Crew Scheduling for Netherlands Railways "Destination: Customer"

Leo Kroon

NS Reizigers and

Erasmus University

P.O. Box 1738

3000 DR Rotterdam, The Netherlands

L.Kroon@fbk.eur.nl

Matteo Fischetti

University of Padova, DEI

Via Gradenigo 6/A

35100 Padova, Italy

Fisch@dei.unipd.it

Abstract: *In this paper we describe the use of a set covering model with additional constraints for scheduling train drivers and conductors for the Dutch railway operator NS Reizigers. The schedules were generated according to new rules originating from the project "Destination: Customer" ("Bestemming: Klant" in Dutch). This project is carried out by NS Reizigers in order to increase the quality and the punctuality of its train services. With respect to the scheduling of drivers and conductors, this project involves the generation of efficient and acceptable duties with a high robustness against the transfer of delays of trains. A key issue for the acceptability of the duties is the included amount of variation per duty. The applied set covering model is solved by dynamic column generation techniques, Lagrangean relaxation and powerful heuristics. The model and the solution techniques are part of the TURNI system, which is currently used by NS Reizigers for carrying out several analyses concerning the required capacities of the depots. The latter are strongly influenced by the new rules.*

1 Introduction

Since a couple of years, there is a separation in the Dutch railway system between the capacity management and the maintenance of the railway infrastructure on one hand, and the exploitation of the railway infrastructure on the other. Capacity management and maintenance of the infrastructure are carried out by the government, and the railway operators are commercial organizations that are allowed to exploit the infrastructure by operating their trains. NS Reizigers is the largest Dutch railway operator, specialized in the railway transportation of passengers.

Improving the quality and punctuality of the train services is currently one of the major issues in the Dutch railway system. Therefore, many projects are being carried out both by the government and by the railway operators. Projects carried out by the government involve the capacity and the quality of the infrastructure. NS Reizigers currently carries out the project "Destination: Customer" ("Bestemming: Klant" in Dutch). This project involves, amongst others, an improvement of the robustness and the stability of the timetable, a decrease of the failure rate of the rolling stock, a redesign of the circulation of the rolling stock in order to increase the seating probability, and a redesign of the duties for the drivers and conductors in order to increase their robustness against the transfer of delays of trains.

A commercially operating railway operator has to be able to react quickly to changes in the transportation market and to be able to adapt his operating plans accordingly. Hence, the railway operator should be able to modify his timetable quickly, and he should be able to modify the corresponding schedules for rolling stock and personnel as quickly as well. Furthermore, he should be able to carry out *what-if* analyses in an easy way, in order to study the long-term consequences of certain future scenarios.

In this situation, information systems that provide insight into the transportation market, as well as “intelligent” information systems that provide active support in the processes of timetabling and scheduling of rolling stock and personnel are indispensable tools. Indeed, timetabling and scheduling are highly complex tasks that are very time consuming when they are carried out manually.

In this paper we describe the “intelligent” information systems that are used by NS Reizigers for supporting the internal planning processes of scheduling drivers and conductors. We focus on the TURNI system which provides real active support to the planners in building up efficient and acceptable duties for the drivers and conductors. We used this system for estimating the consequences of the new rules originating from the project “Destination: Customer” for the required capacities of the depots. The new scheduling rules involve the application of train teams consisting of a driver and a conductor, the restricted assignment of train series to depots, and the required amount of variation per duty. Especially the amount of variation per duty is a key issue for the acceptability of the duties.

2 Scheduling train drivers and conductors

In this section we describe some basic principles of the scheduling processes within the Logistics department of NS Reizigers. In Section 3 we describe the new rules originating from the project “Destination: Customer”.

2.1 Depots

The numbers of drivers and conductors of NS Reizigers are about 3000 and 3500, respectively. Each driver and conductor belongs to one of the depots. There are 29 depots, mainly located at the main nodes of the Dutch railway network (see Figure 1). However, also a number of smaller depots are located near the borders of the country. These depots are necessary because of the trains starting in these locations early in the morning or ending there late in the evening. Each driver or conductor operates from his own depot. That is, each of his duties starts and ends in his own depot and there are no overnight duties. For the drivers there are the additional rules that, depending on his knowledge of the different types of rolling stock and on his knowledge of the infrastructure, each driver is allowed to work on certain types of rolling stock and on certain parts of the infrastructure only.

2.2 Train series, trips and duties

Each duty consists of a sequence of trips to be carried out by one single driver or conductor. Each trip is mainly characterized by a train number, a start and end location and a start and end time. Also a number of other aspects of the trips are relevant, such as the involved train type (Intercity or Stop-train) and train series. Here a train series is a set of trains that provide a regular direct connection between two end stations. The train series play an important role within the project “Destination: Customer”. The following example describes a trip on the 833 Intercity train between Amsterdam (Asd) and Utrecht (Ut) from 9:28 to 10:01. The 833 train is a member of the 800 train series, which provides an hourly connection between Haarlem (Hlm) and Maastricht (Mt) (see also Figure 1).

833	Amsterdam	Utrecht	9:28	10:01	IC	800

The duties should be such that each trip is covered by one driver and a sufficient number of conductors. The number of conductors per trip depends on the length of the corresponding train: the longer the train, the more conductors are required. The latter suggests that the process of scheduling drivers and conductors can start only after the timetable and the rolling stock schedule have been completed. This is also more or less what happens in practice, although in practice there are also some feedback iterations between scheduling personnel and scheduling rolling stock. Recently some research efforts were undertaken to integrate the scheduling of rolling stock and personnel [11, 17].

Each duty has to satisfy certain constraints. To give a few examples (non-exhaustively): the maximum length of a duty is 9:30 hours. Each duty longer than 5:30 hours requires a meal break with a length of at least 30 minutes. The time between the start of the duty and the start of the meal break should be at most 5:30 hours and the same should hold for the time between end of the meal break the end of the duty.

Another important constraint involves the minimum connection time: if a driver or conductor transfers from one train unit to another, then a minimum connection time of 25 minutes is to be taken into account. The mentioned constraints are hard constraints, but also several soft constraints are to be respected.

Besides the constraints that are to be satisfied by the individual duties, there are also several constraints that are not related to the individual duties but that are to be satisfied per depot or by the complete final schedule. For example, per depot at most 5% of the duties can be longer than 9:00 hours. Furthermore, per depot the average length of the duties should be less than 8:00 hours. There are also constraints on the percentage of night duties per depot. These constraints are to be satisfied during the scheduling process in order to facilitate the rostering process afterwards.

