VU Research Portal

An adjustable robust optimization approach for periodic timetabling

Polinder, G. J.; Breugem, Thomas; Dollevoet, Twan; Maróti, Gábor

published in

Transportation Research Part B: Methodological
2019

DOI (link to publisher)

[10.1016/j.trb.2019.07.011](https://doi.org/10.1016/j.trb.2019.07.011)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Polinder, G. J., Breugem, T., Dollevoet, T., & Maróti, G. (2019). An adjustable robust optimization approach for periodic timetabling. *Transportation Research Part B: Methodological*, 128, 50-68.
<https://doi.org/10.1016/j.trb.2019.07.011>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

An adjustable robust optimization approach for periodic timetabling

Gert-Jaap Polinder^{a,e,*}, Thomas Breugem^{b,e}, Twan Dollevoet^{b,e}, Gábor Maróti^{c,d}

^aRotterdam School of Management, Erasmus University, P.O. Box 1738, Rotterdam 3000DR, the Netherlands

^bEconometric Institute, Erasmus School of Economics, Erasmus University Rotterdam, P.O. Box 1738, Rotterdam 3000DR, the Netherlands

^cDepartment of Process Quality and Innovation, Netherlands Railways, P.O. Box 2025, Utrecht 3500HA, the Netherlands

^dSchool of Business and Economics, VU University Amsterdam, De Boelelaan 1105, Amsterdam 1081HV, the Netherlands

^eErasmus Center for Optimization in Public Transport (ECOPT), Rotterdam, the Netherlands

ARTICLE INFO

Article history:

Received 17 December 2018

Revised 7 June 2019

Accepted 22 July 2019

Available online 6 August 2019

Keywords:

Robust optimization

Periodic timetabling

Periodic event scheduling

Adjustable robustness

ABSTRACT

In this paper, we consider the Robust Periodic Timetabling Problem (RPTP), the problem of designing a periodic timetable that can easily be adjusted in case of small periodic disturbances. We develop a solution method for a parametrized class of uncertainty regions. This class relates closely to uncertainty regions known in the robust optimization literature, and naturally defines a metric for the robustness of the timetable. The proposed solution method combines a linear decision rule with well-known reformulation techniques and cutting-plane methods. We show that the RPTP can be solved for practical-sized instances by applying the solution method to practical cases of Netherlands Railways (NS). In particular, we show that the trade-off between the efficiency and robustness of a timetable can be analyzed using our solution method.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

The timetable of a railway operator is the backbone of daily operations. It is therefore no surprise that forced deviations from the planned timetable can be extremely costly: A recent investigation of the Dutch government estimated the annual societal costs of disturbances in the Dutch railway system at 400 to 500 million euros (Ministerie van I&W, KiM, 2017). These costs come from productivity loss, and from the uncertainty that passengers experience in the time their journey takes. Resilience against possible disturbances is therefore a crucial aspect of the operated timetable.

A railway timetable is an assignment of departure and arrival times for a given set of train services. The input to timetabling specifies this set of train services, including all locations where the train shall call. Timetabling is part of tactical planning, which generally is done a few years up to a few months before the actual operations. Scheduling of the resources (namely, rolling stock and crew) is traditionally done in later planning stages. Our focus in this paper is on generating (periodic) timetables, which have to be designed a few months before the actual operation.

We consider periodic timetabling, where the trains follow a repeating pattern: The timetable of a base period (say, one hour) is repeated throughout the day. Periodic timetables are commonly used by passenger railway operators in Europe, although there are examples of non-periodic or partially periodic timetables, as well. Generally, the same periodic timetable is operated for several weeks up to a few months.

* Corresponding author.

E-mail address: polinder@rsm.nl (G.J. Polinder).

The actual operations often differ from what was planned. For example, trains may not be able to run at full speed or have to dwell longer at some stations. These are all *disturbances*, events that can influence the execution of a timetable by causing delays. In this paper, we focus on small disturbances, leading to delays of up to a few minutes. Although the timetable was designed several months before the operations, it should still be operable with a limited amount of delays, no matter when and where the unpredictable disturbances occur. A robust timetable is one for which this is the case. In order to achieve robustness, a timetable allocates time supplements on top of the technically minimal travel times and dwell times. These supplements can absorb the unexpected increase of travel and dwell times and thereby limit the accumulation of delays.

In this paper, we consider *periodic* disturbances, i.e., disturbances that reoccur in every cycle period during one day or a few days. Periodic disturbances are seen, for example, on autumn days when the fallen leaves lead to slippery rails, and thus to increased travel times. When designing the timetable to be operated in autumn, it is known that fallen leaves lead to increased travel times during some days. However, it is not known yet at which days this will happen. Similarly, the dwell times at some stations increase on sunny summer days when many people take the train to the beach. One can expect this to happen, say, in August, but it is unknown which days will be sunny. In both cases, a timetable is announced for a period of several weeks or months, while the periodic disturbances affect a few days only. The timetable must balance the service quality of the normal days and the delays of the disturbed days.

Achieving robustness generally conflicts with achieving efficiency in a timetable. In an efficient timetable, passengers have little waiting time, and can travel quickly from their origin to their desired destination. Robust solutions prefer larger time supplements since they offer more potential to absorb delays. However, the increased time supplements are added on undisturbed days, as well. That is, the passengers will experience longer travel times even if there are no disturbances on that day.

In this paper we aim at finding a periodic timetable which is robust against small periodic disturbances. An adjustable robust periodic timetable is one that can easily be adapted to cope with such disturbances. In particular, the periodic disturbances can be incorporated without changing the structure of the timetable. Note that this leads to a complex optimization problem. The underlying Periodic Timetabling Problem is NP-hard (Serafini and Ukovich, 1989) and known to cause major computational challenges in practice, as well. The robust extension adds to the complexity in that it needs to keep track of the consequences of many possible disturbances.

In this paper, we formulate the Robust Periodic Timetabling Problem (RPTP), an adjustable robust formulation for the Periodic Timetabling Problem. We propose a solution method for a parametrized uncertainty region, and we study the interplay between three values: (i) the severity of the disturbance; (ii) the total amount of timetable adjustments; and (iii) the efficiency of the undisturbed timetable. We apply our approach to practical instances from Netherlands Railways (NS), and show that good solutions can be obtained for challenging real-world instances.

To the best of our knowledge, this paper is one of the first to incorporate periodic disturbances in the Periodic Event Scheduling Problem, and to combine this with standard robust optimization techniques. Furthermore, we are the first to analyze the efficiency and robustness of a timetable with respect to a recovery budget in the context of periodic timetabling in such detail.

The contribution of the paper is twofold. Firstly, we propose an adjustable robust extension of the Periodic Timetabling Problem. Second, we demonstrate the viability of adding adjustable robustness extensions to practically relevant optimization problems. In fact, periodic timetabling is an application where the nominal (i.e., non-robust) optimization problems pose a considerable challenge. Our robust extension turns out to be tractable by using the proper combination of modelling and solution techniques.

The remainder of this paper is organized as follows. In Section 2 we formalize the RPTP. We give an overview of related work in Section 3. In Section 4 we propose a mathematical formulation for the problem. We introduce two assumptions in Section 5, on which we base our solution method in Section 6. We apply this method to practical instances from Netherlands Railways in Section 7. Finally, the paper is concluded in Section 8.

2. Problem description

In this section, we first explain what the Periodic Timetabling Problem is in Section 2.1. After this, we focus on robustness in periodic timetabling and highlight the problem we deal with in this paper in Section 2.2.

2.1. Periodic Timetabling Problem

In the Periodic Timetabling Problem (PTP), the departure and arrival times of trains have to be planned. These departure and arrival events are periodic, i.e., they re-occur every cycle period (e.g., one hour). To assure that the resulting timetable can be (safely) operated, several relations are needed between the event times. These relations are referred to as *activities*. Each activity is a relation between a pair of events, with lower and upper bounds on their time difference. Examples include driving activities (restricting the time difference between a departure and the next arrival), dwell activities (restricting the time a train dwells at a station), transfer activities (to guarantee a good transfer time from one train to another) and safety activities (to guarantee a safe operation of a timetable).

Fig. 1. Example periodic constraint.

Fig. 2. Example periodic constraint with disturbance.

The PTP is naturally modeled as a Periodic Event Scheduling Problem (PESP), as introduced by Serafini and Ukovich (1989). The PESP can be stated as follows. Given the set V of events and the set A of activities, lower and upper bounds ℓ_{ij} and u_{ij} for each $(i, j) \in A$ and a cycle period T , find event times π_i for each $i \in V$ and integer values p_{ij} for each $(i, j) \in A$ such that

$$\pi_j - \pi_i + Tp_{ij} \in [\ell_{ij}, u_{ij}] \quad \forall (i, j) \in A \quad (2.1)$$

is satisfied. That is, assure that the difference between event times π_i and π_j is within the $(T$ -periodic) interval specified by the bounds ℓ_{ij} and u_{ij} , given a cycle period of length T . We assume, without loss of generality, that $0 \leq \ell_{ij} < T$ for all $(i, j) \in A$. The term Tp_{ij} acts as a *modulo*-operator, which is necessary since the interval $[\ell_{ij}, u_{ij}]$ is periodic. The goal of the PESP is to assign a time to each event such that all constraints are satisfied.

The general form (2.1) can be used to express all common timetabling constraints. Consider, for example, a driving activity that states that the time difference between arrival event π_2 and departure event π_1 should be at least 10 min, and at most 13 min. This constraint can be written as $\pi_2 - \pi_1 + Tp_{12} \in [10, 13]$, which is visualized in Fig. 1 (where the cycle time is 60 min). The highlighted area shows the possible event times π_2 , given the scheduled time $\pi_1 = 10$. For this activity, scheduling π_2 at minute 20 is a feasible solution. In fact, this would be the unique optimal solution from an efficiency-perspective.

