University of Leeds

SCHOOL OF COMPUTER STUDIES

RESEARCH REPORT SERIES

Report 96.27

Scheduling Train Drivers and Investigating Alternative Scenarios by Computer

by

Ann S K Kwan, Raymond S K Kwan, Margaret E Parker & Anthony Wren

Division of Operational Research & Information Systems

September 1996

To be presented at the 4th Meeting of the EURO Working Group on Transportation, Newcastle, 9 - 11 September, 1996

INTRODUCTION

The authors are members of a group with extensive experience in scheduling public transport vehicles and their drivers since the early 1960's. The IMPACS bus driver scheduling system (Smith and Wren, 1988; Wren and Smith, 1988) based on heuristics and integer linear programming was developed in the 70's and is still being used by many transport operators. In 1990 we collaborated with the Operational Research Unit (ROR) of British Rail to develop a model which could be used to estimate the costs of producing driver schedules under a range of alternative operating scenarios (Parker, Wren and Kwan, 1995). For this purpose we adapted some parts of the IMPACS system so that they would speedily produce outline schedules meeting British Rail's requirements. In some of the simpler test situations we were able to produce actual operational schedules and so to calibrate our model. The degree of accuracy of the model was more than satisfactory for the purpose of the exercise in hand, but we were aware that much needed to be done in order to incorporate all the needs of a rail system.

The adapted system continued to be evolved alongside researches to enhance its computational speed and robustness. In 1994 a project, sponsored by the Engineering and Physical Sciences Research Council (EPSRC), was launched to take further steps to realise the potentials of the system for rail operators. Many components have been replaced and completely new methods have been developed to resolve some of the rail problems which we had not met in the bus industry. The resulting system, TRACS II, is now being used to produce operable and efficient schedules for many rail companies.

The approach of this project is to undertake scheduling problems from a number of UK train operating companies (TOC's) and gain real scheduling experience with them, thereby further enhancing the system. In all our case studies, TRACS II was able to produce driver schedules that were comparable in quality with, and often better than, those being operated. Many TOC's appreciate the speed and accuracy of TRACS II, and several of them have recently commissioned projects using it for investigating alternative scenarios to aid management decision making. A few TOC's are presently considering installing TRACS II for use by themselves. This paper describes some of the complex issues of train driver scheduling tackled. It then describes the TRACS II system and its applications to a number of real problems.

ASPECTS OF TRAIN DRIVER SCHEDULING PROBLEMS

1. Multi-depot Schedules

In the bus industry drivers are usually restricted to vehicles from their home depots, so that where several depots were involved there were effectively a number of separate scheduling problems. All the train companies with which we have worked have had several (as many as fifteen) depots, with drivers allowed in principle to drive trains from any depot provided that they return to the same starting depot at the conclusion of their shifts. Constructing a set of shifts, each of which satisfies the conditions at some depot, is non-trivial There are various constraints such as route and traction knowledge (see below). Some rules and parameters (e.g. time allowances for signing on) are depot-related, and some of them are adjustable in order to satisfy other rules. For the schedule as a whole, there are shift-type specific capacity constraints to be satisfied for some depots.

2. Problem Size

A driver scheduling problem may be large on two different counts. It may be large because there is much work to be covered requiring many shifts (e.g. in excess of 150). It may also be large, computationally speaking, because there are very many relief opportunities close together in time. In the latter case the number of shifts required to operate the train schedule may not necessarily be large (e.g. 50 shifts). In the former case, the problem will be decomposed intelligently into smaller sub-problems. In the latter case, some of the relief opportunities will be de-selected using heuristics so that the problem size is reduced. Some problems may require both decomposition and relief opportunities de-selection.

Three out of the six case studies below have schedules with the number of shifts in excess of 150. All three required decomposition and one also required relief opportunities de-selection.

3. Drivers Travelling as Passenger

In order to optimise overall efficiency it is often best to transport drivers between locations. In such a circumstance the driver becomes a passenger, either on trains currently being scheduled, or on trains run by other operators, or even on other modes of transport such as taxis. A good scheduling system must be able to consider such possibilities. Early versions of TRACS II catered for drivers travelling between relief points either by assuming a standard travel time, or by assigning more than one driver to a vehicle as part of a shift; the excess drivers would in fact be passengers. Both methods have their shortcomings. Assigning more than one driver means that shifts may overlap and the train work assigned in these circumstances is said to be overcovered. This facility only allows the drivers with the same route and traction knowledge to travel as passenger on certain trains. Using standard travelling times may give an inaccurate travel time which might either be too generous or unachievable in practice.