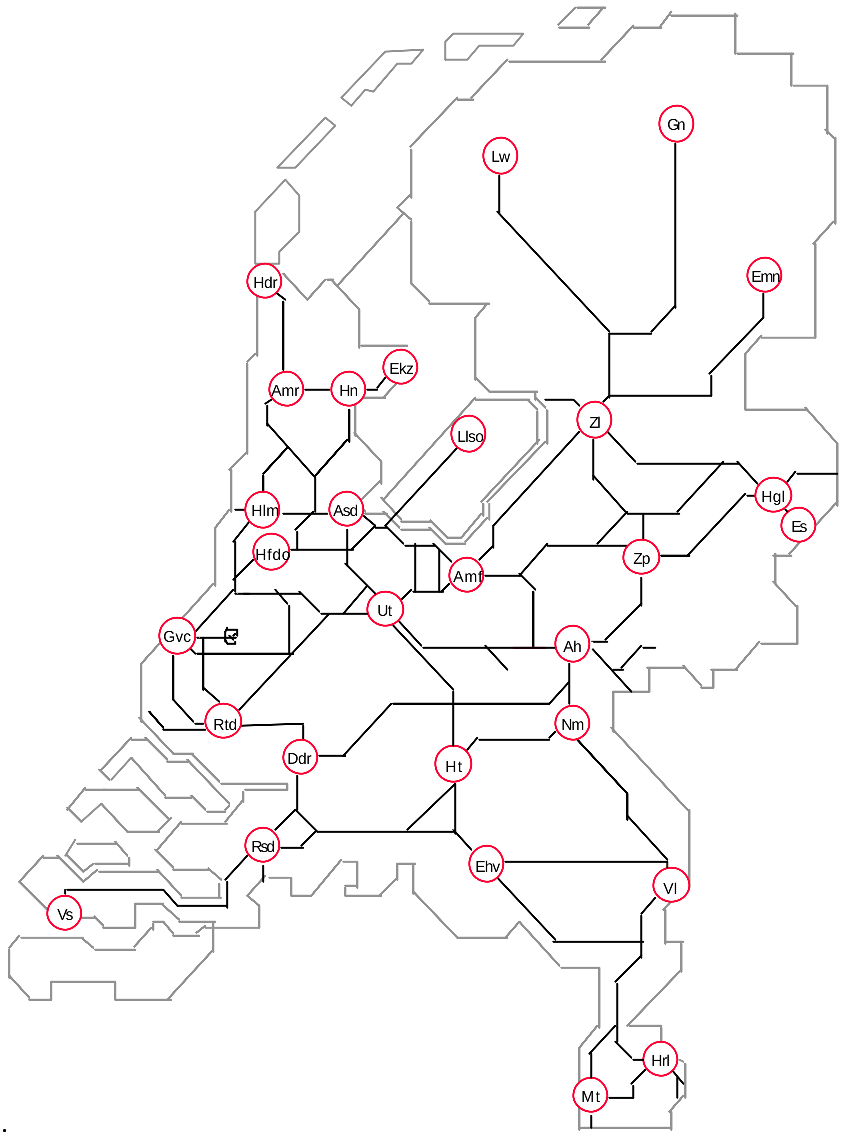

Figure 1: Dutch railway infrastructure and the crew depots of NS Reizigers

The main criteria in the scheduling process of drivers and conductors are *feasibility*, *efficiency* and *acceptability*. Feasibility means that it should be possible to carry out the obtained schedules in practice, and that they are sufficiently robust for outside disruptions and for delays of trains. A key parameter here is the minimum connection time that was mentioned before. The larger the minimum connection time, the more robust the schedule will be. Efficiency means that the percentage of productive time in the duties is high. Non-productive time includes, besides the required pre-time at the start of a duty, the required post-time at the end of a duty and the meal break, also the gaps between the trips, and the so-called P-trips (passenger trips or positioning trips). A P-trip is a trip where a driver or conductor does not act as such, but moves as a passenger in a train from the end location of his previous trip to the start location of his next trip. Acceptability is a qualitative aspect of a schedule, referring to the probability that the obtained schedule will be accepted by the drivers and conductors. Within NS Reizigers, acceptability is highly related to the level of variation in the duties (the more variation, the better).

2.3 Scheduling and rostering

The duties to be carried out are created by the central Logistics department of NS Reizigers. Currently, the operational planning process is supported by the VPT system and by the CREWS system, which are briefly described in Section 4. After the duties have been generated in the Logistics department, they are sent to the depots. Usually, after some iterations of negotiation and adaptation, the duties are accepted by the depots. The generation of the rosters for the depots based on the duties is done locally in the depots. Also in the rostering process many constraints have to be satisfied. For example, after a couple of night shifts a driver or conductor should have a certain number of days off before he may carry out his next duty.

One reason for splitting the central scheduling process from the local rostering process is the fact that both processes are highly complex. Therefore, combining the scheduling and the rostering processes into one process would result in huge combinatorial problems that are practically insolvable. Note that, from a mathematical point of view, the scheduling problem and the rostering problem share a similar structure. Indeed, both problems involve the arrangement of objects into proper sequences. In the scheduling problem the trips have to be sequenced into duties and in the rostering problem the duties have to be sequenced into rosters. For more detailed descriptions of the problems crew scheduling and crew rostering we refer to [6, 7].

3 “Destination: Customer”

As was described earlier, the project “Destination: Customer” is being carried out by NS Reizigers with the objective to obtain an improved quality and punctuality of the train services provided to the customers. One of the aspects to be taken into account within this project is the generation of efficient and acceptable schedules for the drivers and conductors with a maximal robustness with respect to the transfer of delays of trains. For the duties, the following aspects originate from the project “Destination: Customer”:

1. Operating in train teams
2. Variation per duty
3. Assignment of train series to depots

3.1 Operating in train teams

According to the new rules originating from the project “Destination: Customer”, the drivers and conductors will carry out their duties together in so-called train teams. This means that each day a driver and a conductor form a train team and that they carry out the same duty together. Each day the composition of the train teams may be different from the composition of the train teams on the previous day, depending on the rosters for drivers and conductors.

An advantage of the introduction of the train teams is that it will lead to an improved robustness of the schedules. Indeed, in the new situation, only two components will be required for operating a train, namely a unit of rolling stock and a train team. In the traditional situation, three components were required, namely a unit of rolling stock, a driver and a conductor. By reducing the number of components required for operating a train, the probability is increased that all components will be available in time. Furthermore, the rules for the variation in duties (see Section 3.2) aim at keeping also the unit of rolling stock and the train team together as much as possible. This will even further increase this probability.