2.2. Robustness in periodic timetabling

The goal of the PTP is to find a feasible timetable. It is common practice, however, to include an objective which maximizes the efficiency or the robustness of the timetable. These two objectives are conflicting: Efficiency is achieved by minimizing the planned duration of trip and dwell activities, while robustness is achieved by adding time supplements on these activities.

In practice, operations often differ from the plan due to disturbances, thus rendering the planned timetable infeasible. We assume that the uncertainties manifest themselves in the lower bounds of (2.1). This assumption is in line with the literature (Liebchen et al., 2009; Goerigk and Schöbel, 2010; Goerigk, 2015). We note that uncertainties in the upper bounds can be easily incorporated in the models, and that they have no effect on the methodology proposed in this paper.

Consider again the trip time activity shown in Fig. 1. Suppose that a disturbance affects this activity: The travel time will be at least 11 min due to slippery tracks. This implies that the feasibility interval changes from $[10, 13]$ to $[11, 13]$, as shown in Fig. 2. Notice that the highlighted area shrunk by 1 min. As a result, the originally scheduled pair of event times becomes infeasible. A possible solution is to shift event time $\pi_2 = 20$ by one minute to $\pi_2^* = 21$, thereby making the schedule feasible again. The shift of event time π_2 is known as the *recourse action* for the given disturbance.

Given an uncertainty set (i.e., a set of possible disturbances), we determine a timetable that maximizes the efficiency while we require that for every disturbance in the uncertainty set, at least one recourse action should exist such that the timetable remains feasible. This implies that two things have to be taken into account whilst constructing a robust timetable: (i) the set of possible disturbances, and (ii) the set of feasible recourse actions. The latter is restricted by introducing a *budget*. The budget limits the recourse actions: The total amount of adjustments must not exceed the budget. This assures that the adjustable timetable has the desirable property of not being too different from the nominal timetable.

The RPTP can now be stated as follows: Given an uncertainty set, and a recovery budget, create a feasible periodic timetable such that the efficiency is optimized, whilst assuring that, given the budget, there exist feasible timetable adjustments for each possible disturbance. A solution to the RPTP is specified by a nominal timetable, and a recourse action for every possible disturbance. By varying the uncertainty set and recovery budget, different ‘good’ timetables can be obtained.

3. Literature review

The periodic (train) timetabling problem is well-studied in the literature. Generally, the problem is modeled as a PESP. The versatility of the PESP allows to model a broad range of different problems. In this overview we mainly focus on railway timetabling problems. Odijk (1996) and Peeters (2003) give an extensive overview on how to model timetabling

constraints in the PESP framework. A detailed discussion of what can be included in a PESP framework is given in [Liebchen and Möhring \(2007\)](#).

Using a Mixed Integer Programming approach, the PESP is generally solved using one of two well-known formulations. The first is the formulation introduced in [Serafini and Ukovich \(1989\)](#), in which the actual event times are planned, see also [Section 2.1](#). The second formulation is known as the Cycle Basis formulation (see, e.g., [Nachtigall \(1999\)](#), [Liebchen and Peeters \(2009\)](#), [Caimi et al. \(2017\)](#)), in which not the actual event times are planned, but the time differences between events, i.e., the activity durations. The actual timetable can easily be computed afterwards. This latter formulation is generally more efficient than the former when using a commercial solver. The MIP formulations are successfully applied to solve practical timetabling problems (cf. [Liebchen \(2008\)](#)), however, both formulations struggle to solve large scale instances.

Alternative approaches to solve the PESP are constraint programming (see [Kroon et al. \(2009\)](#)), or the modulo-simplex heuristic (see, for example, [Nachtigall and Opitz \(2008\)](#), and [Goerigk and Schöbel \(2013\)](#)). Furthermore, [Großmann et al. \(2012\)](#) developed a sophisticated algorithm for the PESP which combines a state-of-the-art SAT solver with a polynomial transformation from PESP to SAT. The authors show that the resulting algorithm can solve the PESP very effectively.

When operating a periodic timetable in practice, disturbances and disruptions are faced on a daily basis. Common examples of a disruption are the blockage of a track or failures of signs and switches. There is a rich stream of literature on rescheduling the timetable in case of a given disruption, see, for example, [Meng and Zhou \(2011\)](#), [Louwerse and Huisman \(2014\)](#) and [Veelenturf et al. \(2016\)](#). Also, much research has been devoted to rescheduling the timetable in case of known smaller disturbances, such as delays of several minutes, see [D'Ariano et al. \(2007\)](#) and [Pellegrini et al. \(2015\)](#). In both cases, a real-time setting is considered, which means that the disturbance or disruption is known. We refer to [Cacchiani et al. \(2014\)](#) for an overview of recovery models for railway timetabling in such a real-time setting. In contrast, in robust timetabling, the occurrence of disturbances is considered already in the design of the timetable, i.e., in the planning phase.

The field of robust optimization has developed quickly in recent years. We refer to [Ben-Tal et al. \(2009a\)](#) and [Bertsimas et al. \(2011\)](#) for a detailed overview of robust optimization and associated techniques. [Gorissen et al. \(2015\)](#) also give a detailed overview of robust optimization techniques, thereby focusing on applicability in practice. Adjustable robustness is a specific technique in robust optimization, in which a two-stage problem is considered, where in the second stage an additional decision has to be made, according to some decision rule. This idea has been successfully applied to practical problems (see, e.g., [Ben-Tal et al. \(2005\)](#) and [Ben-Tal et al. \(2009b\)](#)). [Yanikoğlu et al. \(2018\)](#) give a survey on several techniques in adjustable robust optimization, thereby providing some examples of decision rules. Interesting work has been done in [Iancu et al. \(2013\)](#) and [Zhen et al. \(2018\)](#), and references therein, which prove optimality for classes of decision rules based on the underlying problem structure.

In recent years, robust timetabling has received much attention (see e.g., [Goerigk and Schöbel \(2010\)](#), [Cacchiani and Toth \(2012\)](#), and [Lusby et al. \(2017\)](#) for detailed overviews). Although much interesting work has been done, there still seems to be a lack of consensus: As noted in [Lusby et al. \(2017\)](#), there is currently no clear and unique definition of robustness for timetabling. We therefore cover only a subset of the approaches in this overview which are most closely related to our research.

Most research regarding robust timetabling has focused on *aperiodic* timetabling. [Fischetti and Monaci \(2009\)](#) introduce light robustness, and apply this robustness concept to the aperiodic timetabling problem. The light robustness framework aims at finding a robust solution, without deviating too much from the optimal nominal objective value. The first papers that apply adjustable robustness in timetabling are [Liebchen et al. \(2009\)](#), [Cicerone et al. \(2009\)](#), and [Cicerone et al. \(2012\)](#). The adjustability of a solution relates to the existence of a rescheduling algorithm, which resolves possible infeasibilities as soon as the disturbance becomes known. Finally, [Goerigk and Schöbel \(2010\)](#) and [Goerigk and Schöbel \(2014\)](#) introduce the concept of Recovery-to-Optimality, which aims at minimizing the recovery cost of the solution (i.e., the cost incurred when making the disturbed solution feasible again).

One of the first approaches to increase the reliability of a *periodic* timetable that we are aware of is by [Kroon et al. \(2008\)](#). The authors apply a stochastic programming approach in order to improve the robustness of an input timetable, by adding time supplements at strategic places. Later on, this method is improved by [Maróti \(2017\)](#). Stochastic programming and adjustable robustness are closely related in that both approaches aim at finding a nominal solution that can be repaired in any disturbance scenario. The main conceptual difference is the objective. Stochastic programming models minimize the expected repair costs subject to a given probability distribution of the disruptions. In contrast, adjustable robustness considers worst-case repair costs; in particular, it does not assume any known probability distribution.

Another approach to increase the robustness of a timetable is presented by [Dewilde et al. \(2014\)](#). Here, a timetable is developed that minimizes the *expected* travel time of the passengers. To this end, the expected travel time is defined, for each edge in the event-activity network, as a function of the buffer allocated to that edge. By minimizing the expected travel time, the trade-off between the planned travel time and the delays that occur in practice is automatically incorporated (see [Sels et al. \(2016\)](#)).

[Liebchen et al. \(2009\)](#) note that the adjustable robustness framework can also be applied to periodic timetabling. They observe that a choice has to be made regarding the (a)periodicity of the disturbances. They, however, do not provide any computational results for the periodic case. [Liebchen et al. \(2010\)](#) consider periodic timetabling, by rolling out a periodic instance to a non-periodic instance. [Goerigk \(2015\)](#) extends the Recovery-to-Optimality approach to the periodic case. For

small cases, a mixed-integer programming approach is taken. For larger instances, for which this approach is intractable as the nominal problem already is notoriously difficult, a local search heuristic is applied. Within a provided time limit, this approach tries to improve the robustness of the current solution.

As is clearly shown in [Lusby et al. \(2017\)](#), current literature has devoted much attention to robust timetabling. The main focus of many of these approaches is on aperiodic timetabling. For periodic timetabling, only few approaches exist, and many of these focus on aperiodic disturbances. This, in general, means that a roll-out has to be done which can be evaluated over different scenarios. In the majority of the cases, a robust timetable is found based only on a limited number of scenarios ([Kroon et al., 2008](#); [Liebchen et al., 2010](#)). These approaches suffer from the curse of dimensionality, i.e., they quickly become intractable if more scenarios are included. Therefore, in practical cases, only few scenarios can be considered. There is only one paper that we are aware of, that considers periodic timetabling with periodic disturbances ([Goerigk, 2015](#)), but it also deals with a finite number of scenarios and a heuristic is used for larger cases.

