

Disruption Management of Rolling Stock in Passenger Railway Transportation

Lars Kjaer Nielsen and Gabor Maroti

ERIM REPORT SERIES <i>RESEARCH IN MANAGEMENT</i>	
ERIM Report Series reference number	ERS-2009-046-LIS
Publication	August 2009
Number of pages	28
Persistent paper URL	http://hdl.handle.net/1765/16557
Email address corresponding author	lnielsen@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	<p>This paper deals with real-time disruption management of rolling stock in passenger railway transportation. We present a generic framework for modeling disruptions in railway rolling stock schedules. The framework is presented as an online combinatorial decision problem where the uncertainty of a disruption is modeled by a sequence of information updates. To decompose the problem we propose a rolling horizon approach where only rolling stock decisions within a certain time horizon from the time of rescheduling are taken into account. The schedules are then revised as the situation progresses and more accurate information becomes available. We extend an existing model for rolling stock scheduling to the specific requirements of the real-time case and apply it in the rolling horizon framework. We perform computational tests on instances constructed from real life cases and explore the consequences of different settings of the approach for the trade-off between solution quality and computation time.</p>
Free Keywords	<p>passenger railway transportation, disruptions, combinatorial decision problem</p>
Availability	<p>The ERIM Report Series is distributed through the following platforms:</p> <p>Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal</p> <p>Social Science Research Network (SSRN), SSRN ERIM Series Webpage</p> <p>Research Papers in Economics (REPEC), REPEC ERIM Series Webpage</p>
Classifications	<p>The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems:</p> <p>Library of Congress Classification, (LCC) LCC Webpage</p> <p>Journal of Economic Literature, (JEL), JEL Webpage</p> <p>ACM Computing Classification System CCS Webpage</p> <p>Inspec Classification scheme (ICS), ICS Webpage</p>

Disruption Management of Rolling Stock in Passenger Railway Transportation

Lars Kjær Nielsen, Gábor Maróti

Rotterdam School of Management, Erasmus University Rotterdam,
P.O. Box 1738, 3063 HD Rotterdam, the Netherlands.

Email: {lnielsen, gmaroti}@rsm.nl

Abstract

This paper deals with real-time disruption management of rolling stock in passenger railway transportation. We present a generic framework for modeling disruptions in railway rolling stock schedules. The framework is presented as an online combinatorial decision problem where the uncertainty of a disruption is modeled by a sequence of information updates. To decompose the problem we propose a rolling horizon approach where only rolling stock decisions within a certain time horizon from the time of rescheduling are taken into account. The schedules are then revised as the situation progresses and more accurate information becomes available. We extend an existing model for rolling stock scheduling to the specific requirements of the real-time case and apply it in the rolling horizon framework. We perform computational tests on instances constructed from real life cases and explore the consequences of different settings of the approach for the trade-off between solution quality and computation time.

1 Introduction

The planning and operation of a busy passenger railway network is a complex task. Rolling stock and crew have to be scheduled to serve the timetable with ever growing demand for capacity. This has led to extensive research on optimizing the utilization of railway resources such as the infrastructure, the rolling stock and the crew (see Caprara et al. [3]). The developed methods have resulted in resource schedules that are highly effective when operations run as planned. However, when operating a dense timetable with many passengers, operations may occasionally have to deviate from the plans. This results in less efficient schedules in the actual operations, and in more serious cases it turns the plans infeasible. Currently such deviations are handled manually. There is a demand for decision support systems for effectively dealing with the challenges posed by the real time operation of a passenger railway system. However, to this date only limited research has been conducted on the real time rescheduling of railway resources.

The challenges in real time decision making differ from the static planning tasks in a number of aspects. The decisions must be made within a tight time frame and information is often uncertain as the environment evolves (Séguin et al. [12] and Grötschel et al. [5]). Furthermore, a highly optimized plan already exists and any changes to that plan will have to be communicated to the involved parties.

Railway resource planning is traditionally divided into four stages; strategic, tactical, operational and short-term (see Huisman et al. [6]). The first three involve different aspects of early planning with time horizons of several weeks up to years. Short-term planning involves decisions with up to a few days horizon and includes real-time planning. The planning of the timetable, rolling stock and crew is divided into separate tasks though the involved resources are interdependent. This way the resulting instances are manageable. Similarly, the real time control and rescheduling of the timetable, rolling stock and crew is performed separately. This paper deals with the real time rescheduling of rolling stock.

The problem and instances considered in this paper come from Netherlands Railways (NS). NS is the major passenger railway operator in the Netherlands. They transport more than 90% of the daily passenger traffic in the country. Figure 1 shows the infrastructure of the railway network of the Netherlands. The solid lines are operated by NS and the dotted lines are operated by other operators.




Figure 1: The infrastructure of NS.

This paper is structured as follows. In Section 2 give a short literature overview and set the terminology we use. Section 3 describes the rolling stock rescheduling process from a practical point of view. The process at NS is illustrated by a detailed example of a disruption.

In Section 4 we formulate a generic framework for the rolling stock rescheduling problem and for its online variant. The framework is generic in that it can

be based on any particular set of constraints in the underlying rolling stock scheduling problem. Also, we analyze the competitiveness properties of the online problem.

The generic framework is specified to the instances of NS in Section 5 by extending the existing rolling stock scheduling model of Fioole et al. [4] for the special needs of real-time planning. The basic model of Fioole et al. [4] itself is a linear integer programming model which has successfully been used by NS for medium-term planning since 2004.

Real life rescheduling has to face uncertainties about the future. To deal with that, we propose a rolling horizon approach in Section 6. In many aspects, it models the way planners at NS currently cope with disruptions.

Section 7 is devoted to the computational results on real-life instances of NS. Finally, we draw some conclusions in Section 8.

2 Literature and terminology

There is a rich list of publications that address rolling stock planning problems with time horizons of a couple of weeks or a few months. We refer to Caprara et al. [3] for a comprehensive overview. We mention here Alfieri et al. [1], Fioole et al. [4] and Peeters and Kroon [11]. The problems considered in these papers have the same underlying structure and assumptions as the rescheduling problems considered in this paper. They model the problem of assigning rolling stock units to timetabled train services while optimizing the capacity, robustness and efficiency of the schedules. The first two mentioned papers formulate the problem as Mixed Integer Programming (MIP) models and apply commercial solvers. The latter uses a branch-and-price approach, but can only solve the problem for simple line structures.

A number of publications address the problem of dispatching trains through a dense railway network. Törnquist [13] provides a recent overview of the research in railway scheduling and dispatching, and classifies the surveyed models with respect to problem type, solution mechanism and type of evaluation. However, the literature is primarily concerned with the timetable services, i.e. dispatching of trains and the management of delays.

Jespersen-Groth et al. [8] offers a comprehensive overview of the practical aspects of disruption management in passenger railway transportation. The literature on real time rescheduling of rolling stock is scarce, we mention here a number of publications that consider problems similar to the one addressed in this paper. Jespersen-Groth and Clausen [7] consider the problem of reinserting canceled train lines after a disruption. They formulate a MIP model which minimizes train cancellations on the disrupted lines.

Walker et al. [15] present an integrated model for train and crew recovery in the event of disruptions. The approach incorporates the complex crew constraints into the timetable modification decisions and minimizes the train idle times. The model does not explicitly consider rolling stock as it does not contain possibilities for changing the assignment of rolling stock to the trains.

Budai et al. [2] consider the rolling stock balancing problem. This is the problem of rescheduling the rolling stock so that it ends up at the correct stations at the end of the day. Their models take the exact order of the rolling stock units in the trains into account and minimize the number of rolling stock units that end up at different stations than planned. The rescheduling is subject to a number of constraints on capacity demand and shunting possibilities. The problem is solved for fairly large instances by iterative heuristics whose short running times are attractive for real time usage. However, the model does not take the stochastic nature of real time operations into account.