The problem is solved by creating automatically a timetable for each relief opportunity. Train times operated by different TOC's or different transport modes can also be specified. A recently added feature includes the ability to allow for a 'leeway' whenever the driver has to change train (either on duty or travel as passenger) as a safety margin against train delays.

4. Shifts on more than three trains

In rail operation a shift may contain small pieces of non-wheel-turning work such as train preparation or disposal. There is therefore a need to create shifts which work on many vehicles (train sets). However, the number of possible shift compositions increases exponentially with the number of train sets allowed in a shift. TRACS II at present has a limit of up to four sets in a shift. Creation of shifts with more than four trains is under investigation. The general strategy is to impose progressively more restrictive heuristics in the shift formation process according to the number of trains allowed in a shift, so that for example only exceptionally efficient four-train shifts would be formed.

On the other hand, shifts without a meal break are not practised in British rail operations but may become an option in future. Although these shifts can be considered as 'part' shifts, they may be transformed into a 'full' shift if a meal break is allocated either at the front, or more usually, at the end of the shift. 'Part' shifts are not as efficient as the usual 'full' shifts, but they can be useful in completing a schedule.

5. Route and Traction Knowledge

Route knowledge considerations are vital in preparing rail schedules. Rail drivers have to be trained to operate on every track on which they may drive, and have to maintain that knowledge by being assigned frequently to work every line within their domain. For this reason, drivers are generally restricted to certain sections of the total operating area dependent on the depots they belong to. These sections usually overlap, so that it is not possible to subdivide the problem into areas of specific knowledge. The same argument applies to traction knowledge where certain drivers may have the knowledge of some, if not all, of the traction types.

To accommodate the route knowledge constraints, a 'route number' is assigned to each piece of work between adjacent relief opportunities. The route numbers are compiled according to the information supplied by the TOC's. The routes known by drivers from particular depots are specified as parameters to TRACS II. Traction knowledge is handled similarly.

6. Constraints on Schedule Composition

The most common constraint on the overall shape of a schedule concerns the proportions of shifts of different lengths. For example, one usual constraint says that there should be no more than 50% of shifts greater than 8 hours and not more than 20% more than 8.5 hours. Examples of other constraints on schedule composition are: bounds on the number of shifts for each depot; average shift length. These constraints can be tackled by classifying shifts into specific types and then placing limits on these shift types during the mathematical phase of TRACS II.

7. Meal Break Rules

There are complicated rules governing when a meal break should occur; e.g. it must occur between the third and fifth hour relative to the start of shift if the shift's length is less than 8 hours; a different set of rules apply if the shift's length is more than 8 hours. Some shifts which are crucial in forming a good schedule might have been excluded because of these rules. Certain intelligence has to be incorporated into the system so that such shifts would become valid by extending either the signing on or the signing off time. Extending the signing on time can place the meal break in the correct time bracket whereas extending the signing off time can force the shift's spreadover into a different spreadover range which might have a slightly more relaxed meal break rule.

8. Other Aspects

Other features of train driver scheduling include:

- Some train work could form the first part of an early shift or the last part of a night shift, depending on how it best contributes to overall efficiency
- Some operators might want schedules with a mixture of shifts of wide ranging spreadovers, e.g. 4 hours to 12 hours

TRACS II

The features of train driver scheduling described above have already been catered for in TRACS II, whose system components are described in detail in a paper presented this year (Kwan, Kwan, Parker and Wren, 1996). The system is driven by three main parameter sets. One describes the trains for which drivers are to be provided, and specifies geographic data, such as when and how drivers can travel between points of the network. Another sets up the scheduling rules, for example, lengths of shifts of different types, position and duration of meal breaks, etc. The third sets up some global constraints and governs the running of the system.

TRACS II consists of a set of inter-linked programs which are normally run serially corresponding to the following major phases:

1. Pre-processing of scheduling data

In this phase, critical relief opportunities are identified and where necessary some of the less important ones are dropped in order computationally to simplify the problem. This phase is also responsible for working out the possibilities for drivers to travel as passengers.