We add to the robust optimization literature by proposing a method to find a periodic timetable, that can be adjusted once a possible disturbance occurs. It deals with periodic disturbances, which has received only limited attention so far. Our model extends a known model, the PESP, for the fairly challenging Periodic Timetabling Problem.

4. Mathematical model

In this section we present a mathematical formulation for the RPTP. We first develop the necessary notation and terminology in [Section 4.1](#). We then present the mathematical formulation in [Section 4.2](#), and we conclude by characterizing the set of admissible decision rules in [Section 4.3](#).

4.1. Notation and terminology

We consider a network of events V and activities A . The set $\bar{A} \subseteq A$ denotes the set of all activities related to the efficiency of the timetable (i.e., the set of activities related to trip, dwell, and transfer activities). Let w_{ij} denote the weight assigned to each of these activities. These weights specify the relative importance (e.g., based on the number of passengers) of each of the activities. Furthermore, let β denote the budget that is available for the timetable adjustments. These adjustments are measured at a subset $\bar{V} \subseteq V$ of the events (e.g., all arrivals at important stations, or at all stations where the trains stop).

Each activity $(i, j) \in A$ has an associated lower bound ℓ_{ij} and upper bound u_{ij} , restricting the time difference between events i and j according to [\(2.1\)](#). Furthermore, a subset $D \subseteq A$ of the activities is prone to disturbances. In the example of the fallen leaves, the set D contains all driving activities that are influenced by fallen leaves, thus, those that traverse regions with many trees. These disturbances manifest themselves in the lower bounds of the activities. Formally, we are given an uncertainty set Z of disturbances, which influence the subset of activities D . Let ζ_{ij} denote the effect of disturbance $\zeta \in Z$ on activity $(i, j) \in D$. It is assumed that ζ_{ij} is in $[0, 1]$. The lower bound of activity $(i, j) \in D$, given that the disturbance $\zeta \in Z$ occurs, is given by

$$\ell_{ij} + g_{ij}\zeta_{ij}. \quad (4.1)$$

Here the non-negative scalar g_{ij} regulates the impact of the disturbance on the lower bound. The found timetable should be robust against all disturbances in Z . That is, for every possible disturbance $\zeta \in Z$, there should exist timetable adjustments such that the solution is feasible for the ‘disturbed’ lower bounds given by [\(4.1\)](#). We assume that Z always includes a ‘nominal disturbance’, i.e., a disturbance ζ^0 with $\zeta_{ij}^0 = 0$, for all $(i, j) \in D$. This assures that the timetable is feasible for the non-disturbed lower bounds.

4.2. Mathematical formulation

We now formalize the RPTP. First, we introduce the following decision variables.

- π_i for all $i \in V$. The variable π_i specifies the event time of event $i \in V$. Due to the periodicity of the timetable, it must hold that $\pi_i \in \{0, 1, \dots, T-1\}$.
- $y_{i\zeta}$ for all $i \in V$ and $\zeta \in Z$. The variable $y_{i\zeta}$ specifies the recovery action (i.e., the shift in event time) for event $i \in V$, given scenario ζ . These variables are the essence of adjustable robustness: The values $y_{i\zeta}$ together indicate how a timetable π is adjusted in order to deal with scenario ζ .
- p_{ij} for all $(i, j) \in A$. The integer variable p_{ij} is used to model the modulo operator in the PESP constraints. In particular, it is used to model the number of cycle periods between events i and j .

The RPTP, for a given budget parameter β , can now be expressed as follows.

$$\min \sum_{(i,j) \in \tilde{A}} w_{ij} (\pi_j - \pi_i + T p_{ij}) \quad (4.2a)$$

$$\text{s.t.} \quad \sum_{i \in V} y_{i\zeta} \leq \beta \quad \forall \zeta \in Z \quad (4.2b)$$

$$\pi_j + y_{j\zeta} - \pi_i - y_{i\zeta} + T p_{ij} \geq \ell_{ij} + g_{ij} \zeta_{ij} \quad \forall (i, j) \in D, \zeta \in Z \quad (4.2c)$$

$$\pi_j + y_{j\zeta} - \pi_i - y_{i\zeta} + T p_{ij} \geq \ell_{ij} \quad \forall (i, j) \in A \setminus D, \zeta \in Z \quad (4.2d)$$

$$\pi_j + y_{j\zeta} - \pi_i - y_{i\zeta} + T p_{ij} \leq u_{ij} \quad \forall (i, j) \in A, \zeta \in Z \quad (4.2e)$$

$$y_{i\zeta^0} = 0 \quad \forall i \in V \quad (4.2f)$$

$$p_{ij} \in \mathbb{Z}_+ \quad \forall (i, j) \in A \quad (4.2g)$$

$$\pi_i \in \{0, 1, \dots, T-1\} \quad \forall i \in V \quad (4.2h)$$

$$y_{i\zeta} \in \mathbb{R}_+ \quad \forall i \in V, \zeta \in Z. \quad (4.2i)$$

The objective (4.2a) expresses that we minimize a weighted sum of the duration of the activities. Generally, this is the minimization of the planned travel times of the passengers (i.e., we optimize the efficiency of the timetable). Constraints (4.2b) assure that the sum of the adjustments to the timetable is at most β , and (4.2c)–(4.2e) assure that the timetable can be made feasible for every possible scenario in Z . In particular, (4.2d) and (4.2e) represent (2.1) for the adjusted timetable. Here, $\pi_i + y_{i\zeta}$ and $\pi_j + y_{j\zeta}$ are the adjusted times when events i and j take place, given that disturbance ζ occurs. For transfer activities, for example, this ensures that the transfers are maintained for all disturbances. Constraints (4.2c) are similar to (4.2d) and take the increased lower bounds for disturbance ζ into account. Constraints (4.2f) enforce that no adjustments are made for the nominal (i.e. undisturbed) scenario, thereby assuring that (4.2a) is correctly specified. Finally, constraints (4.2g)–(4.2i) specify the domains of the decision variables. Note that (4.2) might have infinitely many constraints, depending on the uncertainty region Z .

We emphasize that we do not impose $\pi_i + y_{i\zeta} < T$ for events $i \in V$ and disturbances $\zeta \in Z$. If $\pi_i + y_{i\zeta} \geq T$ for some event i and disturbance ζ , it means that event i is actually rescheduled to $\pi_i + y_{i\zeta} - T$ if disturbance ζ occurs. Also observe that the decision variables p_{ij} are *not* adjustable. This implies that the recourse actions do not change the structure of the timetable. In the next section, we explain in detail what this means.

4.3. Analysis of decision rules

Given a disturbance $\zeta \in Z$, the time when event $i \in V$ takes place is given by $\pi_i + y_{i\zeta}$. Here, the variables $y_{i\zeta}$ describe how the timetable should be adjusted based on the disturbance ζ . In contrast, the variables p_{ij} can *not* be changed in order to cope with the disturbance. In this section, we describe the consequences of the p -variables not being adjustable.

We first explain what we mean by the *cyclic* order of events. Consider three events $i_1, i_2, i_3 \in V$. There are two possible cyclic orderings of the events i_1, i_2 , and i_3 , as visualized in Fig. 3. In the left part of the figure, the cyclic order (i_1, i_2, i_3) is depicted. Here, after event i_1 has occurred, event i_2 takes place before event i_3 . Note that this holds for all three time lines depicted in the left part of the figure. In particular, the cyclic orders (i_1, i_2, i_3) , (i_2, i_3, i_1) , and (i_3, i_1, i_2) are equal. Intuitively, this cyclic order means that after observing event i_1 , we first observe event i_2 and then i_3 . On the right, the cyclic order (i_1, i_3, i_2) ,

Fig. 3. Two cyclic orders of events $i_1, i_2, i_3 \in V$: On the left, the cyclic order is (i_1, i_2, i_3) ; on the right the cyclic order is (i_1, i_3, i_2) .

i_3, i_2) is depicted. In this case, after the occurrence of event i_1 , event i_3 occurs before event i_2 . Observe that the cyclic order of three events is independent of the start time of a cycle and thus invariant under time shifts.

Consider now a set of three events $I = \{i_1, i_2, i_3\} \subseteq V$ that induces a clique in the PESP graph. This means that there is an activity between any two distinct events $i, i' \in I$. More formally, there exists $(i, i') \in A$ for all $i, i' \in I$ with $i \neq i'$. Three events form a clique, for example, if the events take place at the same platform or use a common part of the infrastructure. If $u_{ii'} < T$ for all these (i, i') , the non-adjustability of the p -variables implies that the cyclic order of these three events is not changed as a result of a disturbance. We now prove this claim formally.

Proposition 1. *Let events $i_1, i_2, i_3 \in V$ induce a clique: $(i, i') \in A$ for all $i, i' \in \{i_1, i_2, i_3\}$ with $i \neq i'$. If $u_{ii'} < T$ for all such (i, i') , then every feasible decision rule leaves the cyclic order of events i_1, i_2 , and i_3 intact.*

Proof. Let a feasible solution $(\hat{\pi}, \hat{p}, \hat{y})$ for (4.2a)–(4.2i) be given. Assume, without loss of generality, that $\hat{\pi}_{i_1} \leq \hat{\pi}_{i_2} \leq \hat{\pi}_{i_3}$. It holds by definition that $\hat{\pi}_{i_3} \leq \hat{\pi}_{i_1} + T$. We will prove for every disturbance $\zeta \in Z$ that

$$\hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} \leq \hat{\pi}_{i_2} + \hat{y}_{i_2\zeta} \leq \hat{\pi}_{i_3} + \hat{y}_{i_3\zeta} \leq \hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} + T. \quad (4.3)$$

This implies that the cyclic order remains intact.

