Session 3
Paper (a)

ECONOMIC EVALUATION OF LOG TRUCK SCHEDULING

Tim Robinson
LIRO
P.O. Box 147
Rotorua
New Zealand

ABSTRACT

This paper outlines the theory of log truck scheduling and despatching, and shows why some schedules are better than others. The concept of schedule efficiency is introduced, together with strategy for assessing the quality of manual despatching.

INTRODUCTION

The New Zealand forest industry spends over half a million dollars a day on transporting logs by truck. Even a modest reduction in this cost would save the industry millions of dollars a year and improve its competitiveness in the world market. There are two possible strategies for reducing log transport costs:

1. Use better trucks
2. Make better use of existing trucks

In the last few decades, most of the effort has gone into the first option, and we have seen tremendous advances in truck technology and design in an effort to increase payload and versatility and reduce costs. The second option is to plan routes and schedules so that every truck is used to its full potential. This is not as easy as it sounds!

Many transport managers are wondering whether or not to invest in computer technology to help solve their despatching problems. They are asking **"How much money can a computer save?"** Even more important, **"Just how good is my despatching *now*?"**

SCHEDULING AND DESPATCHING

The Log Truck Cycle

In the course of a day, a log truck will carry several loads. The productive cycle is shown in figure 1: the truck arrives in the forest, is loaded with logs, carries them to a weighbridge, gets weighed, carries the logs to the customer, and then makes the unloaded trip to the next pickup.

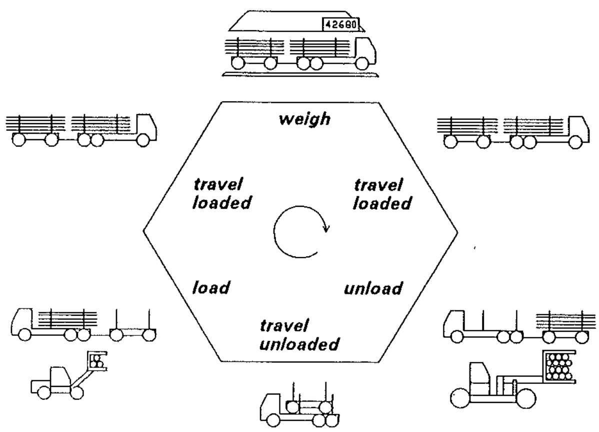

Figure 1: The Log Truck Cycle

Table 1: Sample Schedule for One Truck

Pick Up				Deliver	
Logs		From	At	To	At
8.1m	Kraft	Neptune 43	4:21 am	Technopulp	7:09 am
12m	J grade	Jupiter 29	8:41 am	The Wharf	10:59 am
7.4m	Peeler	Mercury 36	1:44 pm	Cyberboard	5:04 pm

The Schedule

For the purposes of this paper, a *schedule* is a description of each log truck's activities during the day. The essential features are:

1. What logs to carry,
2. Where to pick them up,
3. Where to deliver them, and
4. The time of day for each event.

For example, see table 1.

Despatching

It might be supposed that the schedule is planned well in advance, but in New Zealand, this is rarely true. In practice, the schedule is created on the fly, and can only be published after the event.

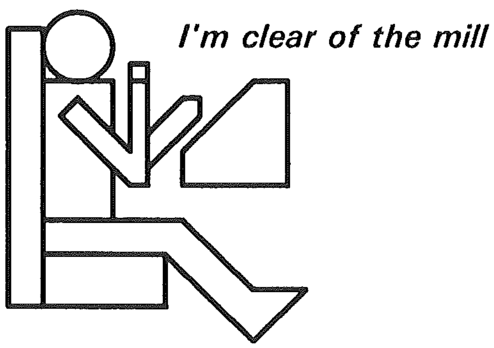

At present, the task of scheduling is performed by human despatchers. The typical New Zealand despatcher operates from a weighbridge or building nearby, and communicates with the truck drivers and loader operators by radio. The day's pickups and deliveries are listed on paper or displayed on white-boards, and the despatcher has to assign these tasks to the trucks as they become available.