2. Building a large shift set

The fundamental approach of TRACS II is to generate a set of all the valid and potentially good shifts, so that at a later phase a subset can be selected to form a schedule. A number of heuristic rules are employed in this phase so that on one hand not all the millions of possible legal shifts would be formed, and on the other hand no shifts vital in an efficient schedule would be missed.

3. Consolidation of shift set

Since the processes in (2) have to be generic for problems from all styles of train operations, the resulting shifts set is often of a size (e.g. 45,000) that the next mathematical phase would find it computationally overwhelming. Hence this phase provides some filtering programs that allow the user to appraise and eliminate some shifts under specific rules, e.g. shifts that are wholly contained within other shifts can be eliminated; shifts whose work components are heavily covered by other shifts are deemed to be of less importance for retaining. A typical size for the shift set to be consolidated to is about 20,000 shifts.

4. Mathematical selection from shift set

This phase employs the set covering integer linear programming model (see, e.g., Wren and Smith, 1988) to select from a large (e.g. 20,000) shift set the most economical subset (the schedule) that covers all the work and satisfies any other constraints present.

5. Adjustments and production of schedule documents

In this phase, some adjustments to the mathematically yielded schedule might be performed. This might be to resolve overlapping work due to the set covering model used. A decision has to be made as to which shift should actually be covering the overlapped work in practice. There might also be some swapping of work pieces between shifts to enhance the intrinsic quality of the schedule, which could not be easily handled mathematically.

In addition, there are a number of programs for managing sub-problems when it is necessary to subdivide a large problem.

CASE STUDIES

The system has been used on behalf of about a dozen train operators in the last two years. The complexity of the problems has varied considerably reflecting the wide range of conditions of the TOC's. The system has been used both to assess the efficiency of existing schedules (TRACS II has always been able to produce schedules at least as efficient as those previously in operation) and to predict the distribution of work amongst depots. Once basic data has been established it is very easy to adjust it to cater for different scenarios, and to re-run the programs. This helps to evaluate the costs of alternative operating conditions by producing full schedules for a range of strategies, e.g.:-

- To consider a depot closure
- To adjust the maximum shift length

- To change the positioning of meal breaks
- To add to the route knowledge of a depot
- To effect a wider range of spreadover of shifts
- To achieve a target average shift length
- To create a new depot, or relief point

We have worked directly with five operating companies, and through Rail Operational Research Ltd (ROR) who have run our programs for other clients. In one case we have taken on board the full scheduling work for one of the clients of ROR. Where ROR has undertaken work themselves we have acted as advisers and have modified the programs as appropriate to deal with particular customer needs. In this paper we present our work for the six companies for whom we have produced the schedules ourselves.

Five of the problems which we have investigated have been undertaken as part of a research contract financed by the Engineering and Physical Sciences Research Council (EPSRC). The others have been tackled under contract to the operating companies, including one company for whom we had initially tackled a problem under the EPSRC research contract. In this paper we cannot detail the conclusions reached as a result of the exercises, but in each case we obtained schedules matching current conditions which were at least as good as those in operation, and often better.

1 InterCity East Coast

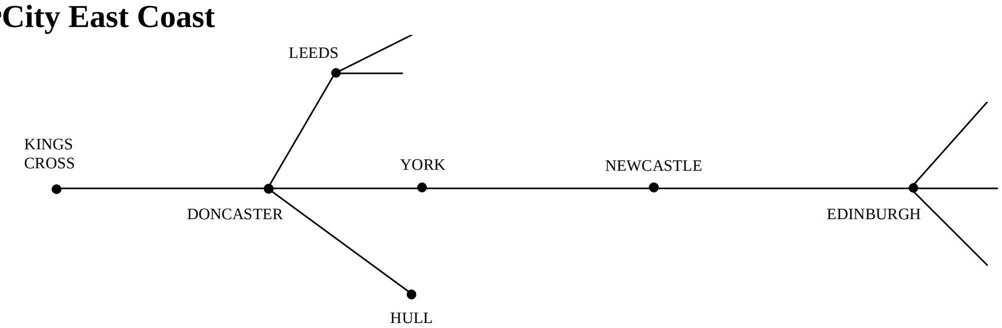

Figure 1 The InterCity East Coast network

The first problem was tackled around the end of 1994, and it is a long-distance fast operation along a 400 mile track from London to Edinburgh. One of the main features of the operation is that two drivers are needed while the electric trains exceed 110 m.p.h between London and Newcastle. There are five driver depots each being one or two mile from train depots.