We first consider the inequality

$$\hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} \leq \hat{\pi}_{i_2} + \hat{y}_{i_2\zeta}. \quad (4.4)$$

By our assumption that $\hat{\pi}_{i_1} \leq \hat{\pi}_{i_2}$, it holds that

$$0 \leq \hat{\pi}_{i_2} - \hat{\pi}_{i_1} = \hat{y}_{i_2\zeta^0} - \hat{y}_{i_1\zeta^0}.$$

We prove by contradiction that $\hat{p}_{i_1i_2} \leq 0$. In order to do so, assume that $\hat{p}_{i_1i_2} \geq 1$. We then obtain from the above that

$$\hat{\pi}_{i_2} + \hat{y}_{i_2\zeta^0} - \hat{\pi}_{i_1} - \hat{y}_{i_1\zeta^0} + \hat{p}_{i_1i_2}T \geq T.$$

Using that $u_{i_1i_2} < T$, this contradicts feasibility for ζ^0 . We conclude that $\hat{p}_{i_1i_2} \leq 0$. To show that (4.4) is satisfied, assume on the contrary that (4.4) were not satisfied for a given $\zeta \in Z$. We then obtain

$$\hat{\pi}_{i_2} + \hat{y}_{i_2\zeta} - \hat{\pi}_{i_1} - \hat{y}_{i_1\zeta} + \hat{p}_{i_1i_2}T \leq \hat{\pi}_{i_2} + \hat{y}_{i_2\zeta} - \hat{\pi}_{i_1} - \hat{y}_{i_1\zeta} < 0 \leq \ell_{i_1i_2} \leq \ell_{i_1i_2} + g_{i_1i_2}\zeta_{i_1i_2}.$$

This contradicts feasibility for ζ . By contradiction, (4.4) must be satisfied. In a similar fashion, it follows from $\hat{\pi}_{i_2} \leq \hat{\pi}_{i_3}$ that

$$\hat{\pi}_{i_2} + \hat{y}_{i_2\zeta} \leq \hat{\pi}_{i_3} + \hat{y}_{i_3\zeta}$$

for all $\zeta \in Z$. It remains to prove that

$$\hat{\pi}_{i_3} + \hat{y}_{i_3\zeta} \leq \hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} + T. \quad (4.5)$$

By our assumption that $\hat{\pi}_{i_3} \leq \hat{\pi}_{i_1} + T$, it holds that

$$0 \leq \hat{\pi}_{i_1} - \hat{\pi}_{i_3} + T = \hat{y}_{i_1\zeta^0} - \hat{y}_{i_3\zeta^0} + T.$$

It follows that $\hat{p}_{i_3i_1} \leq 1$. To see this, assume on the contrary that $\hat{p}_{i_3i_1} \geq 2$. We then obtain

$$\hat{\pi}_{i_1} + \hat{y}_{i_1\zeta^0} - \hat{\pi}_{i_3} - \hat{y}_{i_3\zeta^0} + \hat{p}_{i_3i_1}T \geq T.$$

Using that $u_{i_3i_1} < T$, this contradicts feasibility for ζ^0 . We conclude that $\hat{p}_{i_3i_1} \leq 1$. Now, if (4.5) were not satisfied, then

$$\hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} - \hat{\pi}_{i_3} - \hat{y}_{i_3\zeta} + \hat{p}_{i_3i_1}T \leq \hat{\pi}_{i_1} + \hat{y}_{i_1\zeta} - \hat{\pi}_{i_3} - \hat{y}_{i_3\zeta} + T < 0 \leq \ell_{i_3i_1} \leq \ell_{i_3i_1} + g_{i_3i_1}\zeta_{i_3i_1}.$$

This contradicts feasibility for ζ . By contradiction, (4.5) must be satisfied for all $\zeta \in Z$. It follows that the cyclic order of events is preserved in the solution $(\hat{\pi}, \hat{p}, \hat{y})$. $\square$

The decision rules that we obtain can easily be implemented in practice. Changing the cyclic order of events in order to deal with a disturbance might require dispatchers to set the routing of trains manually. The fact that the cyclic order of events does not change implies that no manual adjustments are necessary to operate the adjusted timetable. It also means that the timetable is adjusted by propagation and absorption of delays only.

5. Modelling assumptions

In this section, we discuss two assumptions we make, which are used in our solution approach for the RPTP. We first discuss the uncertainty region in detail in Section 5.1. Next, we discuss a specific decision rule we use in Section 5.2. The combination of these leads to our solution approach in Section 6.

Fig. 4. Uncertainty region Z_α for two activities. The highlighted areas indicate the uncertainty region for different values of α . Note that there is overlap between the different regions, i.e., $Z_1 \subseteq Z_{1.5} \subseteq Z_2$.

5.1. Parametrized uncertainty region Z_α

We assume the uncertainty set Z as introduced in Section 4 to have a specific structure: We consider a parametrized uncertainty set Z_α , where α regulates the ‘severity of disturbance’ in the system. This is also known as a ‘budget of uncertainty’ (Bertsimas et al., 2011). To give a formal definition of the uncertainty set, Z_α is given by all ζ with $0 \leq \zeta_{ij} \leq 1$, for all $(i, j) \in D$, and

$$\sum_{(i,j) \in D} \zeta_{ij} \leq \alpha.$$

Note that this includes the uncertainty regions considered in Goerigk and Schöbel (2010) and Goerigk (2015). Furthermore, it generalizes the simplex and box uncertainty regions (Ben-Tal et al., 2009a).

Fig. 4 shows the uncertainty region Z_α for $|D| = 2$. The highlighted areas show the uncertainty region for different values of α . Note that the uncertainty region depends only on the ranges of the uncertain parameters. This means that part, but not all, of the correlation is expressed in the uncertainty region: Negatively correlated events will lead to linearly bounded uncertainty (i.e., the simplex region), whereas allowing general correlation (including perfect correlation) is captured by box uncertainty. In this sense, the control parameter α can be seen as a measure for both correlation between the disturbances and the ‘size’ of the disturbance: For two dimensions, $\alpha = 1$ corresponds to the simplex region, and $\alpha = 2$ corresponds to box uncertainty. This approach is more conservative than, e.g., stochastic programming, as only the support of the distribution of uncertain parameters is taken into account: If the solution is feasible for the entire support, then it must be feasible for any distribution over this support.

We conclude by emphasizing that the scenarios ζ are real-valued vectors from the unit hyper-cube. Thus a typical ζ has strictly positive components only. In particular, the sets Z_1 and Z_2 are not confined to vectors with at most one or two non-zero components, respectively.

5.2. Linear decision rule

For the large-scale problems encountered in practice, it is generally impossible to solve (4.2) for arbitrary decision rules. We therefore restrict our attention to linear decision rules, in order to simplify (4.2). This can be seen as a heuristic approach: By only considering linear decision rules, we allow less flexibility when dealing with disturbances. Other, more complex decision rules are not considered. As a result, if we obtain a solution to (4.2), its objective value is an upper bound on the optimal solution value of (4.2).

In particular, we assume that the decision rule $y_{i\zeta}$ can be written as a linear function of ζ . That is, we introduce a decision variable $\delta_{kl}^i \in \mathbb{R}_+$ for each $i \in V$ and $(k, l) \in D$, and assume that $y_{i\zeta}$ can be written as

$$y_{i\zeta} = \sum_{(k,l) \in D} \delta_{kl}^i \zeta_{kl}, \quad (5.1)$$

for all $i \in V$. The functional form (5.1) is generally known as a *linear decision rule*. The linear decision rule assumes that the recovery action is a linear combination (with weights δ_{kl}^i) of the disturbance vectors ζ . The weights express the extent by which a disturbance of an arc influences the realised event times: Each unit of disturbance of arc (k, l) will require a recovery time of δ_{kl}^i at event i . It is important to note that, although the linear decision rule is generally not an optimal policy, it is *always* an admissible recourse function. That is, every feasible linear decision rule is a feasible recourse function.

The resulting model is obtained by substituting the linear decision rule into the original model. That is, we use (5.1) to reformulate constraints (4.2b)–(4.2e). After rearranging terms, we obtain the new set of constraints

$$\sum_{i \in \tilde{N}} \sum_{(k,l) \in D} \delta_{kl}^i \zeta_{kl} \leq \beta \quad \forall \zeta \in Z \quad (5.2a)$$

$$\pi_j - \pi_i + T p_{ij} \geq \ell_{ij} + g_{ij} \zeta_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in D, \zeta \in Z \quad (5.2b)$$

$$\pi_j - \pi_i + T p_{ij} \geq \ell_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in A \setminus D, \zeta \in Z \quad (5.2c)$$

$$\pi_j - \pi_i + T p_{ij} \leq u_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in A, \zeta \in Z. \quad (5.2d)$$

The new model is obtained by replacing (4.2b)–(4.2e) with (5.2a)–(5.2d). Note that (4.2f) is trivially satisfied for (5.1).

6. Solution approach

The key issue when modelling robust optimization problems, is dealing with the intractability of the original model. In the case of linear uncertainty over a polyhedral region, as introduced in Section 5, two methods are most common: The constraint is reformulated as a finite system of inequalities using LP duality (see e.g., Ben-Tal et al. (2009a)), or constraints are separated iteratively using a cutting-plane method (also known as the adversarial approach). Bertsimas et al. (2016) give a detailed comparison of the reformulation technique and the cutting-plane method, where they show there is no clearly dominating method. We found the best-suited method for the RPTP to be a combination of the two. These two techniques that we use to solve the original model are outlined in this section. Finally, a heuristic approach to find a benchmark solution is introduced in Section 6.3.