2.1 Railway rolling stock terminology

In this section we introduce some terminology and features of the rolling stock operations in a line based passenger railway system. For more details see Fioole et al. [4].

Features of the rolling stock operations

The passenger train services at NS are primarily operated by electrically powered *rolling stock units*. A unit consists of a fixed number of carriages and needs no locomotive to be operated. Units are divided into types depending on their characteristics and units of the same type are considered interchangeable during operations.

Rolling stock units can be combined to form *compositions*. This way capacity can better meet demand during peak hours. An essential feature of NS is the fact that, due to efficiency reasons, train compositions are adapted during operations by uncoupling units from or coupling units to trains.

A *trip* is a train service from a departure station to an arrival station at a specific time. A timetable is a set of trips. At NS, the timetable is cyclic which means that the same one hour timetable is repeated.

Each trip has fixed rolling stock connections to the next trip to be carried out. This means that when a train arrives at the arrival station of a trip, it is fixed which trip is to be carried out next. However, the composition of the train may be adapted between trips. The fixed rolling stock connections exist both along the line and at stations at the end of the line. Here an incoming trip is matched with a specific outgoing trip. This is known as *turning* the train. Two trips that form a rolling stock connection are called *consecutive*.

The non-timetabled movements of rolling stock units inside railway nodes is called *shunting*. This includes the uncoupling of units from trains and the coupling of units to trains which is also known as *composition changes*. Usually only certain types of composition changes are allowed. For example the standard allowed composition changes are to couple up to two units to the front end of the train or to uncouple up to two units from the rear end of the train. But this depends on the layout of the station.

The exact details of the shunting operations are locally planned and monitored. This means that the full information on the track occupancy inside

stations and shunting yards is not globally available during operations. As a consequence, any changes to the planned shunting operations required by the dispatchers must be accepted by local planners to ensure local feasibility.

The *duty* of a rolling stock unit is the set of trips that the unit will service along with the position of the unit in the composition of each trip. The assignment of all rolling stock units to duties is called the *rolling stock circulation*. According to the duties of the units each unit will end up at a specific station at the end of the day. The number of units of each type that end their duties at each station is called the *rolling stock balance*. The balance connects the rolling stock circulation to that of the next day. The number of units of each type that is available at a station at any given time is called the *inventory* of the station at that time.

3 Real time control in passenger railways

In this section we describe the practical aspects of the real time monitoring and rescheduling of rolling stock. We discuss the options that are open to dispatchers for rescheduling the rolling stock during operations. All specific examples are from NS. For a comprehensive overview of the disruption management process in passenger railways we refer to Jespersen-Groth et al. [8].

3.1 The real time control process

The resource plans that have been constructed through the various planning steps are highly optimized toward a number of managerial goals. These plans are to be carried out in practice. The execution of the plans involves monitoring the positions and movements of all resources, and reacting to any deviations from the plans.

Such deviations from the operational plans may require rescheduling of some resources. Naturally, some deviations are more serious than others, in practice, there is a distinction between minor and major incidents. The distinction is purely practical and not well defined, but depends on the impact on the operations. Though as a rule of thumb, major incidents require significant changes of the pre-set resource schedules.

A *disturbance* is an event that causes part of the railway operations to deviate from the operational resource plans. For example the boarding of passengers may take longer than planned at a station for a certain train leading to a delay in departure. Disturbances are usually absorbed by the slack in the system or can be handled by minimal changes to the resource allocation.

More serious incidents are known as *disruptions* and these incur major deviations from the planned operations. Disruptions may be caused by various internal or external factors such as a faulty switch on a busy track, broken down rolling stock, or damaged overhead wires. In a disrupted situation the planned resource schedules are no longer feasible and will have to be updated to take the actual situation into account.

The overall rescheduling process in disrupted situations is sequential with respect to the resources involved. First, the timetable is updated – trains may be canceled due to blocked trajectories or failing rolling stock. Second, the rolling stock circulation is rescheduled to cover as many as possible of the trips in the updated timetable. Third, the crew schedules are replanned to ensure that a driver and a conductor are assigned to each train. The involved rescheduling tasks are interconnected in the sense that each step in this process may turn out to be infeasible. This means that the result of the previous step must be changed. Still, the rescheduling problems are solved independently taking the output of the previous step as input. The rolling stock rescheduling thus assumes that an updated timetable is available and the later crew assignment can be carried out independently. For a detailed discussion of real time operations and disruption management in passenger railways we refer to Jespersen-Groth et al. [8].

A complicating issue in the process is the inherent uncertainty related to disruptions: It is usually not known at the occurrence of a disruption how long the involved resources will be unavailable – at best, a forecast is available. This means that the results of any step in the rescheduling process can turn out to be infeasible when more information on the resource availability becomes known.

A further complicating issue is the fact that the exact information on some resources is decentralized. This includes the local shunting possibilities at the stations which is locally planned and monitored. The rescheduling of resources is a task that needs global coordination, so any changes to the resource allocations that have local consequences must be communicated to local operations controllers. Figure 2 shows a simplified schematic view of the entities involved in the real time operations of the rolling stock.

The local operations controllers implement local changes to the shunting processes and regional changes to the rolling stock allocation. The Rolling Stock Managing Center is a centralized unit, it is responsible for real time monitoring and dispatching, and ensures global coordination of the rolling stock usage. The center communicates with the local operators to find feasible solutions when significant changes are needed in rolling stock allocations. Currently the analysis and rescheduling performed by the Rolling Stock Managing Center is done without sophisticated computerized support and primarily relies on the skills and experience of the involved planners.

The rescheduling framework proposed in this paper is meant to be used by the planners at the Rolling Stock Managing Center in an iterative procedure by feeding the system with input and objectives and evaluating the output. If the output has undesirable features or is infeasible according to local planners, the input can be tweaked by penalizing or forbidding these features. For such a system to be accepted by practitioners, the running time must be low – preferably a solution must be available within seconds.




Figure 2: A simplified schematic view of the real time control process.

3.2 Rolling stock rescheduling in practice

During irregular operations it may be impossible to run trains according to the published timetable. A necessary possibility is the option of canceling trips. Note that operators distinguish between two different causes for canceling trips; the first reason is external factors such as failing infrastructure. These force the operators to cancel trips on the involved trajectories and adapt the resource allocations accordingly. Such cancellations are part of the timetable updating process. The adaptation of the timetable is performed according to a *handling scenario* which states how to change rolling stock connections between trips during a given disruption. The second reason is internal factors such as lack of necessary resources to service the trips. These cancellations are caused by the rolling stock rescheduling process or the crew rescheduling process.

When trips are canceled, the rolling stock connections between some trips may be changed. This means that the coupling-uncoupling pattern is infeasible and the planned composition changes must be changed. However, since shunting is planned locally the dispatchers must communicate the changes to the local controllers. Different changes to the shunting plans have different impact on the operational process and have different probabilities of being accepted by the local controllers. For example, suppose that no coupling or uncoupling was originally planned between two consecutive trips. Then introducing a shunting operation between these two trips requires major modifications of the local shunting plans. However, canceling a planned coupling of a unit has less impact on the local shunting plans but will still have to be communicated to the involved local controllers.

The scope of the rolling stock rescheduling is usually at most until the end of the day. At that time the resources are no longer in service and the system must be ready for the operations of the next day. When the operations deviate from the plan some units may end up at different stations than planned. This may have undesirable consequences for the operations of the next day or it may lead to necessary repositioning of trains during night time. A deviation of one unit from the planned number of units of a specific type ending up at a station is known as an *off balance*. Units of the same type are generally interchangeable with respect to end-of-day balances so only deviations in the number of units are counted as off balances.