The driver schedule falls into two parts because of the way the double manning operation is carried out. Some drivers based at one of the depots operate exclusively between London and Newcastle, between which there is a requirement for double manning, and work as pairs throughout the entire shift. The remainder of the drivers pair up on an ad hoc basis and need not be based at the same driver depot. The driver schedules were thus constructed separately according to this criterion with the second category being the larger. In order to produce a schedule where drivers team up on an arbitrary basis it is necessary to specify all the train work requiring double manning twice and to regard the shifts produced as totally independent of each other. Double manning is only required for the wheel turning part of the shift. Two drivers are not required for preparing the train unit for the subsequent journey or for ensuring that the unit is left in a safe manner at the conclusion of the journey.

This problem involved a large geographic area with the opportunities to relieve drivers remotely spaced both in time and distance, a feature that is quite different to intensive urban bus operation which we had previously dealt with. The shift formation process was straightforward and the time required to run the whole process was short. The first schedule was produced before many enhancements had been incorporated in TRACS II. The schedules produced were comparable with those produced manually. Later, the problem was re-run using an up-to-date version of TRACS II and the schedules produced used two shifts fewer than those produced manually. The number of shifts was slightly over 100 in total.

2 Thameslink

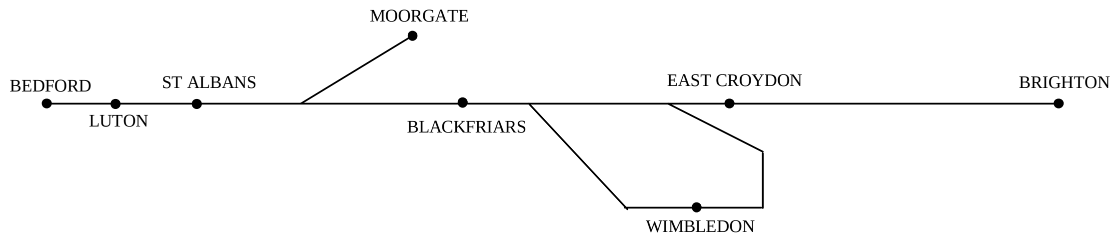

Figure 2 The Thameslink network

The company operates relatively intensive commuter suburban and middle distance services over about 150 miles from Brighton to Bedford passing through the centre of London. There is a short spur leaving the main line to Moorgate, and a loop line branching after Blackfriars through Wimbledon to return towards Bedford. The running time from end to end is about 140 minutes, and there are trains about every 15 minutes over the central section. There are three driver depots; one of which is not on the tracks operated. To get to this particular depot, drivers have to travel as passengers on trains of another company unless they are driving to or from the depot. There are some simple route restrictions applied to drivers of two of the depots.

This problem was tackled at a very early stage before the passenger travelling feature was incorporated. Drivers' travel times between relief points were estimated and put into a table of 'standard' allowances. These allowances might be realistic during usual commuting hours but becoming impracticable during late night or early morning hours when the train service was infrequent. Compared with the manual solution of around 120 shifts, TRACS II was able to use marginally fewer shifts.

This exercise demonstrated that TRACS II's ability to produce schedules of similar quality to manual schedules, especially on the distribution of work amongst depots. The company appreciated that TRACS II could provide a very realistic forecast very quickly and they later requested us to investigate several 'what-if' scenarios. These 'what-if' exercises were completed in a short period of time and they found the results useful in evaluating different options.

3 London Underground

Figure 3 The Piccadilly Line

London Underground is a major metropolitan operator. This exercise covered the Piccadilly Line with trains every two-and-a-half minutes over the busiest period. The end-to-end times are 70 or 80 minutes to Heathrow or Uxbridge. Cockfosters, Oakwood and Arnos Grove are all used as turning points, while at the other end of the line most trains terminate at Heathrow or Rayners Lane. The extension to Uxbridge is only used at peak times. There are train depots near Oakwood and Northfields, while six trains are stabled overnight at South Harrow, restarting in the early morning. Driver depots (i.e., reporting points) are at Oakwood, Wood Green, Acton Town and Northfields. Acton Town and Northfields are only 4 minutes apart, while drivers taking trains to or from South Harrow travel by taxi to Acton Town or Northfields.