6.1. Reformulation

In our approach, we only reformulate (5.2a). That means, (5.2a) is replaced by the system of inequalities

$$\begin{aligned} \alpha \lambda + \sum_{(i,j) \in D} \mu_{ij} &\leq \beta \\ \lambda + \mu_{kl} &\geq \sum_{i \in \tilde{N}} \delta_{kl}^i \quad \forall (k, l) \in D \\ \mu_{ij} &\in \mathbb{R}_+ \quad \forall (i, j) \in D \\ \lambda &\in \mathbb{R}_+. \end{aligned}$$

The decision on whether to reformulate a constraint or not, can be made based on the uncertainty set. In the case of Z_α , the reformulation of one constraint requires a total of $|D|$ new constraints and $|D| + 1$ new variables (see Appendix A for a detailed overview of the reformulation). For the considered instances, D has roughly the same size as A . This implies that reformulating (5.2b), (5.2c), and (5.2d) results in roughly $|A|^2$ additional constraints. As a result, reformulating all of the constraints (5.2b)–(5.2d) might perform well for smaller instances, but would lead to an intractable model (from a practical point of view) for the larger instances. For this reason, we have used the cutting-plane approach for these sets of constraints.

6.2. Cutting-plane method

In order to discuss the cutting-plane method in more detail, we consider the formulation

$$\min \sum_{(i,j) \in \tilde{A}} w_{ij} (\pi_j - \pi_i + T p_{ij}) \quad (6.1a)$$

$$\text{s.t.} \quad \alpha \lambda + \sum_{(i,j) \in D} \mu_{ij} \leq \beta \quad (6.1b)$$

$$\lambda + \mu_{kl} \geq \sum_{i \in \tilde{V}} \delta_{kl}^i \quad \forall (k, l) \in D \quad (6.1c)$$

$$\pi_j - \pi_i + T p_{ij} \geq \ell_{ij} + g_{ij} \zeta_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in D, \zeta \in L_\alpha^{ij} \quad (6.1d)$$

$$\pi_j - \pi_i + T p_{ij} \geq \ell_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in A \setminus D, \zeta \in L_{\alpha}^{ij} \quad (6.1e)$$

$$\pi_j - \pi_i + T p_{ij} \leq u_{ij} + \sum_{(k,l) \in D} (\delta_{kl}^i - \delta_{kl}^j) \zeta_{kl} \quad \forall (i, j) \in A, \zeta \in U_{\alpha}^{ij} \quad (6.1f)$$

$$p_{ij} \in \mathbb{Z}_+ \quad \forall (i, j) \in A \quad (6.1g)$$

$$\pi_i \in \{0, 1, \dots, T-1\} \quad \forall i \in V \quad (6.1h)$$

$$\delta_{kl}^i \in \mathbb{R}_+ \quad \forall i \in V, (k, l) \in D \quad (6.1i)$$

$$\mu_{ij} \in \mathbb{R}_+ \quad \forall (i, j) \in D \quad (6.1j)$$

$$\lambda \in \mathbb{R}_+, \quad (6.1k)$$

where $L_{\alpha}^{ij}, U_{\alpha}^{ij} \subseteq Z_{\alpha}$ for all $(i, j) \in A$. The constraints (6.1d), (6.1e), and (6.1f) are enforced for only a subset of the scenarios. Initially, we have $L_{\alpha}^{ij} = U_{\alpha}^{ij} = \{\zeta^0\}$, and in each cutting-plane iteration, we extend these sets.

For convenience in the description of the algorithm, we only consider (6.1d), noting that the methodology for (6.1e) and (6.1f) is identical. Consider some fixed $(i, j) \in D$. Whenever a solution $(\hat{\pi}, \hat{p}, \hat{\delta})$ is found for the reduced set of constraints, we check whether (6.1d) is violated for some $\zeta \in Z_{\alpha} \setminus L_{\alpha}^{ij}$. If this is the case, then ζ is added to L_{α}^{ij} , and the model is solved again. This procedure is repeated until no more violations can be found. At that point, the found solution is feasible for all possible disturbances. Note that this algorithm is equivalent to applying Benders Decomposition on the *fully* re-formulated model.

To check whether a constraint (6.1d) is violated for some $\zeta \in Z_{\alpha}$, we determine a disturbance $\hat{\zeta} \in Z_{\alpha}$ given by

$$\hat{\zeta} = \arg \max_{\zeta \in Z_{\alpha}} \left\{ g_{ij} \zeta_{ij} + \sum_{(k,l) \in D} (\hat{\delta}_{kl}^i - \hat{\delta}_{kl}^j) \zeta_{kl} \right\}, \quad (6.2)$$

and add $\hat{\zeta}$ to L_{α}^{ij} whenever

$$\hat{\pi}_j - \hat{\pi}_i + T \hat{p}_{ij} < \ell_{ij} + g_{ij} \hat{\zeta}_{ij} + \sum_{(k,l) \in D} (\hat{\delta}_{kl}^i - \hat{\delta}_{kl}^j) \hat{\zeta}_{kl}.$$

Note that if the solution $(\hat{\pi}, \hat{p}, \hat{\delta})$ is feasible with respect to $\hat{\zeta}$, the solution is feasible for all $\zeta \in Z_{\alpha}$.

An optimal solution for (6.2) can be determined efficiently as follows. For a given $(i, j) \in D$, the maximization problem (6.2) boils down to a linear knapsack problem over the activities D , with capacity α and cost coefficients $c_{ij} = g_{ij} + \hat{\delta}_{ij}^i - \hat{\delta}_{ij}^j$ for activity (i, j) , and $c_{kl} = \hat{\delta}_{kl}^i - \hat{\delta}_{kl}^j$ for all activities $(k, l) \in D \setminus \{(i, j)\}$. For this problem, an optimal solution can be found efficiently by determining the $\lceil \alpha \rceil$ activities with largest nonnegative cost coefficient. As noted in Bertsimas et al. (2016), this can be done in $\mathcal{O}(|D| + \alpha \log \alpha)$ time, using an efficient partial sorting algorithm. Gorissen et al. (2015) note that it is not uncommon that the cutting-plane approach converges to optimality in only a small number of iterations.

6.3. Benchmark solution

One way of avoiding delays in a timetable, is to absorb them as soon as possible. Based on this principle, we generate a benchmark solution by requiring that each activity has enough time supplements to absorb any possible disturbance on that activity. That is, we require that

$$\pi_j - \pi_i + T p_{ij} \geq \ell_{ij} + \max_{\zeta \in Z_{\alpha}} g_{ij} \zeta_{ij}. \quad (6.3)$$

This imposes additional restrictions on the durations of activities and leads to solving a standard PESP model, where the efficiency is to be optimized. As an example, assume that all trip activities $(i, j) \in A$ can require up to one additional minute due to a disturbance. In that case, $g_{ij} = 1$ for all trip activities $(i, j) \in A$, and we require that the minimum duration of these activities is at least $\ell_{ij} + 1$. The only thing to be optimized is the additional duration of the trip and dwell time activities.

If a solution can be found, this solution is also feasible for (4.2), with a trivial linear decision rule (i.e., δ equals the zero-function: $\delta_{kl}^i = 0$ for all $i \in V, (k, l) \in D$), and no recovery budget. Such a benchmark solution may not exist if the worst-case disturbances are higher than what the restrictions of the underlying timetabling problem can handle. We note that our case studies never reach such high level, therefore we always find benchmark solutions.

The benchmark solution can be used as a starting solution to be improved upon in the further optimization. Furthermore, it implements a simple and natural idea to cope with the uncertainty, therefore it is a suitable comparison point for the outcome of our robust optimization models. The benchmark solution captures the spirit of current railway practices. Most operators follow the UIC recommendations by setting the time supplements uniformly to 6–8% of the technically minimal travel times. Our benchmark solution slightly deviates from these recommendations because applying them strictly would lead to an infeasible timetabling problem.

To summarize our approach, we propose a solution method for the RPTP that builds upon a linear decision rule. Combining this with the uncertainty region Z_α , this decision rule gives an upper bound for the original problem in (4.2). The 'linearized' model is solved by combining reformulation with cutting planes: Constraints (6.1b) and (6.1c) are the dual reformulation of the budget constraint (5.2a), and Constraints (6.1d)–(6.1f) represent (5.2b)–(5.2d) for only a subset of the disturbances. A simple heuristic is used to generate a feasible solution. Next, the model is solved iteratively, identifying disturbances corresponding to violated cuts in each iteration, until all constraints are satisfied, and hence an optimal solution to (6.1) has been found.

7. Computational experiments

We test the adjustable robust approach on real-life instances of Netherlands Railways (NS), the largest operator of passenger trains in the Netherlands. The instances feature the trains in the so-called 'Kop van Noord-Holland' region in a 1-h period. Note that Kroon et al. (2008) report their computational results for the same region.

The nominal Periodic Timetabling Problem is formulated by using standard methods from the PESP literature, including various pre-processing steps and using cycle periodicity formulation inequalities (Liebchen and Peeters, 2009).

We study the relation between the three aspects of the timetable: recovery budget (β), robustness and efficiency. Robustness is expressed by the parameter α of the uncertainty set, ranging between zero and the number of disturbable arcs $|D|$. Note that D is considered input. In what follows we measure a solution's *efficiency loss* as the total amount of time supplements needed, defined as

$$\sum_{(i,j) \in \tilde{A}} \pi_j - \pi_i + T p_{ij} - \ell_{ij}. \quad (7.1)$$

That is, we consider all trip and dwell activities, and we add up the difference between the planned activity time and the minimal activity time. The best possible efficiency is achieved if the total time supplement is zero.

We consider two instances. Case 1 is a study on a medium-sized, yet non-trivial, sub-network of the region. It includes 20 services connecting 20 stations (Fig. 5a). Case 1 has a PESP model with 589 nodes and 1668 activities. We assume that disturbances can occur on all dwell activities and most of the trip activities related to station Alkmaar (Amr): A trip activity can require up to 1 min more time than in the nominal problem, a dwell activity can require up to 2 min more; there are 22 disturbable activities in total. Thus, $g_{ij} = 1$ if $(i, j) \in D$ represents a trip arriving at or departing from Alkmaar, and $g_{ij} = 2$

Fig. 5. The railway networks considered in the case study.

if $(i, j) \in D$ represents a dwell activity in Alkmaar. The robustness parameter α ranges from 0 to 22. Case 1 has a feasible solution with maximal efficiency where the total travel time of the trains is $\sum_{(i,j) \in \tilde{A}} \ell_{ij} = 622$.

