other changes have been notified, the existing schedule cannot be assumed feasible, and so a new initial solution must be generated. In either case, a modified version of the heuristic is run, which terminates after  $t_2$  minutes.

#### *A Rapid Succession of Changes*

Suppose that while a change is being processed, a further change is notified. The question arises: should the dispatcher finish the first change, then ratify the second, or abort the first change and deal with the two together? Roughly, the 15 minutes allowed for the process can be thought of as 5 minutes to amend the data, 5 to rerun the heuristic, and 5 to redirect drivers. Since the two changes might, if processed separately, require some drivers to be redirected twice, it would seem best to deal with them together if possible, thereby eliminating unnecessary (and perhaps confusing) telephone calls. As a rule of thumb, then, if one change supersedes another, the process should be aborted, and the two changes processed together, unless some drivers have been redirected. In that case, the dispatcher should complete the first change before dealing with the second, and should leave the redirection of drivers as late as is practical.

### **Computational Experiments in Practice (On-line/Off-line)**

In this section we present results using two scenarios namely the off-line (eg., the changes are provided afterward and known apriory) and the on-line (no information is given and the heuristic needs to react to the changes as quickly as a human dispatcher does). These two scenarios use information from a real life scheduling application in the UK. To guarantee the use of the company owned vehicles (or those commonly relied upon contractors) our method is designed to fully use these vehicles first and minimise the number of additional vehicles afterward. Though such a decision may conflict in some occasions with the best schedule the company is happy to adopt such a strategy rather than a schedule which may use other vehicles first due to mileage cost etc. To achieve our goal we assign a large fixed



cost to these additional vehicles. More details on this issue can be found in Currie (2003).

### **Testing Real-Time Changes Off-Line**

The incidence of changes is highly variable, but the following is a rough guide based on the experience of the dispatchers. At Oldbury, there are 12 changes on average per day, of which most (10 or 11) are demand-based. They may occur at regular intervals, or in clusters. Other changes are usually caused by delays due to the road network, which includes Birmingham, or to a vehicle going out of service. Changes arising from difficulties at a works occur more rarely.

At Widnes, on average 5 orders per day are cancelled; one has a change to the quantity, time window or supplying works. One vehicle may also be taken out of service, and one short (30 minute) delay may occur in production. Longer delays (say, 3 or 4 hours) may happen on average once a fortnight. The area covered includes sections of the M60 and A6, and major traffic delays happen occasionally.

Six scenarios, based on the orders for a day at Oldbury and Widnes, have been constructed. Each scenario consists of simulated changes to the orders or available resources. These vary in complexity and impact, but the intention is approximately to reproduce the actual incidence of the various types of change. The exact times of notification are generated randomly. The costs of, and numbers of vehicles used by, the resulting schedules are compared with those that the heuristic would have produced had the changes been known in advance. Tables 1 and 3 summarise the results for Oldbury and Widnes respectively. Examples of scenarios appear in Tables 2 and 4. The results are given in Appendix 2. In the tables, the cost shown includes the deliveries made before the change, and is therefore comparable with the cost of the original schedule. An arbitrary large fixed cost of say £10000 has been assigned to each hired vehicle, to discourage hiring. The costs shown in the tables represent the travelling costs only. Where the cost is omitted, the change has been processed with the following one.

In scenarios O1 and O3 the travelling cost of the schedules produced on the assumption that all changes were known at the outset exceeds the total costs after

Table 1: Summary Results for Oldbury-Based Scenarios

Scenario	Original (No Changes)		Real-Time Changes		Changes Known at Outset	
	Cost	Vehicles	Cost	Vehicles	Cost	Vehicles
O1	4373	7	4441	11	4446	7
O2	4373	7	4410	10	4396	8
O3	4373	7	4446	9	4450	7
O4	4373	7	4389	7	4384	7
O5	4373	7	4446	9	4416	7
O6	4373	7	4871	13	4820	12

