UvA-DARE (Digital Academic Repository)

The new Dutch timetable: The OR revolution

Kroon, L.; Huisman, D.; Abbink, E.; Fioole, P.J.; Fischetti, M.; Maróti, G.; Schrijver, A.; Steenbeek, A.; Ybema, R.

DOI

[10.1287/inte.1080.0409](https://doi.org/10.1287/inte.1080.0409)

Publication date

2009

Document Version

Final published version

Published in

Interfaces

[Link to publication](#)

Citation for published version (APA):

Kroon, L., Huisman, D., Abbink, E., Fioole, P. J., Fischetti, M., Maróti, G., Schrijver, A., Steenbeek, A., & Ybema, R. (2009). The new Dutch timetable: The OR revolution. *Interfaces*, 39(1), 6-17. <https://doi.org/10.1287/inte.1080.0409>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

The New Dutch Timetable: The OR Revolution

Leo Kroon

Department of Logistics, Netherlands Railways (NS), NL-3500 HA Utrecht, and Rotterdam School of Management, Erasmus University Rotterdam, NL-3000 DR Rotterdam, The Netherlands, lkroon@rsm.nl

Dennis Huisman

Department of Logistics, Netherlands Railways (NS), NL-3500 HA Utrecht, and Econometric Institute, Erasmus University Rotterdam, NL-3000 DR Rotterdam, The Netherlands, huisman@few.eur.nl

Erwin Abbink, Pieter-Jan Fioole

Department of Logistics, Netherlands Railways (NS), NL-3500 HA Utrecht, The Netherlands
{erwin.abbink@ns.nl, pieterjan.fioole@ns.nl}

Matteo Fischetti

Department of Information Engineering, University of Padova, and Double-Click sas, I-35131 Padova, Italy,
m.fischetti@turni.it

Gábor Maróti

Rotterdam School of Management, Erasmus University Rotterdam, NL-3000 DR Rotterdam, The Netherlands,
gmaroti@rsm.nl

Alexander Schrijver

Center for Mathematics and Computer Science (CWI), Amsterdam, and University of Amsterdam,
NL-1098 SJ Amsterdam, The Netherlands, lex@cw.nl

Adri Steenbeek

Safiro Software Solutions, NL-1303 AJ Almere, The Netherlands, a.steenbeek@safiro.nl

Roelof Ybema

Department of Logistics, Netherlands Railways (NS), NL-3500 HA Utrecht, The Netherlands, roelof.ybema@ns.nl

In December 2006, Netherlands Railways introduced a completely new timetable. Its objective was to facilitate the growth of passenger and freight transport on a highly utilized railway network and improve the robustness of the timetable, thus resulting in fewer operational train delays. Modifications to the existing timetable, which was constructed in 1970, were not an option; additional growth would require significant investments in the rail infrastructure.

Constructing a railway timetable from scratch for about 5,500 daily trains was a complex problem. To support this process, we generated several timetables using sophisticated operations research techniques. Furthermore, because rolling-stock and crew costs are principal components of the costs of a passenger railway operator, we used innovative operations research tools to devise efficient schedules for these two resources.

The new resource schedules and the increased number of passengers resulted in an additional annual profit of €40 million (\$60 million); the additional revenues generated approximately €10 million of this profit. We expect this profit to increase to €70 million (\$105 million) annually in the coming years. However, the benefits of the new timetable for the Dutch society as a whole are much greater: more trains are transporting more passengers on the same railway infrastructure, and these trains are arriving and departing on schedule more than they ever have in the past. In addition, the rail transport system will be able to handle future transportation demand growth and thus allow cities to remain accessible to more people. Therefore, we expect that many will switch from car transport to rail transport, thus reducing the emission of greenhouse gases.

Key words: transportation: rail, scheduling, vehicles, crew; programming: large-scale systems, integer; network graphs: multicommodity.

In a small and densely populated country, such as The Netherlands, public transport plays an important mobility role and is indispensable to the economy and the public welfare. The backbone of the Dutch public transport network is the national passenger railway system. In this paper, we will describe how operations research techniques supported the development of a new timetable that facilitates operating a higher number of trains and fewer train delays.

Until 1995, Netherlands Railways (Nederlandse Spoorwegen, or NS) was a state-owned company, operating passenger and freight services and building and maintaining the railway infrastructure. Because of European Union regulations and liberalization of the railway market, NS was split into several companies during the period 1995–2002. The state maintains ownership of the infrastructure because of its strategic value. ProRail, a nonprofit organization owned by the state, is responsible for maintaining and allocating the infrastructure. Several companies, including NS, now operate the passenger services; there are also a number of freight operators. Approximately 95 percent of the trains are passenger trains, which transport approximately 1.1 million passengers on an average workday. In 2006, almost nine million different passengers traveled about 15.8 billion passenger kilometers by train. On average, each Dutch citizen travels approximately 1,000 kilometers by train per year. NS, which owns the license to operate passenger trains on all main lines until 2015 (Figure 1), is still by far the largest passenger-train operator.

