Control and Cybernetics

vol. **22** (1993) No. 4

A decision support system approach to milk tanker routing

by

C.B. Basnet, D.N. Clark, L.R. Foulds

Department of Management Systems
The University of Waikato
Hamilton
New Zealand

Vehicle scheduling is concerned with the efficient routing of a fleet of vehicles that must visit, from a central depot, a collection of clients to perform some service. Milk tanker routing is an example, where milk tankers need to be assigned to dairy farms for collecting milk and transporting to factories for processing. Traditionally, vehicle scheduling scenarios have been modelled by operations researchers as mixed integer programs to be solved by advanced mathematical programming techniques. Although attempts have been made to create realistic models, experienced schedulers have often found them hard to use. In this paper, we present a decision support system (DSS) approach to the milk tanker routing scenario in the New Zealand dairy industry. This DSS is designed to help the transport managers of the dairy companies build the schedules using their experience and preferences.

1. Introduction

Throughout New Zealand, dairy companies are faced with the question of how to collect milk from their supplier farms efficiently, using road milk tankers, and have it delivered to their factories for processing. Many companies own more than one factory, which must be supplied with milk each day, and operate fleets of milk tankers over two shifts per day: with some of the tankers based at locations other than factories. Each company has contracts with several farmers who supply milk regularly, often (but not always) daily to the company. To provide an indication of the problem size, a typical application may have two factories, 15 milk tankers, and up to 1500 suppliers.

In this paper, we discuss the milk tanker routing scenario in the New Zealand dairy industry and present a decision support system designed to aid the transport managers of the dairy companies. We begin by briefly reviewing the research on vehicle scheduling. We then introduce the milk tanker routing problem in the New Zealand dairy industry. Next, the requirements for a decision support system approach to this problem are identified. This is followed by a description of an actual decision support system called **FleetManager**. Finally, concluding remarks are presented.

2. Vehicle Scheduling

Milk tanker routing problem is a variation of the classical vehicle scheduling problem (VSP), which is concerned with the efficient routing of a fleet of vehicles that must visit, from a central depot, a collection of clients to perform some service. Each vehicle's route begins at the depot, includes serving several clients, and ends at the depot. Each vehicle's capacity for service must not be exceeded. Variations of this problem include multiple depots, multiple runs/day, and time window constraints.

One of the first studies of the VSP was reported by Dantzig and Ramser (1959) who formulated it as a linear programming model. Since then many others, including Clarke and Wright (1964), Gillet and Miller (1974), Foulds et al. (1974, 1977a, 1977b), Christofides et al. (1981), Fisher and Jaikumar (1981), and Bodin et al. (1983) have all contributed to the area. More recent contributions include a column generation approach to solve the VSP with time windows by Desrochers et al. (1991), and a network models approach by Carraresi and Gallo (1984) and by Soumis et al. (1991). Dell'amico (1989) has devised an assignment algorithm for the multiple depot VSP and Desrochers et al. (1990) have devised a classification scheme for VSP models. The mathematical programming literature on the VSP has been surveyed by Golden et al. (1977) and Watson–Gandy and Foulds (1981). Bibliographies have been given by Bodin et al. (1983), Lawler et al. (1985), and Golden and Assad (1988). An excellent survey of the VSP has been given by Bodin (1990) who discusses important components, approaches for solving practical problems, and other issues.

Ferland and Fortin (1989) introduced a heuristic approach for the VSP in which time windows of subsets of tasks are allowed to slide. Golden and Assad (1986) and Soloman (1987) have also considered time windows. Applications of VSP models include: school bus scheduling, public bus scheduling, beer delivery, milk collection, printing press scheduling, cement delivery, garbage collection, newspaper distribution, medical specimen collection, retail goods delivery, fuel oil delivery, bulk mail conveyance, meal delivery, mass transit crew scheduling, postal truck scheduling, ambulance service, disabled transportation, police patrols, and the delivery of industrial gases to hospitals. We now go on to examine just one of those applications that is of particular importance within the New Zealand agricultural sector.

