CLOSE ENCOUNTERS WITH VEHICLE SCHEDULING

Bruce Benseman
Industrial Research Limited
PO Box 31-310, Lower Hutt

INTRODUCTION

This paper is about real vehicle scheduling experiences in the dairy, oil, and logging industries. Vehicle scheduling concerns servicing a group of customers, with known demands or supplies, using a set of vehicles from one or more depots. Vehicle scheduling problems differ from industry to industry, in their assumptions about customers, vehicles and objectives.

- Some customers may be restricted to particular vehicles or particular depots; some customers may need both pickups and deliveries; some customers may get priority over other customers; some customers may allow load splitting between successive deliveries; some customers need different servicing frequencies from others (this is period routing); some customers have stocks that need replenishing. We consider an example of this inventory routing problem later.

- Vehicles may have different sizes, speeds, costs, and loading rates. Vehicles (and customers) may have different availability periods. Some vehicles may have compartments that are either interchangeable or dedicated.

- Most vehicle scheduling approaches minimise the time or the distance or the cost of travel by the transport fleet. Some include penalty functions for extra vehicles or driver overtime, or for not meeting customer service levels.

We use computers because they are fast and because the world is complex and keeps changing. Computer models help us make better decisions because they consider more information than we can. We get new customers, new loads, new vehicles. We need to respond to the unexpected things in life, like plant breakdown, supplier failure, and new market opportunities. Computerised vehicle scheduling should be interactive to involve the user's intelligence. If we can successfully marry the computer speed with our human intelligence we will get better schedules in a shorter time.

Vehicle scheduling models started appearing in the 1970s. Today the food industry uses computer models to distribute products from warehouses to shopping outlets. The oil industry has the inventory routing problem of keeping their stocks everywhere sufficient. The New York City Sanitation Department routes street cleaners and garbage trucks. The Ameri-Gas Corporation uses vehicle scheduling to distribute fuel oil. US Coastguards route ships to check their marine buoys. And the Florida and Toronto transport authorities have experimented with dial-a-ride bus services. Other published models include school busing, packet switching in computer networks, and delivering newspapers, soft drinks, and mail.

ROUTING DAIRY TANKERS

In the 1980s, DSIR was invited by a large New Zealand dairy company to help them develop a vehicle routing programme on their VAX computers. Today this company uses about 200 vehicles to collect whole milk from over 6000 suppliers. Our Computer Aided Tanker Scheduling programme (CATS for short) is a visual interactive vehicle scheduling programme that seeks to combine human intelligence and computer speed. CATS uses simple commands like "DATA", "BUILD", "EDIT", and "MAP" to let the user create and modify loads and evaluate alternative routes; and the novice user has an on-line context sensitive "HELP". Overseas research shows novice computer users prefer menu-driven systems but experienced users want command-driven systems because they are much faster.

- The DATA command reads and checks the ROADS, FLEET, and SUPPLY data. The ROADS data changes infrequently and comprises suppliers and locations, distances between farms and depots and neighbours, farm access restrictions, and map coordinates for farms, roads and towns. The FLEET data comprises the number, capacity, and class of each vehicle type, various road speeds, hook-up times, pump rates, driver availability, and load safety factors. Plus vehicle-supplier pairs to seed the first N loads, and groups of suppliers that must be collected on the same load. The SUPPLY data says what wholemilk volumes are expected; this changes significantly from week to week; and is forecast using historical trends and user judgment.

- The BUILD command starts new routes by sending the largest tanker available to the farthest uncollected supplier. Then BUILD uses a modified Clark and Wright algorithm, looking for the best uncollected neighbour that can be fitted onboard the tanker, without any access problems. BUILD also handles groups of farms that must be collected on the same load. Running on a 386 PC, CATS generates up to 50 routes per minute. BUILD also checks each route for supplier crossovers, and will automatically stop for the user to vet any route created with a low vehicle utilisation.

- The MAP command gives a global picture of the wholemilk collection to date. See Figure 1. It displays all the assembled routes, highlights any that are below standard, and paints waiting suppliers with an "abc" code indicating supply volume. MAP also lets the user reduce and magnify detail, or pan to neighbouring areas.