The main lines represent approximately 90 percent of the total passenger demand; the NS operating revenues on these lines are approximately €1.5 billion per year. In 2006, its operating income on these lines was approximately €200 million.

Railway services are especially attractive for traveling between the large cities. For example, the NS market share of passenger movement between the four largest cities—Amsterdam, Rotterdam, The Hague, and Utrecht—is above 50 percent during peak hours, a period in which the road network between these cities is highly congested. If all commuters who currently travel by train switched to traveling by car, these cities would become almost inaccessible. This would have a dramatic negative impact on the Dutch economy. Moreover, railway transport is environmentally

Figure 1: This map shows The Netherlands passenger railway network as of January 2008. The total network, which is mostly double-tracked and electrified, is approximately 2,800 kilometers long. The four large cities—Amsterdam (Asd), Rotterdam (Rtd), The Hague (Gvc), and Utrecht (Ut)—are all in the western part of the country.

friendly (CO₂ emission per train-passenger kilometer is about one-third of the emission of an average car-passenger kilometer; see Nederlandse Spoorwegen 2008). Therefore, reducing greenhouse pollution is a benefit of reducing automobile travel.

For these reasons, the Dutch government wants to stimulate the growth of railway transport. However, it does not want to invest additional billions of euros in new infrastructure, although its infrastructure is already very highly utilized compared with other countries (Figures 2 and 3).

Figure 2: This graph shows the top 10 European countries that have the highest ratio of train kilometers over track kilometers. The Netherlands railway system is one of the most intensively used networks within Europe (International Union of Railways 2007).

Figure 3: This graph shows the top 10 European countries that have the highest ratio of passenger kilometers over line kilometers. NS transports the most passengers on a kilometer of railway line in Europe, illustrating its efficient use of railway infrastructure (International Union of Railways 2007).

Moreover, building infrastructure is time-consuming and requires considerable land, which is unavailable in most places in The Netherlands.

In 2002, all parties involved in the Dutch railway sector (ProRail, NS, and the freight operators) wrote a report, *Benutten en Bouwen* (Utilize and Build), which studied these issues. This report was the catalyst for launching a project to construct a new timetable with the primary goal of managing future growth on an already highly utilized railway infrastructure. Because NS is the largest of several operators that use the same infrastructure, NS constructed one timetable for all operators; in doing so, it considered the wishes of the other operators. When there was a conflicting wish, ProRail, an independent organization, made the final decision to resolve the conflict. On December 10, 2006, NS began operating this new timetable, marking a new era for all railway passengers in The Netherlands. In the remainder of this paper, we will discuss why a new timetable was necessary, how we developed it, and the challenges we had to conquer. We will focus on the key role that operations research tools played in the construction of this new timetable.

Background

The last major change in the Dutch timetable occurred in 1970. Since then, the amount of passenger transport

on the Dutch railway network has nearly doubled, with 8 billion passenger kilometers in 1970 growing to 15.8 billion in 2006. During the same period, freight transport increased by 285 percent. To facilitate this growth, NS scheduled more and larger trains without changing the basic structure of the timetable. Moreover, since World War II, the infrastructure has only been extended slightly.

This had two important consequences: (1) Further growth based on the existing timetable was impossible without significant investments in the infrastructure (as mentioned above, this was not an option), and (2) the buffers in the system became smaller and smaller, resulting in train delays; delays are the most frustrating aspect of train travel. Therefore, one of the NS primary objectives was to improve the punctuality of the railway system. The definition of punctuality in The Netherlands is the percentage of trains that arrive at one of its 35 main stations with a delay of less than three minutes.

Thus, NS had to run more trains on the network and improve the punctuality of the railway system. In principle, these two goals conflict; it is easy to add more trains to an existing timetable if we allow lower punctuality. Conversely, it is also simple to improve the punctuality by reducing the number of trains in an existing timetable. The only way we could fulfill both goals was to develop a new timetable from scratch.

When NS decided to develop a new timetable, it realized that it could also meet several other objectives. For example, introducing a timetable that is easy to remember on the most important lines in the western part of the country led to new commercial opportunities. In the new timetable, a long-distance train that stops at the main stations only should arrive every 15 minutes, and a regional train that stops at all stations should also arrive every 15 minutes. Moreover, the timetable should improve connections with the neighboring countries of Germany and Belgium.

The only characteristic of the 1970 timetable that remained intact was its cyclic nature. In the Dutch case, this means that each hour a train leaves at the same minute in the same direction. Passengers find this timetable property, which many other European countries also successfully apply, to be very convenient. Therefore, we considered only cyclic variants of the new timetable. One disadvantage of a cyclic

timetable is that during certain periods of the day, there may be many trains serving a relatively low demand. We alleviate this by allowing some exceptions; for example, not all trains run in the late evenings.