3. Milk Tanker Routing Problem

The essential problem in milk tanker routing is to assign vehicles to start from a factory depot, collect milk from dairy farms (suppliers) with given milk output, and return to the same depot. The tanker vehicles have uniform capacity limitations and the factories have minimum demands to meet. In this problem, there are no time windows, travel distance, or travel time restrictions. There can be multiple depots (processing facilities). A single commodity (milk) is to be transported. The milk available with each client is estimated. A single tanker can make multiple runs in a shift. A mathematical formulation of this problem is given below:

Let: m = the number of vehicle runs

n = the number of locations, including factory depots (the first p locations are the depots)

C = vehicle capacity

a_i = the milk available at location $i (i = 1, 2, \ldots, n)$, $a_i = 0$ for $i = 1, 2, \ldots, p$

d_{ij} = the distance from location i to location j $(i = 1, 2, \ldots, n;\ j = 1, 2, \ldots, n)$, $d_{ii} = \infty$

x_{ijk} = 1, if vehicle k travels directly from location i to location j 0, otherwise, $(i = 1, 2, \ldots, n;\ j = 1, 2, \ldots, n;\ k = 1, 2, \ldots, m)$

D_i the demand of factory depot i for milk, $i = 1, 2, \ldots, p$

Objective:

To devise depot–to–depot routes for the vehicles so that all milk is collected and demand is met within vehicle capacity, with the minimum distance travelled.

Minimize

$$\sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{m} d_{ij} x_{ijk} \tag{3.1}$$

[Total cost is minimized]

Subject to:

$$\sum_{i=1}^{n} \sum_{k=1}^{m} x_{ijk} = 1, \qquad j = p+1, p+2, \ldots, n \tag{3.2}$$

[One vehicle visits each location]

$$\sum_{i=1}^{n} x_{ipk} - \sum_{j=1}^{n} x_{pjk} = 0, \qquad k = 1, 2, \ldots, m; \quad p = 1, 2, \ldots, n \tag{3.3}$$

[Tours must be continuous: if a vehicle run enters a location, it must leave the location]

$$\sum_{i=1}^{n} a_i \left(\sum_{j=1}^{n} x_{ijk} \right) \leq C, \qquad k = 1, 2, \ldots, m \tag{3.4}$$

[Vehicle capacity must not be violated]

$$\sum_{i=1}^{p} \sum_{j=p+1}^{n} x_{ijk} \leq 1, \qquad k = 1, 2, \ldots, m \tag{3.5}$$

[There are at most m vehicle runs]

$$q_i - q_j + n \sum_{k=1}^{m} x_{ijk} \leq n - 1, \tag{3.6}$$

for some real numbers q_i, $i = p+1, p+2, \ldots, n$; $j = p+1, p+2, \ldots, n$, and $i \neq j$.
[There cannot be any subtours]

$$\sum_{k=1}^{m} \sum_{i=p+1}^{n} a_i \left(\sum_{j=1}^{n} x_{ijk} \right) \left(\sum_{h=p+1}^{n} x_{rhk} \right) \geq D_r, \qquad r = 1, 2, \ldots, p \tag{3.7}$$

[The demands of factory depots must be met]

$$x_{ijk} = 0 \text{ or } 1, \quad i = 1, 2, \ldots, n;\ j = 1, 2, \ldots, n; \text{ and } k = 1, 2, \ldots m. \tag{3.8}$$

[Vehicle run k either travels from location i to location j or it does not]

Constraints (3.7) make this problem different from the usual multidepot vehicle routing problem that assume depots to have infinite capacity for service. Here, each processing facility has a demand for milk that needs to be filled.