At the time this was the largest problem we had ever tackled, with 2500 relief opportunities. There were a number of special conditions which we had never met before; the ability to sign on or sign off at points other than the depots at certain times of day, and a limit on the earliest possible starting time at one of the depots. Travelling between places during day time could be estimated because the train service was so frequent. During late night, extra consideration had to be given for drivers to travel by taxis. All these special circumstances were handled for this demonstration exercise by hard coding into the routine which checks the validity of shifts.

There were 72 train sets in use daily, of which 18 returned to a depot between the peaks. There were thus 90 train blocks, of which 18 operated for up to five hours over the morning peak, and 18 started before the afternoon peak. The earliest train started at around 0445, and the latest finished around 0130; trains did not necessarily finish work at the depot from which they had started. The current schedule used 169 shifts.

Hitherto the largest problem we had tackled, without sub-dividing it, had had about 700 relief opportunities, so it was clear that steps would have to be taken to reduce the problem size. In our bus driver scheduling work we had developed heuristics to eliminate unlikely relief opportunities and to decompose large problems into sensible smaller units (Wren and Smith, 1988), but both had been developed from our specialised knowledge of bus driver scheduling from a single depot and were unlikely to work in the current complex situation. Manual analysis resulted in a reduction of about 300 relief opportunities that were unlikely to be useful in the early morning and late evening.

We were now faced with decomposing a problem of about 2200 relief opportunities. We formed the train blocks into two groups; blocks which started in the morning at Oakwood or started in the afternoon anywhere and finished at Oakwood, and the rest (which were associated similarly with Northfields or South Harrow). These groups were of approximately equal size, and it was observed that about half the current shifts were entirely on trains from a single group, indicating no significant correlation between driver depot and train depot. We then formed four subproblems which were of decreasing size; each consisted of similar ratios of peak, off-peak and evening trains.

Following the strategy used in bus driver scheduling the first subproblem was solved and the work contained in the least efficient shifts was removed from the solution and merged with the second subproblem. This cascading process was continued automatically by the system through to the fourth subproblem, with about 30% of the work being carried forward each time (the reason for the original subproblems being made progressively smaller). Finally the retained shifts from the first three subproblems were added to the result of the fourth subproblem, yielding a schedule with 167 shifts. This saving of two shifts was accepted in principle by the company, who, however, pointed out that three shifts required the driver to travel as passenger back to the other end of the line to sign off, which would be unpopular. We therefore added a constraint forbidding this type of shift and obtained a solution with 168 shifts; the company then accepted that our original solution with 167 was preferable.