Table 2: Scenario O3: 12 request changes occurring in clusters

Change	Time	Nature of Change	Cost	Hired
1	07:45	B101 20→17 tonnes 6W	4286	7
	07:47	RY203 cancelled		
	07:56	S50002 required 10:00-13:00		
	08:01	W101 cancelled		
2	08:43	W20004 product 1444→2300	4300	7
	08:50	RG2006W 17→20 tonnes (2x10 tons) 6W		
3	12:30	H20015 +20 tonnes 11:45-16:00	4324	8
	12:36	A10004 cancelled		
4	13:15	RY203 reinstated	4446	9
	13:17	A10004 reinstated		
	13:20	W101 reinstated		
	13:29	H21202 product 1444→1824		
Original schedule at Oldbury			4373	7
If changes known at outset			4450	7

the last amendment (£4441 and £4446 respectively). These results, which at first sight look odd, are due to the necessity to hire more vehicles when the changes are *not* known in advance. The arbitrary fixed cost of a hired vehicle, which is imposed whenever a vehicle is hired, causes the program to hire as few vehicles as possible. During rescheduling, vehicles already hired are replaced with in-house vehicles whenever possible. Thus, in each case, the total 'cost' (including the fixed costs) is smaller when changes are known from the outset.

The dispatcher at Oldbury points out that among his secondary objectives is the wish to equalize jobs among the company's owner-drivers during slack times. He would like to give each of them at least 3 jobs during the day. Occasionally it happens that, without warning, a vehicle cannot access a customer site, so that the load has to be returned to the works or diverted to a different customer. The program as written does not deal with this situation directly, as it assumes that vehicles en route to a customer at the time of the change successfully complete that delivery. The dispatcher would have to re-input the 'failed' delivery, mark as 'completed' the delivery that replaced it (if any) and adjust the starting time

and place of the vehicle. The dispatcher also comments that it would be helpful if changes to scheduled deliveries were highlighted on the output.

Table 3: Summary Results for Widnes-Based Scenarios

Scenario	Original (No Changes)		Real-Time Changes		Changes Known at Outset	
	Cost	Vehicles	Cost	Vehicles	Cost	Vehicles
W1	11343	1	10619	1	10628	0
W2	11343	1	10652	1	10655	0
W3	11343	1	10412	2	10621	0
W4	11343	1	10653	1	10655	0
W5	11343	1	11189	1	11183	1
W6	11343	1	11168	2	11160	2

The dispatcher at Widnes similarly wishes to equalize jobs among the company's owner-drivers during slack times. The comments made about Oldbury therefore apply here. Bearing in mind the larger fleet size at Widnes, the dispatcher might be more inclined to consider reducing the number of in-house vehicles kept, in the interest of economy.

### Testing Real-Time Changes On-Line

This section presents the results of a day's testing at Widnes. A schedule was prepared for that day based on data available the previous evening. During the day, changes were recorded whenever the Dispatch Office was notified of them, and the schedule recast for the rest of the day. Table 5 records each change, the time it was notified, and the cost, and number of hired vehicles used, according to each schedule. For instance in row #6 (change #6), there is a new request denoted by H056 and which requires a 500 tons of product number 1506 that needs to be delivered to customer whose coordinates are (213,229). This change occurs at 8.25 in the morning. At the suggestion of the dispatcher, the customers that

Table 4: Scenario W1: 7 request changes occurring at regular intervals

Change	Time	Nature of Change	Cost	Hired
1	07:26	D30814 cancelled	11257	1
2	08:14	F70601 16→20 tonnes	11270	1
3	09:30	D32501 required 11:45-12:30	11270	1
4	09:55	F70902 cancelled	11247	1
5	10:39	D30807 cancelled	11135	1
6	11:08	D30602 cancelled	10730	1
7	11:39	D30810 cancelled	10619	1
Original schedule at Widnes			11343	1
If changes known at outset			10628	0