The objective of minimization of the total distance travelled by all the milk tankers per shift is a surrogate for minimizing the cost per kilogram delivered to the factories. However, this is only one of many considerations in the comparison of possible scheduling options. Often it is only a secondary consideration of the busy planner who is under pressure to produce a satisfactory schedule. Combinations of other factors such as: the level of customer service, equity of route generation including driving and visiting times, rostering arrangements for drivers, efficient vehicle and driver utilization, the total schedule time, company financial strategies, access problems including certain vehicle–customer combinations, road inclinations in winter conditions, queuing at factories, accidents and breakdowns, geographical obstacles which complicate distance and time estimation, vehicle cleaning and servicing, union rules, labour and traffic codes, company and customer policies, and unpredictable human behaviour. Thus the dairy company scheduler is typically faced with a multitude of ill-defined objectives and constraints, the relative priorities of which may change markedly in a short time. These factors are seldom captured in computer packages that are based on mathematical programming.

We now go on to discuss decision support systems in general and how we propose to use them to address the milk tanker routing scenario in New Zealand.

4. The Decision Support System approach to milk tanker routing

Although the definition of **Decision Support Systems** (DSS) has been elusive [Davis (1988), Keen (1987) Er (1988)], the field has flourished with developments in technology. In his review of a decade of DSS development Keen (1987) stressed the need for balance among the three DSS elements: **decision, support** and **system**. He felt that more research effort on the decision component was required to restore this balance, as the technology for the system component was no longer the bottleneck. To achieve "the mission of DSS – helping people make better decisions", Keen stressed the need for an active supporting role for "decisions that really matter".

Sprague (1987) suggested extending the concept of the DSS "to include tracking-monitoring-altering and communicating as well as the more traditional intelligence-design-choice view of decision making." This would enable decision support systems to provide better support for ill-structured tasks. Clearly, focusing on the data and logic models of the DSS, without attention to the computer-human interface, will be inadequate. Acknowledging the influence of personality and cognitive style on an individual's decision style adds yet another dimension to the type of decision support that is appropriate for a specific situation [Er, 1988]. We next discuss a particularly fruitful area for the implementation of decision support systems — in milk tanker scheduling.

4.1. Requirements for a DSS for Milk Tanker Scheduling

We start from the premise that a vehicle scheduling computer system should be developed in order to aid, rather than to replace, the scheduler who is part of the scenario described in Section 3. New Zealand dairy company transport managers have traditionally used a large–scale map together with coloured pins displaying the farms with which they have supply contracts. This is a very useful visual aid and, to gain acceptance, it is usually productive for a decision support system to contain a computerized map. Geographical features and all relevant locations can be represented on the screen as a result of digitizing their coordinates. Relevant factory, milk tanker, and supplier information can be represented by using a colour graphics display system with windows and pull–down menus. This will allow a complementary combination of skills. The scheduler has the skill, superior to that of the digital computer, to recognise patterns in the location of suppliers and routes. These patterns will suggest possible options when routes have to be modified. Before these patterns can be translated into new routes, the scheduler needs information on predicted individual supplier output totalled for any cluster of suppliers. The generalisation of this information is a task to which the digital computer is ideally suited. Such a scheduler–computer combination marries the pattern recognition skills and the specialized knowledge that the experienced scheduler usually has, along with the numerical and recall ability of the computer.

One of the keys to the design and successful implementation of a DSS within the New Zealand dairy industry is first to discover the behaviour and strategies of a typical experienced milk tanker scheduler. One must then devise ways in which a DSS can make this person more efficient. We now discuss the tasks that such a scheduler typically faces on a daily basis.