4 Regional Railway North East

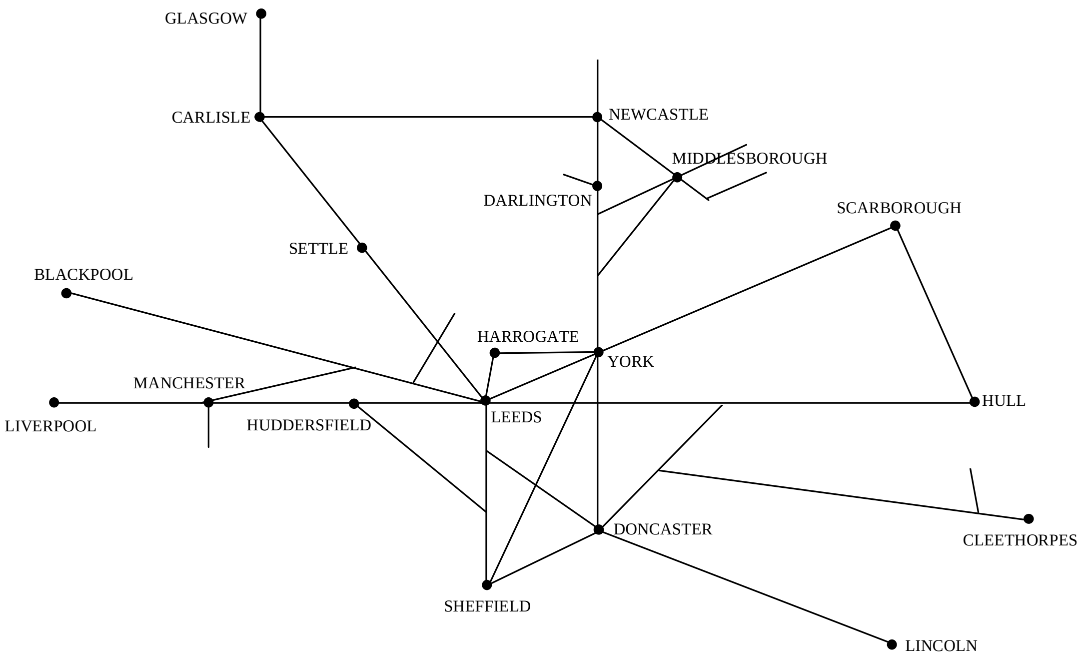

Figure 4. The Regional Railway North East Network

Regional Railways North East (RRNE) first assisted our research in April 1995 by providing us with driver shifts for Monday to Friday based on three depots, together with the rules on which the shifts were based. The initial study combined an urban commuting area with a rural area in the north of England. Relief places are close together both in time and distance. However the route knowledge restrictions help in restricting the size of the problem since the drivers from certain depots can drive only on certain parts of the network. Certain assumptions had to be made about times allowed to travel as passenger between points on the system, as TRACS II could not at that stage deal with drivers travelling other than as drivers on trains which were being scheduled. At this stage we were not yet able to create shifts working on more than three different traction units. Despite these limitations, the schedules produced were satisfactory after minor adjustments, and actually used 4 fewer shifts than the existing operation.

RRNE approached us in October 1995, asking us to repeat the exercise we had already carried out, but using a range of different scenarios. This work was undertaken in the period up to March 1996, and assisted the company to narrow the choice of scenario to a single set of rules affecting such matters as signing on and off, preparation and disposal allowances and meal breaks, with three options for other rules.

By the spring of 1996, TRACS II had been extended to allow the formation of shifts involving more than three units, and to incorporate provision for passenger travel on trains of any appropriate company or indeed by taxis, provided that these had all been coded into the data. The following work was undertaken based on this new version.

In April we were commissioned by RRNE to apply alternative scenarios to their whole Monday to Friday operation. This was by far the largest and most complicated problem we had ever tackled, with around twenty depots or driver links, and over 400 daily driver shifts. There were restrictions on route and traction knowledge for each depot, so that any individual train could only be driven from a small number of depots. The route knowledge of neighbouring depots did however overlap considerably, so that the whole operation was a single interacting process rather than a set of individual depot-based problems. The numbers of shifts to be assigned to certain depots were also limited.

Ideally we should have liked to have had time to develop a solution strategy based on the TRACS II capability of subdividing a large problem, solving the first subdivision and carrying forward badly scheduled work to the next subproblem. This is done automatically in our earlier bus driver scheduling system. However, in bus driver scheduling, the subdivision is carried out by consideration of how fragments of work from a single depot best hang together, and a different set of rules would have had to be developed for train operation. We could of course have followed the same strategy, but carried out the subdividing processes manually as we did with London Underground. We might first have considered the northernmost part of the operation. After creating shifts based on this, we could have removed the poorer shifts, automatically adding their work to later subproblems and continued the process over the whole operation.

However, RRNE required us to obtain schedules as quickly as possible, and we knew that in this first exercise over the whole network we would meet a wide range of new situations in coding data and in deciding how some local difficulties should be tackled. We would not therefore have had time to carry out the work in sequence as above. Instead, we first coded and developed shifts for a set of relatively self-contained work. This enabled us to assess the magnitude of the whole problem and to make some decisions as to how to proceed.

We then subdivided the whole operation into five subproblems based on suitable combinations of depots, route and traction knowledge. These subproblems were then solved in parallel. For each subproblem many tens of thousands of possible shifts were generated according to the most restrictive of the possible scenarios. The shifts generation program provides information on any train work which cannot be included in shifts according to the rules being followed. This work was examined carefully, and dealt with by a number of actions.

Sometimes work was uncovered due to errors in coding or to inconsistencies in allocation of work to subproblems. In other cases, work which was in theory suitable for one subproblem could not be fitted together with other work in that subproblem to make an efficient shift according to the parameters being followed. Our next task was therefore to correct the errors and to exchange problematic work between subproblems.

However, in some cases the problems had arisen because of real difficulties at certain points. For example, in order to prevent the generation of very large numbers of unlikely shifts, TRACS II places maximum lengths on the amount of idle time between components of a shift. At certain places it was not possible to find matches for some incoming trains according to these rules. We therefore had to allow the rules to be bent in certain circumstances. In other cases, examination of the problem situation revealed that there was no way of covering certain work according to the given operating scenario, and some of the parameters originally stipulated had to be altered by agreement with RRNE.