Although for the majority of the trains one conductor is sufficient, some trains should be serviced by two conductors. In general, the second conductors cannot be part of a train team since there is no matching driver. Therefore, the introduction of the train teams will also require the introduction of special duties for second conductors.

3.2 Variation in duties

For the drivers and conductors, the amount of variation in the duties is a key issue for the acceptability of the duties. In fact, in the current situation most duties are quite varied, and the drivers and conductors would like to keep it like this as much as possible. On the other hand, from the point of view of the robustness of the schedules, it would be optimal to keep the units of rolling stock and the train teams as much as possible together. As was described above, this would even further increase the probability that all components required for operating a train are available in time. However, keeping a unit of rolling stock and a train team together during a complete duty is inefficient from the point of view of the rolling stock, because the train team needs a meal break after a certain amount of time. Furthermore, the units of rolling stock are usually scheduled in an up-and-down fashion between the two end stations of the involved train series. Thus, keeping the rolling stock and the train teams together for a long period of time would result in duties without any variation. Therefore, the following rules were introduced in order to create an acceptable compromise between the objective of keeping the rolling stock and the train teams as long as possible together and the objective of having an acceptable amount of variation in the duties.

> Each duty longer than 5:30 hours not involving a long (>150 km) Intercity train series should contain at least two out of the following three elements:
>
> - Two train series
> - Two train types (Intercity or Stop-train)
> - Two corridors of the infrastructure

Note that these rules are considered as hard rules, which means that they must be satisfied by each of the duties. In general, the variation rules will allow for the generation of duties with the following simple structure: the first part of the duty before the meal break is carried out on one unit of rolling stock servicing a certain train series, and the second part of the duty after the meal break is carried out on another unit of rolling stock servicing another train series. Together with the rules for the assignment of train series to depots, which will be described in the next section, the majority of the duties will have the meal break in the home depot, which has additional managerial advantages.

3.3 Assignment of train series to depots.

The “Destination: Customer” rules also describe that, in principle, each train series may be assigned only to those depots where the train series has a starting and/or ending train. For example, according to this rule the work on the 800 train series between Haarlem (Hlm) and Maastricht (Mt) that was described earlier may be assigned to the depots Haarlem and Maastricht. Furthermore, in the early morning there are some trains in this train series which start in Utrecht (Ut) or Eindhoven (Ehv) and in the late evening there are some trains that end in these locations. Therefore, the work on the 800 train series may also be assigned to these depots.

The rationale behind this assignment rule is that each driver and conductor will operate more often on the same train series than before. Indeed, in the traditional situation each train series could basically be assigned to each of the 29 depots, in particular the work for the conductors. The application of the assignment rule will lead to a higher exposure to and thereby more experience with the peculiarities of the train series and the timetable. Hopefully, this will lead to an improved quality and punctuality of the train services. Obviously, the assignment rule has a large impact on the duties for the drivers and conductors and on the required capacities of the depots, hence the need for an effective automatic crew scheduling system for making an a priori estimate of these consequences.

There are a few exceptions to the assignment rule. Indeed, for some of the depots a strict application of the rule would result in a number of allowable train series that is too small for satisfying the variation rules described in the previous section. In such a case, the depot is considered as a satellite depot of a nearby larger depot. This means that the satellite depot may be assigned also to (part of) the train series of the larger depot in order to be able to satisfy the variation rules. Satellite depots mainly occur near the borders of the country where only a limited number of train series is available.

4 Planning systems

4.1 VPT

Until recently, the scheduling of drivers and conductors was done mainly manually. That is, the duties were created manually by the planners and after their creation the duties were recorded in the information system VPT ("Vervoer Per Trein" in Dutch; "Transport By Train" in English). This system was developed especially for NS, not only for supporting the logistic planning processes but also for supporting the planning and control processes in the daily operations. For supporting the scheduling of drivers and conductors, VPT carries out several checks in order to see whether all trains have been supplied with a driver and with a sufficient number of conductors, whether the duties are feasible, and whether the relevant additional constraints are satisfied. However, VPT itself does not provide real active support in creating the schedules, in the sense that VPT does not come up with suggestions for schedules. On the other hand, VPT certainly provides relevant support in the process of scheduling drivers and conductors.

In many cases, new schedules were created with the support of VPT by modifying previous plans. However, in the current situation, where NS Reizigers operates commercially, this more or less manual planning process is considered as insufficient since it is too time consuming and since it does not allow for the organization to react as flexibly as desired to changes in the transportation market. Furthermore, this manual way of planning is prohibitive for developing complete redesigns of the operational plans and for carrying out "what-if" analyses efficiently.

4.2 CREWS

Therefore it was decided some years ago by the Logistics department of NS Reizigers to use the CREWS system [15]. CREWS supports the scheduling of drivers and conductors and can be used in one of three modes: (*i*) the manual mode, (*ii*) the semi-automatic mode, and (*iii*) the automatic mode. In the manual mode, the system does not provide any active scheduling support to the planner, but it continuously checks whether all relevant constraints are satisfied by the schedule generated by the planner. In the semi-automatic mode, the system provides the planner with suggestions for appropriate next steps in the planning process, which he may or may not follow. In the automatic mode, the system creates a schedule automatically, without any interaction of the planner.

In comparison with the former situation, where schedules were created mainly manually, the use of CREWS is certainly a step forward. CREWS provides several functions which are useful in the planning process, even in the manual mode. For example, the system contains a large database with extensive facilities to generate trips, and to sort and filter the data, which supports the planners in looking for appropriate next steps in the planning process. The database contains a locking mechanism in order to prevent the same trips to be handled by several planners at the same time. The system also provides a large number of graphical representations of the data and the obtained subschedules, such as Gantt charts and time-space diagrams. Furthermore, the system has a user-system interface that is rather user-friendly. Such aspects of the system are useful in supporting the planners in their planning activities.

The semi-automatic and the automatic modes of CREWS are based on the A*-algorithm [5]. To some extent, this algorithm tries to mimic the planning methods used by the planners when they manually build up a schedule: the algorithm creates a schedule by starting with an empty schedule and by successively inserting one trip after another.

The algorithm builds up a search tree, where each node of the search tree is a representation of a corresponding subschedule. The costs of each node are calculated based on two components: (*i*) the costs of the subschedule, and (*ii*) an estimate of the costs that are to be made for completing the subschedule to a final schedule.