Most dairy companies in New Zealand operate two shifts per day during the season when most suppliers are actively producing milk for collection. A master schedule is produced which is based on the demand for milk by each factory requiring milk from the suppliers of the company. Suppliers are grouped together into what is called a **run**, which is a sequence of suppliers that are visited by a tanker in a specified order. (The complete set of runs for a shift is called a **schedule**.) The object of this exercise is to allocate runs to factories so as to satisfy the demand of each factory. This involves assigning a tanker to each run. The initial run for each tanker begins at its base, visits the suppliers of the run in the order specified, and then ends at the factory for that run (which may be at a different location from the initial base). Subsequent runs of the tanker will begin at that factory and may well end at another factory. It is usually efficient to attempt to orchestrate the final run of the shift for each tanker to end at its base (assuming the base is a factory with positive demand) to minimize empty running. Because some suppliers have a relatively low output at certain times of the season, it is not considered worthwhile to visit them daily. Thus part of the allocation problem is the identification of which suppliers are to be visited for the current shift. There is also the question of the accurate prediction of supplier output. Also a judgement has to be made as to when frequency of visiting a particular supplier is to be either increased or decreased, due to a change in output.

Tanker schedulers usually approach these questions by first establishing which suppliers are to be visited on a given shift and then allocating them to the different factories to satisfy factory demand. If there has been no significant change from the previous similar shift the runs of that last shift are usually modified to generate the runs for the present shift. In examining previous runs, schedulers often ask themselves two key questions:

(i) What are the requirements of a new schedule that differ from the previous schedule?

(ii) How should the previous schedule be modified to create a satisfactory new schedule?

The first question usually involves constraints governing the feasibility of any new schedule. The second question involves not only these constraints but also the measurement of how satisfactory the new schedule is, in terms of various objectives. We now identify some possible main-menu options of a DSS that the scheduler may find useful in the search for answers to these questions.

The schedule list option

In order to begin the process of new schedule generation, based on the previous schedule, the scheduler must first be able to access the previous schedule. The DSS should have a listing of all the runs for any previous schedule, along with various of its summary statistics that will be used in the generation of a new schedule. Examples of such statistics are for each run: total supplier output, access, tanker capacity, distance travelled, and time taken.

The schedule checking option

There should be some means by which the scheduler can ascertain how well a schedule will meet the requirements of the scheduler. A schedule may be checked for: tanker capacity, amount of milk collected from individual suppliers, amount of milk remaining in the vats of suppliers, and distance travelled.

The schedule modification option

Having pinpointed where a schedule is deficient, the scheduler must then devise modifications to it that produce a satisfactory new schedule. Thus, in answering the second question, the system must provide for modifying existing schedule by such means as:

Adding a new supplier to a run,
Deleting a supplier from a run,
Transferring a supplier from one run to another,
Interchanging suppliers between different runs, and
Creating a new run.

The selection of these options should be guided by the provision of relevant statistics associated with them, such as: tanker capacity utilization, and run duration. Naturally there must be a mechanism by which the new schedule can be recorded, usually by overwriting the schedule that is being modified.

The schedule creation option

It is also desirable that the DSS can be used when, for instance, questions concerning changes to the size of the tanker fleet base, the acquisition of the new suppliers or the location of a new factory are to be addressed. In these cases it is not simply a question of merely modifying an existing schedule, but rather one of constructing a complete schedule from scratch. This is appropriate in a start up situation, or when there are significant changes in the conditions or data that trigger a rationalisation of resources. This task can be accomplished by carrying out systematically the clustering of suppliers into subregions. Despite what has been stated earlier about the undesirability of automatic run creation by mathematical programming methods, it is useful for a DSS to contain such an option.

We now introduce a new DSS that has been devised for the New Zealand dairy industry and which is based on the above considerations.

5. FleetManager

FleetManager is a decision support system that has been developed at the University of Waikato for use by New Zealand milk tanker schedulers. It has been written in Turbo Pascal version 6.0, and is designed to be run on an IBM compatible personal computer with a high resolution colour monitor. The basis of the system is a digitized map of the area of operations of the dairy company that shows all relevant locations and roads. It has pull–down menus, windows which display factory, supplier, and milk tanker information, and is user–friendly and mouse–driven. It can be used to create new, or to improve existing, milk tanker routes.