A specific reason for launching the new timetable in December 2006 was that, at the start of the development process in June 2003, everyone had expected that three major infrastructure extensions would be completed by December 2006: (1) A high-speed line between Amsterdam and Belgium, (2) the Betuwe freight line between the port of Rotterdam and Germany, and (3) four parallel tracks between Amsterdam and Utrecht, which would allow fast (long-distance) trains and slow (regional and freight) trains to have their own tracks. Unfortunately, the construction work took longer than anticipated, causing many challenges for the timetabling project. In the *Implementation and Challenges* section, we will discuss these challenges and our solutions.

Modeling the Timetabling Problem and Related Planning Problems

To construct a new timetable and its related resource schedules, we modeled and solved a sequence of planning problems. As input, we had to define a line system. Then, we calculated the timetable, including the detailed routings through the stations. Finally, we needed to construct rolling-stock and crew schedules. Note that we solved these different planning problems sequentially.

However, to avoid going back and forth between the three planning phases, at each phase we considered specific characteristics of the subsequent problem. For example, the required number of conductors on a train depends on the length of the train, which is determined during the rolling-stock scheduling. Therefore, minimizing the required number of conductors is part of the objective function in the rolling-stock scheduling problem.

In the remainder of this section, we will give a short overview on how we solved these planning problems. Huisman et al. (2005) and Caprara et al. (2007) provide extensive overviews of these solutions.

Line System

The line system, which is the collection of all train lines, is the key input for the timetabling process, in

which each *line* has an origin station and a final destination station; a frequency and a certain stopping pattern indicate the stations at which the trains on the line call: at all stations (regional trains) or at major stations only (long-distance trains). For example, there is a line from The Hague to Groningen that runs once per hour and stops only at major stations.

Timetabling

The timetable describes the *planned* departure and arrival times of every train at every station. These time instants are called *events*. Other relevant events are the time instants at which the trains pass junctions, bridges, and other locations where coordination of train movements is required. Bridges, which need to be opened frequently for ships, are a particular bottleneck in the Dutch railway system because of the numerous waterways in The Netherlands.

Trains running from one station to another or trains dwelling inside a station define the relationship between these events. Similarly, the headway time between two consecutive trains on the same route also defines a relationship between two events.

To increase the robustness of the timetable, we increased the running times, dwell times, and headway times by time supplements based on experience and expert opinions. Time supplements in the running and dwell times absorb small disturbances in the real-time operations, allowing trains to recover from delays. Time supplements in the headway times, also called buffer times, reduce the propagation of delays from one train to another.

Thus, time supplements and buffer times add to the predictability of the railway system. However, the downside is that extending time supplements might require more rolling stock and crew and thus influence the costs of operating the timetable. As we mentioned earlier, we considered only cyclic variants of the timetable. Therefore, we only developed decision-support tools to generate cyclic timetables. The conversion of a cyclic timetable to a timetable for a whole week is a relatively simple copy and paste operation.

The model that we used for timetable generation describes the cyclic timetabling problem in terms of the periodic event scheduling problem (PESP) constraints. This is a generic model for scheduling a set of periodic events, such as the event times in a cyclic

timetable (Serafini and Ukovich 1989). All PESP constraints are expressed as differences of event times. For example, the running time of a train from one station to another is the difference between the arrival time and the departure time. We give a mathematical expression in the appendix. Similarly, the headway time is the difference between the departure times of two consecutive trains on the same track. The headway times, in particular, make the timetabling problem an extremely complex combinatorial problem because the orders of the trains on the tracks are not known a priori. Moreover, we computed all time differences modulo 60 to reflect the timetable's cycle of one hour. For example, the time intervals between minutes 10 and 35 of an hour and between minutes 50 and 15 have a length of 25 minutes.

We can solve small instances of the cyclic timetabling problem as mixed-integer programming (MIP) problems. However, typical NS instances contain approximately 8,400 events to schedule and about 70,000 PESP constraints. Because each PESP constraint results in a binary variable (representing the previously mentioned orders of the trains), and the LP relaxation of such a formulation is very weak, the MIP approach does not work in this situation.

To solve the PESP, we developed the CADANS module. CADANS is based on constraint programming techniques, primarily because we are interested in finding a feasible solution. It includes several special-purpose constraint programming techniques. For example, we paid special attention to the procedure for backtracking from infeasible branches of the search tree. Therefore, the required computing time is often only a few minutes. Schrijver and Steenbeek (1993) provide additional details of the techniques that underlie CADANS.

CADANS provides a feasible solution that satisfies all the PESP constraints, if such a solution exists. If a feasible solution does not exist, then CADANS indicates the conflicting constraints and shows how to solve the infeasibility. The system user can then relax one of the constraints (e.g., a transfer possibility) and run CADANS again.

The implemented constraint programming techniques do not provide facilities for direct optimization of the timetable. However, once CADANS obtains a feasible timetable, a postoptimization procedure improves the timetable by modifying the arrival and

departure times of the trains, without changing the basic structure of the timetable. In this step, CADANS improves the transfer times for the passengers at specific stations.