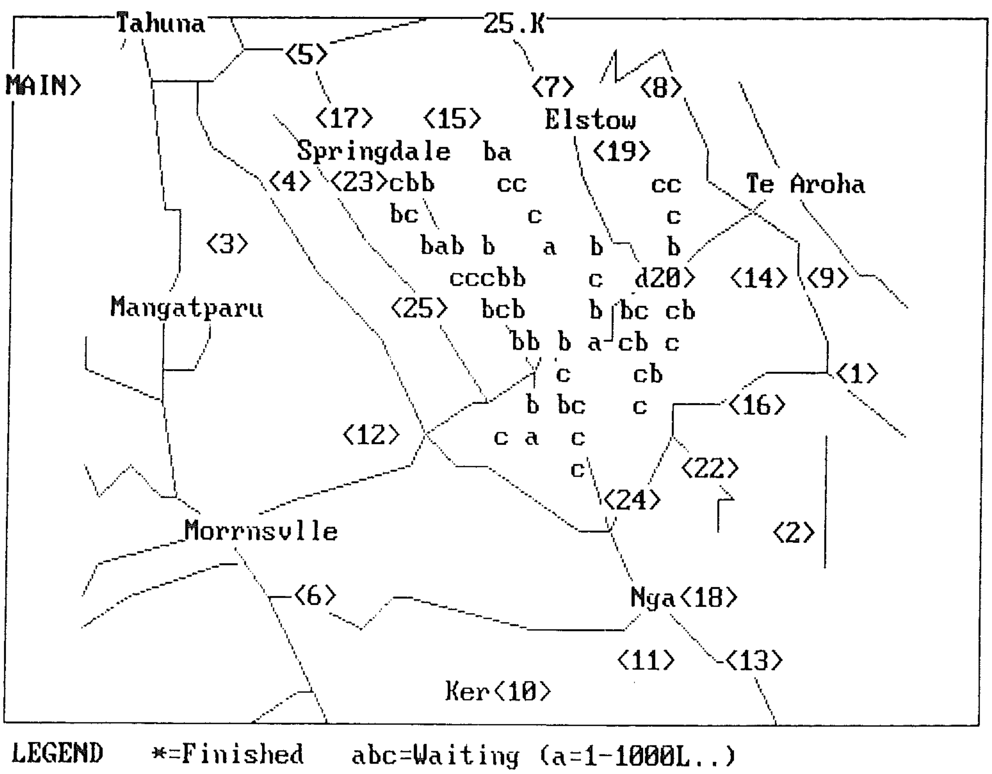

<RTE>	UTIL	SIZE
1	100	23000
2	100	23000
3	97	23000
4	100	23000
5	98	23000
6	100	23000
7	99	23000
8	99	23000
9	98	23000
10	99	23000
11	100	23000
12	98	23000
13	98	23000
14	99	23000
15	100	23000
16	98	23000
17	98	23000
18	98	23000
19	98	23000
20	98	23000
21	97	23000
22	100	23000
23	100	23000

Figure 1. Map of Current Routes and Uncollected Suppliers

- The LIST command provides supplier route and vehicle summary reports with options for more details in each category.
- The EDIT command lets the user check and modify any generated route. Generally CATS builds good routes, but the Clark and Wright algorithm is greedy and shortsighted, sometimes leaving small groups of farms stranded. EDIT encourages the user to check low scoring routes for overlooked opportunities. EDIT displays supplier volumes, vehicle statistics, and a map of suppliers on or near the specified route. EDIT also lets us add, drop and move suppliers, reverse the pick-up order, join this route to other routes, and change the current vehicle.
- The RUN command assembles truck loads into runs taking one tanker type at a time, once we have collected all the suppliers and are happy with the routes. CATS shares out the loads done by each vehicle type, to equalise the time spent by each driver.