The first sets of shifts covering two scenarios were sent to RRNE about six weeks after the start of the exercise. Following examination by RRNE some errors in the data were corrected. There was also a new requirement for the system to allow for a specified leeway whenever a driver has to change train either to carry on their duty or to travel as passenger. This leeway feature was used as a safety measure for possible delays and was later implemented. The revised schedules for the whole network were despatched to the company within three weeks later.

5 West Anglia and Great Northern

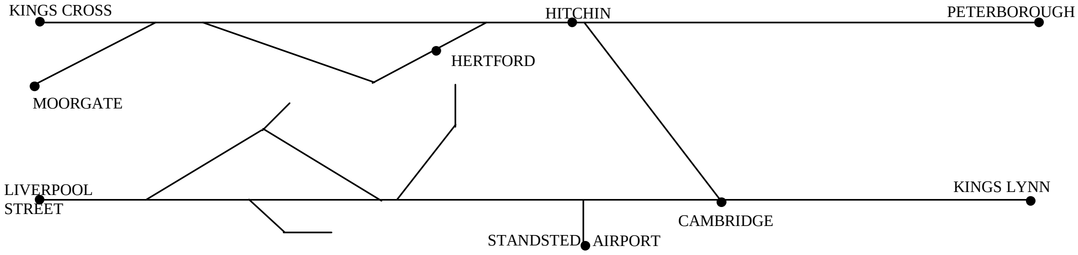

Figure 5 The WAGN network

We are studying a portion of their network which comprises 4 driver depots, and we have undertaken a preliminary exercise in which we have produced schedules for two depots for calibration purposes. The company operates an intensive service over a system which is broadly H-shaped, with several additional spurs and loops. The total number of weekday shifts is about 150, from 4 depots. Currently we have scheduled shifts based at two of the depots, of which there are about 100 shifts.

The intensive service gives rise to very many relief opportunities which are closely packed together. In order to reduce the problem to a manageable size, a program which eliminates many relief opportunities using heuristics has to be used. One main feature of the manual schedule is the extensive use of passenger travelling. A few of the passenger journeys involved travelling on three different trains. At the moment, our system could only cope with passenger travelling on a maximum of two different vehicles plus further walking to or from places. Hence, some of the passenger journey times have to be adjusted manually.

The solution produced apparently saves one shift. However, there are a few shifts whose sign on time falls into the 'unsocial hours' period of 0000 - 0459 inclusive. This will be further investigated in future.

6 South Wales and West

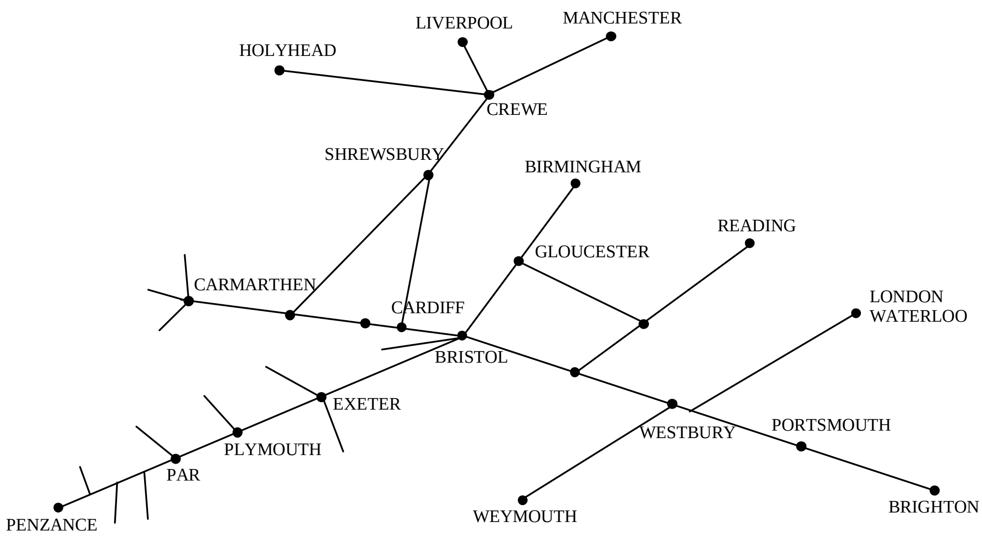

Figure 6 The South Wales and West network

The exercise was undertaken in the winter of 1995 to 1996 through Rail Operational Research, although the work has been done by our team. The company operates over a wide geographic area with 13 driver depots. The rail network has a mixture of long distance, commuter, rural and branch line traffic. A particular feature here is that there is a very significant amount of travel as passenger within shifts, and that trains of other companies are sometimes used for this purpose. We therefore had to extend our arrangement for passenger travel to allow the system to read details of trains which it was not going to schedule but might be used if appropriate for movement of drivers.