5.1. FleetManager as a Daily Scheduling Tool

The option of user-created runs being generated manually is based on a computerized map of the area containing the features described earlier. Runs may be created by the user, who clicks with a mouse on the locations of suppliers in the sequence they are to be visited. Various statistics are automatically generated, such as the percentage of vehicle capacity utilized, distance and time travelled. These are available via windows. Alterations are easily handled. **FleetManager** can accommodate the following:

The possibility of more than one milk tanker base,

Suppliers that are picked up on less than daily basis and the capability of judging how often a supplier should be visited,

Multiple shifts,

Supplier vat size limitation,

Tanker capacity,

The ability to identify suppliers of various output sizes not yet visited,

The ability to identify tankers with spare capacity during run construction,

The ability to modify and fine tune existing schedules including the transfer of suppliers between runs and shifts, a summary of all runs of a schedule for a given shift,

The allocation of the total output of each shift among various destinations,

Warnings concerning illogical outcomes, such as unvisited suppliers and overloaded tankers.

FleetManager also has an interface that communicates with the company mainframe computer to import and export the information necessary to support it, including the updating of supplier output. There is also a suppliers' output forecasting option.

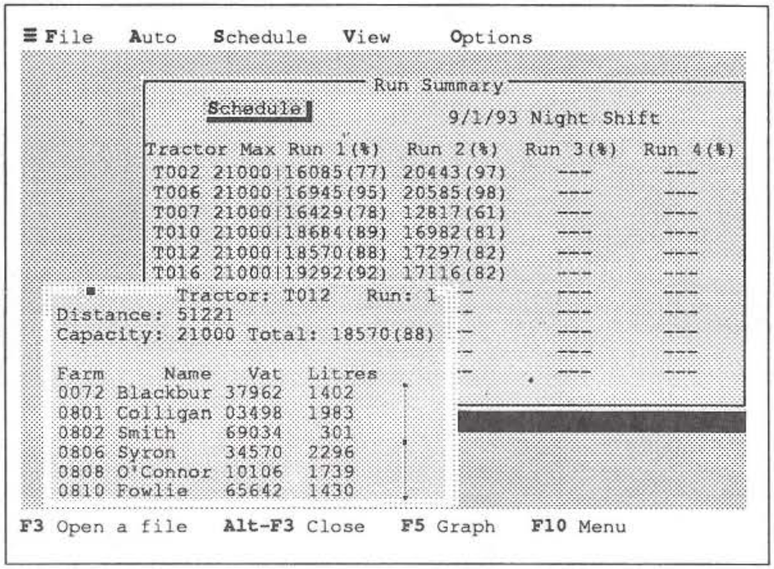

Figure 1. Screen of an Interactive Session

5.2. The Window Based User Interface

FleetManager provides multiple, resizable, overlapping windows to assist the scheduler in his/her tasks. The user interacts through a mouse with pull–down menus, dialogue boxes, list views, buttons, and hot keys. The screen of an interactive session with FleetManager is shown in Figure 1.

The menu bar at the top displays available menus. The status line at the bottom displays keystrokes or shortcuts to carry out pertinent tasks. Here, the user is shown a summary of the runs for a particular shift. The user has double–clicked on a particular run in the list, and the particulars of the run are being displayed in the smaller window.

Figure 2 presents the menu chart of the pull—down menus. The **File** menu provides file manipulations, exit, and shell (to DOS) functions. The Auto menu provides for the use of **FleetManager** to generate automatically all schedules and to look at the effects of changes in milk yield. The **Schedule** menu provides most of the schedule manipulation options. There are provisions for looking at a summary of the schedule, to edit a run (including swapping of suppliers between runs) and to check the performance of a schedule.