Our users responded enthusiastically to CATS. They said CATS was the company's best computer system because it was the only one designed to help them. We also compared the CATS schedule to the manually prepared ones very early in the project. We found CATS could reduce the fleet distance by up to 7% and the number of tanker trips by up to 14%. Similar savings have been shown by overseas studies in Northern Ireland and Ontario, Canada. Bruce Golden, one of the USA's leading schedulers, says it is reasonable to expect annual savings of 10% in distribution costs on the implementation of a computerised vehicle routing system. And of course, vehicle scheduling packages have other spin-offs, such as helping transport managers decide what new vehicles to buy.

Unfortunately in the late 1980s, the dairy industry went through an export slump, and the company suspended all computer development while it undertook a strategic systems review. Recently it investigated the leading vehicle routing package from the UK, but the company's conclusions were not encouraging. Potential savings anticipated by the consultants seemed uncertain; no actual schedulers could be released; too many computer routes had to be hand-checked; and the costs of altering the UK software were too expensive. Two NZ dairy companies have tried Australian routing software since then. One company has abandoned their consultants because of insufficient progress. The other company is still hopeful. The rest of the industry is waiting, and using computers to just store historical milk weights, and to update their set of man-made routes. In fact it seems no dairy company in the world is actually using route creating software for collecting wholemilk.

SCHEDULING COASTAL TANKERS

Recently IRL was contracted by Coastal Tankers Ltd (CTL) to help them schedule their vessels. Ship scheduling algorithms consider how to combine ships and ports to form voyages, and what product volumes should go on each voyage. Our aim is to minimise shipping and port costs, reduce product stockouts at the ports, and avoid surplus stock at the refinery. We have found nothing in the Operations Research literature about the simultaneous optimisation of both voyages and cargoes. Other models generally optimise routes, given fixed cargoes and specific origin and destination dates.

Ship scheduling problems are difficult because of their inherent requirement for 0-1 variables on the possible vessel trips. When we must decide the departure and arrival dates, the problem size can quickly expand and make the problem unsolvable. Most approaches tend to improve an initial trial schedule and then progressively sub-optimise the routes or the cargoes, but not both together. Initially we had a similar plan, but we found we could optimise both the voyages and the cargoes together, using some clever preprocessing.

Our COAST software helps CTL plan the next 30-40 voyages for the vessels to minimise total distribution costs and keep all New Zealand ports stocked up with gasolines, diesels, jet kerosines, marine fuels, and bitumens. The model has to consider factors like the daily accrual rates of components at the refinery, the current stock levels and usage rates of oil products at the major ports, the varying capacities of the coastal ships, and the port tank farms, and the ship and port costs. COAST also remembers flows down the refinery pipeline to Auckland, and planned oil imports and exports at all ports. COAST can quickly produce new schedules to respond to changes in supply and demand around the country. The latest model handles user-committed voyages, ship survey outages, feedstock liftings from New Plymouth, and special drops to Tiwai Point and the Devonport Naval Centre.

The 3500 decision variables in COAST represent the yes/no decisions on each possible ship voyage, and what volumes of each grade to drop at each port. The model has 2000 constraints that ensure volumes are loaded properly onboard at the refinery and subsequently unloaded at the ports; that weight restrictions on each ship and draught requirements at each port are obeyed; and that user limits on service frequency and voyage complexity are respected.

When refinery stocks fall below minimum levels or exceed tank tops, COAST invokes various penalties. Over the 90 day period, port stocks get continually depleted and replenished as new ship voyages are created. The user defines minimum stock levels (in terms of day's cover) for the ports, and when grade stocks fall below this level the model penalises its behaviour.

This initial COAST optimisation is based on average voyage lengths. This solution gets audited so loading and delivery days can reflect the actual time taken to travel between the ports chosen. The overall variation in days between the initial and the targeted schedule is small. Because some delivery dates will be moved slightly with the adjustments, the adjusted schedule is re-optimised to fine-tune the drops for each port. This scheduling adjustment considers such things as available refinery berths, ship cleaning times, loading and discharging times, steaming and berthage times, and daylight restrictions at some ports.