Case 2 is a large study that includes all 44 services of the region, connecting 38 stations (Fig. 5b). It covers about 8% of all passenger services of NS. Disturbances of up to 1 min can occur on all trip activities. Here, the set $D \subseteq A$ contains all driving activities and $g_{ij} = 1$ for all $(i, j) \in D$. The PESP model has 1321 nodes and 4862 activities of which 276 activities are disturbable. Case 2 has a feasible solution with maximal efficiency, too, where the total travel time of the trains is $\sum_{(i,j) \in \tilde{A}} \ell_{ij} = 1406$.

We emphasize that our case studies are non-trivial timetabling problems. The nominal optimization problem itself takes 6.2 seconds for Case 1 and about 5 min for Case 2. In fact, Case 2 approaches the limits of what periodic timetable optimization models can handle in a non-robust setting. Adding just a few services to Case 2 raises the typical computation times dramatically: Near-optimality cannot be reached after several hours of CPU times, the optimality gaps often remain above 90%.

Our computations are carried out on a machine with an Intel Xeon E5-2650 v2 2.60Ghz processor and with 64 GB RAM. The integer programs are solved by Cplex 12.8.0 with 15 parallel threads. The cutting-planes are implemented via a Lazy Constraint Callback.

7.1. Kop van Noord Holland: Case 1

In this section, we discuss the computational results for Case 1. We first investigate the relation between the robustness and the recovery budget for a maximally efficient timetable, and then the trade-off between the efficiency, robustness, and recovery budget.

7.1.1. Relation between robustness and recovery budget

Our first experiments aim at finding the relevant values of the recovery budget β . For each value $\alpha \in \{0, \dots, 22\}$ we adjust our robust optimization model to determine the smallest possible recovery budget under the assumption that all time supplements are zero (that is, that the solution is maximally efficient).

In order to speed up these calculations, we observe that an optimal solution for a higher value of α is always a feasible solution for a lower value of α . Therefore we solved the problem in a decreasing order of α , using the best solution found so far as a starting solution by means of Cplex's MIP start feature.

The solid line in Fig. 6 indicates the optimal budget for each value of α . Mind that $\alpha = 0$ is the nominal scenario (i.e., there is no disturbance), and $\alpha = 22$ is the scenario where each disturbable activity is indeed maximally disturbed. It turns out that each α admits a maximally efficient timetable. Therefore the obtained values indicate the largest amount of useful recovery budget for each α . We note that the recovery budget can rise to 360 min which is more than 57% of the total nominal travel times.

Fig. 6. Trade-off between robustness and recovery budget. Solid line: The optimal budget for each α . Other lines: Maximally efficient timetables, optimized for one specific value of α and evaluated under various levels of disturbances.

Fig. 7. Trade-off between efficiency and robustness for various recovery budgets β .

For each value of α , we have obtained a timetable that is maximally efficient and minimizes the required budget β . We now evaluate three of these timetables for all possible robustness values α . We consider the timetable found without considering robustness ($\alpha = 0$) and those found with a low and a high uncertainty level ($\alpha = 1$ and $\alpha = 22$ respectively). The dotted line in Fig. 6 shows the budget needed to recover the timetable optimized for $\alpha = 0$. The recovery budget of this particular solution starts growing rapidly with α , and no recovery is possible for $\alpha > 3$.

The dashed and dash-dotted line in the figure show the optimal recovery costs of two timetables that are optimized for a low and a high level of robustness, respectively. The benefit of taking a suitable level of robustness into account is clearly visible in the figure. Optimizing a timetable for a small level of robustness leads to recovery costs for high values of α that are higher than the optimum. For $\alpha = 22$, the recovery costs are 392 which is roughly 9% from the optimal costs (360). If a high level of uncertainty is taken into account, the recovery costs are close to the optimum. Only for medium values of α the costs are slightly higher than the optimum. It is interesting that for none of the timetables that we have found the recovery budget equals the optimal budget for every value of α .

In this figure we see that maximally efficient solutions give rise to very different recovery costs. The result demonstrates our method's ability to find the most robust solution among all efficient solutions.

7.1.2. Trade-off between efficiency, robustness, and recovery budget

In the second set of experiments we apply the adjustable robust optimization model as described in Sections 4–6: we seek the lowest efficiency loss for given combinations of robustness level α and recovery budget β . We set a time limit of 4 CPU hours for each optimization run. Again, we carry out the computations, for each given β , in a decreasing order of α , and we use the previous round's solution as a feasible starting solution.

Fig. 7 summarizes our results for $\alpha \in \{0, \dots, 22\}$ and $\beta \in \{0, 10, 20, 30, 50, 75, 100, 150\}$. The vertical axes indicate the efficiency loss, both in absolute and relative terms. The left axis shows the amount of added time supplements, as defined in (7.1). The right axis displays the relative objective value increase with respect to a maximally efficient solution. Not all combinations are solved to optimality within the time limit; the dashed lines indicate the best found lower bounds. In fact, the computation times vary heavily, both for finding good solutions, and for proving optimality.

Fig. 7 shows the intuitively clear relation between the recovery budget and the steepness of the curves. The time supplements trade delay propagation for delay absorption: A high budget allows for more propagation, and hence less time supplements are needed in the nominal timetable. The efficiency rapidly decreases for small β as the robustness level α increases, whereas large values of β admit a milder efficiency loss.

The curves for the individual budget values β reach their maximum height at surprisingly low values of α ; this effect becomes increasingly stronger as β decreases. For example, the efficiency loss for $\beta = 50$ is identical for each $\alpha \geq 13$. The uncertainty parameter $\alpha = 13$ requires such an amount of time supplements that is sufficient for all scenarios, even with $\alpha = 22$. That is, there exists an optimal solution for $\alpha = 13$ which is optimal for $\alpha = 22$, as well.

More detailed information about the results can be gained from Fig. 8. Here, heatmaps are shown for the solutions found with $\alpha \in \{1, 13, 22\}$ and $\beta \in \{0, 30, 100\}$. The darker the color on the segments between stations, the more time supplements are added here.

Fig. 8. Graphical results for Case 1.

Clearly, most of the time supplements are added around the station of Alkmaar (Amr), as is expected. Some are added at the tracks right next to Alkmaar, and some also a bit further away, even though the disturbance could not directly affect these tracks. Secondly, we see that time supplements are added if there is no recovery budget available ($\beta = 0$) and that the solution hardly changes if more uncertainty is considered. If more budget is available, we see that a higher level of uncertainty is needed before time supplements are necessary. This confirms that our model indeed adds time supplements at reasonable locations and first uses the available recovery budget whenever possible.

The main benefit of our approach is the mere fact that it enables one to carry out detailed studies on the trade-off between robustness, efficiency and recovery budget. The decision makers gain a valuable insight into the consequences of sacrificing some efficiency for the sake of improved robustness.

7.2. Kop van Noord Holland: Case 2

Case 2 contains all services of the 'Kop van Noord Holland' region, and disturbances are allowed on each trip activity. This is a non-trivial study; the nominal timetabling problem, without any robustness consideration, needs about 5 min to be solved to optimality. We want to explore the capabilities and limitations of our adjustable robust approach by considering an underlying problem that is close to the limitations of the non-robust model.

Table 1
Results for the ‘Kop van Noord Holland’ – Case 2.

α	$\beta = 10$				$\beta = 100$			
	UB	LB	Opt.	Impr.	UB	LB	Opt.	Impr.
1	38	2	95%	86.2%	0	0	0%	100.0%
2	190	6	97%	31.2%	201	0	100%	27.2%
3	199	13	94%	27.9%	201	0	100%	27.2%
4	211	33	85%	23.6%	201	0	100%	27.2%
5	219	40	82%	20.7%	204	0	100%	26.1%
10	236	117	50%	14.5%	204	0	100%	26.1%
25	266	164	38%	3.6%	204	0	100%	26.1%
50	266	216	19%	3.6%	204	56	72%	26.1%
100	266	244	8%	3.6%	212	31	85%	23.2%
150	266	259	3%	3.6%	212	123	42%	23.2%
200	266	262	1%	3.6%	212	198	6%	23.2%
250	266	266	0%	3.6%	225	180	20%	18.5%
276	266	266	0%	3.6%	225	180	20%	18.5%

(a) The case $\beta = 10$.

(b) The case $\beta = 100$.

Fig. 9. Trade-off between efficiency and robustness in Case 2 for various recovery budgets β . Solid line: objective value; dashed line: lower bound.

We consider uncertainty parameters $\alpha \in \{1, \dots, 5, 10, 25, 50, 100, 150, 200, 250, 276\}$, recovery budgets $\beta \in \{10, 100\}$, and we minimize the efficiency loss for each combination of these parameters. Recall that $\alpha = 0$ is the nominal scenario, and $\alpha = 276$ is the scenario where each disturbable activity takes a disturbance of 1 min. Each optimization run has a time limit of 5 h.

The budgets of 10 and 100 min comprise 0.71% and 7.1% of the total train travel time of an optimally efficient solution, respectively. For a comparison, the timetables of NS feature travel times where the planned time supplement is 7–8% of the technically minimal travel times. We report our computational results with a wider range of budget values in [Appendix B](#).