The **View** menu provides a list of suppliers that meet certain criteria (such as exceeding 80% of their vat capacities) chosen by the user in a dialogue box. The **Options** menu permits customising the display. With the use of a mouse, by pointing and clicking or double–clicking on the lists and views displayed by

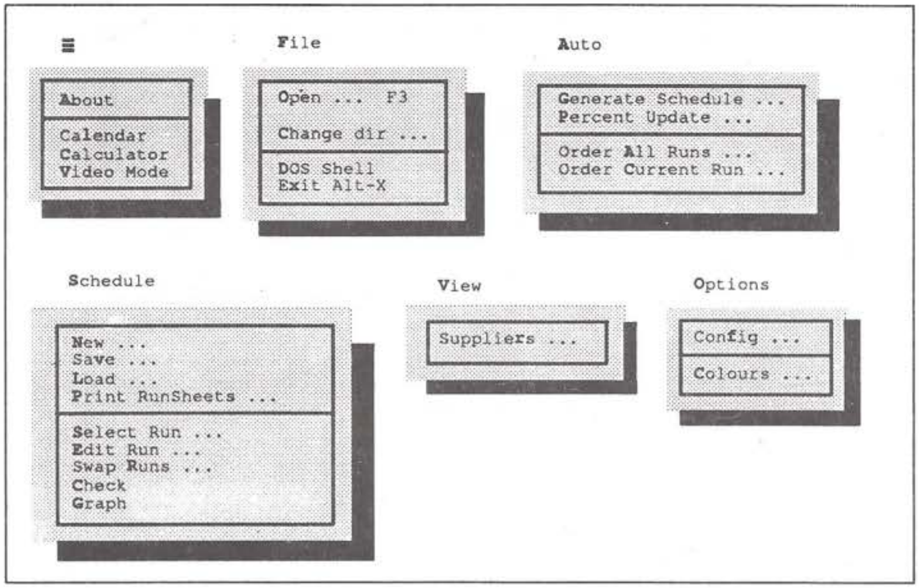

Figure 2. Menu Chart

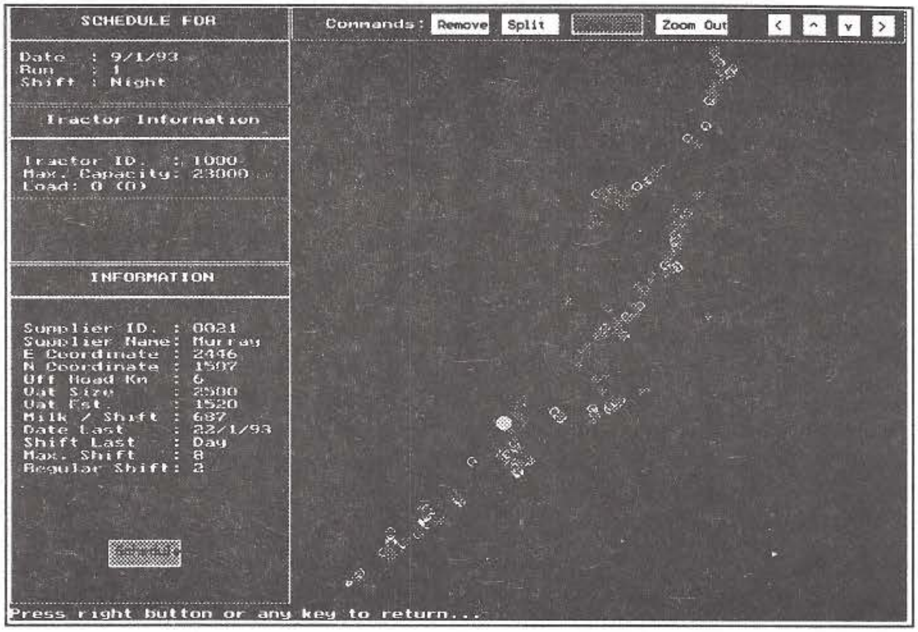

Figure 3. Graphical User Interface

the various menu commands, all the functions mentioned in Section 4.1 can be achieved easily.