The model reports on the costs of port visits, stockouts and overflows; the destinations, arrivals, and cargoes for each voyage, and the stock levels and cover at each port. The model uses graphics to summarise stock levels either by grade or by port. See Figure 2. The graphs highlight potential stock problems, and lets the user respond appropriately. For example, if planned imports coincide with high stock levels nation-wide can be delayed or cancelled, and the model re-run to lessen pressure on storage.

Results to date have been very encouraging. The ability to optimise cargoes and routes simultaneously means the user avoids any local optimums, and gets the best overall solution. To produce the 90-day schedule takes 2-3 hours on a 486 PC, compared to a week doing the schedule by hand. COAST reduces stockout and port costs by about 90% and 15% respectively, because the model considers so many more factors. Test runs indicate COAST can respond well to changing import levels and tankage availabilities.

Manual scheduling is a difficult and time-consuming problem. Too many trade-offs between different products and ports and arrival dates confuse the poor human scheduler. Optimisation models, however, can cope with this complex set of options, achieving the best delivery pattern and volumes over the entire planning horizon. As well as considering port shortages, COAST also considers shipping and port costs, and minimises the overall costs of the schedule. The speed of computer models means the scheduler can respond faster and better to new demands and schedule disruptions. Models can also provide reliable information on distribution costs and enable users to study the effects of proposed changes in port facilities and future shipping capacity.

DESPATCHING LOG TRUCKS

Log truck schedulers are concerned with the daily hauling of timber from different forest stands to destinations such as sawmills, sorting yards, and ports. Different timber