This system is preferred to a rigid schedule for several reasons:

1. Log supply and demand can change from hour to hour, as new loads accumulate in the forests and customers change their minds (sad but true!);

2. Trucks and other machines can break down, forcing tasks to be reassigned;

3. Queues at weighbridges and customers make it difficult to predict exactly how long any particular trip will take.

Despatching Aids

The despatcher must make scores of despatch decisions every day; at peak times, he or she may have half a dozen incoming calls a minute. The despatcher must not only record what has happened, but must try to anticipate what is going to happen later in the day. Despatch decisions made in the morning will determine the distribution of trucks and loads available in the afternoon. Very few despatchers have anything more than pencil and paper to help them.

One notable exception is the despatch team at CHH Kinleith, which has an electro-mechanical Truck Control Board to help keep track of truck location [1]. In one form or another, this device has been operating since the sixties.

Many fleet operators in other industries have turned to computers to assist with despatching. Milk tanker fleets have been scheduled by computer for many years. St John's Ambulance (Auckland), the Automobile Association Emergency Breakdown Service (Auckland) and the Police (Wellington) all have various kinds of Computer Aided Despatching. Wellington Combined Taxis has gone a step further: their computer makes all the despatch decisions and operates the radio as well!

The New Zealand forestry industry has been relatively slow to adopt computer technology, particularly in its transport operations. For example, a 1985 LIRA report [2] assumed that transport operators would use computers solely for costing and accounting purposes. There was no hint that computers could be used for scheduling and despatching. The one serious attempt at computer aided log truck despatching in the 80's merely reinforced the industry's distrust of computers.

The 1990's have brought renewed interest in the use of computers as a despatch aid. Success in other countries [3] and other industries [4] has encouraged New Zealand forestry companies to rexamine this technology. Many transport managers are asking themselves what kind of improvement they can expect from a computer system.

Theoretical Work

Vehicle scheduling has long been a subject of intense theoretical scrutiny [5], although this has tended to concentrate on other industries. In forestry, an American team has investigated log distribution in a way that can be used for cost estimation and route planning, but not truck scheduling [6]. Papers concerning log truck scheduling are surprisingly few [7], [8], [9].

Schedule Quality

What is so difficult about despatching anyway? And what makes one schedule better than another?

Roger! Head out to Saturn Forest

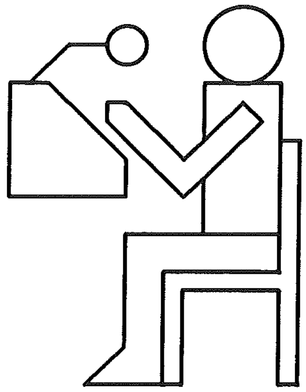

The aim of every transport manager is to ensure that all the required work gets done within the law and at the smallest possible cost. If every truck could be doing useful work at every minute of the day, there would be no problem. In practice, there are a number of ways a fleet can work at less than total efficiency:

Unloaded Trips
Log trucks are paid to haul logs from forests to customers, but every loaded trip requires an *unloaded* trip from the customer back to the forest. In some schedules, a truck will cover half its daily distance in an unloaded (and unprofitable) condition. By contrast, it is often possible to arrange for the truck to visit several different forests and several different customers in the course of a day, and to have the total loaded distance greater than the total unloaded distance. The success of this strategy (*back loading*) is measured by the *loaded running percentage*, which is the total loaded distance divided by total overall distance. Transport managers place great store in loaded running percentages, but they don't tell the whole story - after all, they *increase* whenever the driver takes a detour on the way to the customer.

Slack Time
The length of a driver's shift is limited by law to fourteen hours a day. (There are other restrictions too, but this is the most important one in practice.) If a driver works only 13 hours, the final hour is wasted (*slack time*). As we shall see, the standing costs of a log truck are at least as great as the running costs, and slack time and queueing time are very expensive.