3.3 Objectives of rolling stock rescheduling

The primary goal in the real time operations is to keep the circulation feasible. In addition to that, there are several objectives to be taken into account in the process. The objectives represent different aspects of the overall managerial objective and are often conflicting and difficult to quantify. The following lists the different perspectives of the real time process and the objectives associated with them:

The service perspective: The inconvenience for passengers should be minimized which means that canceling trips should be avoided. Assigning too low capacity on trips with high expected demand is undesirable as well.

The robustness perspective: The propagation of the effects of the disruption should be kept local and thus the changes to the rolling stock assignment should be minimized. The shunting operations i.e. coupling and uncoupling decrease the robustness of the system so the number of shunting operations should be kept moderately low.

The process perspective: The deviation from the original plan should be minimized as changes must be communicated to the involved parties. Especially rolling stock that ends up at a different station at the end of the day than planned may propagate the disruption to the next day.

The efficiency perspective: The number of carriage kilometers driven is closely related to the operational cost of the rolling stock schedule. Nightly repositioning of train units incurs a cost, as well.

These objectives cannot be achieved directly. Some of the objectives are conflicting like the robustness considerations and the efficiency. Other criteria are difficult to quantify due to lacking real time information. For example, the precise numbers of passengers on trains are not available in real time – at best a forecast based on historical data is available. Therefore only a subset of the mentioned criteria are taken into account here. We use three objective components that represent the central aspects of the mentioned objective perspectives.

First, cancellations are penalized. This criterion has a very high priority in practice.

Second, changes to the coupling-uncoupling pattern need to be communicated to the local shunting operators and there is a chance they cannot be implemented due to local resource constraints. Some of such changes require a substantial modification of the shunting process. Suppose, for example, that no coupling or uncoupling was planned originally between two consecutive trips. Then, introducing coupling or uncoupling has a fair chance to turn out impossible, therefore this modification of the shunting plan is penalized heavily. On the other hand, canceling a planned shunting operation or (un)coupling a different type of unit is unlikely to have major consequences at the local level, so this is only lightly penalized.

Third, off balances may have consequences for the circulation of the next day or they may be necessary to resolve by repositioning a train during the night. In either case, off balances are undesirable and are thus penalized.

These three criteria together, minimize canceling trips, changes to the shunting plans and off balances, capture the essence of the overall managerial goal for the disruption management.

3.4 An example of a disruption

In this section we describe an example of the rolling stock rescheduling problem. The example involves the 3000 line of NS.

3.4.1 The 3000 line

The 3000 line is an intercity line of NS from Den Helder to Nijmegen. The trains operating the line call at a number of major stations on the line including Den Helder (Hdr), Schagen (Sgn), Alkmaar (Amr), Amsterdam (Asd) and Nijmegen (Nm). The line is operated with units that are available in lengths of 4 and 6 carriages.

The timetable for the 3000 line is cyclic and has trains in both directions every half hour. The *time-space diagram* for the timetable of the northern part of the line is seen in Figure 3. The diagram shows stations in the vertical dimension and time in the horizontal dimension. The trips in the timetable are shown as diagonal lines between stations. The diagram also shows how incoming trains are matched with outgoing trains at the end of the line in Den Helder.

3.4.2 Disruption on the 3000 line

In this example, a disruption occurs at 12:45: The tracks cannot be used in either direction between Schagen and Alkmaar. At the occurrence, the infrastructure is estimated to be unavailable for one and a half hours. The trips between Schagen and Alkmaar are canceled and the timetable is updated according to a handling scenario. In the updated timetable trains are turned at each side of the disrupted region (i.e. in Schagen and Alkmaar). The updated timetable for the estimated duration of the disruption is shown in Figure 4.

However, at time 13:30 it becomes clear that the necessary repairs to the infrastructure will take an hour longer than first estimated. This leads to a new updated timetable with two more canceled trips in each direction (see Figure 5). The new update of the timetable is a forecast of the development of the disruption, and it may turn out to need adjustment later.

3.4.3 Challenges for the rescheduling

The original circulation of rolling stock units determines where and when units are coupled to or uncoupled from trains during the day. However, the original circulation is not feasible for the updated timetable when the disruption occurs

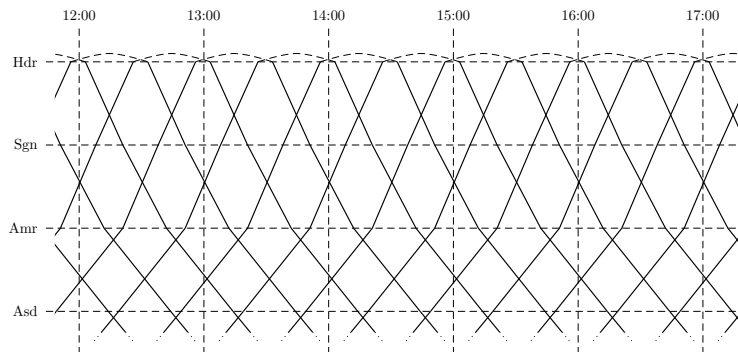


Figure 3: Time-space diagram for part of the timetable for the 3000 line.




Figure 4: Time-space diagram for part of the timetable for the 3000 line. There is a disruption between Schagen and Alkmaar. The estimated length of the disruption is one and a half hours. The bold line is a train whose path differs from the original timetable.




Figure 5: Time-space diagram for part of the timetable for the 3000 line. The length of the disruption is two and a half hours. The bold line is a train whose path differs from the first updated timetable.

and consequently a number of the planned shunting operations may not be possible to carry out as originally planned.

Originally, the train departing from Den Helder at 12:33 is supposed to arrive in Alkmaar at 13:09. But due to the disruption it is turned at Schagen station. In Figure 4 the path of the train is marked with a bold line and as it can be seen, the train now serves different trips in the timetable. At the same time, other trains take over the original trips of this one and some of the couplings and uncouplings meant for the original train must be changed or canceled. Thus any planned coupling or uncoupling of units may not be possible to carry out at the planned time instants.

The rolling stock circulation is rescheduled according to the updated timetable, yet it is still uncertain if the disruption will last shorter or longer than first estimated. When it becomes clear at 13:30 that the repairs to the infrastructure will last longer than expected, the recently updated circulation must be changed again. In Figure 5 the before mentioned train now has had its path changed again.

4 The Rolling Stock Rescheduling Problem

In this section we present a generic framework for modeling disruptions in the rolling stock schedules. The rolling stock rescheduling process is considered separate from the updating of the timetable, so the timetable changes are input for the rolling stock rescheduling. We first define the Rolling Stock Rescheduling Problem (RSRP). It is the problem of rescheduling the current circulation of rolling stock to a changed timetable. We then extend the problem to the online version called Online RSRP. Here several changes to the timetable may occur over time and the circulation has to be rescheduled accordingly every time.

4.1 RSRP

An instance of RSRP consists of the following elements:

- The original timetable \mathcal{T}_0 which consists of a set of trips including the rolling stock connections between them.
- The original rolling stock circulation \mathcal{C}_0 which is an assignment of the available rolling stock units to the trips in \mathcal{T}_0 .
- A time instant t at which it becomes known that the timetable is changed.
- The updated timetable \mathcal{T}' which becomes known at time instant t and which is to be served from t till the end of the day.

The problem is to modify circulation \mathcal{C}_0 to serve timetable \mathcal{T}' using the available rolling stock. The rolling stock assigned to trips in \mathcal{T}_0 that depart before time t cannot be changed. The resulting circulation is called \mathcal{C}' . The same constraints as in early planning apply to the rolling stock assignment.

The objective of the RSRP measures the deviation of the rescheduled circulation \mathcal{C}' from the original circulation \mathcal{C}_0 by weighing the objective criteria for canceled trips, changes to the shunting processes and off balances. The definition of the objective function is intentionally vague as it is very application dependent. The objective function is discussed in detail in Section 5.2.