5.3. Graphical User Interface

A user can interact with **FleetManager** through a graphical interface as well. The graphical interface and window based interface can be accessed from each other at the touch of a button. The graphical interface displays a road map of the area and the location of the suppliers. This map is based on digitised coordinates of the supplier locations and the roads supplied by the Department of Survey and Land Information of the New Zealand Government. The graphical interface was created using the graphics subroutines conveniently available in Turbo Pascal. The interface has provisions for zooming and panning. The user successively clicks on the supplier locations to build up a run. Information on current run and current supplier is continually updated and displayed. Figure 3 illustrates the graphical interface.

5.4. Automatic generation of schedules

FleetManager contains the option of automatically creating runs. If new schedules are needed, such as in a start–up situation, **FleetManager** provides suggested schedules, which can be modified by the user. Automatic generation of schedules is also useful as a planning tool: to examine the effects of changes in the milk yield, factory demands, shifts, and tanker capacities. The automatic creation of routes involves the following steps: the assignment of suppliers to factories, the assignment of tankers to factories, the creation of an initial run for each tanker and the assignment of second and subsequent run to each tanker from a factory.

The milk tanker routing problem differs from the usual VSP in the recognition of the demands of the factories (constraints (3.7) in Section 3). **FleetManager** handles this by using the transportation algorithm. The program sums up the available milk supply and distributes the milk to the factories on the basis of the demand. Any excess milk is distributed to the factories by user input. This milk distribution also permits allocation of the milk tankers to the milk depots. Using the factories as demand points and the dairy farms as supply points, the dairy farms are allocated to the factories by the transportation algorithm. The distance from a factory to the suppliers is treated as the cost.

The original problem (3.1) is thus converted to a single depot VSP, and, as mentioned earlier, there are many efficient algorithms available to solve this problem. The sweep algorithm of Gillet and Miller (1974) is employed in **FleetManager** to automatically generate the runs. The locations to be visited are then ordered by the Farthest Insertion procedure [Syslo et al., 1983, p 363].

Given what was stated earlier about the utility of such models, it must be stressed that they do not generate schedules satisfying all the various constraints while recognizing the complicating factors that were mentioned earlier. Rather, the use of this model is to provide an initial schedule that is optimum in some sense. This schedule may then be modified by the scheduler in the light of those complicating factors. Automatic schedule generation may also be used as a strategic planning tool in which various options can be tested and the costs of those can be compared.

6. Implementation

An earlier version of the DSS approach described herein has been used by a number of New Zealand dairy companies successfully. Recently a similar approach was adopted in an Irish dairy company. As a result of using the DSS, milk collection costs have been reduced at that company by US $ 1,300,000 annually; implementation of this project is detailed in Foulds et al., (1993). The DSS described in this paper is currently being introduced into a New Zealand dairy company.

7. Summary and conclusions

We have introduced the milk tanker routing scenario that commonly occurs within the New Zealand dairy industry. We have discussed a decision support system approach to this task. We have also described an actual DSS that has been developed for the New Zealand milk tanker scheduling scenario. This DSS approach appears to be the practical, and acceptable to the transport managers for actual daily schedule generation. The traditional algorithmic approach was incorporated for automatic generation of schedules, and for use as a strategic planning tool.

References

L.D. Bodin, B.L. Golden, A. Assad, and M. Ball (1983) Routing and Scheduling of Vehicles and Crews: The State of the Art, *Computers and Operations Research*, **10**, 163–211.

L.D. Bodin (1990) Twenty Years of Routing and Scheduling, *Operations Research*, **37**, 571–579.

P. Carraresi and G. Gallo (1984) Network Models for Vehicle and Crew Scheduling, *European Journal of Operational Research*, **16**, 139–151.

N. Christofides, A. Mingozzi, and P. Toth (1981) Exact Algorithms for the Vehicle Routing Problem, Based on Spanning Tree and Shortest Path Relaxations, *Mathematical Programming*, **20**, 255–282.

G. Clarke, and J.W. Wright (1964) Scheduling of Vehicles from a Central Depot to a Number of Delivery Points, *Operations Research*, **12**, 568–581.

M.W. Davis, (1988) Applied Decision Support, Prentice Hall, New Jersey.

G.B. Dantzig and J.H. Ramser (1959) The Truck Dispatching Problem, *Management Science*, **10**, 80–91.

M. Dell'amico (1989) A New Assignment Algorithm for the Vehicle Scheduling Problem, *Ricerca Operativa*, **19**, 3–21 (In Italian).

M. Desrochers, J.K. Lenstra and M.W.P. Savelsbergh (1990) A Classification Scheme for Vehicle Routing and Scheduling Problems, *European Journal of Operational Research*, **46**, 322–332.

M. Desrochers, J. Desrosiers, and M. Soloman (1991) Using Column generation to Solve the Vehicle Routing Problem with Time Windows, *Operational Research '90*, (UK) 411–419.

M.C. Er (1988) Decision Support Systems: A Summary, Problems, and Future Trends, *Decision Support Systems*, **4**, 355–363.

J.A. Ferland and L. Fortin (1989) Vehicle Scheduling with Sliding Time Windows, *European Journal of Operational Research*, **38**, 141–268.

M.L Fisher, and R. Jaikumar (1981) A Generalized Assignment Heuristic for Vehicle Routing, *Networks*, **11**, 109–124.

L.R. Foulds, P.M. Ghare and W.C. Turner (1974) Transportation Routing Problem: A Survey, *AIIE Transactions*, **6**, 288–301.

L.R. Foulds, E.G. Read, and D.F. Robinson (1977a) A Manual Procedure for the School Bus Routing Problem, *Australian Road Research*, **7**, 21–24.

L.R. Foulds, J.P. Tham, and L. O'Brien (1977b) Computer Based Milk Tanker Scheduling, *Journal of Dairy Science and Technology*, **12**, 141–145.

L. Foulds, D. Clark, M. Butler, and P. Doyle (1993) A Decision Support System Reduces Milk Collection Costs In An Irish Dairy, *New Zealand Journal of Business*, **14**, 74–80.

B. Gillet and L. Miller (1974) A Heuristic Algorithm for the Vehicle Dispatch Problem, *Operations Research*, **22**, 340–349.

B. Golden, T. Magnanti, and K. Nguyen (1977) Implementing Vehicle Routing Algorithms, *Networks*, **7**, 113–148.

B.L. Golden and A. Assad (eds), (1986) Special Issue on Time Windows, *American Journal of Mathematical Management Analysis*, **6**, 251–339.

B.L. Golden and A. Assad (1988) Vehicle Routing: Methods and Studies, North Holland, Amsterdam.

P.G.W. Keen (1987) Decision Support Systems: The Next Decade, *Decision Support Systems*, **3**, 253–265.

E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys (eds), (1985) The Travelling Salesman Problem, Wiley–Interscience, New York.

M. Soloman (1987) Algorithms for the Vehicle Routing and Scheduling Problem with Time Window Constraints, *Operations Research*, **35**, 254–265.

F. Soumis, M. Sauve and L. Le Beau (1991) The Simultaneous Origin-destination Assignment and Vehicle Routing Problem, *Transportation Science*, **25**, 188–200.

R.H. Sprague, Jr. (1987) DSS in Context, *Decision Support Systems*, **3**, 197–202.

M.M. Syslo, N. Deo, and J.S. Kowalik (1983) Discrete Optimization Algorithms, Englewood Cliffs, New Jersey: Prentice Hall Inc.

C.D.T. Watson–Gandy and L.R. Foulds (1981) The Vehicle Scheduling Problem: A Survey, *New Zealand Operational Research*, **9**, 73–92.