Queues
Forestry has been expanding rapidly for over thirty years, and as a result, many of its facilities are overloaded. Log trucks can form queues when waiting for loading, weighing and unloading operations; by far the worst queues occur at weighbridges of major log yards and at tally sheds of export ports. Queues build up because trucks arrive faster than the facility can deal with them. Some queues can be avoided by careful despatching so that trucks arrive at regular intervals; other queues are structural and the only solution is to upgrade the facilities.

Night Time
A driver may only work for fourteen hours a day, but the *truck* could work for another ten. Even allowing time for routine maintenance, there is clearly room for a second shift if customers can be persuaded to receive deliveries at night. (Of course, this means extra shifts of loader and stacker operators, too.) For example, a clear opportunity exists for night running between central processing yards and export ports.

For many transport operations, night running is impossible. The transport manager will therefore confine his or her attention to reducing queues, slack time and unloaded trips.

SCHEDULING CASE STUDY

As mentioned earlier, before anyone can estimate how much money a computer system will save them, it is necessary to form an estimate of the efficiency of the existing manual despatching operation. Curiously, very little is known about this vital question; many transport managers have expressed an opinion that their despatching could be improved, but no concrete evidence has yet emerged. (Some managers may know more than they are telling, but this is no help to those who don't!)

LIRO is conducting an investigation to determine the typical efficiency of manual despatching of log trucks in New Zealand. The pilot study focuses on the log transport operation of Pan Pacific Forest Industries in Hawkes Bay. This was a good starting point as the fleet is of medium size, and very advanced in the way it already uses computers. They have an automated weighbridge and some of the trucks are equipped with on-board computers. Neil Weber (Log Transport Supervisor) has been exceptionally helpful and has provided large amounts of high quality data.

The study method was to take schedules that Pan Pac had already used, and to compare them with optimal (cheapest) schedules created later. The efficiency of any schedule is then:

$$efficiency = \frac{cost\ of\ optimal\ schedule}{cost\ of\ actual\ schedule} * 100\%$$

This leaves just two problems:

1. Determining the cost of a schedule;
2. Finding the optimal schedule.

Trucking Costs

The costs of running a log truck may be broken down as shown in figure 2. Running costs (fuel, tyres, maintenance and Road User Charges) make up less than half the total. For more detailed truck costing, see [10].

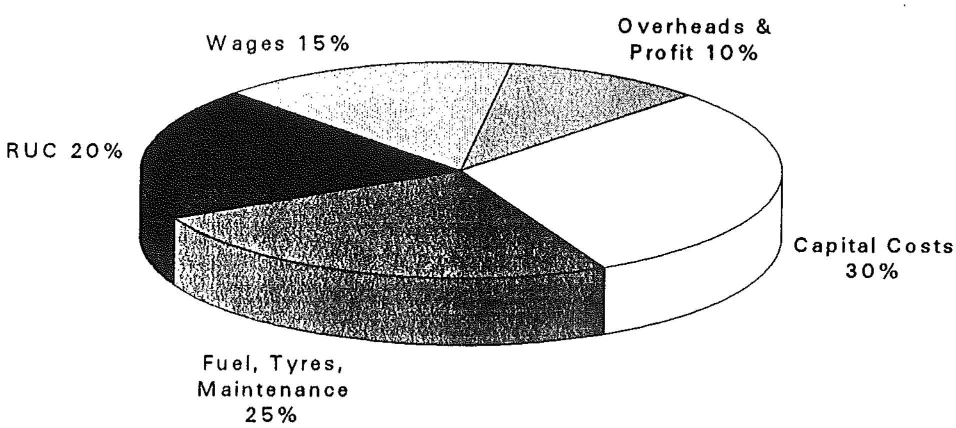

Figure 2: Cost Breakdown for Log Truck Operation (Indicative Values Only)

So it is a simple matter to determine the cost of any schedule: just add up the running costs of each trip, and add a day's standing cost to the result.