Table 5: On-line Test at Widnes: 28/11/02

No.	Nature of Change Comments	Occurrence of a change	Travelling Cost	Vehilces Hired
1	S403: 16ton→14ton 8:15-9:00 + 14ton 12:15-13:00	07:05		
2	S312: 48ton→38ton	07:19	15927	24
3	New D335: 48ton (2268) 10:30-11:00 to (237,246)	07:37		
4	D318: 100ton→60ton	07:37	15856	30
5	D035: 400ton→600ton	08:25		
6	New H056: 500ton (1506) to (213,229)	08:25		
7	New H057: 10ton (1514) to (220,240)	08:27	17646	41
8	New H060: 40ton (1514) to (220,236)	08:55		
9	New H061: 20ton (1514) to (212,250)	09:04		
10	New H062: 20ton (1597) to (212,250)	09:04	17917	43
11	New H064: 16ton (1514) to (228,241)	09:17		
12	H304 and H305: Cancelled	09:18		
13	S404: Advance each load by 2 hours not possible for 1st 2 loads	09:21	17891	44
14	S307: Cancelled 1st load already on way	09:27		
15	S406: 7ton→9ton	09:36	17792	46
16	New D104: 20ton (1088) to (217,239)	09:55	18043	47
17	New H067: 20ton (1514) to (219,241)	10:06	17981	49
18	S418: 8ton→9ton but load already on way	10:26		
19	S307: reinstate 2nd load, 8ton→5ton 11:45→12:30	10:48	17969	49
20	New H074: 20ton (1514) to (203,233)	11:11	17991	49
Original schedule at Widnes			15915	23
If changes known at outset			18521	33
Actual itinerary			14901	48
Schedule created by program based on deliveries actually made			16967	29

required coated product are constrained to be supplied by the works specified by the company, whereas the heuristic has been allowed to determine the works supplying dry product.

Two items in Table 5 require comment. Changes 3 and 4 have the effect of advancing the customers' time windows for the arrival of 17-tonne trucks. The schedule was recast, with the result that the remaining loads of some hired vehicles, which had been due to be delivered in the afternoon, were cancelled. These hired vehicles were no longer required, but new hired vehicles were needed for other loads. The number given for hired vehicles is the total number that have been hired at any time, regardless of whether or not they are in use simultaneously, since the arbitrary fixed cost is assumed non-refundable. Similarly, change 17, which consists of one new customer, causes one existing hiring to be cancelled, and replaced with 2 hired vehicles.

Two differences between the schedule created by the heuristic and the actual itinerary are apparent. Firstly, in the actual itinerary, vehicles usually pick up from

the same works throughout their tour. Only in 8 cases did a vehicle return to a works other than the one from which it had made its last pickup. By contrast, the heuristic caused that to happen 48 times. Secondly, as mentioned above, the heuristic was allowed to change the works specified by the company, for dry goods only. It used this flexibility on 27 occasions.

The proposed system has some limitations as some changes can not necessarily be implemented. For example, according to the schedule created by the program, request S418 was already loaded by the time change 18 (increase from 8 to 9 tonnes) was notified. The program could not therefore accommodate the change, even if it could be made in practice on the day. This consideration limits the rigour of the exercise as a system test and hence further development is needed.

Furthermore, discussions with the dispatcher at Widnes after the test revealed that not all events had been communicated to the office on the day. For example, one vehicle had been detained at a customer site, and another had broken down during the day. These events would have vitiated the schedule created by the heuristic, but, had they been known, the program would have allowed for them.

The actual itinerary does not cover all the customers on the list of jobs for the day, and time windows cannot always be achieved. This means that comparisons between schedules generated by the heuristic and actual itineraries must be viewed with caution. In this case, actual times at customer sites were not available, but, making the same assumptions about travelling and servicing times as are made by the heuristic, 31 deliveries out of 235 would have been made after the time window.

The output from the program was shown to the dispatcher, who thought it satisfactory. This fact, and the results shown in Table 5, support the view that the system as written is viable.

### **An Additional Investigation**

For purposes of comparison, the program was also used to create and cost a schedule including only those deliveries that were actually made. This was run 3 times, assuming (i) arbitrarily large fixed costs of hiring vehicles and uniform haulage rates (line *S1*), (ii) zero fixed costs and uniform haulage rates (line *S2*) and (iii) zero

Table 6: Varying Cost Assumptions

			Travelling Cost	Vehicles Hired	Total minutes late
S0		Actual Itinerary 28.11.02	14901	48	4431
S1	Program:	Fixed Costs for hired vehicles	15562	20	
S2		No Fixed Costs	15561	29	
S3		Triple rates for hired vehicles	18962	24	

fixed costs and high (triple) haulage rates for hired vehicles (line S3). The dramatic reduction in the number of vehicles hired, yielded by all runs, demonstrates the scope for saving. Run *S2* results in more vehicles being hired, but yields no reduction in costs (as compared with *S1*). In run *S3*, as might have been expected, fewer vehicles are hired, but the overall cost increases.