We compare the model's outcome to the benchmark solution described in [Section 6.3](#). This benchmark timetable adds (at least) 1 min time supplement to each disturbable activity; its efficiency loss is 276 which is 19.6% of 1406, the total train travel time of an optimally efficient solution. We take various measures to reduce the solution time of the adjustable robust approach. We use the benchmark timetable as a feasible starting solution. Also, similarly to [Section 7.1](#), we start the computations with $\alpha = 276$ and we proceed toward $\alpha = 1$ using the solutions of the previous rounds as a starting point.

[Table 1](#) and [Fig. 9](#) summarize our results with the robust optimization approach using a linear decision rule in terms of the best objective value and the best lower bound of the MIP models. In addition, we report an upper bound on the optimality gap and the improvement upon the benchmark solution.

Our model is able to find feasible solutions that significantly improve the benchmark timetable. For $\alpha \geq 2$, our solutions have 3.6 to 31.2% less efficiency loss. The improvement is limited under the tight recovery budget $\beta = 10$. However, the more realistic budget $\beta = 100$ allows for an improvement of 18.5% even under the highest disturbance level, and for slightly larger improvements for lower disturbance levels. For $\alpha = 1$ we find dramatic improvements; the choice of $\beta = 100$ admits a maximally efficient robust timetable. We note that the cases with $1 < \alpha < 275$ are intractable by scenario-based robust optimization approaches. Indeed, the uncertainty regions have, for integer α , at least $\binom{276}{\alpha}$ vertices; this results in more than 10^{10} scenarios for $\alpha = 5$, and more than 10^{81} scenarios for $\alpha = 150$.

Fig. 10. Graphical results for Case 2.

In addition, our model provides meaningful lower bounds on the efficiency loss (see again Fig. 9). They provide a valuable insight to decision makers for evaluating the candidate timetables. The proven optimality gap is remarkably low for $\alpha \geq 100$ (with a budget of 10) and for $\alpha \geq 200$ (with a budget of 100). It turns out that the small improvement percentage 3.6% for $\beta = 10$ (see Table 1) is often paired with a lower bound that prevents substantially better solutions. We do admit, though, that the optimality gaps for $\alpha < 100$ are high. This is also reflected in the fact that the solutions for $\alpha = 2$ and $\alpha = 3$ have a lower objective value with $\beta = 10$ than with $\beta = 100$, even though a higher budget allows for a better solution. The superior solutions for $\alpha = 1$ suggest that the large gaps are likely caused by not finding near optimal solutions. For situations with moderate disturbance, better solutions might be found with a different solution methodology.

Fig. 10 displays the heatmaps of the solutions that are found with two different combinations of α and β . In the case with small disturbances, we see that the majority of the time supplements are added at the west side of the network and at the tracks in the diagonal direction. The reason for this is that the majority of the trains use these tracks and thus a correct execution of the timetable is very important here. Furthermore, time supplements are added on the very north of the network, as there is only single track available here. Hence, the timetable is very sensitive to any disturbance and delays easily propagate to other trains using the same tracks. With an increased recovery budget and more uncertainty, time supplements are added throughout the network, but the largest increase is seen on the southern tracks.

8. Conclusion

In this paper, we proposed an adjustable robust extension of the Periodic Timetabling Problem (PTP): the Robust Periodic Timetabling Problem (RPTP). The RPTP is a challenging optimization problem, as the non-robust PTP is already known to be computationally difficult. We developed a solution method for the RPTP for a class of uncertainty regions. This class relates closely to uncertainty regions known in the robust optimization literature, and allows for a natural quantification of the robustness of a timetable.

The developed solution method consists of two parts. We first reduce the model size by assuming a linear decision rule, i.e., the recourse actions depend linearly on the disturbances. The size of the resulting model is only a fraction of the original model's size, albeit at a loss of freedom in the decision rule. As a result, this approach can be considered to be heuristic. The new model is solved using dual reformulation techniques and cutting-plane methods, both well-established techniques in the robust optimization literature.

We showed that the RPTP can be solved for practical-sized instances by applying the solution method to instances of Netherlands Railways (NS). In particular, we studied the interplay between the severity of the disturbances, the total amount of timetable adjustments, and the efficiency of the undisturbed timetable. Our experiments show that the RPTP remains computationally tractable: Good solutions for challenging instances can be found in reasonable time. We also saw, however, that for large scale instances with moderate disturbance there remains room for improvement.

Interesting further research would include the analysis of different, more complex, decision rules. In particular, it would be interesting to quantify the possible efficiency loss due to a linear decision rule. Note, however, that this would most likely greatly complicate any possible solution method. Furthermore, it would be interesting to consider a different class of uncertainty regions, based on, e.g., historical data.

Further research is needed to refine the solution methodology in order to find better solutions in situations with moderate disturbances.

Acknowledgments

We thank Prof. Dick den Hertog for his valuable comments and suggestions, which were helpful in improving this work.

Appendix A. Robust Counterpart

We now show how to derive the robust counterpart of (5.2a). We first reformulate the general constraint

$$a + \mathbf{b}^\top \boldsymbol{\zeta} \leq u \quad \forall \boldsymbol{\zeta} \in Z, \quad (\text{A.1})$$

with $a \in \mathbb{R}$, and $\mathbf{b}$ and $\boldsymbol{\zeta} \in \mathbb{R}^d$, following the methodology of Ben-Tal et al. (2009a). We assume that the uncertainty set Z is given by

$$Z = \{\boldsymbol{\zeta} \in \mathbb{R}^d : \mathbf{D}\boldsymbol{\zeta} \leq \mathbf{d}\}, \quad (\text{A.2})$$

for some $\mathbf{D} \in \mathbb{R}^{p \times d}$ and $\mathbf{d} \in \mathbb{R}^p$. Constraint (A.1) can equivalently be written as

$$a + \max_{\boldsymbol{\zeta} \in \mathbb{R}^d} \{\mathbf{b}^\top \boldsymbol{\zeta} : \mathbf{D}\boldsymbol{\zeta} \leq \mathbf{d}\} \leq u. \quad (\text{A.3})$$

By applying LP-duality, this transforms to

$$a + \min_{\boldsymbol{\theta} \in \mathbb{R}_+^p} \{\mathbf{d}^\top \boldsymbol{\theta} : \mathbf{D}^\top \boldsymbol{\theta} = \mathbf{b}\} \leq u. \quad (\text{A.4})$$

It is easily seen that (A.4) can only be satisfied if we can find a $\boldsymbol{\theta}$ such that

$$a + \mathbf{d}^\top \boldsymbol{\theta} \leq u \quad (\text{A.5a})$$

$$\mathbf{D}^\top \boldsymbol{\theta} = \mathbf{b} \quad (\text{A.5b})$$

$$\boldsymbol{\theta} \in \mathbb{R}_+^p. \quad (\text{A.5c})$$

The system of inequalities (A.5) is known as the *tractable reformulation* of (A.1).

We now consider the special case where $Z = Z_\alpha$. This implies that $\mathbf{D}^\top = [\mathbf{1}, -\mathbf{I}, \mathbf{I}]$, with $\mathbf{I}$ the identity matrix of dimension d , and $\mathbf{d}^\top = [\alpha, \mathbf{0}^\top, \mathbf{1}^\top]$. Hence, if we write $\boldsymbol{\theta}^\top = [\lambda, \boldsymbol{\phi}^\top, \boldsymbol{\mu}^\top]$, where $\lambda \in \mathbb{R}_+$ and $\boldsymbol{\phi}, \boldsymbol{\mu} \in \mathbb{R}_+^d$, and use the definition of $\mathbf{D}$ and $\mathbf{d}$, we obtain the constraints

$$a + \alpha\lambda + \mathbf{1}^\top \boldsymbol{\mu} \leq u \quad (\text{A.6a})$$

$$\lambda \mathbf{1} - \boldsymbol{\phi} + \boldsymbol{\mu} = \mathbf{b} \quad (\text{A.6b})$$

$$\lambda \in \mathbb{R}_+ \quad (\text{A.6c})$$

$$\boldsymbol{\phi}, \boldsymbol{\mu} \in \mathbb{R}_+^d. \quad (\text{A.6d})$$

The $\boldsymbol{\phi}$ variables function solely as slack variables, and hence the model can be written in the more compact form

$$a + \alpha\lambda + \mathbf{1}^\top \boldsymbol{\mu} \leq u \quad (\text{A.7a})$$

$$\lambda \mathbf{1} + \boldsymbol{\mu} \geq \mathbf{b} \quad (\text{A.7b})$$

$$\lambda \in \mathbb{R}_+ \quad (\text{A.7c})$$

$$\boldsymbol{\mu} \in \mathbb{R}_+^d. \quad (\text{A.7d})$$

Appendix B. Additional Numerical Results for Case 2

Table 2 summarizes our computational results on Case 2, using the recovery budget values {5, 10, 20, 50, 100, 200}.

Table 2

Results for the 'Kop van Noord Holland' – Case 2.