4.2 Online RSRP

Here we define the online version of RSRP. In this version the timetable may be changed several times. Online RSRP is the problem of updating the current assignment of rolling stock to trips in the timetable whenever the timetable is changed. More formally, an instance of online RSRP consists of the following elements:

- The original timetable \mathcal{T}_0 .
- The original rolling stock circulation \mathcal{C}_0 .
- A finite list of changes to the timetable,

$$\langle t_1, \mathcal{T}_1 \rangle, \dots, \langle t_n, \mathcal{T}_n \rangle.$$

An element in the list is a pair consisting of a time instant t_i and an updated timetable \mathcal{T}_i . We assume that $t_1 < \dots < t_n$. Each element in the list represents a time instant where new constraints are added to the problem which possibly renders the current solution infeasible. The list thus represents the uncertainty of the real time operations.

The task is then at time instant t_i to reschedule circulation \mathcal{C}_{i-1} to serve timetable \mathcal{T}_i with the available rolling stock. The resulting circulation is called \mathcal{C}_i . At time instant t_i , the rolling stock assigned to trips in \mathcal{T}_i departing before t_i is fixed. At this time there is no knowledge of future changes in resource availability, i.e. the rest of the list $\langle t_{i+1}, \mathcal{T}_{i+1} \rangle, \dots, \langle t_n, \mathcal{T}_n \rangle$ is not known.

The objective of online RSRP measures the deviation of the final circulation \mathcal{C}_n from the original circulation \mathcal{C}_0 by weighting the objective criteria.

We note that it can be proven that no competitive online algorithm exists for generic Online RSRP. The proof is straightforward but we omit it in this paper.

This way of modelling the progression of a disrupted situation is also known as a wait-and-see approach. In this approach, no assumptions are made on the probability distribution of the length of the availability of resources.

4.3 The disruption example revisited

We here return to the disruption example from Section 3.4. The disruption can be modeled as an instance of the online RSRP the following way: The original timetable \mathcal{T}_0 for the 3000 line is given; part of it is shown as a time-space diagram in Figure 3. The available rolling stock is given: 11 units of length 4

and 20 units of length 6 are used. The units are assigned to trips in \mathcal{T}_0 . The assignment of units is given as the circulation \mathcal{C}_0 .

The disruption between Alkmaar and Schagen starts at 12:45 and the information on the duration of the disruption changes at 13:30. This is modeled as a list of changes to the resource availability containing the following two elements:

- First element is $\langle t_1, \mathcal{T}_1 \rangle$ where $t_1 = 12:45$, \mathcal{T}_1 is equal to \mathcal{T}_0 except from the canceled trips according to Figure 4.
- Second element is $\langle t_2, \mathcal{T}_2 \rangle$ where $t_2 = 13:30$, \mathcal{T}_2 is equal to \mathcal{T}_0 except from the canceled trips according to Figure 5.

The task is then to reschedule the rolling stock at time t_1 to serve timetable \mathcal{T}_1 without the knowledge of the later update $\langle t_2, \mathcal{T}_2 \rangle$. That information becomes available at time t_2 , where the rolling stock must be rescheduled to serve \mathcal{T}_2 .

4.4 Theoretical competitiveness properties

The following definition by Karlin et al. [9] measures the performance of an online algorithm by comparing it to the optimal offline algorithm, OPT, which knows the entire input sequence in advance.

Definition 1. *Let \mathcal{I} be the set of all instances for an online minimization problem P . The competitive ratio of an online algorithm \mathcal{A} for P is defined as the smallest k such that for some constant α ,*

$$\mathcal{A}(I) \leq k \cdot \text{OPT}(I) + \alpha \quad \forall I \in \mathcal{I}.$$

An online algorithm \mathcal{A} is said to be k -competitive if the competitive ratio is k . If the algorithm is k -competitive for some constant k , the algorithm is said to be competitive.

We shortly note the following proposition for the online RSRP.

Proposition 2. *No online algorithm for the online RSRP is competitive.*

An informal sketch of the proof is given here: Consider a timetable with three trips from station A to station B where one rolling stock unit is assigned to each trip. If the timetable is changed so that the second trip is canceled then the three rolling stock unit must be assigned to the first and the third trip to avoid off balances. If at least one is assigned to the third trip, then a new timetable update with a cancellation of that trip will lead to off balances. If no unit was assigned to the third trip it leads to a canceled trip due to rolling stock shortage. The offline algorithm, which knows the timetable updates beforehand, will distribute the rolling stock units on the trips that are not canceled thus leading to a solution with no off balances or cancellations due to lack of rolling stock. To overcome the additive term α in Definition 1, the construction can be copied an arbitrary number of times. The result holds for randomized online algorithms too as either the expected number of off balances or the expected number of cancellations due to rolling stock shortage will be positive.

The timetable and the updates used in the proof of Proposition 2 are rather unrealistic and only serve a theoretical purpose. Nevertheless, the result does show that there exist instances where any online algorithm may be arbitrarily worse than the optimal offline algorithm.

5 A model for RSRP at NS

RSRP, as it is described in Section 4, is an *abstract* framework for modeling disruptions in the rolling stock schedules. In this section we describe an existing rolling stock circulation model which has been developed for problems specific for NS. Then we extend this model so that it can serve as an underlying model for a *concrete* RSRP application.

5.1 A model for rolling stock scheduling at NS

The model for rolling stock scheduling is described by Fioole et al. [4] and Maróti [10]. It is used by NS to set up generic rolling stock circulations at an early planning stage with planning horizons of a few weeks/months. It is out of the scope of this paper to describe the complex model in detail. We here give a less formal description of the model and then describe the extensions implemented for the RSRP context.

Let R be the set of trips, let M be the set of rolling stock types, let C be the set of allowed compositions and let S be the set of stations. Let $T \subseteq R \times R$ be the set of pairs of consecutive trips.

The units in the inventories at the stations can be coupled to departing trains, while uncoupled units are added to the inventory. The inventories thus form a simple model of the shunting process. The inventories of each station at the end of the day connect the rolling stock assignment to the next day. If specific inventory sizes are wanted, they are known as *target inventories*.

The model is based on a constrained multi-commodity flow in a graph constructed from the time-space diagram of the timetable. The core of the model is the decisions on what compositions are assigned to each trip. The model then links the compositions on consecutive trips through the possible composition changes at the stations and the units in the inventory there. The assignment of compositions to trips for a train follows a path in the transition graph shown in Figure 6 where the possible composition changes are represented by arcs between compositions on consecutive trips. The model is summed up in the following list:

- Binary decision variable $X_{r,c}$ decides whether composition $c \in C$ is used on trip $r \in R$.
- Let $(r, r') \in T$ be a pair of consecutive trips and let $(c, c') \in C \times C$ be a pair of compositions. Then binary variable $Z_{r,r',c,c'}$ denotes whether composition c on trip r is changed to composition c' on trip r' . We have

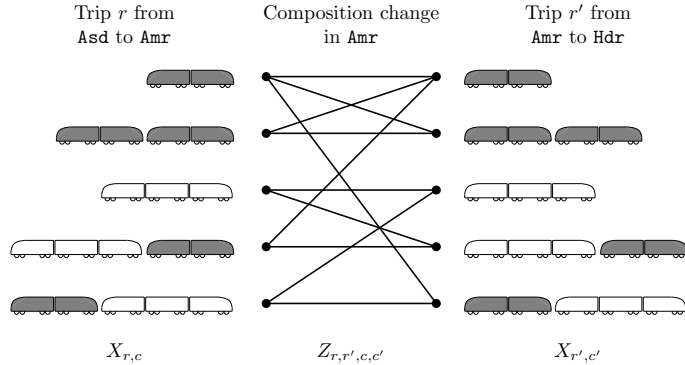


Figure 6: Transition graph for two consecutive trips. In the example it is possible to couple one unit to the front of the train or uncouple one unit from the rear of the train at station Amr.

such a variable only if the composition change from c to c' is technically possible.