## Conclusions and Suggestions for Further Research

In this study a tabu search heuristic that was developed earlier by the authors to solve a variant of the Pickup and Delivery Problem with Time Windows has been adapted to deal with changes to demands and resources that occur while the schedule is in operation. These adaptations have been tested, on-line and off-line, at two centres of a construction company in the UK, and the results are presented. The current implementation, though it may need additional practical adjustment, was considered satisfactory by the dispatchers.

It is clear that each vehicle's schedule is recast, not from the moment that the change takes place, but from the time the vehicle is next due to be empty, at a node. A more efficient use of vehicles might be possible if a vehicle that is about to leave a customer site at the time of the change could be made to await instructions. This could be done if all drivers could simultaneously be told to wait while the schedule was recast. This is not possible at the moment, as the only way of communicating with drivers is by telephone (or by leaving a message at the next destination). If the technology allowed direct instant communication, the possibilities introduced by having all drivers wait for a recast schedule would be a worthwhile subject for further research.

It was noted that both dispatchers would like as a secondary objective to balance

the drivers' loads and to guarantee a certain minimum load for all drivers. A bi-objective approach could be a useful way forward, especially from the practical point of view of the user.

Schedules could be better revised if empty vehicles could be diverted at, say, the next road junction on their route, rather than keeping to the original schedule as far as their destination works. Such a research avenue would obviously require the knowledge of the entire road network which can be generated via GIS. The implementation of such a methodology, using parallel computing strategies if possible, would speed up the process of the re-optimisation and renders the computerised system more flexible and more reliable.

*Acknowledgments-* The authors would like to thank the epsrc for the case studentship.

## References

- Currie, R. & Salhi, S. (2003), 'Exact and heuristic methods for a full-load, multi-terminal, vehicle scheduling problem with backhauling and time windows', *Journal of the Operational Research Society* **54**, 390–400.
- Currie, R. & Salhi, S. (2004), 'A tabu search heuristic for a full-load, multi-terminal, vehicle scheduling problem with backhauling and time windows', *Journal of Modelling and Mathematical Algorithms* **3**, 225–243.
- Currie, R.-H. (2003), *Distribution Management: An investigation into a full-load, multi-terminal vehicle scheduling problem with backhauling and time windows*, PhD Thesis, School of Mathematics and Statistics, University of Birmingham, UK.
- Du, T., Li, E. & Chou, D. (2005), 'Dynamic vehicle routing for online b2c delivery', *Omega* **33**, 33–45.
- Eglese, R., Maden, W. & Slater, A. (2006), 'A road timetable to aid vehicle routing and scheduling', *Computers and Operations Research* **33**, 3508–3519.

- Fu, L. (2001), ‘An adaptive routing algorithm for in-vehicle route guidance systems with real-time information’, *Transportation Research B* **35**, 749–765.
- Gendreau, M. & Potvin, J.-Y. (2004), ‘Issues in real-time fleet management’, *Transportation Science* **38**, 397–398.
- Hanshar, F. & Ombuki-Berman, B. (2007), ‘Dynamic vehicle routing using genetic algorithms’, *Applied Intelligence* **27**, 89–99.
- Huisman, D., Freling, R. & Wagelmans, A. (2004), ‘A robust solution approach to the dynamic vehicle scheduling problem’, *Transportation Science* **38**, 447–458.
- Ichoua, S., Gendreau, M. & Potvin, J.-Y. (2000), ‘Diversion issues in real-time vehicle dispatching’, *Transportation Science* **34**(4), 426–438.
- Ichoua, S., Gendreau, M. & Potvin, J.-Y. (2003), ‘Vehicle dispatching with time-dependent travel times’, *European Journal of Operational Research* **144**(2), 379–396.
- Larsen, A. (2001), *The dynamic vehicle routing problem*, PhD Thesis, Technical University of Denmark.
- Lee, T. (2004), *Real Time Optimisation for Road Stone Production*, MPhil dissertation, School of Mathematics and Statistics, University of Birmingham, UK.
- Li, J., Mirchandani, P. & Borenstein, D. (2007), ‘The vehicle rescheduling problem: Model and algorithms’, *Networks*, DOI 10.1002/net **47**, 211–229.
- Montemanni, R., Gambardella, L., Rizzoli, A. & Donati, A. (2005), ‘Ant colony system for a dynamic vehicle routing problem’, *Journal of Combinatorial Optimization* **10**, 327–343.
- Potvin, J.-Y., Kervahut, T., Garcia, B.-L. & Rousseau, J.-M. (1995), ‘The vehicle routing problem with time windows part 1: Tabu search’, *INFORMS Journal on Computing* **8**(2), 158–184.



- Potvin, J.-Y., Xu, Y. & Benyahia, I. (2006), ‘Vehicle routing and scheduling with dynamic travel times’, *Computers and Operations Research* **33**, 1129–1137.
- Salhi, S. (2002), ‘Defining tabu list size and aspiration criterion within tabu search methods’, *Computers and Operations Research* **29**, 67–86.
- Spiney, M. & Powell, W. (2004), ‘The dynamic assignment problem’, *Transportation Science* **38**, 399–419.
- Taillard, E., Badeau, P., Gendreau, M., Guertin, F. & Potvin, J.-Y. (1997), ‘A tabu search heuristic for the vehicle routing problem with soft time windows’, *Transportation Science* **31**(2), 170–186.
- Thangiah, S. & Awan, S. (2004), ‘Real-time split delivery pickup and delivery time window problems with transfers’, *Technical report: SRT90-2004, Artificial Intelligence and Robotics Lab, Slippery Rock University* pp. 1–31.
- vanWoensel, T., Kerbache, L., Peremans, H. & Vadaele, N. (2003), *An ants colony optimisation approach to VRP models with dependent travel times*, Proceeding of the 4<sup>th</sup> Aegan Conference on Manufacturing Systems.
- Wallace, M. (2007), *The dynamic vehicle routing: A meta-heuristic based investigation*, OhD dissertation, Institute of Mathematics, Cardiff University, UK.

## Appendix 1: The TS heuristic

```

Step 1 Obtain an initial solution using a constructive heuristic.
Step 2 Use SHIFT and Or-Opt to improve the solution and call the improved solution S1.
Step 3 Define the tabu list size and stopping limit. Set C(best), C(current) to the cost of S1. Make
      all moves non-tabu.
Step 4 Do until the diversification limit is reached:
  4.1 Initialise iteration number.
  4.2 Do until stopping limit is reached:
    4.2.1 Use CROSS and Or-Opt moves to search the neighbourhood.
    4.2.2 If feasible moves exist for which solution cost < C(best), then
      Perform move, set C(best) = C(current) = the least cost of such neighbours.
      Make the reverse move tabu.
    4.2.3 Else
      If non-tabu moves or moves that pass aspiration level exist, then
        For infeasible moves, calculate the penalty.
        Perform move, set C(current) to the least cost of such neighbours.
        Make the reverse move tabu.
      Else indicate stopping level reached.
      End if
    End if
  End do
  4.3 Apply diversification strategies
End do

```

Figure 5: The Tabu Search Heuristic

## Appendix 2: Results of Off-Line Tests

In Tables 7 to 18, the cost shown includes the deliveries made before the change, and is therefore comparable with the cost of the original schedule. The costs shown in the tables represent the travelling costs only. Where the cost is omitted, the change has been processed with the following one.

Table 7: Scenario O1: 12 request changes occurring at regular intervals

Change	Time	Nature of Change	Cost	Hired
1	07:33	B101 20→17 tonnes 6W	4364	7
2	08:14	RY203 cancelled	4335	7
3	08:58	S50001 required 10:00-13:00	4338	7
4	09:22	W101 cancelled	4282	7
5	10:10	W20003 product 1444→2300	4285	7
6	10:43	RY202 17→20 tonnes (2x10 tons) 6W	4301	8
7	11:45	H20015 +20 tonnes	4347	9
8	12:28	A10004 cancelled	4322	9
9	12:55	RY203 reinstated	4355	10
10	13:50	A10004 reinstated	4389	10
11	14:14	W101 reinstated	4441	11
12	14:30	H20017 product 1444→2300	4441	11
Original schedule at Oldbury			4373	7
If changes known at outset			4446	7

Table 8: Scenario O2: 11 request changes and one vehicle breakdown, occurring at regular intervals

Change	Time	Nature of Change	Cost	Hired
1	08:09	RY2026W → 241 159 6W	4387	7
2	08:24	18228 out of service after current job	4387	8
3	08:42	RY203 +3 tonnes	4392	8
4	09:06	H20009 14:00-16:00	4392	8
5	09:21	A10001 20→17 tonnes (6W)	4394	8
6	09:40	A10004 cancelled	4370	8
7	10:52	H20014 +20 tonnes	4421	9
8	11:21	H208 cancelled	4389	9
9	12:19	W20004 cancelled	4341	9
10	13:03	H20018 Berkswell, Hopwas	4345	10
11	13:33	W20004 reinstated 20→17 tonnes (6W)	4391	10
12	13:56	A105 product 1823→2300	4410	10
Original schedule at Oldbury			4373	7
If changes known at outset			4396	8

Table 9: Scenario O3: 12 request changes occurring in clusters

Change	Time	Nature of Change	Cost	Hired
1	07:45	B101 20→17 tonnes 6W	4286	7
	07:47	RY203 cancelled		
	07:56	S50002 required 10:00-13:00		
	08:01	W101 cancelled		
2	08:43	W20004 product 1444→2300	4300	7
	08:50	RG2006W 17→20 tonnes (2x10 tons) 6W		
3	12:30	H20015 +20 tonnes 11:45-16:00	4324	8
	12:36	A10004 cancelled		
4	13:15	RY203 reinstated	4446	9
	13:17	A10004 reinstated		
	13:20	W101 reinstated		
	13:29	H21202 product 1444→1824		
Original schedule at Oldbury			4373	7
If changes known at outset			4450	7

Table 10: Scenario O4: 11 request changes and one vehicle breakdown, occurring in 2 clusters

Change	Time	Nature of Change	Cost	Hired
1	07:54	RY203 → 241 159	4383	7
	08:00	R101 +17 tonnes		
	08:11	H20009 14:00-16:00		
	08:12	H102 20→17 tonnes (6W)		
	08:24	18228 out of service after current job		
	08:26	A10004 cancelled		
2	12:30	H20017 +20 tonnes	4389	7
	12:36	A50001 cancelled		
	12:43	W20505 cancelled		
	12:44	W201 Hereford or Hopwas		
	12:49	BH102 20→17 tonnes (6W)		
	12:55	H20013 product 1444→2300		
Original schedule at Oldbury			4373	7
If changes known at outset			4384	7

Table 11: Scenario O5: 10 request changes, 2 road network disruptions

Change	Time	Nature of Change	Cost	Hired
1	07:45	B101 20→17 tonnes 6W	4336	7
	07:47	RY203 cancelled		
	07:56	S50002 required 10:00-13:00		
	08:01	Add 30 minutes to zone 1 times at Rugeley		
2	08:43	W20005 product 1444→2300	4352	8
	08:50	RG2006W 17→20 tonnes (2x10 tons)6W		
3	12:30	H20012 +20 tonnes	4377	9
	12:36	A10001 cancelled		
4	13:15	RY203 reinstated	4446	9
	13:17	A10001 reinstated		
	13:20	Add 50% to zone 3 times at Hopwas		
	13:29	H20003 product 1444→1824		
Original schedule at Oldbury			4373	7
If changes known at outset			4416	7

Table 12: Scenario O6: 12 request changes and one breakdown in production

Change	Time	Nature of Change	Cost	Hired
1	08:45	A10001 needed before 12:00	4381	7
	09:07	W20004 +60 tonnes		
2	09:17	H21201 1444→1514	4519	9
	09:44	WN100 relocate→ 258 189		
3	09:52	Weeford works breaks down	4554	11
	10:28	R108 20→17 tonnes (6W)		
4	10:35	S50002 +100 tonnes	4764	12
	13:02	W20004 Ber, Her, Hop, Rag, Rug, Wil		
5	13:09	BH103 cancelled	4729	12
	13:15	H20015 needed after 15:30		
	13:19	H20018 needed after 15:30		
	14:11	BH103 reinstated		
6	14 16	H20003 +40 tonnes	4871	13
Original schedule at Oldbury			4373	7
If changes known at outset			4820	12

Table 13: Scenario W1: 7 request changes occurring at regular intervals

Change	Time	Nature of Change	Cost	Hired
1	07:26	D30814 cancelled	11257	1
2	08:14	F70601 16→20 tonnes	11270	1
3	09:30	D32501 required 11:45-12:30	11270	1
4	09:55	F70902 cancelled	11247	1
5	10:39	D30807 cancelled	11135	1
6	11:08	D30602 cancelled	10730	1
7	11:39	D30810 cancelled	10619	1
Original schedule at Widnes			11343	1
If changes known at outset			10628	0

Table 14: Scenario W2: 6 request changes, 1 vehicle breakdown, occurring at regular intervals

Change	Time	Nature of Change	Cost	Hired
1	07:36	D30814 cancelled	11262	1
2	08:01	F70701 16→20 tonnes	11277	1
3	08:17	D32501 required 11:45-12:30	11277	1
4	08:54	T904SBL out of service after current job	11281	1
5	09:11	D30809 cancelled	11170	1
6	09:57	D30602 cancelled	10765	1
7	10:25	D30810 cancelled	10652	1
Original schedule at Widnes			11343	1
If changes known at outset			10655	0

Table 15: Scenario W3: 7 request changes occurring in clusters

Change	Time	Nature of Change	Cost	Hired
1	07:35	D30814 cancelled		
	07:39	F70701 16→20 tonnes		
	07:45	D32501 required 11:45-12:30	11279	2
2	09:30	F71001 cancelled		
	09:40	D30809 cancelled	11130	2
3	10:39	D30602 cancelled		
	10:54	D30810 cancelled	10412	2
Original schedule at Widnes			11343	1
If changes known at outset			10621	0

Table 16: Scenario W4: 6 request changes, 1 vehicle breakdown, occurring in clusters

Change	Time	Nature of Change	Cost	Hired
1	08:29	D30814 cancelled		
	08:35	F70701 16→20 tonnes		
	08:41	D32501 required 11:45-12:30		
	08:48	T904SBL out of service after current job	11282	1
2	09:31	D30809 cancelled		
	09:34	D30602 cancelled		
	09:42	D30810 cancelled	10653	1
Original schedule at Widnes			11343	1
If changes known at outset			10655	0

Table 17: Scenario W5: 3 request changes, vehicle and works breakdown

Change	Time	Nature of Change	Cost	Hired
1	08:06	S30701 cancelled	11242	1
2	08:40	S40002 cancelled	11190	1
3	09:03	G897APL out of service after current job	11191	1
4	10:09	D30814 required 15:00-16:00	11192	1
5	12:01	Salford (coated) 30 minutes delay	11189	1
Original schedule at Widnes			11343	1
If changes known at outset			11183	1

Table 18: Scenario W6: 5 request changes, 2 road network disruptions

Change	Time	Nature of Change	Cost	Hired
1	07:37	D09102 required from Salford (dry)		
	07:42	D30814 cancelled		
	07:46	D32303 cancelled	11226	2
2	08:58	D09102 required 10:00-12:00		
	09:04	D30812 add 60 minutes (time zones 2 and 3)	11225	2
3	11:25	Forest Hill add 30 minutes (time zone 2)		
	11:31	D08301 cancelled	11168	2
Original schedule at Widnes			11343	1
If changes known at outset			11160	2

**University of Kent**

<http://www.kent.ac.uk/kbs/research-information/index.htm>