We first scheduled a part of the company's operation covering the West Country which is a large, but relatively isolated, geographic area consisting of a main line from Exeter to Penzance with many branches. This part of the operation was covered principally from seven depots, but we included the work of six shifts which operated into the area in question from three other depots. One further depot in the area was expected to cause difficulties in meeting the relevant conditions, and its work was initially excluded from the schedule. After also removing from consideration the work of two night shifts which consisted of many very short shunting movements which could be well defined in advance, there were 49 shifts in the current schedule; and TRACS II produced a satisfactory schedule also using 49 shifts. Experiments were then carried out with a number of alternative scheduling rules.

It was then decided to extend the work to the whole of the company, and initially this was to be divided into several areas as subproblems, one of which was the West Country area which was already investigated and relatively self-contained. The remaining areas have a total of around 150 shifts. Although the size of the remaining subset is big in terms of number of shifts, the huge geographic area it involved and the route restrictions helped in restricting the number of possible combinations. It was then decided to divide the problem into two subproblems only; one contained the West Country area and the other contained the remaining areas. The schedule of the whole network produced by TRACS II has three fewer shifts than that of the manual process. We then tested a number of scenarios on the whole operation. The whole exericse was completed within about one month.

OTHER EXERCISES USING TRACS II

ROR themselves have carried out five further exercises using the system on behalf of their client operating companies. These are :

- InterCity West Coast (on catering staff)
- Mersey Rail
- North London Rail
- InterCity East Coast
- Great Western

CONCLUSION

We have demonstrated that it is possible using TRACS II to produce efficient schedules for train drivers according to widely differing circumstances. In all cases these have been acceptable to management, and usually cheaper than existing schedules. The system can be used both to produce production schedules and to investigate alternative operating scenarios, such as using a different set of depots to cover the operations, or adopting different working rules.

We believe that the RRNE exercise is the first time anywhere in the world where train driver schedules for non-urban situations have been produced on such a large scale. We have demonstrated that the system is also suitable for intensive urban operations. It is a very powerful tool for producing schedules under alternative operating scenarios. Unlike a manual scheduler, the system has no preconceived ideas of how the portions of train work should be put together and this might lead to better shift formulation.

REFERENCES

Fores, S., Wren, A. and Proll, L.G. (1996). A column generation approach to bus driver scheduling. Submitted to 4th Meeting of the EURO Working Group on Transportation, Newcastle.

Kwan, A.S.K., Kwan, R.S.K., Parker, M.E. and Wren, A. (1996). Producing train driver shifts by computer. In: *Computers in railways V* (J. Allen, C.A. Brebbia, R.J. Hill, G. Sciutto and S. Sone, eds.), Vol. 1: *Railway Systems Management*, pp. 427-435. Computational Mechanics Publications.

Parker, M.E., Wren, A. and Kwan, R.S.K. (1995). Modelling the scheduling of train drivers. In: *Computer-aided transit scheduling* (J.R. Daduna, I. Branco and J.M.P. Paixao, eds.), Springer-Verlag, pp.359-370.

Smith, B.M. and Wren, A. (1988). A bus crew scheduling system using a set covering formulation. In: *Transportation Research*, vol.22A, pp. 97-108.

Wolfenden, K. and Wren, A. (1966). Locomotive scheduling by computer. In: *Proc. British Joint Computer Conference*, pp. 31-37, London, IEE Conference Publication no. 19.

Wren, A. and Kwan, R.S.K. and Parker, M.E. (1995). Scheduling of rail driver duties. In: *Computers in Railways IV* (T.K.S. Murthy, B. Mellitt, C.A. Brebbia, G. Sciutto, S. Sone, eds.), Vol. 2 : *Railway Operations*, pp. 81-90. Computational Mechanics Publications.

Wren, A. and Smith, B.M. (1988). Experiences with a crew scheduling system based on set covering. In: *Computer-aided transit scheduling* (J.R. Daduna and A. Wren, eds.), pp. 104-118, Springer-Verlag.