- Integer variables $I_{t,s,m}$ count the inventory of units of type $m \in M$ at station $s \in S$ at time t . Integer variables $I_{s,m}^\infty$ count the end-of-day inventory of units of type m at station s .

A rolling stock assignment is feasible if certain upper and lower bounds on train lengths hold, the composition changes between consecutive trips are legal, and the inventories are always non-negative.

We note that the model does not take the duties of rolling stock units explicitly into account. Instead it works with the units at an aggregated level and divides them into groups of units of interchangeable types. However, it is always possible to convert the compositions resulting from model to duties and experience shows that the duties created from the compositions are also feasible in practice.

5.2 Extending the model for RSRP

Now we extend the model to the rescheduling context by adding a number of simple features. First of all, the original rolling stock circulation has to be taken into account. We assume that \mathcal{T}_0 and \mathcal{C}_0 are known. In particular, we know which composition is assigned to each trip in the original circulation.

The rescheduling takes place at time t where \mathcal{T}' is the new timetable. All compositions on trips departing before time t are fixed in the model so that the trips get the same composition as in \mathcal{C}_0 since the past cannot be changed. For the remainder of the day compositions may be modified freely.

In order to accommodate the possibility of canceling trips, a new binary variable $X_{r,0}$ is added for each trip $r \in R$ that departs after time t . The variable denotes the assignment of a special *empty composition* to the trip which is the

equivalent of assigning zero rolling stock units to the trip: In other words, the trip is canceled.

For measuring off balances in the end-of-day inventories some additional variables are needed. For each type of rolling stock $m \in M$ and for each station $s \in S$ the parameter $i_{s,m}^\infty$ denotes the target inventory and the variable $D_{s,m}$ denotes the deviation from the target. The deviation is measured by adding constraints

$$\begin{aligned} D_{s,m} &\geq i_{s,m}^\infty - I_{s,m}^\infty && \forall s \in S, m \in M \\ D_{s,m} &\geq I_{s,m}^\infty - i_{s,m}^\infty && \forall s \in S, m \in M \end{aligned}$$

A surplus of one unit at a station incurs a deficit of one unit at another station which in total results in two deviations. Thus two deviations constitute one off balance.

An important adjustment of the original model is in the use of the objective function. It is composed of the following terms.

$$\sum_{r \in R} w_r X_{r,0} + \sum_{r,r' \in T} \sum_{c,c' \in C \times C} \gamma_{r,r',c,c'} Z_{r,r',c,c'} + \sum_{s \in S} \sum_{m \in M} \beta_m D_{s,m}$$

The first term incurs a penalty of w_r for canceling trip r . The second term incurs a cost of $\gamma_{r,r',c,c'}$ for the composition change from c to c' between consecutive trips r and r' . The values γ depend on the composition change that was originally planned between the trips as explained in Section 3.3. The third term penalizes a deviation of rolling stock type m in the end-of-day inventory at station s by β_m .

The model can be solved by commercial MIP solvers (e.g. CPLEX) for realistic instances in reasonable time. Smaller instances take seconds or minutes and larger instances take minutes or hours. Experience shows that the running time is rather unpredictable and heavily depends on the structure of the objective function parameters. For some instances the running time is unacceptably long for use in a real time setting. This motivates a heuristic approach to reduce problem size and thereby decrease running time.

6 Rolling horizon framework for the online RSRP

The extended model described in Section 5 models RSRP where the exact duration of the disruption is known. However, in the online version each updated timetable is related to some uncertainty as the situation may develop further as time progresses. Also, the running times can be a bottleneck.

We use an approach inspired by observations from practice. During a disruption dispatchers only plan a certain time ahead and then revise the solution as time progresses and more information becomes available. This is possible since, in practice, only the most immediate decisions are executed and the remaining ones can be rescheduled again later. This approach is known as planning with a

rolling horizon. We note that Törnquist [14] uses a similar approach for railway traffic disturbance management.

The rolling horizon framework presented in this section does not include a specific algorithm for rescheduling the circulation. In principle, any model for rolling stock rescheduling can be used as the core of the solution approach. The rolling horizon framework only states how to decompose the problem into a series of subproblems without losing the overall structure of the problem. In the computational tests presented in this paper the model described in Section 5 is used for rescheduling the circulation.

6.1 Rescheduling with rolling horizon

The rolling horizon approach works by considering only those trips that are within a certain time horizon from the time at which rescheduling takes place. Whenever new information on the timetable and rolling stock becomes available, the circulation is rescheduled for a further time horizon.

To formalize the above mentioned idea, let h be a time horizon parameter, i.e. how far ahead we wish to take the current information into account. At time instant t_i when the information $\langle t_i, \mathcal{T}_i, \rangle$ becomes known, trips in \mathcal{T}_i that depart no later than $t_i + h$ are taken into account and the remaining are ignored. If a new timetable update $\langle t_{i+1}, \mathcal{T}_{i+1} \rangle$ arrives, then the circulation is rescheduled accordingly for the time interval from t_{i+1} to $t_{i+1} + h$.

For a time interval without any timetable update, the current timetable is still the best available forecast on the development of the disruption. Still, a feasible plan only exists from the last update point until the end of the horizon. So we introduce a new parameter p , called the period, which denotes how often the circulation should be updated when the available information does not change. Then if no new information arrives for p minutes create an artificial information update $\langle t_i + p, \mathcal{T}_i \rangle$ and reschedule accordingly.

The idea of rolling horizon rescheduling is shown in Figure 7: A time-space diagram is shown for a timetable with trips between stations A to D. At time t_1 a disruption occurs which leads to some cancellations of trips. The circulation is rescheduled to serve the updated timetable until $t_1 + h$. However shortly later, at time t_2 new information becomes available and again the circulation is rescheduled until $t_2 + h$. Then no new information arrives for some time which means that rescheduling is performed at time $t_2 + p$ and again at time $t_2 + 2p$ and so on.

Realistic values of h are around two to three hours as that allows dispatchers to react to the most immediate conflicts. Parameter p should be somewhat smaller than h to allow a smooth roll of the horizon, it could be for example one hour if h is two hours.




Figure 7: Rescheduling using rolling horizon. When new information is available, the circulation is rescheduled h time steps ahead.

6.2 Accounting for the objectives when rescheduling with rolling horizon

The objectives for cancellation of trips can be taken into account in the same way as in Section 5.2, namely by penalizing the assignment of the empty composition to a trip. This is of course only possible for trips within the horizon but the rolling stock will be rescheduled for all trips as time progresses and the horizon rolls on. Similarly, the penalties for changes to the shunting plans can be taken into account when using the horizon.

The end-of-day balances are a different matter. In a generic iteration, the rolling stock circulation model is applied for the time interval $[t_i; t_i + h]$. Therefore the end-of-day off balances cannot be taken explicitly into account during the day with a rolling horizon approach. We now propose a heuristic approach to account for off balances.

6.2.1 Heuristic for accounting for off balances with rolling horizon

When the circulation is rescheduled at time t_i , the deviation variables $D_{s,m}$ measure the deviation of the inventories at the end of the horizon at time $t_i + h$. The question is which target inventories are to be selected. We propose to use the original (undisrupted) rolling stock circulation \mathcal{C}_0 as a guideline.

One can explicitly compute the intermediate inventories $i_{s,m}^t$ according to \mathcal{C}_0 : these numbers describe how many units of type m are located at station s at time instant t . We declare the values $i_{s,m}^{t_i+h}$ to be the target inventories when planning for the time interval $[t_i; t_i + h]$.