α	$\beta = 5$		$\beta = 10$		$\beta = 20$		$\beta = 50$		$\beta = 100$		$\beta = 200$	
	UB	LB	UB	LB	UB	LB	UB	LB	UB	LB	UB	LB
1	103	6	38	2	11	0	0	0	0	0	2	0
2	213	37	190	6	202	0	213	0	201	0	117	0
3	228	53	199	13	202	0	213	0	201	0	178	0
4	236	64	211	33	211	1	213	0	201	0	178	0
5	237	123	219	40	212	10	213	0	204	0	178	0
10	271	143	236	117	220	65	213	1	204	0	178	0
25	271	185	266	164	256	49	213	55	204	0	181	0
50	271	249	266	216	256	166	232	55	204	56	183	0
100	271	263	266	244	256	221	232	157	212	31	183	61
150	271	269	266	259	256	239	233	182	212	123	192	85
200	271	270	266	262	256	249	233	213	212	198	192	57
250	271	264	266	266	256	255	233	221	225	180	218	147
276	271	271	266	266	258	255	234	225	225	180	217	138

References

- Ben-Tal, A., El Ghaoui, L., Nemirovski, A., 2009a. *Robust Optimization*. Princeton University Press.
- Ben-Tal, A., Golany, B., Nemirovski, A., Vial, J., 2005. Retailer-supplier flexible commitments contracts: a robust optimization approach. *Manuf. Serv. Oper. Manage.* 7 (3), 248–271.
- Ben-Tal, A., Golany, B., Shtern, S., 2009b. Robust multi-echelon multi-period inventory control. *Eur. J. Oper. Res.* 199 (3), 922–935.
- Bertsimas, D., Brown, D., Caramanis, C., 2011. Theory and applications of robust optimization. *SIAM Rev.* 53 (3), 464–501.
- Bertsimas, D., Dunning, I., Lubin, M., 2016. Reformulation versus cutting-planes for robust optimization. *Comput. Manage. Sci.* 13 (2), 195–217.
- Cacchiani, V., Huisman, D., Kidd, M., Kroon, L., Toth, P., Veelenturf, L., Wagenaar, J., 2014. An overview of recovery models and algorithms for real-time railway rescheduling. *Transp. Res. Part B* 63, 15–37.
- Cacchiani, V., Toth, P., 2012. Nominal and robust train timetabling problems. *Eur. J. Oper. Res.* 219 (3), 727–737. doi:10.1016/j.ejor.2011.11.003.
- Caimi, G., Kroon, L., Liebchen, C., 2017. Models for railway timetable optimization: applicability and applications in practice. *J. Rail Transp. Plann. Manage.* 6 (4), 285–312. doi:10.1016/j.jtrpm.2016.11.002.
- Cicerone, S., D'Angelo, G., Di Stefano, G., Frigioni, D., Navarra, A., 2009. Recoverable robust timetabling for single delay: complexity and polynomial algorithms for special cases. *J. Comb. Optim.* 18 (3), 229–257. doi:10.1007/s10878-009-9247-4.
- Cicerone, S., Di Stefano, G., Schachtebeck, M., Schöbel, A., 2012. Multi-stage recovery robustness for optimization problems: a new concept for planning under disturbances. *Inf. Sci.* 190, 107–126. doi:10.1016/j.ins.2011.12.010.
- D'Ariano, A., Pacciarelli, D., Pranzo, M., 2007. A branch and bound algorithm for scheduling trains on a railway network. *Eur. J. Oper. Res.* 183 (2), 643–657.
- Dewilde, T., Sels, P., Cattrysse, D., Vansteenwegen, P., 2014. Improving the robustness in railway station areas. *Eur. J. Oper. Res.* 235 (1), 276–286. doi:10.1016/j.ejor.2013.10.062.
- Fischetti, M., Monaci, M., 2009. Light robustness. In: Ahuja, R., Möhring, R., Zaroliagis, C. (Eds.), *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 61–84. doi:10.1007/978-3-642-05465-5_3.
- Goerigk, M., 2015. Exact and heuristic approaches to the robust periodic event scheduling problem. *Public Transp.* 7 (1), 101–119. doi:10.1007/s12469-014-0100-5.
- Goerigk, M., Schöbel, A., 2010. An empirical analysis of robustness concepts for timetabling. In: Erlebach, T., Lübbecke, M. (Eds.), *10th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS'10)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 100–113. doi:10.4230/OASICS.ATMOS.2010.100.
- Goerigk, M., Schöbel, A., 2013. Improving the modulo simplex algorithm for large-scale periodic timetabling. *Comput. Oper. Res.* 40 (5), 1363–1370. doi:10.1016/j.cor.2012.08.018.
- Goerigk, M., Schöbel, A., 2014. Recovery-to-optimality: a new two-stage approach to robustness with an application to aperiodic timetabling. *Comput. Oper. Res.* 52, 1–15. doi:10.1016/j.cor.2014.06.025.
- Gorissen, B., Yanikoğlu, İ., den Hertog, D., 2015. A practical guide to robust optimization. *Omega* 53, 124–137. doi:10.1016/j.omega.2014.12.006.
- Großmann, P., Hölldobler, S., Manthey, N., Nachtigall, K., Opitz, J., Steinke, P., 2012. Solving periodic event scheduling problems with SAT. In: Jiang, H., Ding, W., Ali, M., Wu, X. (Eds.), *Advanced Research in Applied Artificial Intelligence*. In: *Lecture Notes in Computer Science*, 7345. Springer Berlin Heidelberg, pp. 166–175.
- Iancu, D., Sharma, M., Sviridenko, M., 2013. Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61 (4), 941–956.
- Kroon, L., Huisman, D., Abbink, E., Fiole, 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.
- Kroon, L., Maróti, G., Retel Helmrich, M., Vromans, M., Dekker, R., 2008. Stochastic improvement of cyclic railway timetables. *Transp. Res. Part B* 42 (6), 553–570. doi:10.1016/j.trb.2007.11.002.
- Liebchen, C., 2008. The first optimized railway timetable in practice. *Transp. Sci.* 42 (4), 420–435. doi:10.1287/trsc.1080.0240.
- Liebchen, C., Lübbecke, M., Möhring, R., Stiller, S., 2009. The concept of recoverable robustness, linear programming recovery, and railway applications. In: Ahuja, R., Möhring, R., Zaroliagis, C. (Eds.), *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1–27. doi:10.1007/978-3-642-05465-5_1.
- Liebchen, C., Möhring, R., 2007. The modeling power of the periodic event scheduling problem: railway timetables—and beyond. In: Geraets, F., Kroon, L., Schöbel, A., Wagner, D., Zaroliagis, C. (Eds.), *Algorithmic Methods for Railway Optimization: International Dagstuhl Workshop, Dagstuhl Castle, Germany, June 20–25, 2004, 4th International Workshop, ATMOS 2004, Bergen, Norway, September 16–17, 2004, Revised Selected Papers*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 3–40. doi:10.1007/978-3-540-74247-0_1.
- Liebchen, C., Peeters, L., 2009. Integral cycle bases for cyclic timetabling. *Discrete Optim.* 6 (1), 98–109. doi:10.1016/j.disopt.2008.09.003.
- Liebchen, C., Schachtebeck, M., Schöbel, A., Stiller, S., Prigge, A., 2010. Computing delay resistant railway timetables. *Comput. Oper. Res.* 37 (5), 857–868. doi:10.1016/j.cor.2009.03.022.
- Louwerse, I., Huisman, D., 2014. Adjusting a railway timetable in case of partial or complete blockades. *Eur. J. Oper. Res.* 235 (3), 583–593.
- Lusby, R., Larsen, J., Bull, S., 2017. A survey on robustness in railway planning. *Eur. J. Oper. Res.* doi:10.1016/j.ejor.2017.07.044.
- Maróti, G., 2017. A branch-and-bound approach for robust railway timetabling. *Public Transp.* 9 (1), 73–94. doi:10.1007/s12469-016-0143-x.
- Meng, L., Zhou, X., 2011. Robust single-track train dispatching model under a dynamic and stochastic environment: a scenario-based rolling horizon solution approach. *Transp. Res. Part B* 45 (7), 1080–1102.

- Ministerie van I&W, KiM., Mobiliteitsbeeld2017. <https://www.kimnet.nl/mobiliteitsbeeld#bereikbaarheid-article6>. Accessed: 2018-10-22.
- Nachtigall, K., 1999. Periodic network optimization and fixed interval timetables.. Universität Hildesheim Habilitation thesis.
- Nachtigall, K., Opitz, J., 2008. Solving periodic timetable optimisation problems by modulo simplex calculations. In: Fischetti, M., Widmayer, P. (Eds.), 8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS'08). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany doi:[10.4230/OASICS.ATMOS.2008.1588](https://doi.org/10.4230/OASICS.ATMOS.2008.1588).
- Odijk, M., 1996. A constraint generation algorithm for the construction of periodic railway timetables. Transp. Res. Part B 30 (6), 455–464. doi:[10.1016/0191-2615\(96\)00005-7](https://doi.org/10.1016/0191-2615(96)00005-7).
- Peeters, L., 2003. Cyclic Railway Timetable Optimization. Erasmus University Rotterdam PhD thesis.
- Pellegrini, P., Marlière, G., Pesenti, R., Rodriguez, J., 2015. Recife-milp: an effective milp-based heuristic for the real-time railway traffic management problem. IEEE Trans. Intell. Transp. Syst. 16 (5), 2609–2619.
- Sels, P., Dewilde, T., Cattrysse, D., Vansteenwegen, P., 2016. Reducing the passenger travel time in practice by the automated construction of a robust railway timetable. Transp. Res. Part B 84, 124–156. doi:[10.1016/j.trb.2015.12.007](https://doi.org/10.1016/j.trb.2015.12.007).
- Serafini, P., Ukovich, W., 1989. A mathematical model for periodic scheduling problems. SIAM J. Discrete Math. 2 (4), 550–581. doi:[10.1137/0402049](https://doi.org/10.1137/0402049).
- Veelenturf, L., Kidd, M., Cacchiani, V., Kroon, L., Toth, P., 2016. A railway timetable rescheduling approach for handling large-scale disruptions. Transp. Sci. 50 (3), 841–862. doi:[10.1287/trsc.2015.0618](https://doi.org/10.1287/trsc.2015.0618).
- Yanikoğlu, İ., Gorissen, B., den Hertog, D., 2018. A survey of adjustable robust optimization. Eur. J. Oper. Res. doi:[10.1016/j.ejor.2018.08.031](https://doi.org/10.1016/j.ejor.2018.08.031).
- Zhen, J., den Hertog, D., Sim, M., 2018. Adjustable robust optimization via Fourier–Motzkin elimination. Oper. Res. 66 (4), 1086–1100.