The algorithm aims at finding a node corresponding to a final schedule with minimal total costs. To some extent, it is possible to backtrack to another node in the search tree. Since, in the worst case situation, this might result in a more or less complete enumeration of the solution space requiring too much time and memory, the search tree is bounded in two ways: (*i*) each node may have only a certain maximum number of successors (by default: 3), and (*ii*) each level of the search tree may have only a certain maximum number of nodes (by default: 10). In building up the search tree, the algorithm proceeds in a more or less greedy way, where nodes with lower costs are preferred over nodes with higher costs. Nevertheless, often the computing time and the amount of memory required for solving even moderately sized instances is quite large.

Furthermore, the nature of the A*-algorithm allows for the checking of the constraints of the individual duties during the creation of a schedule, but it prohibits the checking of the additional constraints per depot, such as the constraints on the average length of the duties per depot. Moreover, due to the greedy approach of the A*-algorithm, the algorithm does not look forward in order to consider the later consequences of certain choices, and, due to the described limitations of the search tree, it backtracks only to a limited extent. This may sometimes result in schedules that are rather inefficient due to a high number of P-trips. Therefore, the planners are not always completely satisfied with the schedules produced by the automatic mode of CREWS.

4.3 TURNI

In order to be able to estimate the consequences of the “Destination: Customer” rules for the required capacities of the depots as well as possible, in particular the consequences of the new rule involving the assignment of train series to depots, the Logistics department of NS Reizigers decided to look for alternative algorithmic approaches for supporting the crew scheduling process. One of the available alternatives turned out to be the TURNI system. This system has a relatively simple user-system interface, but a powerful scheduling module. Initially, only a few customizations of the system to the situation at NS Reizigers were necessary for obtaining useful results.

4.3.1 Model description

In contrast with the automatic mode of CREWS, which uses techniques originating from Artificial Intelligence, the TURNI system uses mathematical programming techniques originating from Operations Research. Since it is difficult to describe the constraints to be satisfied by the individual duties explicitly in a mathematical programming model, the model underlying the TURNI system is a set covering model with a number of additional constraints.

If a set covering model is applied, then the solution mechanism consists of a duty generation module and a duty selection module. After a certain set of feasible duties has been generated, the duty selection module aims at selecting a subset of feasible duties in such a way that each of the trips is covered by at least one of the selected duties, that the relevant additional constraints are satisfied, and that the total involved costs are minimal. The set of feasible duties may be generated a priori, or it may be generated “on the fly” during the solution process.

Set covering models have been popular in the airline industry for many years [12, 16]. However, in the railway industry the sizes of the crew scheduling instances are, in general, a magnitude larger than in the airline industry, which prohibited the application of these models in the railway industry until recently. But due to the increase in the computational power of nowadays’ computers and algorithms, set covering models become more and more applicable in the railway industry as well [6, 7, 13, 14]. We mainly handled instances with up to 2500 trips. However, we also carried out a number of experiments with much larger instances (about 8000 trips corresponding to all trips of a single day), which turned out to be promising for the near future.

Using the notation $t = 1,\dots,T$ for the trips to be covered, $d = 1,\dots,D$ for the potential duties, and $c = 1,\dots,C$ for the additional constraints to be satisfied, the set covering model with additional constraints can be formulated as follows:

$$\min \sum_{d=1}^{D} c_d x_d \tag{1}$$

subject to

$$\sum_{d=1}^{D} a_{t,d} x_d \geq 1 \qquad \forall t = 1,..., T \tag{2}$$

$$\sum_{d=1}^{D} b_{t,d} x_d \leq u_c \qquad \forall c = 1,..., C \tag{3}$$

$$x_d \in \{0,1\} \qquad \forall d = 1,..., D \tag{4}$$

The meaning of the binary decision variables x_d is as follows:

$$x_d = \begin{cases} 1, \text{ if duty } d \text{ is selected} \\ 0, \text{ otherwise} \end{cases}$$

In the 0-1 matrix $a_{t,d}$ each row represents a trip and each column represents a feasible duty, and $a_{t,d} = 1$ if and only if trip *t* is covered by duty *d*. Besides the regular trips that have to be covered, the TURNI system also allows one to include a number of *suggested* or *inadvisable* trips into the model. These trips may be covered by a duty, but they need not be covered. Adding such trips to an instance may be helpful for finding a feasible schedule or for improving the overall efficiency of the schedule. There is no constraint (2) corresponding to the suggested or inadvisable trips. The input for the TURNI system concerning the trips that need to be covered and on the other available trips is obtained from the CREWS system via a simple data interface.

The additional constraints (3) are not related to the individual duties (such constraints are handled at the duty generation level), but to certain forbidden combinations of duties. These constraints may be related to e.g. the maximum number of duties per depot, the maximum average length of the duties per depot, or the maximum percentage of long duties per depot.

The cost coefficients represent the fact that the main objective is to minimize the number of duties required to cover all trips. However, also other cost aspects, such as additional costs for P-trips, uncovered suggested trips, covered inadvisable trips, or penalties for discouraging undesirable characteristics or violated soft constraints may be included in the cost coefficients.

4.3.2 Solution technique

The resulting model is solved by applying dynamic column generation [3,9], Lagrangean relaxation and several heuristics. Dynamic column generation means that the duties are not generated a priori, but "on the fly" during the solution process. The application of this dynamic approach is required by the fact that the complete set of feasible duties is extremely large in general, hence generating all feasible duties a priori is not a feasible approach. We next give a brief sketch of the solution framework used by TURNI.

The overall execution is organized into a sequence of *passes,* in each of which the system tries to obtain better and better solutions. The initialization pass is aimed at determining an initial set of duties covering all the input trips. The first pass finds the first efficient solutions for the current problem and is designed to reach soon a hopefully good solution by using an aggressive solution strategy. The next passes improve upon the solutions found in previous passes, by using a more accurate solution strategy. Typically, the best solutions are determined already during the very early passes, but in some hard cases also later passes can find improved solutions.

Within each pass, the algorithm iterates between a *duty generation* module and a *d* *uty selection* module. The duty generation module generates new prospective feasible duties based on a dynamic programming heuristic. As customary [3], the prospectiveness of a feasible duty is evaluated based on dual information related to the Linear Programming relaxation of the underlying model. However, TURNI uses Lagrangean relaxation and subgradient optimization instead of Linear Programming for calculating the required dual information, in a vein similar to the one recently proposed in [8] for the solution of pure set covering problems. Feasible duties generated in earlier stages whose prospectiveness turns out to be low during later stages of the algorithm, may be deleted in order to keep the number of active duties manageable.