This guideline may be fairly inaccurate for early time instants t_i , and intuitively it becomes more and more precise as t_i approaches the end of the day. Therefore the importance of the off balances in the objective function should

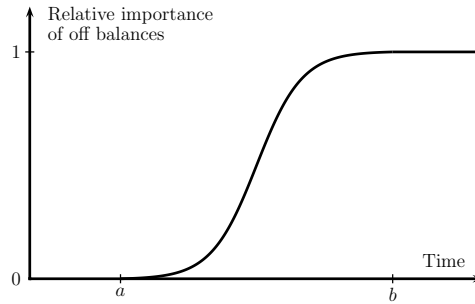


Figure 8: Increasing the relative importance of off balances by a scaled logistic function.

increase with t_i . We achieve this by multiplying the third term in the objective function by a parameter $\rho(t)$ so it reads

$$\sum_{s \in S} \sum_{m \in M} \rho(t) \beta_m D_{s,m}$$

where the parameter $\rho(t)$ depends on the time t at which the horizon ends.

There are several ways to increase $\rho(t)$ over time. We have restricted this study to a particular class of functions that increase the relative importance of the intermediate inventories from 0 to 1 over time. The function used is a scaled logistic curve. We introduce two parameters a and b that guide when the intermediate inventories start to be taken into account and when they are taken into account with full weight. The function is as follows

$$\rho(t) = \begin{cases} 0, & t < a \\ \frac{f(t) - f(a)}{f(b) - f(a)}, & a \leq t < b \\ 1, & t \geq b \end{cases}$$

where

$$f(t) = \frac{1}{1 + e^{-\frac{1}{30}(t - \frac{b-a}{2})}}$$

The function is scaled so that it maps the interval $[a, b]$ onto the interval $[0, 1]$. The curve is sketched in Figure 8.

6.3 Parameters for rescheduling with rolling horizon

With the described rolling horizon framework, there are four parameters that determine the performance of the rescheduling process. The parameters h and p determine for how long the circulation is rescheduled and how often the solution is revised. Parameter a determines when the intermediate balance is taken into account and parameter b determines when the intermediate inventories are as important as at the end of the day.




Figure 9: The Noord-Oost lines.

7 Computational results

In this section we present some computational results of the extended model from Section 5 used in the framework from Section 6. We explore the relationship between the horizon parameters and the solution quality and characteristics using a set of test instances. We first describe the instances and then discuss the results.

7.1 Test instances

To analyze the relationship between parameters and solutions we use a realistic set of instances of NS. These instances involve a disruption on the so-called Noord-Oost lines, a system of interconnected lines with a closed rolling stock circulation (see Figure 9). The Noord-Oost lines form the most challenging cases for rolling stock scheduling at NS as they have a complicated structure (e.g. trains are split or combined underway). We note that Fioole et al. [4] report computations times of several hours for tactical and operational scheduling problems on the Noord-Oost lines when the planning horizon is an entire day.

The disruption occurs at 12:00 between Utrecht (Ut) and Amersfoort (Amf). During this period no trains can run between the two stations. The disruption lasts two and a half hours. However, the actual length of the disruption is not known at the occurrence of the disruption but only an estimated length is available. An instance consists of the original timetable with the original circulation and a list of timetable updates that become available at certain times. Table 1 shows the timetable updates that are available in each instance. The timetable $\mathcal{T}_{t_1-t_2}$ denotes the timetable for the Noord-Oost lines with all trips canceled between Utrecht and Amersfoort in the time interval from t_1 to t_2 . All trains are turned in Utrecht and Amersfoort according to the handling

Instance	Timetable updates
1	$\langle 12:00, \mathcal{T}_{12:00-13:00} \rangle, \langle 12:30, \mathcal{T}_{12:00-14:30} \rangle$
2	$\langle 12:00, \mathcal{T}_{12:00-13:00} \rangle, \langle 13:00, \mathcal{T}_{12:00-13:30} \rangle, \langle 13:30, \mathcal{T}_{12:00-14:30} \rangle$
3	$\langle 12:00, \mathcal{T}_{12:00-13:00} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 13:30, \mathcal{T}_{12:00-14:30} \rangle$
4	$\langle 12:00, \mathcal{T}_{12:00-13:00} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 14:00, \mathcal{T}_{12:00-14:30} \rangle$
5	$\langle 12:00, \mathcal{T}_{12:00-13:00} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:30} \rangle$
6	$\langle 12:00, \mathcal{T}_{12:00-13:30} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:30} \rangle$
7	$\langle 12:00, \mathcal{T}_{12:00-13:30} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 13:30, \mathcal{T}_{12:00-14:30} \rangle$
8	$\langle 12:00, \mathcal{T}_{12:00-13:30} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 14:00, \mathcal{T}_{12:00-14:30} \rangle$
9	$\langle 12:00, \mathcal{T}_{12:00-13:30} \rangle, \langle 13:30, \mathcal{T}_{12:00-14:30} \rangle$
10	$\langle 12:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 13:00, \mathcal{T}_{12:00-14:30} \rangle$
11	$\langle 12:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 13:30, \mathcal{T}_{12:00-14:30} \rangle$
12	$\langle 12:00, \mathcal{T}_{12:00-14:00} \rangle, \langle 14:00, \mathcal{T}_{12:00-14:30} \rangle$
13	$\langle 12:00, \mathcal{T}_{12:00-14:30} \rangle$

Table 1: Instances with a disruption on the Noord-Oost lines between Utrecht and Amersfoort.

scenario. For example, in instance 6, the disruption is first estimated to last one and a half hours, but at time 13:00 the estimated length of the disruption is changed to two and a half hours in total.

The difference between the instances is when information becomes available and the accuracy of that information. The instances all have the same optimal offline solution.

7.1.1 Objective parameters

The overall objective is to balance the three objectives; changes to the shunting plans, off balances and canceled trips. The following weights are used in the tests: Introducing new shunting operations costs 100, and minor changes to already planned shunting operations cost either 1, 2 or 5 depending on their nature. Off balances cost 200. The cost of canceling a trip is set to 10,000 which outweighs all other objective parameters. This reflects the practitioners' preference to use train cancellations as a last option. The same objective function is used for each test instance.

7.2 Tuning the horizon parameters

In this section we investigate different settings of the parameters of the rolling horizon approach. We then discuss their consequences and implications for practical use.

7.2.1 Parameters for the intermediate inventories

As mentioned in Section 6, off balances are handled heuristically by increasing the importance of balancing the intermediate inventories. The parameters a

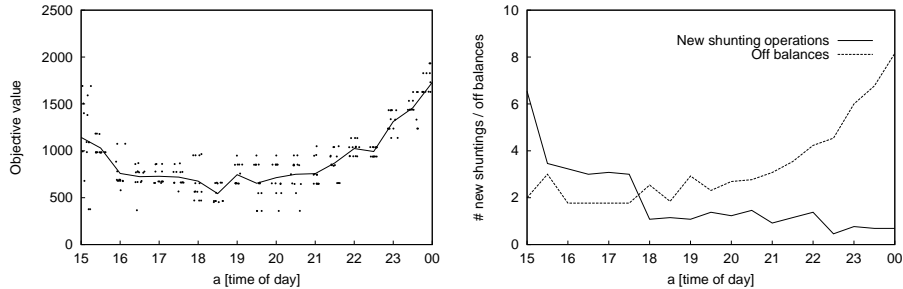


Figure 10: *Left*: Objective function of each instance as a function of parameter a . The line represents the average. *Right*: Average number of new shunting operations and off balances.

and b determine when the importance of the intermediate inventories starts to be taken into account and when it is taken into account with full weight respectively.

In the first set of experiments, we set the horizon to $h = 3$ hours and the update parameter to $p = 1$ hour. We then applied the rolling horizon framework with a at different values between 15:00 and 0:00 (midnight) with 30 minutes between values. The parameter b which controls when off balances in the intermediate inventories are taken into account with full weight is set to 6 hours after a in all tests.

The left diagram of Figure 10 shows the objective values of all instances as a function of the parameter a . The points have been scattered slightly in the horizontal direction to ease the visualization. The average outcome over all instances is plotted as a line in the diagram as well. It turned out that none of the obtained solutions requires cancellations of trips. In the right diagram of Figure 10 the average number of new shunting operations and off balances are shown as they are the major contributors to the objective function.