An alternative costing scheme based on haulage rates paid by the forestry company to its trucking contractors was rejected because:

1. Companies don't like to see their contract haulage rates published;
2. Haulage contracts vary according to company policy, and rates cannot be compared directly;
3. Just under half of Pan Pac's drivers are direct employees, not contractors;
4. Haulage rates tend to be based on *nominal* amounts of standing time and unloaded running, and this conceals the true cost of these items.

Timing

Although the costing aspects determine which schedule is *best*, it is the timing factors that determine which schedules are *possible*. The regulations that restrict the timing of a schedule can be briefly summarised:

1. A driver may not work more than fourteen hours a day;
2. A driver may not drive for more than eleven hours a day;
3. A driver may not drive for more than five and half hours without a half hour rest break.

The third rule is a subject of some controversy. Many transport managers allow their drivers to claim queueing time and loading time as rest breaks, but Pan Pac insists on a definite break of at least half an hour. Rule two seldom has any impact on log truck drivers, as so much time is taken up with loading, weighing and unloading operations.

From a scheduling point of view, driving time and running time (see figure 3) are the same thing.

Shift Length Calculation

The total shift length for any truck is just the sum of the times required to do each task in the schedule. Shift time can be broken down as follows:

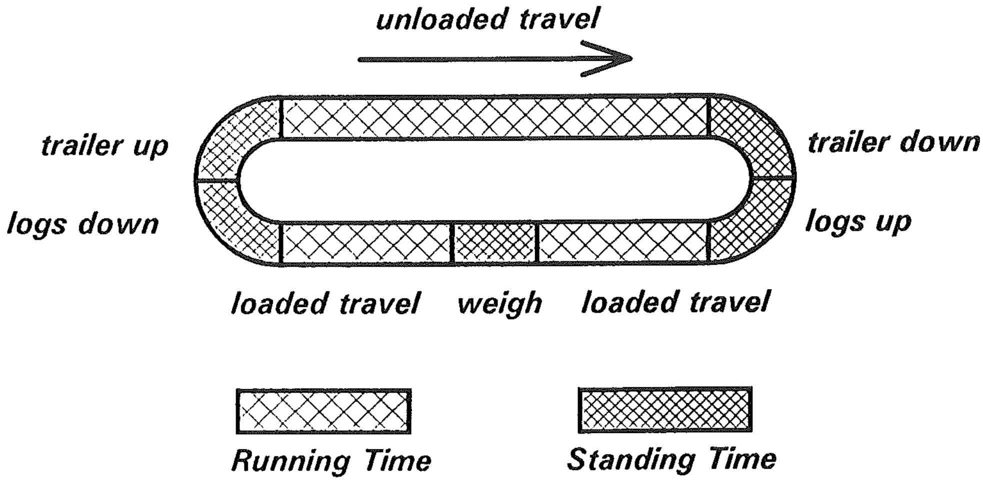

Figure 3: Another way of looking at the log truck cycle - Running Time and Standing Time

Loading Operation...
  Trailer Down...
   - Lift down trailer
  Logs Up...
   - Lift up logs
   - Collect docket
   - Secure load with chains
Loaded Travel
   - To weighbridge
   - From weighbridge
Weighing Operation
Unloading Operation...
  Logs Down...
   - Unsecure load
   - Lift down logs
  Trailer Up...
   - Lift up trailer
Unloaded Travel

The loading and unloading operations are subdivided into parts that depend on truck configuration (lifting the trailer), parts that depend on the log length (lifting the logs), and parts that are much the same for all operations (chains, dockets). The weighing operation is assumed to be the same for all configurations and log sizes.

Unfortunately, timing information for these elements was not available, so a student (Steven Bates) was hired over the summer of 1993-94 to conduct a time study of loading an unloading operations [11]. It was not necessary to conduct a similar study into log truck travel times, as these could be deduced from truck computer records supplied by Neil Weber. LIRO hopes to publish a report later in 1994 which will contain all the relevant timing information used in the study.