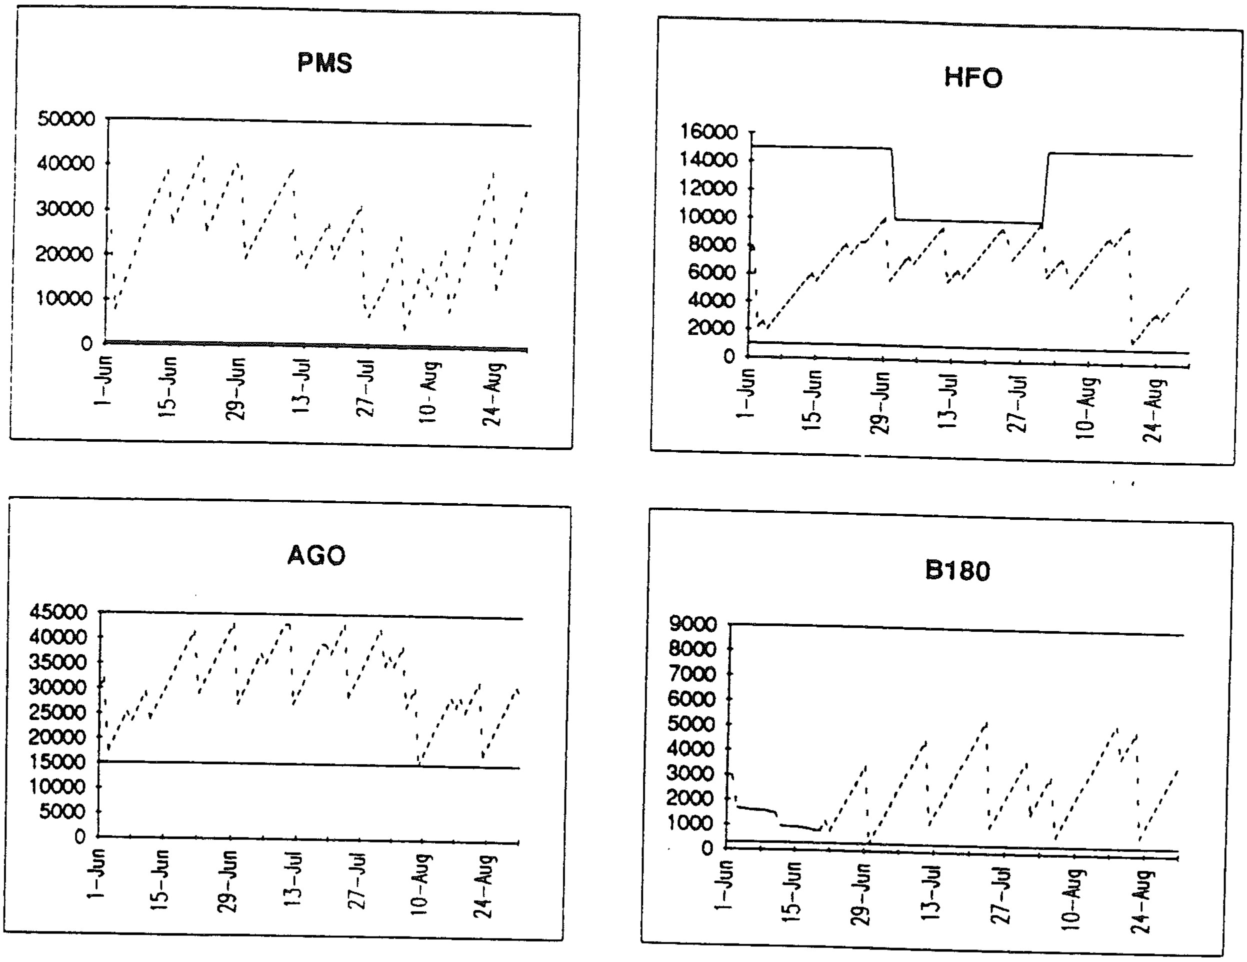

Figure 2. Product Stock Levels at Marsden Point Refinery

products are defined basically by length and diameter. Trucks are typically subcontracted and differ in their hauling and forest access capacities. The basic company objective is to satisfy the demands for different products at each destination while optimising the use of equipment within the technical and labour constraints.

The ideal scheduling software needs to know the volume of wood available at each forest origin and the demands at the destinations, the number of available trucks, their capacity and travel times, and the teams available for loading and unloading, and the times required for this work. The model must decide which origin satisfies each demand, what their truck fleet requirements are, and what the work schedule for each truck is. The model should prevent delays, look for the shortest possible routes, regularly supply the destinations throughout the day, and adhere to the meal breaks and rest periods for drivers and workers. It must programme each workday, assigning jobs, schedules, and specific origin and destination locations to each truck. The system should also programme the workday so that the truck drivers can complete their shifts near their homes. At the end of the day this information will be given to the drivers and despatchers so that they will know their activities for the following day.

One of the few systems found in the Management Science literature is the ASICAM system developed at the University of Chile by Andres Weintraub and Rafael Epstein. ASICAM has been running since 1990, and is based on a simulation model driven by heuristic rules to assign trips to trucks. The system runs daily and takes data according to the supply of timber products at origins, demand for destinations, and truck fleet characteristics, costs and times. The model seeks to recreate an entire day's operation and for each departure assigns trucks in order to satisfy all the despatching requirements and to minimise the total cost of wood transport. The system outputs truck requirements, a schedule for each, and basic statistics to calculate the performance. The system has been allegedly implemented in 7 Chilean forestry companies, and trialled in 4 South African ones. The authors report substantial improvements in efficiency, with some companies reducing annual transport costs or their truck fleet by up to 15%. Weintraub also claims that trucks have less empty runs, and drivers can programme their repairs on time because they know their exact itinerary.

Scientists from the Logging Industry Research Organisation have actually visited Chile to investigate ASICAM. However LIRO did not get a copy of the software, and has now suspended their interest in this Chilean product. The ASICAM authors warn that their application will not show any benefits unless the entire company is structured to operate in accordance with their software. Success, they say, depends on having good communication systems, and trucks that don't break down. The main disadvantages of ASICAM seem to be:

- ongoing support for the software is difficult from Chile;
- New Zealand companies have many more trucks than their Chilean counterparts.

CONCLUSIONS

What lessons can the logging industry learn from the vehicle scheduling experiences of the dairy and the oil industries?