The duty selection module looks for a heuristic solution for the overall model, based on the currently available set of feasible duties. The applied heuristic is also driven by the Lagrangean dual information. That is, the heuristic evaluates and selects duties from the set of active duties based on the Lagrangean dual information.

Within each pass, the generation and selection phases are cyclically applied for a certain number of iterations, hopefully updating the best solution found. Thereafter a *fixing procedure* is activated to select some duties that appear to be particularly efficient and to fix them as belonging to the final solution. The choice of the duties to be fixed depends on a score computed as a function of the current Lagrangean cost and of the current *duty efficiency*, the latter being defined as the ratio between the overall duration of the trips covered by the duty (skipping all trips covered by the previously-fixed duties as well as the inadvisable trips) and the paid working time. The overall process is then repeated on the trips not covered by the fixed duties: the duty generation and selection phases are iterated for a while on the subproblem resulting from fixing, new duties are fixed, etc. In this way, typically several millions of possible duties are evaluated within the duty generation module, and thousands of alternative solutions are constructed and evaluated during the selection phase.

Fixing typically leads to better and better solutions, up to a point where the current solution cannot be improved any further without releasing some of the fixed duties. In this situation the program ends the current pass, applies a refinement procedure to hopefully improve the best solution by means of trip exchanges among duties, and begins a new pass: the pass counter is incremented, some or all of the fixed duties are released according to their efficiency (as defined earlier), and the program is reapplied on the uncovered trips.

The elaboration continues until some stop condition (time limit, maximum number of passes, etc.) is met. Computational experience has shown that this solution approach has the advantage of finding quickly good solutions for the problem at hand (already after the first fixings, the solutions are typically comparable or better than the best manual ones), while using the remaining processing time for improving them by exploring different fixing patterns. For more details of the duty selection mechanism, we refer the reader to [8]. Other approaches for solving large set covering problems can be found in [1, 2, 3, 4, 9, 10, 16].

4.3.3 Output and other functions.

The results of a run of the algorithm can be shown by the TURNI system in tabular format, where the successive trips in each duty are listed along with some duty statistics, or in a graphical format in the form of a Gantt chart or a time-space diagram. An example of such a Gantt chart is shown in Figure 2 which represents some stop-train duties for train teams from the depot Maastricht (Mt) in the southern part of the Netherlands. This Gantt chart clearly shows that these duties indeed have the structure imposed by the “Destination: Customer” rules. That is, each duty longer than 5:30 hours consists of one part before and one part after the meal break, each one involving a different train series. In this Gantt chart, the train series can be recognized from the first two digits of the train number.

An interesting feature of the TURNI system is the fact that it allows one to observe the best obtained solutions while the algorithm is still running. This is enabled by the multi-tasking facilities of the underlying operating system. At the same time, the system continuously provides an optimistic estimate (that is, a Lagrangean lower bound) of the optimal solution costs. In this way, the user of the system may decide for himself whether or not he is satisfied by the obtained solutions and whether or not he wants to interrupt the run of the algorithm.

Other useful functions provided by the TURNI system are the following. First, the system allows the generation of so-called deadheading trips. That is, the system may generate automatically some additional trips if this is necessary for obtaining a feasible schedule. This may be particularly important if one has only data available on the train movements of commercial passengers trains and not on the additional train movements for positioning the rolling stock. A generated deadheading trip triggers the fact that some of these positioning trips are missing in the available data set. Also useful is the fact that the system allows one to check whether a certain manual solution satisfies all constraints. Here the obtained log-file points at all the constraints that are violated by the given manual solution.

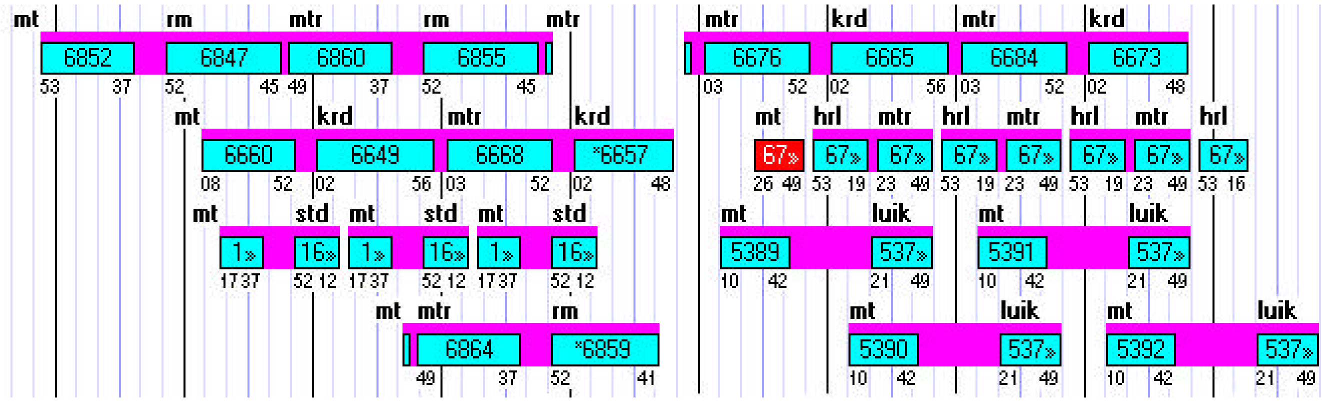

Figure 2: Gantt chart with duties with variation for the depot Maastricht (Mt)

5 Modeling "Destination: Customer" and other rules

5.1 "Destination: Customer" rules

The TURNI system contains a mechanism for specifying for each trip the depots that may be used for servicing the trip. This mechanism could be used to handle the "Destination: Customer" rules concerning the assignment of train series to depots in a straightforward way.

Handling the variation rules was less straightforward. First, we recoded the train series indicators of the trips in such a way that trips involving train series that should not appear in the same duty obtained the same train series indicator. For example, two stop-train series covering (for a large part) the same corridor of the infrastructure obtained identical train series indicators. Furthermore, in order to guarantee that the train teams will not transfer too often from one train series to another, we used an upper bound on the number of transfers per duty from one train series to another. Finally, we used a mechanism for stimulating the generation of duties with the required amount of variation. That is, duties longer than 5:30 hours involving only one train series could be penalized or forbidden in several gradations. Although this was not part of the official "Destination: Customer" rules, we also used a mechanism to prevent the generation of duties with so-called "artificial variation". That is, duties for which the majority of the work is carried out on one train series and only a very minor part of it is carried out on another train series. This mechanism counts a train series in a duty only if the amount of work on this train series exceeds a certain lower bound.

Although from a technical point of view the mechanism for handling the variation rules operated as expected, it had to be handled with some care in practice. In particular, if the generation of duties without the required amount of variation is strictly forbidden, then the instances should be chosen carefully. If an instance does not contain a sufficiently large number of different train series, or if the amounts of work on the different train series do not sufficiently match with each other, then it is obviously impossible to generate duties satisfying the variation rules. Although a careful selection of the instances to be handled indeed usually leads to useful results, we intend to improve the mechanism for handling the variation rules by making it somewhat more subtle. We intend to achieve this by making the variation constraints dependent on the involved depots. This will allow one to focus on the generation of duties satisfying the variation rules initially for only a subset of the depots, while satisfying the variation rules for the other depots can be postponed to later instances.

5.2 Underway train changes

From the point of view of the robustness of the duties, it is desirable that a train team remains as long as possible on the same train and does not get off the train somewhere underway. In order to achieve this, we used the following approach: each trip record was extended with an additional attribute describing whether the trip is the First (F) or the Last (L) trip of the train. Now an underway train change is defined as a transfer of a driver or conductor from trip t_1 to trip t_2, where

TRAIN(t_1) <> TRAIN(t_2) and

t_1 is NOT the Last trip of TRAIN(t_1) or

t_2 is NOT the First trip of TRAIN(t_2).

The duty generation module of the TURNI system was modified in such a way that the number of underway train changes per duty can be bounded or penalized. In the above definition, the first condition implies that the trips t_1 and t_2 are related to trains with different train numbers. In general, this implies that the trains are different. The first part of the second condition means that the train corresponding to trip t_1 has not yet come to his end station. Hence, if the driver or conductor leaves this train, then the train has to be supplied with a different driver or conductor on his remaining route to his end station. This is undesirable, since this may lead to a transfer of a delay. The second part of the second condition means that, if the driver or conductor boards the train corresponding to trip t_2, then it is somewhere underway the route of this train, which is again undesirable.

This mechanism for structuring the duties turned out to be quite useful, in particular when dealing with long Intercity train series. The mechanism could be used to make a trade-off between maximum robustness of the duties (by not allowing any underway train change) and maximum efficiency (by allowing underway train changes).

5.3 Additional rules for train teams

The fact that drivers and guards will operate in train teams implies that for the train teams at least the traditional rules for the drivers have to be applied. These traditional rules for the drivers are somewhat more restrictive than the traditional rules for the conductors. In most cases, these more restrictive rules for the drivers are related to the combining and splitting of trains, and then they are due to the fact that a driver has a fixed position in front of the train. Therefore a driver is less flexible in his movements than a conductor who can move more or less freely throughout the train. Splitting and combining of trains occurs in several train series of NS Reizigers. Although the additionally required synchronization of trains may have a bad impact on the punctuality, it also allows for more direct connections for passengers and at the same time an efficient use of the infrastructure.

The following example involves the combining and splitting of trains in the 6000 train series (Utrecht - Tiel) and the 16000 train series (Geldermalsen - ’s Hertogenbosch) in Geldermalsen. A train in the 6000 train series arriving from Utrecht (Ut) in Geldermalsen (Gdm) is split into two parts: the front part of the split train continues to Tiel (Tl) as a 6000 train and the rear part of the split train continues to ’s Hertogenbosch (Ht) as a 16000 train. As a consequence, the train team of the arriving train should continue to Tiel and the part of the 16000 train bound for ’s Hertogenbosch should be supplied with another train team. In the reverse direction, the 6000 train from Tiel and the 16000 train from ’s Hertogenbosch are combined into one single 6000 train from Geldermalsen to Utrecht. The train from ’s Hertogenbosch arrives first in Geldermalsen and becomes the front part of the combined train. The train from Tiel arrives second and becomes the rear part of the combined train. Thus the train team arriving from ’s Hertogenbosch may continue on the combined train to Utrecht, and the train team from Tiel has to leave its train unit.

We modeled these additional rules by recoding the train numbers of the trips where necessary. For example, for the trains 6040 and 16040, there are three trips, namely

6040	Tiel	Geldermalsen	12:36	12:46	Stop-train	6040
6040	Geldermalsen	Utrecht	12:49	13:11	Stop-train	16040
16040	’s Hertogenbosch	Geldermalsen	12:20	12:38	Stop-train	16040

Here the first column represents the original train number and the last column represents the recoded train number that we really used. The fact that the recoded train number of the Tiel-Geldermalsen trip is different from the recoded train number of the Geldermalsen-Utrecht trip, together with the minimum connection time of 25 minutes that is to be respected between two trips with different recoded train numbers, prohibits the train team on the Tiel-Geldermalsen trip to continue on the Geldermalsen-Utrecht trip. This train team may continue e.g. on the next 16000 trip Geldermalsen-'s Hertogenbosch. Furthermore, due to the recoding of the trips, the train team on the 's Hertogenbosch-Geldermalsen trip may continue on the Geldermalsen-Utrecht trip.

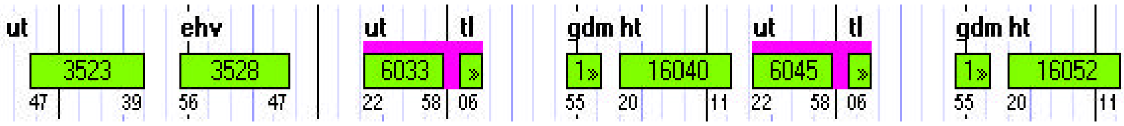

Figure 3: Duty involving the 6000/16000 train series.

That this recoding of the trips indeed has the required effect, is shown by the duty in the Gantt chart in Figure 3. This duty includes two times the route Utrecht – Tiel – Geldermalsen – 's Hertogenbosch – Utrecht, which is correct according to the given description. Note that the Gantt chart shows the recoded train numbers. The variation in this duty is provided by the trips from Utrecht (Ut) to Eindhoven (Ehv) and vice versa on the Intercity train series 3500.

5.4 Duties for second conductors

As was described earlier, some trains must be serviced by two conductors. Thus the introduction of the train teams requires the generation of separate duties for these second conductors since for the second conductors there is no matching driver. There are some exceptions from this rule, since for the duties on a couple of long Intercity lines, train teams consisting of one driver and two conductors will be used.

A problem that we encountered here is the fact that the involved set of trains is only a relatively small subset of the set of all trains. This implies that the trips for the second conductors do not fit nicely together, as the trips for the train teams do. This problem is solved partly by the fact that it was decided that most of the "Destination: Customer" rules need not be applied to the duties for the second conductors. Hence, the trips for the second conductors may be assigned to all depots and the variation rules need not be satisfied here: only the traditional rules, involving aspects such as the average duty length per depot and the meal break, need to be satisfied. Nevertheless, the instances of the scheduling problem containing only the trips for the second conductors turned out to be infeasible. In particular, it turned out to be impossible to cover each of these trips by a feasible duty.

We solved this problem by first using the deadheading mechanism of the TURNI system. This mechanism generates additional trips at appropriate positions in time and space in such a way that a feasible covering of all trips is obtained. This mechanism uses a complete "distance" matrix containing the time required to travel from one location to another. We will not go into the details of the deadheading mechanism here. After having run the instance with the activated deadheading mechanism, we added a large set of additional timetabled inadvisable trips matching the generated deadheading trips as well as possible. In this way, we finally obtained a feasible solution covering all required trips.

Nevertheless, it turned out that in general the duties for the second conductors will be intrinsically less efficient than the duties for the train teams. However, within NS Reizigers the advantages of the introduction of the train teams is preferred over the lower efficiency of the duties for the second conductors.

6 Computational results

In this section we present computational results based on the experiments that we carried out. All experiments were performed on a PC with a Pentium II processor with 450 MHz and an internal memory of 128 Mbytes. Note that currently much faster PC's are available.

6.1 Instances

Since the timetable of NS Reizigers and its circulation plan for the rolling stock are both more or less cyclic with a cycle length of one week, the crew schedules can be cyclic with a cycle length of one week as well. Furthermore, although there are some differences between the timetables and the circulation plans for the different workdays of the week, these are in general so small that, from a crew scheduling point of view, the workdays can be considered as identical. Only during the weekends the timetables and the circulation plans are significantly different. In particular, during the weekends there are less early trains, there are no additional trains for rush hours, and for some train series the frequency is reduced. Traditionally, each workday counts for about 15% of the work, Saturday counts for about 13% of the work and Sunday counts for about 12% of the work.

In accordance with the foregoing, we split our set of trips into instances for a single workday (in particular, Tuesday), instances for Saturday, and instances for Sunday. Furthermore, we split these instances further into instances for train teams and instances for second conductors. The numbers of trips per day for train teams varied between about 8000 for Sunday and 10000 for Tuesday.

6.2 Train teams/separate train series

We started the analysis of each individual day by handling the involved train series separately from each other. Obviously, this will prohibit the generation of duties satisfying the variation rules. However, this gave a lot of insight into the peculiarities of the different train series, which proved to be quite useful in the analysis of the combined instances. For example, some train series have a complicated circulation of the rolling stock. An example is the 5500 train series Utrecht – Baarn. In principle this train series should be serviced from the depot Utrecht (Ut), since this is the only depot along this train series. However, in the early morning, the first train from Baarn to Utrecht is supplied with a unit of rolling stock from Amersfoort (Amf). Therefore, this train series implies at least one early duty for the depot Amersfoort. Similarly, there is also at least one late duty for this depot. Furthermore, handling the train series separately from each other also learned how easy it would be to move part of the involved duties from one depot to another without increasing the total number of duties.

In general, each of the separate train series is handled very fast, since the corresponding numbers of trips were small. The largest of these instances contained about 400 trips. The computing times for solving these instances varied from seconds for the smaller instances to minutes for the larger ones. According to the planners, the obtained solutions for these instances were of very good quality in all cases.

6.3 Train teams/combined train series

For solving the real scheduling problem, we combined the instances for the separate train series in such a way that the final duties would also satisfy the variation rules. This sometimes required a careful composition of the instances. The latter was particularly true for the areas of the satellite depots, since in these areas with only few train series it is hard to satisfy the variation rules. Here we handled instances with up to 1000 trips. We limited the size of these instances in order to have a better control over the final result. In the areas of the large depots with many train series (Amsterdam, Rotterdam, The Hague, Utrecht, Eindhoven), we generated instances that were as large as possible. In these areas it is relatively easy to satisfy the variation rules, and therefore the obtained duties were almost always immediately acceptable. The largest instances that we handled for these areas contained about 2500 trips.

We pre-processed these instances as much as possible by applying the available forcing mechanism, which can be used to force certain pairs of trips to be scheduled directly after another in one duty. Applying this mechanism may improve the quality of the solution, since certain required combinations of trips can not be broken anymore. At the same time it may lead to a reduced computing time.

In almost all of these experiments, it appeared that after the second pass of the algorithm (see Section 4.3.2) only marginal improvements of the quality of the solution were obtained. In particular, the number of duties required to cover all trips hardly ever decreased after the second pass. Obviously, the time required to get to this point depends on the number of trips in the involved instance. For an instance with about 1000 trips it required usually about 15 to 20 minutes on our hardware. For an instance with up to 2500 trips some more time would be necessary, but in all cases this point was reached within one hour of computing time.

In these experiments, we also had to manipulate the numbers of duties for the depots in order to obtain a maximal match of capacity and demand in the depots. In order to achieve this, we used both the mechanism for bounding the number of duties in a certain depot and the mechanism for attracting some extra duties to another depot by applying a cost reduction for the duties there. The first mechanism is handled by the additional constraints (3) and the second mechanism is handled by modifying the objective function (1) appropriately. Unfortunately, it turned out that a perfect match could not be obtained, due to the strict rules for the assignment of train series to depots.

6.4 Duties for second guards

The numbers of trips per day for second conductors varied between 1200 for Sunday and 1500 for Tuesday. For each day, these numbers of trips could be handled in one instance. Although these instances were not extremely large, it took somewhat more time to obtain nearly optimal solutions here. This was mainly due to the fact that we had to use the deadheading mechanism both for obtaining feasible solutions and for obtaining better ones. The heuristic deadheading mechanism sometimes succeeded only after a couple of passes to find a crucial deadheading trip, thereby improving the quality of the solution. Nevertheless, the final results after a couple of hours of computing time for these instances were quite in accordance with the results that were obtained manually by the planners in parallel.