Routing Trains Through Stations

Stations form a bottleneck in the Dutch railway system because many trains from different directions come together at the stations, and the stations have limited infrastructure. Moreover, to provide good transfer opportunities for the passengers, trains preferably arrive (and depart) more or less at the same time and at adjacent platforms. Therefore, trains can easily hinder each other inside the stations, and finding appropriate routes for the trains through the stations is as important as determining their arrival and departure times. As long as these routes are undetermined, the timetable is incomplete.

Each passenger train requires an inbound route from the point it enters the station area to a platform, and a matching outbound route from this platform to the train's point of departure. The routes of a freight train need not pass along a platform track. The routes must consider the train arrival and departure times that we determined in the timetabling step. The most important requirement is that two trains do not claim the same part of the infrastructure at the same time.

To solve the problem of routing trains through a station, we developed the STATIONS module. Within STATIONS, we first list the potential routes through the station for each train. In the model, we represent each routing possibility by a binary decision variable that indicates whether or not the model selects the routing possibility. The set of selected routes is such that it is optimal according to the prespecified preferences and such that the selected routes fit with each other. That is, no two selected routes claim any part of the infrastructure at the same time, and the selected routes provide the required cross-platform connections for the passengers. Compared to cyclic timetabling, the routing problem is relatively easy. We solved the problem by CPLEX after several preprocessing steps. Zwaneveld et al. (2001) provide additional details.

DONS

The two modules CADANS and STATIONS are the kernel of the automatic timetabling system, *designer*

Figure 4: This figure illustrates two single-deck rolling-stock units with three and four carriages, respectively. A train consists of several of these units.

of network schedules (DONS). DONS also contains a database for storing the data and the obtained timetable, as well as a graphical user interface for the communication between the system and the user. Moreover, DONS provides an interface to the *simulation model for networks* (SIMONE), which allows us to evaluate the robustness of a cyclic timetable that we developed using DONS (Middelkoop and Bouwman 2001). Hooghiemstra et al. (1999) provide information about DONS.

Rolling-Stock Scheduling

Electrical train units, which can drive in either direction without a locomotive, operate most trains in The Netherlands. These units exist in different types (e.g., single-deck or double-deck) and subtypes. Figure 4 shows an example.

Train units of the same type can be combined to form longer trains with more seating capacity. The different subtypes within a type allow for much flexibility in the seating capacities of the trains. For example, with three- or four-carriage single-deck train units, all train compositions with a length of more than three carriages can be formed, except for compositions with five carriages. However, all trains have a specific length limit (typically, the maximum is 12 or 15 carriages).

The goal in scheduling rolling stock is to allocate an appropriate amount of the appropriate rolling-stock type to each train in the given timetable. In this context, the term *appropriate amount* means that the capacity of each train should be sufficient to transport the expected numbers of passengers. For example, an NS criterion is that each passenger must have a seat if the travel time is more than 15 minutes because seat availability is an important factor that passengers use in deciding whether to travel by train or not. However, if the capacity of a train exceeds the demand by too great an amount, this creates an inefficient situation. There is also a constraint on the number of train units

available; during peak hours, most trains will simultaneously require more units. A further complexity is that demand varies substantially during the day and on a line. For example, workdays have two peaks, one in the morning and one in the afternoon, with high travel demands in opposite directions. The rest of the day is the off-peak period with a lower, yet considerable, travel demand. To operate the rolling stock efficiently, NS addresses this demand-variation problem by adjusting the lengths of the trains during the day. This results in many shunting movements, which have a negative impact on the robustness of the real-time operations, primarily because the shunting movements use the same infrastructure as the regular trains.

Therefore, NS must always find a balance between three conflicting objectives in rolling-stock scheduling: (1) service, (2) efficiency, and (3) robustness. In this context, service means offering as many passengers as possible a seat. Efficiency aims at minimizing the amount of rolling stock and the number of rolling stock kilometers. NS addresses robustness by reducing the number of shunting movements and by having a line-based rolling-stock circulation. Preferably, the rolling stock per line is of a single type because this simplifies recovery if there is a disruption in the real-time operations.

Determining the order of the different train units in a train (Figure 5) is a difficult aspect of rolling-stock scheduling.

In this example, two trains depart from The Hague and Rotterdam with compositions AB and BA, respectively. A indicates a unit with three carriages, and B a unit with four carriages; the right character is the front unit of the train. Upon arrival in Utrecht, the trains are coupled onto each other within a few minutes; this results in a single train with composition ABBA. In Zwolle, the train is split again into one train bound for Groningen and one for Leeuwarden. Moreover, because the travel demand in the Northern part of the line is relatively low, the last train unit of the train is uncoupled in Zwolle. Thus, it is necessary to know the order of the units in the train. The train arrives in Groningen and Leeuwarden with compositions BA and B.