- Using CATS showed us that it is possible, given an accurate road database, to build good dairy tanker routes very quickly. Computer speed lets despatchers adapt schedules to changing supply and demand conditions, and real-time emergencies like delays and breakdowns. The CATS users could avoid under-loaded tankers and excessive work hours. And they responded enthusiastically, saying CATS was the company's best computer system because it was the only one designed for them. CATS users especially liked its interactive commands to see and modify the routes. Fewer customers and trucks in the logging industry means this route supervision can be minimised.

- Scheduling the coastal oil tankers is much more difficult. We have to plan 10-15 routes per vessel, with up to 6 different products each, to 12 ports all with different demands and capacities and priorities. Because port visits cost thousands of dollars per day, each visit must be effective, and optimal solutions are important. Little has been published on the inventory routing problem. Companies jealously guard any system they have already developed. However our COAST software shows it is possible to build friendly packages that both really save costs and improve service. Despatching logging trucks seems easier, because time horizons are shorter, load configurations are simpler, and poor decisions are not as costly.

- No good universal vehicle routing system exists. Neither may off-the-shelf packages for the logging industry exist. Consultants' attempts to adapt packages from one industry to another are often unsuccessful. Every problem is different and requires its own tailor-made system.

- Logging despatchers should not expect too much. Consultant representatives usually over-sell their products, and their backroom developers cannot deliver on time or price. Talk to the developer direct. I prefer to start with small, simple models which develop gracefully.

- No vehicle scheduling system can be fully automatic. We always need a person to handle the subjective and unanticipated things. The logging industry needs a quick, visual, interactive, effective scheduling tool. The user must be able to see where the trucks are and when they will arrive at their destinations. The despatchers must be able to update the computer schedules, and create their own from scratch if necessary.

- Quality products result from good designers working with concerned customers. I am confident that our IRL management and computer scientists, in a joint venture with some logging company, can develop an effective visual interactive system that makes log-truck scheduling both fun and profitable.

REFERENCES

"Computerized bulk milk collection systems", International Dairy Federation Bulletin, Document 176, 1984.

Bell, P C. "Visual interactive modelling in operational research: Successes and opportunities", Journal of the Operational Research Society, Vol 36, No. 11, 1985.

Benseman, B R. "Production Planning in the New Zealand Dairy Industry", Journal of the Operational Research Society, Vol 37, No. 8, 1986, pp 747-754.

Benseman, B R. "Visual Interactive Vehicle Scheduling", Asian-Pacific Operations Research APORS'88 Conference Proceedings, Seoul. North Holland 1990, pp 425-430.

Brown, Gerald G; Graves, Glenn W; and Ronen, David. "Scheduling ocean transportation", Management Science, 33(3), 335-346, March 1987.

Clarke, G; and Wright, J W. "Scheduling of vehicles from a central depot to a number of delivery points", Journal of the Operational Research Society, Vol 12, 1964, pp 568-581.

Millar, Harvey H; and Gunn, Eldon A. "Dispatching a fishing trawler fleet in the Canadian Atlantic groundfish industry", European Journal of Operational Research, 55, 148-164, 1991.

Miller, David M. "An interactive, computer-aided ship scheduling system", European Journal of Operational Research, 32(3), 363-379, December 1987.

Weintraub, A; Morales, R; Seron, J; Epstein, R; y Traverso, R. "Managing Operations in the Line Forest Industries", Proceedings of the 1991 Symposium on Systems Analysis in Forest Resources, March 3-6, 1991, USDA, Forest Service General Technical, Report SE-74, pp 31-34.

ABOUT THE AUTHOR

Bruce Benseman works as a Operations Research Consultant for Industrial Research Limited (formerly part of DSIR). He has two Honours degrees in Mathematics and Economics, and was President of the Operational Research Society of New Zealand for a number of years. Bruce has specialised in location, production and distribution studies for over 20 years; and is currently supporting large planning models in the meat, dairy, oil and power industries, both in New Zealand and Australia. You can contact Bruce on 04-569-0000, or write to:

Bruce Benseman
Industrial Research Limited
P O Box 31-310
Lower Hutt