We note that when the balancing is initiated rather late, more off balances remain in the solutions. However, when balancing is initiated rather early, more new shunting operations are introduced. If the process starts too early many shunting operations are introduced without resolving more off balances. The best balanced solutions seem to be found with a between 18:00 and 20:00. This means that the intermediate inventories should be taken into account in the rescheduling from around 15:00 to 17:00 as the horizon is 3 hours.

The characteristics of all results with a specific choice of parameters are shown in Table 2. The parameters are $h = 3$ hours, $p = 1$ hour, $a = 18:30$ and $b = 0:30$. In all instances 1 or 2 off balances occur and up to 2 new shunting operations are introduced. In the optimal offline solution it is possible to reschedule the rolling stock in such a way that no off balances occur and no new shunting operations are introduced. However, when the balancing heuristic is not used and off balances are not taken into account in the rescheduling, it

Instance	New shunting operations	Minor changes to shunting plans	off balances	Objective
1	2	60	2	66000
2	2	57	2	65700
3	1	52	2	55200
4	0	65	2	46500
5	0	65	2	46500
6	0	65	2	46500
7	2	53	1	45300
8	0	65	2	46500
9	2	53	1	45300
10	2	60	2	66000
11	2	52	2	65200
12	0	65	2	46500
13	2	60	2	66000
(*)	0-2	22-26	11-13	-
offline	0	37	0	3700

Table 2: Results of all instances with parameters $h = 3$ hours, $p = 1$ hour, $a = 18:30$ and $b = 0:30$. (*) shows the results of all instances when rescheduling with the same horizon parameters except that off balances are not taken into account.

results in 11-13 off balances. The large reduction in off balances comes at a cost of some minor changes to the existing shunting operations. In real-life railway practice a reduction in potential off balances of this magnitude would be an excellent result. This suggests that the heuristic guidance of rolling stock to the target inventories works well in these instances.

We observe that parameter a represents a trade-off between off balances and changes to the shunting plans. The tests have shown that a being around 18:00 to 20:00 yields balanced results for the Noord-Oost lines. However, tests involving other lines have shown that the best choice of parameter a depends on the structure of the involved lines. For longer lines the balancing should be initiated earlier to allow units to end at far away stations.

7.2.2 Horizon length

A set of experiments have been conducted to explore the relationship between the length of the horizon and the solution quality. The length of the horizon represents how far ahead in the current timetable trips are taken into account. Naturally, looking further ahead gives a larger solution space and thus potentially better solutions. However, this comes at a cost of computation time and potentially poor decisions due to the online nature of the problem.

The instances have been tested with values for the horizon parameter h between 2 and 5 hours with 15 minutes between values. Parameter a is set to 22:00 with b being 6 hours later. The update parameter p is set to 1 hour. The objective values have been plotted as a function of horizon length in figure 11.




Figure 11: Objective function of each instance as a function of parameter h . The line represents the average.

There is a relationship between solution quality and horizon length; though longer horizons do not yield strictly better solutions, there is a tendency toward getting better results when using longer horizons.

7.2.3 Update parameter

The parameter p determines how often the circulation is rescheduled when no updated timetables arrive. This way it controls the roll of the horizon. Revising the situation more often potentially yields better results as non-executed decisions can be revised when later trips come into the horizon. However, in a real-life setting the frequency of updates is limited by the time it takes to discuss possibilities with local planners and communicate the changes to the involved parties.

We have tested the instances for values of p from 30 to 120 minutes with 15 minutes between values. The remaining parameters are fixed at $h = 4$ hours, $a = 21:00$ and $b = 03:00$. The objective values are plotted as a function of parameter p in Figure 12. It shows that updating more often potentially yields better solutions on average. The effect does not strictly increase with the value of p though.

7.3 Running times

The tests have been conducted on a Pentium 4 2.8GHz desktop with 1 GB of RAM using CPLEX 10.1. As discussed earlier, the size of the solution space depends on the horizon length so the relationship between solution time and horizon length is explored. The computation times are presented in Figure 13 as a function of horizon length. Values of h from 2 hours to 5 hours were tested with 15 minutes between values. As mentioned in Section 5, the computation time of the model can be somewhat unpredictable. This is clearly visible




Figure 12: Objective function of each instance as a function of parameter p . The line represents the average.

in the diagram where some computations take up to one minute while other computations with the same parameters only take a few seconds. There is a significant correspondence between the longest experienced computation times and the length of the horizon. When using a horizon of more than three and a half hours unpredictably long computation times may occur. We note that when setting h large enough to reschedule the rest of the day, computation times of up to 10 minutes occurred.

8 Conclusions

This paper deals with real-time disruption management of railway rolling stock. We defined the Rolling Stock Rescheduling Problem (RSRP). The main assumption is that the timetable update is explicitly given. The goal is to adjust the original rolling stock schedules for the updated timetable, taking various objectives into account.

We also defined the online variant of RSRP where the uncertainty about the duration of the disruption is modeled by a sequence of timetable updates. In order to deal with such uncertainties (and in order to reduce problem size), we proposed a rolling horizon framework as a solution approach. In this framework we only consider rolling stock decisions within a certain horizon of the time of rescheduling. The schedules are then revised as the situation progresses and more accurate information becomes available.

The (online) RSRP is an abstract framework which needs to be adapted for the concrete specifications of real-life railway scheduling problems. In this paper we extended an existing rolling stock scheduling model that has been applied successfully by the major passenger railway operator NS for its medium term planning since 2004. The most important adjustments concern the objec-

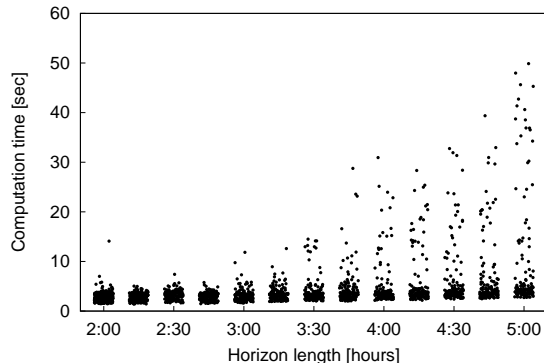


Figure 13: Computation time of each step of all instances as a function of parameter h .

tive function. We penalize cancellations of train, modifications of the shunting process as well as the end-of-day rolling stock off balances.

We note that the rolling horizon approach takes the off balances into account in a heuristic way. Based on the undisrupted rolling stock circulation, we define target inventories, and we minimize the deviation from these targets. Arguably, the heuristic target becomes more accurate as the current horizon approaches the end of the day. Therefore the objective coefficients of the off balances increase as the horizon shifts ahead in time.

The performance of the rolling horizon approach and the heuristic for balancing the end-of-day inventories depends on a number of parameters. These parameters include the length of the horizon and parameters that control the relative importance of the off balances compared to the other criteria. Though the best values are instance dependent, the test results indicate that starting the balancing too early results in many changes to the shunting plans without much effect on the final balances. However, starting too late leaves little time for balancing the rolling stock and leads to many off balances. Good results are achieved when taking off balances into account once the horizon starts after 21:00, though the best time depends on the structure of the involved lines.

The length of the horizon offers a trade-off between solution quality and computation time, though longer horizons do not offer strictly better solutions, mainly because of the uncertainty related to the disruptions.

The tests show that the method can be used to reschedule the rolling stock during a disruption with minor effects for the shunting plans. At the same time the number of off balances can be reduced by a few changes to the planned shunting operations. The obtained solutions may have a large competitive ratio, but the values themselves are quite appealing in practice. This along with the short computation times indicates that the approach is a good candidate for the core of a decision support system for rolling stock rescheduling.

The currently implemented model leaves a number of practically relevant issues out of account such as maintenance of rolling stock. These require further extensions of the model. Also, further discussions with planners are needed to refine the objective function for a better match to the preferences of the decision makers.