The model for *rolling stock allocation* (ROSA), which we developed for generating rolling-stock circulations, is basically an integer, multicommodity flow

Figure 5: This figure shows a rolling-stock plan on the line The Hague/Rotterdam–Utrecht–Zwolle–Leeuwarden/Groningen. A white wagon in the left picture corresponds to a unit with three carriages (A) and a black wagon to a unit with four carriages (B). In Utrecht, the train units from The Hague and Rotterdam are coupled to create one longer train. In Zwolle, they are uncoupled again. Moreover, one unit remains in Zwolle.

model. Each commodity represents a single subtype of train units. However, as we indicated earlier, the order of the units in a train is relevant; therefore, the model must consider it. Because a standard multicommodity flow model cannot handle this situation, we extended the model using the concept of the *transition graph*. For each trip of a train and for each feasible composition, this graph describes the allowed compositions on the next trip of that train. For example, if the composition is ABA, then the composition AB may be feasible on the next trip, but the composition AA is definitely not because it would require complex shunting operations. For each train, we must find a path in its transition graph.

We implemented ROSA using ILOG OPL Studio, a modeling environment for mathematical programs, and we solved it using CPLEX. Correctly describing the transition graph in terms of the decision variables was crucial for obtaining acceptable running times. Therefore, we could solve most instances within minutes, and the most complex instances within about one hour. These most complex instances involve situations in which complete trains are coupled or uncoupled, as the example above describes. Fiiole et al. (2006) provide additional details of the applied model.

Crew Scheduling

Each train in the timetable requires a train driver and a number of conductors. The latter depends on the

rolling-stock composition of the train. In the remainder of this paper, we will use the general term *crew* rather than the terms *train driver* and *conductor*. Approximately 6,000 crew members operate from 29 crew bases throughout the country. Each crew member belongs to a specific crew base. A *duty* starts and ends in a crew base and describes the consecutive trips for a single crew member. For each day, we generate a number of anonymous duties. *Rosters* prescribe how to assign the anonymous duties to individual crew members on consecutive days.

Within NS, the duties are planned centrally for all crew bases, while the crews at the bases generate the rosters themselves. *Crew scheduling* is the first step. It is inevitable that the crew scheduling step must be carried out centrally because most of the trips could be assigned to several crew bases. Moreover, the crew costs are determined in this step.

NS considers three important goals in crew scheduling: (1) efficiency, (2) acceptability, and (3) robustness. Efficiency means that the total number of duties is as small as possible. The objective of acceptability is enhanced via labor rules and company agreements, for example, on the amount of *variation* in the duties. Furthermore, specific rules within NS focus on a fair allocation of the *sweet* and *sour* workload among the bases. An example of the sweet part of the workload is a trip on a long-distance train, while work on a regional train is considered as sour. Abbink et al. (2005) provide additional details. Robustness of the crew duties, i.e., preventing propagation of delays via the crew schedule, depends on several elements, including the transfer times of the crews when transferring from one train to another.

We generate the crew schedules using the TURNI system; TURNI (Kroon and Fischetti 2001) includes an algorithm specifically developed to solve large crew-scheduling problems. The model behind TURNI is a set-covering model. In such a model, there is a binary decision variable for each potential duty (one if the potential duty is selected and zero otherwise). The problem is then to select a subset of duties from a predetermined set of feasible potential duties such that it covers each trip by at least one duty, it satisfies all additional constraints at the crew-base level, and the total costs of the selected duties are as low as possible. A duty is feasible if it satisfies all constraints at the

duty level, for example, the maximum duty length, the location, and the duration of the meal break. All duties of a base taken together must satisfy the additional constraints at the crew-base level. NS has too many of these constraints to mention them all; however, a maximum average duty length of eight hours is an example of such a constraint. Constraints related to the fair allocation among the bases of the sweet and sour amounts of work are also considered as base constraints.

Because trains stop at many stations, there are many options for changing the crew; this is possible at any major station on the route. This results in many trips and in many trips per duty; these numbers are typically much higher than in airline applications. Therefore, the number of feasible potential duties is extremely large. To manage this complexity, TURNI uses column generation, where a pricing model generates the feasible potential duties *on the fly* whenever they are needed. Thus, only the pricing model must consider the complex constraints at the duty level. Lagrangean relaxation, subgradient optimization, and several heuristics solve the resulting extended set-covering problem in a manner similar to that proposed by Caprara et al. (1999) for the pure set-covering problem. Kroon and Fischetti (2001) provide details.

A typical workday at NS includes approximately 15,000 trips for drivers and 18,000 for conductors. The resulting number of duties is approximately 1,000 for drivers and 1,300 for conductors. This leads to extremely difficult crew scheduling instances. Nevertheless, because of the highly sophisticated applied algorithms, TURNI solves these cases in 24 hours of computing time on a personal computer. Therefore, we can construct all crew schedules for all days of the week within just a few days.

Implementation and Challenges

Although the project to construct the new timetable started in June 2003, the development of the models and algorithms in the decision-support tools had begun several years earlier.

Implementing the OR Systems

In the 1990s, NS management recognized the great potential of applying operations research in the planning process. Initially, it preferred to buy automated

timetabling software. However, because no off-the-shelf packages were available, and external IT companies had failed to develop a prototype, NS decided to pursue an innovative approach. Its objective was to stimulate the development of new methods to solve the timetabling problem in close cooperation with the scientific community, in particular with the Centre for Mathematics and Informatics (CWI) and Erasmus University Rotterdam. After several years of research, NS (and following the NS split, NS in partnership with ProRail) implemented the methods developed as part of the DONS system. It has been in use for timetabling studies since the end of the 1990s.

After this successful project, NS management decided to establish an internal OR group within its Logistics Department. We started several new research projects, some together with the scientific community and others internally. The ROSA system was one such project. For crew scheduling, we performed a benchmark of commercially available systems; this resulted in the selection of TURNI. However, making it operational required a large amount of joint R&D work between NS and Double-Click sas, the supplier of TURNI. The system has been in operation since 2002.

There were three crucial factors in the success of this approach. The first was that some of us worked in the same department in which the planning problems were solved manually. Thus, it was easy for us to get data and knowledge about the real problems. The second factor was the existence of a central database containing the timetable and rolling-stock schedules. This database is an integrated system for registration and distribution of the manually created plans. For crew planning, we already had a manual planning system, CREWS, developed by the company SISCOG (Morgado and Martins 1998). We could easily use the central database and CREWS to find the right data for developing and testing our OR methods. Moreover, once the OR methods were proven to be successful, we could easily connect them to the central database and CREWS; thus, we could distribute the generated plans to operations via these systems. The third success factor was management support in challenging the OR experts to develop sophisticated solution methods and in investing in these projects without any guarantee of final success.

The New Timetable Construction

In the process of constructing the new timetable, we used DONS, ROSA, and TURNI intensively and, for the first time, together. In developing this timetable, we considered approximately 10 line systems; each varied radically from every other line system. For all these line systems, we used DONS to generate a one-hour timetable. Because the timetable is cyclic, a one-hour timetable can be repeated throughout the day. Therefore, we can base many evaluations of the complete timetable on the one-hour timetable. For example, we simulated the 10 one-hour timetables to determine their expected punctuality. Furthermore, the NS marketing department evaluated the consequences of these timetables for the passengers. Relevant criteria were the number of direct connections and the travel times. The outcome was the attractiveness of the timetables in terms of the expected passenger growth or decline. Finally, for each timetable, we estimated the operating costs related to rolling stock and crew.

In March 2005, the NS board decided to use 2 of the 10 developed timetables and asked for a timetable that combined the best aspects of both. The result was an 11th timetable.

During 2005, it became clear that two of the three earlier mentioned major infrastructure projects (the high-speed line and the Betuwe freight line) would not be finished in December 2006. Because these are separate lines, their lack of completion would be relatively minor: the timetable would require modification in only a few sections of the country. We performed these changes manually.

At the end of 2005, the NS board made the final decision to introduce the resulting timetable in December 2006. The process of planning the detailed rolling-stock schedules started in January 2006. The crew schedules were subsequently constructed.

Challenges During the Process

During 2006, many unexpected events happened. When the timetable plans were communicated to the public in spring 2006, many people reacted negatively to parts of the timetable. The first challenge was to address this negative reaction. It is inevitable that when one modifies the timetable of a complete country, the connections of some passengers worsen. In particular, politicians in the northern

provinces were unhappy because the travel times to and from these provinces increased by a few minutes, partly because of the opening times of a specific bridge. Therefore, the politicians began to lobby the Dutch parliament for a faster connection with the western part of the country. However, significant changes were not possible. Therefore, the NS board decided to decrease the travel times to the north by temporarily reducing some time supplements in that area; however, this had negative punctuality consequences. In June 2007, the opening times of the bridge were changed, thus providing a better solution. The planners directly adjusted these changes in the central database.

A second, more difficult challenge arose in August 2006, when it became apparent that the construction of the four tracks between Amsterdam and Utrecht would not be completed until four months after the start of the new timetable. At that time, we had already completed the rolling-stock schedules, and we were about to begin the crew scheduling process. This was a serious problem because the line between Amsterdam and Utrecht belongs to the kernel of the Dutch railway network (Figure 1). Postponing the introduction of the new timetable was not an option because of the limited time remaining and the national debate already underway regarding the timetable. We had to find another solution. Therefore, NS decided to introduce a temporary timetable between Amsterdam and Utrecht by canceling one regional train per hour and modifying some other regional trains. These modifications had a significant negative impact on the robustness of the timetable for this important part of the network, and also on the punctuality of the trains in the entire country. To include such significant modifications at such a late time in the planning process was only possible by using TURNI to construct the crew schedules. Previously, when the crew schedules were made manually, the crew scheduling process started at least a year ahead; it now became possible to start the crew-scheduling phase just three months before the introduction of the new timetable.

The third major challenge involved the expected shortage of crew capacity. In January 2006, we foresaw that the available crew capacity would not be sufficient to operate the increased number of trains in

the new timetable. The best solution was to construct more efficient crew schedules. We believed that this could be done by modeling some rules in a different way. For example, one rule prescribes that over an entire week the average length of all duties of a crew base should not exceed eight hours. In our previous computations, we always took this rule into account by limiting the average duty length per crew base for each day of the week to eight hours. Obviously, this is a tighter constraint than the actual rule. When we conducted a few experiments, we noticed that we could improve the crew schedules by applying the actual rule instead of the tighter rule. Therefore, we developed an extension to TURNI to address week instances. However, such instances are so huge that they cannot be solved in a single run. Therefore, we developed a procedure in which we used TURNI iteratively to solve instances of up to 15,000 trips (Abbink et al. 2007). This allowed us to reduce the number of duties by an additional 2 percent; this was sufficient to operate the new timetable.

The introduction of the new timetable on December 10, 2006 went smoothly. Initially, NS permanently monitored the operations to detect any initial problems with the timetable and the rolling-stock circulation, and, in particular, to verify that each train had a good match between seat supply and demand. Furthermore, based on a detailed analysis of the operation, NS applied several relatively minor modifications to the timetable. The final timetable, with the four parallel tracks between Amsterdam and Utrecht, started in April 2007.

Portability

Other railway companies in Europe face challenges that are similar to those NS faced in The Netherlands. We suggest that countries with highly utilized railway infrastructures consider the construction of a new timetable from scratch. Because many of these countries use a cyclic timetable, they can also use the approach we described in this paper: the models within DONS are generic models for solving cyclic timetabling problems.

Furthermore, because of the liberalization of the European railway market, many railway companies are interested in tools for optimizing their resource

schedules. New rolling-stock scheduling systems based on the ideas described in this paper are currently under development. Moreover, another European railway company is using TURNI, the crew scheduling system. Finally, NS has used the tools described to bid for contracts to operate some lines abroad; several of its bids have been successful.

Impact and Success

As we mentioned above, railway transport is a critical transport mode in The Netherlands. Hundreds of thousands of daily commuters use the train, and millions of people use the railway system regularly. It is obvious that when a modification in the railway timetable changes so many daily-life patterns, it will cause much media attention and many discussions in the Dutch parliament. Communicating the need for this dramatic change was not easy. Therefore, NS asked Johan Crujff, the most distinguished Dutch soccer player ever, to discuss the advantages of the new timetable in television commercials. In one commercial, he talked about how frequently the trains in the railway system run: "...when you just missed your train, you are always in time for the next one." He also compared the railway timetable with soccer-game tactics.

After more than a year of operating the new timetable, we can measure its success. NS had an all-time high number of passengers in 2007. When we made a detailed analysis of all individual routes, we discovered that the routes on which we put more trains into service had a much higher increase (as high as 15 percent) in passengers than the average (which was 2.8 percent). Overall, we expect that approximately two percentage points of the long-run passenger increase will be because of the new timetable.

Moreover, in 2007, train punctuality reached a record high: 87 percent of the trains arrived within three minutes of their scheduled arrival time; in both 2005 and 2006, this percentage was 84.8, more than two percentage points less. This is even more remarkable because the punctuality over the first four months of 2007 was almost at the same level as in 2006 because of the delayed opening of the four tracks between Amsterdam and Utrecht. If we replace the first four months of 2007 by the same months of 2008, the punctuality is 87.5 percent.

The NS marketing department conservatively estimated that the changes in the new timetable generated an additional annual profit of €10 million in 2007, which it expects will increase to €20 million in 2009 and later years. One factor in both estimations was an expected punctuality increase of 1.5 percentage points.

Several years ago, NS made an agreement with consumer organizations that would allow a bonus fare increase of 2 percent if NS could achieve a record annual average punctuality level (86.8 percent). Normally, NS is only allowed to increase fares on a par with the inflation rate. NS achieved this high punctuality record in 2007 and was permitted the additional fare increase as of February 2008. The result was an additional annual profit of about €20 million from 2008 onwards.

Moreover, ROSA and TURNI led to more efficient resource schedules. We estimate the savings of the optimized rolling-stock schedules over manually constructed schedules to be 6 percent. NS achieved these savings by the introduction of ROSA on part of the long-distance network in 2005. At that time, we compared the schedules that ROSA generated with those that were manually generated. Six percent corresponds to an annual savings of €18 million. NS invested these savings primarily in improving the seat availability for the passengers, resulting in higher customer satisfaction.

We estimate the benefits of TURNI and its extensions at another €12 million per year. TURNI has been in use since 2002. In its first year of operation, we compared the automatically generated crew schedules with the manual ones. Applying exactly the same rules, we obtained an improvement of 2 percent. By improving the algorithm in the following years, we gained another 2 percent improvement. Overall, we have been able to reduce the number of drivers per train-kilometer by approximately 15 percent because we were able to construct schedules that permitted adjustments to labor standards and regulations. It is clear that without using TURNI we could not have achieved this effect; however, it is hard to measure which part of the additional 11 (15 – 2 – 2) percent is the TURNI contribution.

Adding up all the quantifiable benefits, we find that our total annual additional profit is approximately €70 (20 + 20 + 18 + 12) million (\$105 million).

Finally, the new timetable is having a positive impact on the Dutch society as a whole. The independent advisor of the Dutch government, Centraal Planbureau (CPB), estimates the direct benefits to the Dutch economy to be at €8 million per year for every percentage-point increase in punctuality. More importantly, because of the new timetable, an additional increase in railway transport will be possible without further significant infrastructure extensions, which would require huge investments. For example, doubling the number of tracks on a route of approximately 40 kilometers would cost about €1 billion (\$1.5 billion). The new timetable will be able to accommodate additional railway transport in the future with only limited infrastructure additions; therefore, the railway system will be able to facilitate the growing demand for transportation to the main cities during rush hours. This will help to reduce the pressure on the roads into and inside these cities and to replace car traffic by rail traffic, thereby reducing the pollution from greenhouse gases.

Appendix

In this appendix, we give an example of a PESp constraint. Suppose that the running time between stations A and B is 19 or 20 minutes. Then, denote the departure time in A as d_A and the arrival time in B as a_B . The PESp constraint reads then as follows:

$$19 \leq a_B - d_A + 60 * q_{AB} \leq 20.$$

In this formula, q_{AB} denotes the binary decision variable indicating whether the departure in A and the arrival in B are in the same hour or in consecutive ones.

Acknowledgments

Many people participated in the construction of the new timetable. Although we thank all of them, there are too many to name individually. However, we would like to thank one person in particular: Tjeu Smeets, head of the Logistics Department of NS from 1992 to 2004. He initiated the cooperation between NS and the scientific community, and he always encouraged the development of new operations research models for supporting the NS logistic railway processes. Without his strategic vision, the new timetable and all its benefits would not have been realized.

References

- Abbink, E., J. Van't Wout, D. Huisman. 2007. Solving large scale crew scheduling problems by using iterative partitioning.

- Retrieved September 9, 2008, <http://drops.dagstuhl.de/opus/volltexte/2007/1168/pdf/07001.AbbinkErwin.Paper.1168.pdf>.
- Abbink, E., M. Fischetti, L. G. Kroon, G. Timmer, M. J. C. M. Vromans. 2005. Reinventing crew scheduling at Netherlands Railways. *Interfaces* 35(5) 393–401.
- Caprara, A., M. Fischetti, P. Toth. 1999. A heuristic algorithm for the set covering problem. *Oper. Res.* 47(6) 730–743.
- Caprara, A., L. G. Kroon, M. Monaci, M. Peeters, P. Toth. 2007. Passenger railway optimization. C. Barnhart, G. Laporte, eds. *Transportation*. Elsevier, Amsterdam, 129–187.
- Fioole, P.-J., L. G. Kroon, G. Maróti, A. Schrijver. 2006. A rolling stock circulation model for combining and splitting of passenger trains. *Eur. J. Oper. Res.* 174(2) 1281–1297.
- Hooghiemstra, J. S., L. G. Kroon, M. A. Odijk, M. Salomon, P. J. Zwaneveld. 1999. Decision support systems support the search for “win-win” solutions in railway network design. *Interfaces* 29(2) 15–32.
- Huisman, D., L. G. Kroon, R. M. Lentink, M. J. C. M. Vromans. 2005. Operations research in passenger railway transportation. *Statistica Neerlandica* 59(4) 467–497.
- International Union of Railways (UIC). 2007. *International Railway Statistics 2005*. UIC, Paris.
- Kroon, L. G., M. Fischetti. 2001. Crew scheduling for Netherlands Railways “Destination: Customer.” S. Voß, J. R. Daduna, eds. *Computer-Aided Scheduling of Public Transport*. Springer, Berlin, 181–201.
- Middelkoop, D., M. Bouwman. 2001. SIMONE: Large scale train network simulations. B. Peters, J. Smith, eds. *Proc. 2001 Winter Simulation Conf.*, Institute of Electrical and Electronic Engineers Press, Piscataway, NJ, 1042–1047.
- Morgado E. M., J. P. Martins. 1998. CREWS_NS: Scheduling train crews in The Netherlands. *AI Magazine* 19(1) 25–38.
- Nederlandse Spoorwegen (Netherlands Railways—NS). 2008. Retrieved September 9, 2008, <http://www.ns.nl/cs/Satellite/ns2007/nl/artikel/include/1198251860645/energie+en+klimaat?p=1160724828590>.
- Schrijver, A., A. Steenbeek. 1993. Spoorweg Dienstregelingontwikkeling (Railway timetable development). Retrieved September 9, 2008, <http://homepages.cwi.nl/~lex/files/rapp1a.pdf>.
- Serafini, P., W. Ukovich. 1989. A mathematical model for periodic event scheduling problems. *SIAM J. Discrete Math.* 2(4) 550–581.
- Zwaneveld, P. J., L. G. Kroon, C. P. M. Van Hoesel. 2001. Routing trains through a railway station based on a node packing model. *Eur. J. Oper. Res.* 128(1) 14–33.