7 Conclusions

Within NS Reizigers, it is strongly believed that a proper structuring of the duties for the drivers and the conductors will lead to an improvement of the quality and the punctuality of the train services. The new rules originating from the project “Destination: Customer” aim at the specification of proper structures for the duties. On one hand, the duties should be as robust a possible against the delays of trains, and on the other hand they should allow for an acceptable amount of variation per duty. The new rules also describe that each train series may be assigned to a limited number of depots only. In this paper we described how we used the TURNI system to estimate the consequences of these new rules for the required capacities of the depots. Whether the “Destination: Customer” rules will indeed lead to an improvement of the robustness and stability of the duties for drivers and conductors will have to be proved in practice. However, we also intend to use a simulation model for assessing the effects of the new rules a priori.

For the drivers, the practical conclusion of the experiments is that the total number of available drivers will be sufficient for covering all the work. However, the drivers are sometimes wrongly distributed over the depots. For the conductors, the practical conclusion is that an additional number of conductors will have to be hired for covering all the work. This is first due to the fact that the more restrictive rules for the drivers will also be applied for the conductors, because drivers and conductors will operate in train teams. It is also due to the intrinsically inefficient duties for the second conductors. Also the conductors are sometimes wrongly distributed over the depots. Nevertheless, in practice no drivers nor conductors will be obliged to move to other depots.

In order to obtain schedules that will be feasible in practice, the “Destination: Customer” rules will initially be relaxed to some extent in order to guarantee that all the work can be covered by the available drivers and conductors. For example, one option for dealing with this is to apply the new rules initially only on weekdays and not during the weekends. Anyway, based on the experiments that we carried out, the ideal capacities of the depots have been determined. These will be used in the long run as a target for the capacities of the depots.

Based on the experiments that we carried out with the TURNI system, it can be concluded that the algorithmic part of the system is quite powerful. In our experiments, we have mainly solved instances with up to about 2500 trips. In general, these instances could be handled effectively within a computing time of about one hour. Furthermore, we also carried out a number of experiments with much larger instances (about 8000 trips), which turned out to be promising. Based on the results of these experiments, we are confident that in the near future it will be possible to handle all trips of NS Reizigers corresponding to one single day in one instance.

Although the new rules had not been settled yet completely when we started with our experiments, we managed to generate duties with (more or less) the desired structures. Initially this required some ingenuity sometimes. In particular, satisfying the new rules for the required amount of variation per duty was rather complex. However, after some adaptations to the model and the system, it became more and more easy to obtain duties with the desired structures. Nevertheless, we intend to implement some further adaptations to the model and the system in order to further simplify this by allowing more control over the final result. Anyway, the planners and the management of the Logistics department of NS Reizigers were rather satisfied already with the results obtained so far.

An important conclusion from our experiments is also the fact that practical crew scheduling problems are highly complex. Therefore, the kernel of an intelligent system that is intended to provide real active support for the solution of these problems should be a powerful algorithm that does not only consider the feasibility of the individual duties, but also the feasibility of the schedule as a whole. A more or less greedy heuristic that successively builds up the individual duties and that has only limited backtracking facilities is deemed to be incapable of capturing the complexities of scheduling drivers and conductors.

References

[1] E. Balas, and M.C. Carrera, "A dynamic subgradient-based branch-and-bound procedure for set covering," Operations Research, 44 (1990) 25-42.

[2] E. Balas, and A. Ho, "Set covering algorithms using cutting planes, heuristics and subgradient optimization: a computational study," Mathematical Programming Study, 12 (1980) 37-60.

[3] C. Barnhart, E.L. Johnson, G.L. Nemhauser, M.W.P. Savelsbergh and P.H. Vance, "Branch-and-Price: Column Generation for Solving Huge Integer Programs," in J.R. Birge and K.G. Murty (eds.), Mathematical Programming: State of the Art 1994, University of Michigan, 1994, 186-207.

[4] J.E. Beasley, "A Lagrangian heuristic for set covering problems," Naval Research Logistics, 31 (1990) 151-164.

[5] L. Bolc, and J. Cytowski, "Search Methods for Artificial Intelligence," Academic Press, 1992.

[6] A. Caprara, M. Fischetti, P.L. Guida, P. Toth, and D. Vigo, "Solution of large scale railway crew planning problems: the Italian experience," In N.H.M. Wilson (ed.), Computer Aided Transit Scheduling, Springer-Verlag, 1999.

[7] A. Caprara, M. Fischetti, P. Toth, D. Vigo, and P.L. Guida, "Algorithms for railway crew management," Mathematical Programming, 79 (1997) 125-141.

[8] A. Caprara, M. Fischetti, P. Toth, "A heuristic method for the set covering problem," Operations Research, 47 (1999) 730-743.

[9] J. Desrosiers, Y. Dumas, M.M. Solomon and F. Soumis, "Time constrained routing and scheduling," in M.O. Ball et al. (eds.), Handbooks in OR & MS, Vol. 8, Elsevier Science, 1995, 35-139.

[10] M.L. Fisher, and P. Kedia, "Optimal solutions for set covering/partitioning problems using dual heuristics," Management Science, 36 (1990) 674-688.

[11] R. Freling, "Models and techniques for integrating vehicle and crew scheduling," Ph.D. Thesis, Erasmus University Rotterdam, 1997.

[12] K.L. Hoffmann, and M.W. Padberg, "Solving airline crew scheduling problems by branch-and-cut," Management Science, 39 (1993) 657-682.

[13] L.G. Kroon and M. Fischetti, "Scheduling drivers and guards: the Dutch Noord-Oost case", Proceedings of the 34th HICSS conference, Hawaii, January 2000.

[14] A.S.K. Kwan, R.S.K. Kwan, M.E. Parker, and A. Wren, "Producing train driver schedules under different operating strategies," In N.H.M. Wilson (ed.), Computer Aided Transit Scheduling, Springer-Verlag, 1999.

[15] E.M. Morgado, and J.P. Martins, "Crews-NS: scheduling train crews in the Netherlands," AI Magazine, Spring 1998, 25-38.

[16] D. Wedelin, "An algorithm for large scale 0-1 integer programming with application to airline crew scheduling," Annals of Operations Research, 57 (1995) 283-301.

[17] A. Wren, and N.D. Fares Gualda, "Integrated scheduling of buses and drivers," In N.H.M. Wilson (ed.), Computer Aided Transit Scheduling, Springer-Verlag, 1999.