With this timing information, it is a simple matter to add up the element times (and the compulsory 30 minute rest break) and come up with total shift length.

OPTIMAL SCHEDULES

Armed with timing and costing information, it is easy enough to assign a cost to each schedule, and to determine its legality. The optimal schedule is just the *cheapest legal* schedule. But how to find it?

Finding optimal schedules is far from easy. For a transport operation with 60 loads a day distributed over 20 trucks (eg. Pan Pac on a slow day), there are over 200 000 possible schedules for each truck and over 10^{100} possible schedules for the whole fleet. Clearly there are too many to find the optimum by guesswork or exhaustive search!

We therefore present:

1. A quick method for estimating the cost of the optimal schedule;
2. A method for finding the optimal schedule for a single truck;
3. A method for scheduling a fleet of trucks;
4. A simulation method for testing the schedule in the presence of queues and random variations in timing;
5. Some computational results.

In the descriptions that follow, a lot of technical detail has been left out, but the main ideas remain. For a more technical description, see [12] and [13].

ESTIMATING THE OPTIMAL COST

To determine the efficiency of a schedule, it is not necessary to determine the optimal schedule, merely the *cost* of the optimal schedule. It turns out there are ways to estimate this cost without too much trouble.

So far, we have regarded the truck cycle as having a large number of components: loading, loaded running, weighing, unloading, unloaded running. For this costing exercise, however, the cycle is reduced to two parts, each of which includes some running time and some standing time.

Loaded Trip
For this exercise, a loaded trip from a forest pickup point to a customer begins with the truck empty, the trailer down, and the loader ready to begin. The loaded trip thus includes lifting the logs on the truck and trailer, securing the load, driving to a weighbridge, weighing, driving on, unchaining the load, and lifting down the logs. The loaded trip stops there, with the truck at the customer but with the trailer still down.

Unloaded Trip
The unloaded trip carries on from there. The trailer is piggy-backed onto the truck, the truck drives to the next pickup point, and the trailer is lowered again.

The loading and unloading operations have each been split into two parts, the trailer handling part belonging to the unloaded trip and the log handling part belonging to the loaded trip.

This division ensures that each loaded and unloaded trip can be given a time and a cost that is completely independent of the preceding or following trips.

This being so, the loaded and unloaded trips of a schedule may be considered separately, and thus become *two entirely independent problems*:

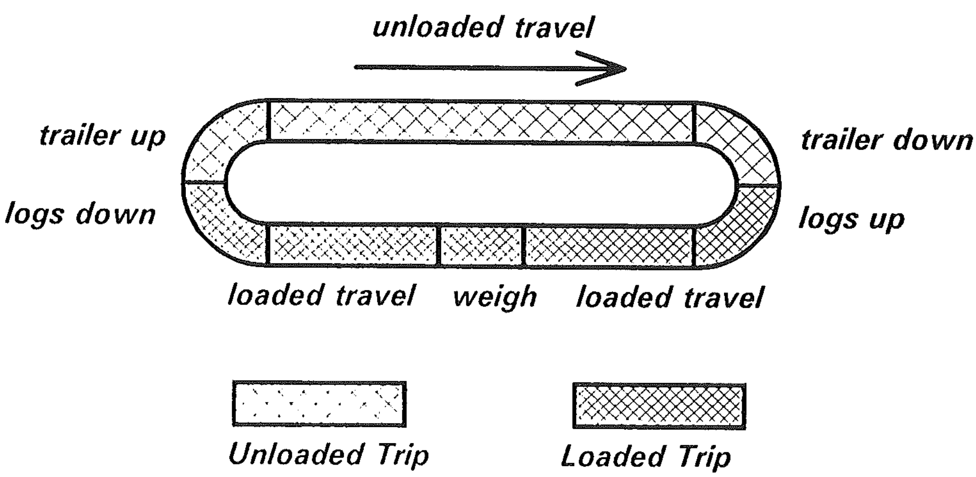

Figure 4: Yet another view: Loaded and Unloaded Trips

1. How can we move the log loads from the forests to the customers as cheaply as possible;

2. How can we get the empty trucks back from the customers to the forests as cheaply as possible?

Since the individual trucks are no longer a part of the formulation, the trip costs must include both running costs and standing costs on a pro rata basis, thus:

trip travel cost = trip running cost
*+ trip duration * minute standing cost*

Here, the *trip duration* is the trip running time together with the standing time associated with the trip (see figure 4), and the *minute standing cost* is the daily standing cost divided by the shift length (fourteen hours minus a compulsory half hour break).

Effects of Problem Separation

By separating the loaded and unloaded problems, all semblance of scheduling has been lost. Ideas such as time of day, trip sequence, and even individual trucks have been cast aside, leaving only route planning and distribution issues. From a mathematical point of view, this is a very good thing, because the problem that remains is very easily solved using textbook techniques (the Stepping Stone Algorithm [14]).

For the loaded trips, often there is no choice at all - each load in the forest is already earmarked for a particular customer. So the cost of all loaded trips is just a matter of adding up the costs for each trip.

For unloaded trips, the situation is only slightly more complicated. Each customer receives a known number of loads, and so parts with an equal number of empty trucks. Likewise, each pickup point despatches a known number of loads, and thus requires an equal number of empty trucks. Using the Stepping Stone Algorithm to match up supplies and demands of empty trucks in the cheapest possible way, the total cost of unloaded trips can be found. Combining this with the cost of loaded trips gives an estimate of the total cost.

A great deal of detail has been discarded in order to reach this estimate, and it would be well to review what has been left out. The estimate assumes:

1. Each truck has its base on a log route;
2. Any truck can carry any load;
3. All trucks have the same costs;
4. The total slack time is zero;
5. The total queueing time is zero;
6. The time required for any operation can be predicted exactly.

The estimate can be improved by including truck bases and a variety of truck configurations, but it still lacks conviction because there is no actual schedule.

FINDING A GOOD SCHEDULE

Finding a good schedule is best tackled in two parts:
1. Finding a good schedule for one truck;
2. Fitting single-truck schedules to cover all the required loads.

Tree Search Procedure

A good method for scheduling a single truck is to use a *tree search* procedure. (This is a *mathematical tree* and has nothing to do with forestry!) Imagine a truck starting out its base in the early morning. Which forest to visit first? Make imaginary duplicates of the truck and send one copy out to each forest. Now wait until one of the imaginary trucks has reached its destination. What load to pick up? Once again, make more copies of the truck and load each copy with one of the loads. Then send the trucks out to their respective customers.

Things start to get a bit hectic as the day progresses! Every time a truck reaches a forest, it makes further copies of itself, each copy picks up a load and heads for its particular customer. Every time a truck reaches a customer, it drops its load and makes further copies of itself, and sends them out to the various forests. One copy of the truck knocks off work and is sent home. When the required shift length is up, all the trucks stop.

Now examine the base. It will be full of imaginary trucks, every one of which will have done a legal shift, and every legal shift will be represented. All the costs are known, so it is trivial to pick the best schedule.

The tree search is a systematic way to enumerate *every legal shift* for a single truck - and there may be millions of them! Really, we only want *one* schedule. The *best* one.

Pruning the Tree Search

Fortunately, there are a number of tricks that can be used to keep down the number of phantom trucks in the system.

Suppose each log load has a price attached. Then each single-truck schedule can be assigned a profit, which is the total value of loads delivered, minus the total running cost of the trips made. (The standing cost can be ignored for the moment, since it is the same for each.) As each phantom truck progresses through the day, it can keep a track of its *partial profit*, which is the total value of loads delivered so far, minus the running cost of trips made so far.

Now imagine a phantom transport manager at each forest and each customer. As each phantom truck comes past, he records the partial profit of the truck on his clipboard. If a truck arrives which has less partial profit than the previous truck at the same site, the phantom transport manager says to the driver:

"I had a truck here just now that has already made more profit than you. Even if you went on from here to do the best you could with the time you have left, that truck could do exactly the same work and get a bigger final profit. And he'd be home before you, too. So you might as well give up now."

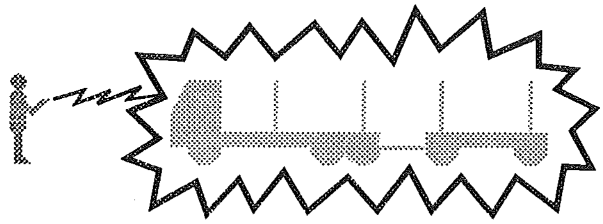

Figure 5: Removing a Phantom Truck

This removes one imaginary truck from the system, and all future duplicates made from it. Experiments show that removing slow, unprofitable phantoms from the system is sufficient to keep the number of trucks under control.

There are a couple of other tricks that can also help, like checking that the truck has time to make his current delivery and still get home before the end of the shift. If he hasn't, he might as well stop at once.

The pruned tree search can be modified to allow for customers opening and closing, and to ensure that no load gets delivered twice by the same truck. Both these modifications are relatively easy to do.

Computer experiments show that removing phantom trucks in this way can reduce their number from half a million to under ten thousand.

Assembling a Fleet Schedule

Using the technique above, or one like it, it is possible to create optimal schedules for single trucks. Curiously, it is a great deal more difficult to bring these schedules together in a way that covers all required deliveries.

One method that is easy to use, is the so called "Greedy Method". This produces solutions in a reasonable time, but the final fleet schedule is not necessarily the best that could be achieved.

The Greedy Method

Start with the whole truck fleet and the complete set of required log loads. Use the pruned tree search to create the best possible single-truck schedule. Cross this truck off the list, and also cross off the log loads carried by that truck. Now create the best possible single-truck schedule from the fleet and loads that remain. Cross *that* truck and *those* loads off the list. And so on. Eventually you must run out of loads or trucks. If you ran out of loads, then you have created a fleet schedule that carries all the required loads. If you ran out of trucks, you'll have to try something else or get some more trucks!

Obviously this is a more satisfactory method of estimating the optimal transport cost, because it doesn't just come up with a cost, but it also proves that the logs can be carried at the costs. Furthermore, this method could be used as a despatching aid, not just an auditing tool, because it provides a schedule that the despatcher can try to put into practice.

At the end of the cost estimation section, there was a list of six false assumptions that had to be made to get the cost estimate. Using the methods shown here, these six are reduced to two:

1. The total queueing time is zero;
2. The time required for any operation can be predicted exactly.

SIMULATION

The only way to deal with queueing and random factors is *simulation.*

In the previous section, the reader was asked to imagine phantom trucks travelling from place to place, picking up and delivering logs. Later, it was revealed that a computer can work with phantoms too. To simulate queues, we need imaginary loaders, weighbridges, and unloading machinery as well.

When a truck arrives at a forest, it will be loaded with logs, *provided the loader is not already busy*. If the loader *is* busy, the truck will have to wait in line. Queues can also occur at weighbridges and customers.

A simulation can be used in two ways:

1. A scheduling simulation;
2. A despatch simulation.

Scheduling Simulation

In a scheduling simulation, it is assumed that the schedule has already been worked out in advance. In a way, all the simulator is doing is *testing* the schedule to make sure it will work in practice, by including the queues and random factors. A five minute wait or too many red lights can turn a tight schedule into an illegal one. Simulation is a useful tool to help determine how much slack time must be included in the schedule to deal with these annoying realities.

Despatching Simulation

Deliberately including slack time in a schedule seems like a waste, but this is the best that can be done if the schedule is to be created in advance. On the other hand, if the schedule could be modified as the day goes on, it would be possible to cope with queues and random factors as they occur. After all, this is the reason why the New Zealand forest industry still uses despatchers.

A despatch simulation can be used in three ways:

1. To compare different despatch strategies with each other;

2. To train dispatchers in the same way that flight simulators are used to train pilots;

3. As an *electronic score-board*, to remind the despatcher of what trucks and loads are still available, and to warn the despatcher of impending events.

In fact, the third of these options is a pretty good description of the Despatching Aids used by the AA Emergency Breakdown Service, St John's Ambulance, the police, and most taxi companies. The truck control board at Kinleith can also be viewed as a despatch simulator, even though there is no electronics involved.

COMPUTATIONAL EXPERIENCE

LIRO has written experimental computer programmes implementing these methods in Turbo Pascal and running on a IBM compatible 486 PC running at 33 MHz. These programmes were applied to a sample problem with 81 log loads, four truck configurations, 16 pickup points and 9 delivery points.

The Stepping Stone algorithm produced an estimated cost in 12 seconds, while the Greedy Method produced a full fleet schedule in 45 minutes. The initial cost estimate matched the cost of the Greedy schedule to within 4%.

The simulation took less than a minute to run.

CONCLUSION

This paper has presented a strategy for evaluating and potentially improving the quality of despatching and scheduling of log trucks. The main forms of unproductive time have be laid out and discussed. The concept of schedule efficiency has been introduced and defined, together with a collection of tools for its estimation:

1. A computational efficient estimate of optimal schedule cost;

2. Techniques for creating schedules for single trucks and complete fleets, yielding an improved estimate of optimal cost that includes slack time;

3. A simulation technique to determine the effect of queues and random time variations.

LIRO hopes to determine the efficiency of several New Zealand log transport operations, and to present these results in a report later this year.

References

[1] *Reducing Log Truck Delays*, LIRA Report Vol. 4 No. 11 (1979).

[2] J.A. Stulen, *Computers for Log Transport*, LIRA Report Vol. 10 No. 2 (1985).

[3] P. Cossens, *Evaluation of ASICAM for Truck Scheduling in New Zealand*, LIRO Report Vol. 18 No. 7 (1993).

[4] B. Benseman, *Computer Scheduling Tanker Transport*, DSIR Physical Sciences Report 15 (1991).

[5] C.D.T. Watson-Gandy and L.R. Foulds, *The Vehicle Scheduling Problem: A Survey*, NZOR Vol. 9, No. 2 (1981).

[6] E.S. Miyata, Z.B.Zabinsky, D.B. Rideout and C.N.Mann, *Least-Cost Analysis of Log Transportation*, Proceedings of the 1986 Winter Meeting

of the American Society of Agricultural Engineers.

[7] Z. Shen and J. Sessions, *Log Truck Scheduling by Network Programming*, Forest Products Journal Vol. 39, No. 10 (1989)

[8] R. McCormack, *Coordinating Contractor Trucks Through Computer Aided Despatch of Log Trucks*, Proceedings of the IUFRO Centennial Meeting, Berlin 1992, Project Group P3.07.

[9] M.F. Dodd, *Reducing Transport Costs Through Truck Scheduling with the aid of a Linear Programme*, B.For.Sc. Thesis 1990.

[10] R. Goldsack, *Costing Handbook for Log Truck Contractors*, NZ Logging Industry Research Association 1988. A revised version of this manual is in preparation.

[11] S. Bates and T. Robinson, *Log Truck Unloading*, LIRO Report (in preparation).

[12] T. Robinson, *Approaches to Log Truck Scheduling and Despatching*, Proceedings of the 29th Annual Conference (Auckland 1993) of the Operational Research Society of New Zealand, pp 193-200.

[13] T. Robinson, *Tour Generation for Log Truck Scheduling*, Proceedings of the 30th Annual Conference (Palmerston North 1994) of the Operational Research Society of New Zealand (in preparation).

[14] H.A. Taha, *Operations Research, an Introduction*, Macmillan Publishing Co., Inc.