More tests need to be conducted to establish the best choice of parameters for different classes of instances. Also, it would be useful to investigate stochastic variants of the involved models.

Acknowledgments

We would like to thank the planners from the Rolling Stock Managing Center at NS for their input on the practical details of rolling stock operations and real time control.

This work was partially supported by the Future and Emerging Technologies Unit of EC (IST priority - 6th FP), under contract no. FP6-021235-2 (project ARRIVAL).

References

- [1] A. Alferi, R. Groot, L. Kroon, and A. Schrijver. Efficient Circulation of Railway Rolling Stock. *Transportation Science*, 40(3):378, 2006.
- [2] G. Budai, G. Maróti, R. Dekker, D. Huisman, and L. Kroon. Rescheduling in railways: the rolling stock balancing problem. Technical Report ARRIVAL-TR-0120, ARRIVAL project, January 2008. Submitted to *Journal of Scheduling*.
- [3] A. Caprara, L. Kroon, M. Monaci, M. Peeters, and P. Toth. Passenger Railway Optimization. In C. Barnhart and G. Laporte, editors, *Handbook in OR & MS vol. 14*, pages 129–187. Elsevier, 2007.
- [4] P.J. Fioole, L. Kroon, G. Maróti, and A. Schrijver. A Rolling Stock Circulation Model for Combining and Splitting of Passenger Trains. *European Journal of Operational Research*, 174:1281–1297, 2006.
- [5] M. Grötschel, S.O. Krumke, J. Rambau, T. Winter, and U. Zimmermann. Combinatorial Online Optimization in Real Time. Technical report, Konrad-Zuse Zentrum für Informationstechnik Berlin, 2001.
- [6] D. Huisman, L.G. Kroon, R.M. Lentink, and M. Vromans. Operations Research in Passenger Railway Transportation. *Statistica Neerlandica*, 59(4):467–497, 2005.
- [7] J. Jespersen-Groth and J. Clausen. Optimal Reinsertion of Cancelled Train Lines. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, August 2006.

- [8] J. Jespersen-Groth, D. Potthoff, J. Clausen, D. Huisman, L.G. Kroon, G. Maróti, and M.N. Nielsen. Disruption Management in Passenger Railway Transportation. Technical report, Econometric Institute, Erasmus University Rotterdam, 2007.
- [9] A.R. Karlin, M.S. Manasse, L. Rudolph, and D.D. Sleator. Competitive snoopy caching. *Algorithmica*, 3(1):79–119, 1988.
- [10] G. Maróti. *Operations Research Models for Railway Rolling Stock Planning*. PhD thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, 2006.
- [11] M. Peeters and L.G. Kroon. Circulation of Railway Rolling Stock: a Branch-and-Price Approach. *Computers & Operations Research*, 35(2): 538–556, 2008.
- [12] R. Séguin, J.Y. Potvin, M. Gendreau, T.G. Crainic, and P. Marcotte. Real-Time Decision Problems: An Operational Research Perspective. *The Journal of the Operational Research Society*, 48(2):162–174, 1997.
- [13] J. Törnquist. Computer-based decision support for railway traffic scheduling and dispatching: A review of models and algorithms. *ATMOS2005 (Algorithmic MeThods and Models for Optimization of RailwayS)*, Palma de Mallorca, Spain, October, 2005.
- [14] J. Törnquist. Railway traffic disturbance management - An experimental analysis of disturbance complexity, management objectives and limitations in planning horizon. *Transportation Research Part A*, 41(3):249–266, 2007.
- [15] C.G. Walker, J.N. Snowdon, and D.M. Ryan. Simultaneous disruption recovery of a train timetable and crew roster in real time. *Computers and Operations Research*, 32(8):2077–2094, 2005.

Publications in the Report Series Research* in Management

ERIM Research Program: "Business Processes, Logistics and Information Systems"

2009

How to Normalize Co-Occurrence Data? An Analysis of Some Well-Known Similarity Measures

Nees Jan van Eck and Ludo Waltman

ERS-2009-001-LIS

<http://hdl.handle.net/1765/14528>

Spare Parts Logistics and Installed Base Information

Muhammad N. Jalil, Rob A. Zuidwijk, Moritz Fleischmann, and Jo A.E.E. van Nunen

ERS-2009-002-LIS

<http://hdl.handle.net/1765/14529>

Open Location Management in Automated Warehousing Systems

Yugang YU and René B.M. de Koster

ERS-2009-004-LIS

<http://hdl.handle.net/1765/14615>

VOSviewer: A Computer Program for Bibliometric Mapping

Nees Jan van Eck and Ludo Waltman

ERS-2009-005-LIS

<http://hdl.handle.net/1765/14841>

Nash Game Model for Optimizing Market Strategies, Configuration of Platform Products in a Vendor Managed Inventory (VMI) Supply Chain for a Product Family

Yugang Yu and George Q. Huang

ERS-2009-009-LIS

<http://hdl.handle.net/1765/15029>

A Mathematical Analysis of the Long-run Behavior of Genetic Algorithms for Social Modeling

Ludo Waltman and Nees Jan van Eck

ERS-2009-011-LIS

<http://hdl.handle.net/1765/15181>

A Taxonomy of Bibliometric Performance Indicators Based on the Property of Consistency

Ludo Waltman and Nees Jan van Eck

ERS-2009-014-LIS

<http://hdl.handle.net/1765/15182>

A Stochastic Dynamic Programming Approach to Revenue Management in a Make-to-Stock Production System

Rainer Quante, Moritz Fleischmann, and Herbert Meyr

ERS-2009-015-LIS

<http://hdl.handle.net/1765/15183>

Some Comments on Egghe's Derivation of the Impact Factor Distribution

Ludo Waltman and Nees Jan van Eck

ERS-2009-016-LIS

<http://hdl.handle.net/1765/15184>

The Value of RFID Technology Enabled Information to Manage Perishables

Michael Ketzenberg, and Jacqueline Bloemhof

ERS-2009-020-LIS

<http://hdl.handle.net/1765/15412>

The Environmental Gains of Remanufacturing: Evidence from the Computer and Mobile Industry
J. Quariguasi Frota Neto, and J.M. Bloemhof
ERS-2009-024-LIS
<http://hdl.handle.net/1765/15912>

Economic Modeling Using Evolutionary Algorithms: The Effect of a Binary Encoding of Strategies
Ludo Waltman, Nees Jan van Eck, Rommert Dekker, and Uzay Kaymak
ERS-2009-028-LIS
<http://hdl.handle.net/1765/16014>

Language Selection Policies in International Standardization – Perception of the IEC Member Countries
Hans Teichmann and Henk J. de Vries
ERS-2009-031-LIS
<http://hdl.handle.net/1765/16038>

Dominant Design or Multiple Designs: The Flash Memory Card Case
Henk J. de Vries, Joost P.M. de Ruijter and Najim Argam
ERS-2009-032-LIS
<http://hdl.handle.net/1765/16039>

Standards Education Policy Development: Observations based on APEC Research
Donggeun Choi, Henk J. de Vries and Danbee Kim
ERS-2009-033-LIS
<http://hdl.handle.net/1765/16040>

Scheduling deliveries under uncertainty
Adriana F. Gabor, Rommert Dekker, Timon van Dijk, and Peter van Scheepstal
ERS-2009-040-LIS
<http://hdl.handle.net/1765/16236>

A simple alternative to the h-index
Ludo Waltman and Nees Jan van Eck
ERS-2009-043-LIS
<http://hdl.handle.net/1765/16556>

Disruption Management of Rolling Stock in Passenger Railway Transportation
Lars Kjaer Nielsen and Gabor Maroti
ERS-2009-046-LIS
<http://hdl.handle.net/1765/16557>

* A complete overview of the ERIM Report Series Research in Management:
<https://ep.eur.nl/handle/1765/1>

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship