ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Master Thesis Operations Research and Quantitative Logistics

An iterative framework for rolling stock rescheduling with railway infrastructure availability constraints

Name student: Jia Hui Zhu

Student ID number: 498700

Supervisor: Prof. dr. D. Huisman

Second assessor: Dr. T.A.B. Dollevoet

Company supervisor: P.J. Fioole

Date final version: August 26, 2022

The content of this thesis is the sole responsibility of the author and does not reflect the view of the supervisor, second assessor, Erasmus School of Economics or Erasmus University.

Abstract

Disruptions on the railway network lead to reduced availability of the railway infrastructure. This affects passengers that travel by train as it can cause delays and even complete inaccessibilities of certain stations due to cancellations of train services. In the face of such disruptions, rolling stock dispatchers are tasked with adjusting the rolling stock schedule in real-time such as to minimize the inconveniences for passengers whilst sticking as close as possible to the original plan. Previous research has often made simplifying assumptions regarding the practical possibilities of changing train compositions and performing shunting movements. In this thesis, we focus on developing a rolling stock rescheduling method which ensures feasibility with respect to the availability of the railway infrastructure. In particular, we research the possibility of performing shunting movements at stations which in practice do not allow for shunting, due to the large number of trains that pass through or due to the complexity of the station layout. We introduce an iterative rolling stock rescheduling algorithm which alternates between creating an interim rolling stock schedule and extracting the suggested shunting movements which can be performed whilst simultaneously forbidding the shunting movements which cause conflicts with other trains due to minimum headway time restrictions. We test our solution approaches with disruption instances that contain complete railway blockages throughout the Netherlands. To counter the disruptions, we allow for shunting at the station Utrecht Centraal and model the exact infrastructure of the station to evaluate the feasibility of the suggested shunting movements. Our algorithm succeeds in adjusting the rolling stock schedule within running times of around a few minutes. We successfully identify feasible shunting movements and therefore improve upon the rolling stock schedule that would otherwise be obtained in the case that performing shunting movements at Utrecht Centraal is prohibited.

Preface

This thesis makes up my graduation project to obtain the master's degree in Operations Research and Quantitative Logistics and therefore marks the end of my status as a student at Erasmus University Rotterdam. I performed this research in collaboration with Netherlands Railways (NS) during my time as an intern in the Performance Management and Innovation (PI) department.

I would first like to thank my daily supervisor from NS, Pieter-Jan Fioole, for the guidance and for the weekly meetings which helped me stay on track and helped me complete my project. I would also like to thank my supervisor from Erasmus University Rotterdam, Dennis Huisman, for all his input and feedback during our (almost) bi-monthly sessions. Additionally, I would like to thank everyone from team BOS and from the department PI who supported me during my journey as an intern. In particular, I would like to thank Gábor and Maaike for providing me with useful data and insights and my fellow interns Camiel, Marja and Susan for turning the long office days into convivial get-togethers.

Finally, I would like to thank my friends and family for the continuous support and curiosity regarding (the course of) my thesis. My dearest Anni, Anna, Mert and Graham, thank you for being you.

Jia Hui Zhu
August 26, 2022

Contents

1 Introduction

The Dutch railway network plays a big role in the mobility system of the Netherlands and is mainly operated by Netherlands Railways (NS). Several elements of the existing train network are intertwined and simultaneously implemented and maintained, namely the timetable, the planning of rolling stock and the scheduling of crew. Disruptions on the railway network, such as a missing rolling stock unit in a train trip or a complete blockage of a railway track as a result of an accident or a train breakdown, require each of the plans to be updated to minimize the deviations from the original plan and to minimize the inconveniences for passengers.

We will focus our research on the rescheduling of rolling stock units in the case of disruptions in the railway system. It is necessary that the rolling stock is efficiently operated and utilized to maintain a good balance between the service to passengers and the incurred operational costs. Unavailability of rolling stock influences passenger satisfaction, as passengers possibly have to stand or even wait for the next train in case not enough trains are available to facilitate the service. An efficient rescheduling process reduces such problems and also reduces the number of trains that are needed as a buffer to catch disrupted train services, leading to a reduction in the capital costs for the train operator.

Currently, the rescheduling of rolling stock after disruptions is mostly done by hand. As a result, railway planners manually try to find an updated rolling stock allocation without the use of any mathematical methods or solvers. This can largely be attributed to the difficulty of the problem; previous research has succeeded in finding a rolling stock schedule for a simplified version of the railway network, leaving finer details with respect to shunting out of the picture. For a rolling stock plan to work in practice, feasible train compositions need to be assigned to each train trip, the associated shunting movements need to be executable and shunting duties need to be created for the limited number of present shunting drivers such that all shunting movements are performed. Past papers have often tackled these problems separately, meaning that there have always been some practical requirements that are not considered in their solution approaches. One of these requirements is the availability of train tracks for the shunting movements to take place. To the best of our knowledge, no previous research has laid their focus on ensuring feasibility of a rolling stock schedule with regard to this aspect. In practice, this is an important part of rolling stock rescheduling; if the rolling stock movements are not possible due to unavailability of railway tracks, the entire plan could fall apart and could need to be adjusted.

We introduce an iterative rolling stock rescheduling algorithm which ensures that all shunting movements are feasible with respect to the available infrastructure. The basis for our algorithm is the composition model by Fioole et al. (2006), which outputs a rolling stock circulation for the passenger train timetable. To check whether each of the suggested trips and shunting movements can be performed, we implement an infrastructure model which is based on the model by Van Aken et al. (2017) for solving the Train Timetable Adjustment Problem (TTAP), which aims at finding an alternative timetable that minimizes the deviation from the original timetable in terms of the incurred delays and cancellations of train services. Our tailored version of this model represents every shunting movement as an event that takes place on a specific track at a specific station and outputs whether a shunting movement is possible and whether other trains need to be shifted to make the shunting movement possible. Generally, rolling stock dispatchers do not have the authority to delay other trains to perform shunting movements. We therefore allow trains to be shifted by at most a few minutes, such as to stay close to the measures that are taken in practice. Our algorithm iteratively alternates between the composition model and the infrastructure model to find a feasible rolling stock reschedule with respect to the available infrastructure.

We test our algorithm on five disruption instances which consist of full blockages that take place at and around Utrecht Centraal, which is the station whose exact infrastructure we model and analyze. In practice, shunting at Utrecht Centraal is assumed to be impossible due to the large number of trains that pass through the station and is therefore not considered a viable option during the creation of the rolling stock schedule. Our results show that for disruption instances that take place adjacent to Utrecht Centraal and disruption instances that take place somewhere else in the country, our algorithm successfully identifies feasible shunting movements that can take place at Utrecht Centraal to counter the disruption and to improve the rolling stock schedule with respect to several operational, passenger service level and stick-to-the-plan objectives compared to the situation in which shunting at Utrecht Centraal is not considered. We furthermore find that our algorithm performs well in case we put emphasis on any one of the aforementioned objectives and that our approach also provides added value in the case that all trains are fixed and cannot be shifted in time.

The remainder of this thesis is organized as follows. Section 2 provides a detailed description of the problem we address. Literature covering similar problems is discussed in Section 3. Section 4 covers the models and methods we use for the generation of our inputs. Section 5 provides a mathematical formulation for our two solution components and introduces the iterative rolling stock rescheduling algorithm. Computational results and a sensitivity analysis of our algorithm are presented in Section 6. Finally, Section 7 gives conclusions and suggestions for future research.

2 Problem description

In this section, we will elaborate on the planning process of NS, introduce the used terms and concepts and give a formal problem description. Section 2.1 describes the different phases of the sequential planning process. Section 2.2 discusses the occurrence of disruptions and the susceptibility of the different phases of the planning process to such disturbances; fast real-time rescheduling needs to take place to minimize the inconvenience for passengers. Section 2.3 covers several aspects of existing rolling stock rescheduling methods and introduces the infrastructure framework that will be researched in this thesis. Section 2.4 formalizes the problem that we study in this research and specifies our objectives and assumptions.

2.1 The planning process at NS

Currently, the planning of the railway operations by NS is divided into several problems which are solved separately and sequentially. First, the passenger demands for each pair of origin and destination stations are estimated. This is used to create an origin-destination matrix, which forms the input for the line planning process. The train lines which are operated are created next. A train line consists of an origin and destination station and the route along which it travels, i.e. the intermediate stations that it visits along the way. A trade-off is present here between the convenience for passengers and the operational costs. Passengers prefer direct connections, which, to a certain extent, would lead to geographically long train lines if the number of direct travellers is maximized. Since the number of passengers can differ substantially across the various regions that long train lines would visit, however, this can lead to inefficient utilisation of rolling stock.

A rolling stock type is then assigned to the train line, which specifies at which stations the train line will stop. For passenger trains, we distinguish between regional trains, or *sprinters*, which stop at all passenger stations, long-distance trains, or *intercities*, which only stop at major stations and high-speed trains, or *intercity direct*, which make use of high-speed routes. A frequency is also chosen for each train line, which determines how many times the train line is operated per hour.

Based on the created line plan, the timetable is constructed. The Dutch railway timetable, called the *basic hour pattern*, consists of a cyclic timetable with a cycle time of one hour which forms the basis for every hour of every day. The basic hour pattern contains all passenger train lines, as well as freight trains and empty trains going to and from turning locations. We will refer to all train services in the timetable as *planned trains*. Each train service typically makes use of one rolling stock type and contains the tracks, or *infra lines* and stations or junctions, or *infra points*, that are part of its service, along with the associated running and dwell times. The infra lines and infra points are connected through train activities, which can be driving between infra points, passing through an infra point and having a short or long stop. Several requirements need to hold for the created timetable, such as a minimum headway time between two trains that make use of the same track, a minimum dwell time at stopping stations to allow passengers to get on and off the train, minimum travel times including buffers between two subsequent stations and a minimum turnaround time at terminal stations for each train. Furthermore, a timetable which satisfies transfers as much as possible and which distributes trains on the same route as equally as possible is desired.

Rolling stock is then assigned to the services in the timetable. A rolling stock schedule specifies which rolling stock units are assigned to which trip in the timetable, as well as all movements that the rolling stock units make whilst empty. Different types of rolling stock units are available for the different train types. Each train unit is self-propelled and can be controlled and steered by a driver from both sides. In practice in the Netherlands, the scheduling of rolling stock for the intercity trains and for the sprinter trains is split up into separate processes. We will only create a rolling stock schedule for the intercity train lines. Table 1 shows the six different rolling stock types that are used for intercity trains and their associated lengths, number of seats, number of carriages and the number of units available on the railway network.

Table 1: Rolling stock types for intercity trains.

Name	**Length (m)**	**Number of seats**	**Number of carriages**	**Number of units available**
ICM-3	80.6	228	3	76
ICM-4	107.1	299	4	45
VIRM-4	108.6	405	4	83
VIRM-6	162.1	597	6	72
DDZ-4	101.8	373	4	25
DDZ-6	154.0	607	6	18

The sequence of train units that is used for a train service is called a *composition.* The assignment of rolling stock units to train services needs to satisfy several restrictions, such as the rolling stock capacities of NS and the maximum length of the train compositions. Train units of the same type can be combined (split up) to create longer (shorter) trains. This process is called *(un)coupling* and is generally performed with rolling stock of the same type. Such composition changes can be performed at stations where the train stops. The order of the train units which make up the composition is of importance. For instance, if a composition consists of a rolling stock unit of type a and b, then the two possible compositions are ab and ba, with the first letter denoting the back of the train and the last letter denoting the front of the train. The (un)coupling of trains units can often only happen at either the front or at the back of the train. Therefore, it is relevant in which order the composition arrives at the station. Uncoupled train units are parked at shunting yards at the respective station until further use. Shunting drivers are responsible for performing these shunting movements.

Finally, a crew schedule is made. Drivers and conductors are assigned a duty, which is a sequence of tasks that typically corresponds to being on board of several consecutive train trips. Enough time needs to be available between two consecutive tasks to get to the starting location of the next task and each duty must contain a break roughly in the middle of the service.

Each of the different phases of the planning process is dependent on the previous phase(s) and feedback loops are present between the different stages. It is for instance possible to adjust the timetable to create a more beneficial rolling stock or crew schedule. Additionally, an important task of the operation control department is the real-time adjustment of the made plannings in the case of disruptions. This thesis will focus on the adjustment of the rolling stock planning after the occurrence of a disturbance, which will be described in more detail in the next section.

2.2 Rescheduling after disruptions

Disruptions lead to reduced availability of railway tracks, stations and train units. Such disruptions can for instance be the mechanical failure of certain sections of the railway infrastructure, due to which all services that use this specific section are cancelled. A train unit can also experience a mechanical failure, which causes this train unit to be unavailable for the services it is assigned to. This can either lead to partial blockages, in which certain infra lines and/or infra points have a reduced number of tracks and/or platforms to their disposal, or complete blockages, in which entire infra lines and/or infra points become completely inaccessible. As a result of these disruptions, the train units that were planned to operate the cancelled services are unable to reach the stations at which they were supposed to arrive at the end of these services. This means that they are also unable to start other services that they were potentially assigned, leading to more cancellations. Furthermore, maintenance requirements also make up part of the instances in which rescheduling needs to be done. Following a misplacement, a rolling stock unit that is planned to undergo maintenance may need to be extracted from the middle of an active composition, which requires several more rolling stock movements than initially planned.

The timetable and the rolling stock plan need to be adjusted to account for these disturbances. In practice, the timetable is usually adjusted first. This is done in such a way that further cancellations are minimized whilst also minimizing the operational costs and staying close to the original timetable. Several contingency plans exist in the current rescheduling procedures that specify which measures can be taken to update the timetable. This includes the retiming of trains, changing the route a train traverses, adding or removing stops and fully cancelling trains. The goal is then to create a feasible timetable which is as close as possible to the original timetable. Existing methods for rolling stock rescheduling will briefly be discussed in the next section. Finally, a new crew schedule needs to be determined given the updated timetable and the updated rolling stock plan.

2.3 Rolling stock rescheduling models

This thesis focuses on the procedures that are in place for the rescheduling of rolling stock. Obstructions on the railway network can have major consequences for the rolling stock planning; since train services are cancelled, rolling stock units can be stranded at their intermediate stations and this can lead to delays for passengers. Other types of disturbances, such as the shortcoming of a train unit in a passenger train composition, can lead to seat shortages and therefore decreased passenger satisfaction.

Current rolling stock rescheduling models focus on salvaging the schedule as much as possible with the available resources. Adjustments are typically made by incorporating the disturbance information in the existing solution methods for the rolling stock planning and generating a new rolling stock schedule with the reduced availability of rolling stock units as a result of the disruption. The resulting schedules are feasible on a zoomed-out network level. In practice, this can lead to infeasibilities on a local level, due to the negligence of local details. Several key aspects have been investigated in the literature that contribute to modelling the problem in a more realistic manner:

- The availability of shunting drivers. Each shunting yard has a limited number of shunting drivers available during a day. For each of these shunting drivers, a duty is created which consists of a sequence of shunting movements that needs to be performed. Sufficient time needs to be available to get from one task to the next and several labor regulations, such as an agreed starting and ending time and a meal break roughly in the middle of the duty, need to be satisfied. Hoogervorst (2021) introduced the Rolling Stock and Shunting Driver Rescheduling Problem (RSSDRP), which incorporates this aspect into their rolling stock rescheduling method in such a way that a set of feasible duties can be created for all present shunting drivers with each rolling stock planning that is produced.

- The availability of shunting tracks and the allocation of rolling stock units on shunting tracks. If rolling stock units are uncoupled from an incoming train composition, they are parked at a shunting yard until further use. Only the units that are parked at the front of shunting tracks are directly available. This can be troublesome if the parked rolling stock units are of a different type; generally, units can only be coupled to units of the same type. This means that a rolling stock unit of a specific type can be unreachable because another unit is parked in front of it on the same track. Freling et al. (2005) introduced the Train Unit Shunting Problem (TUSP), which models the exact shunting movements that take place by matching incoming movements to outgoing movements of the same unit type and on the same track. Haahr and Lusby (2017) further integrate this problem into the rolling stock rescheduling framework by creating a rolling stock plan that takes into account the lengths and capacities of the available shunting tracks.

- The availability of the railway infrastructure. Most disruptions take place during the day, whilst all planned trains are up and running. The viability of the rescheduling measures, such as changing train compositions and performing the required shunting movements, is significantly impaired by the busyness of the surrounding infrastructure. Safety regulations are in place regarding the spacing of trains; a minimum headway time needs to be present between two trains that make use of the same track or crossing. Additionally, the exact infrastructure is of importance; shunting is often only possible on one side of a track and not every track is connected to the shunting yard. Therefore, the exact track that a train arrives on and the order of its composition are relevant in determining a workable rolling stock schedule. To the best of our knowledge, previous research regarding the rescheduling of rolling stock has made assumptions regarding these aspects, leading to serious simplifications.

Our research will primarily focus on incorporating elements of the railway infrastructure into the creation of an adjusted rolling stock schedule. We will formally describe our problem in the next section.

2.4 Problem definition

The aim of this research is to extend the models for rolling stock rescheduling to ensure feasibility with respect to the available railway infrastructure. The input for our method consists of the original timetable, the original rolling stock schedule, information on the railway network layout, the different rolling stock units and types, the allowed compositions and composition changes for each trip and the disruption circumstances. The method then outputs a rolling stock schedule which specifies for each trip in the timetable the composition that is utilized, as well as the shunting movements that are performed. Moreover, our method ensures that all rolling stock movements satisfy all infrastructure related constraints of the stations that we consider, such as the availability of the railway tracks and the minimum headway times between all planned trains and all rolling stock movements.

Our method will create a rolling stock schedule that takes into account several objectives simultaneously:

- Minimizing the operational costs. We aim to minimize two different cost inducing processes:
  - → The number of carriage kilometers. To maximize the efficiency of a rolling stock schedule, it is beneficial to allocate the available rolling stock units as cost effectively as possible. For this purpose, the number of carriage kilometers driven is an accurate benchmark; driving a carriage costs traction power and after a certain number of driven kilometers, train units need to be inspected and possibly repaired.
  - → The number of shunting movements. Changing the composition of a train can be a cumbersome process; it can require a number of shunting movements from the platform to the shunting yard and vice versa. Shunting crew and the railway infrastructure also need to be available to perform these movements. Therefore, minimizing the number of performed shunting movements can contribute to lowering the incurred operational costs.
- Maximizing the service levels for passengers. For this purpose, we aim to minimize two criteria:
  - → The number of seat-shortage hours. One of the key performance indicators of NS is the probability of obtaining a seat in a train, with a target figure of 95 percent during peak hours. During disruptions, passenger flows are likely to accumulate at the stations adjacent to the disturbance. In these situations especially, is it important to use sufficient rolling stock for the planned train services. We will use the number of *seat-shortage* hours as our benchmark, which is the sum over all trips of the expected number of passengers without a seat multiplied by the duration of the respective trip.
  - → The number of cancellations from lack of rolling stock units. Disruptions can lead to decreased availability of the railway infrastructure. As a consequence, passengers experience nuisances in the form of train delays and cancellations. It is possible that, on top of such cancellations, cancellations occur because no rolling stock units are available to drive a train service. A rolling stock schedule should avoid incurring such cancellations as much as possible.

- Stick-to-the-plan. It is beneficial if the adjusted plan is similar to the original plan, since this limits the number of extra shunting movements that need to be performed at night to get each rolling stock unit to its starting position for the next day. Additionally, an advantage for both passengers and crew members is that they are familiar with the original plan and may prefer the rolling stock units to be operated as before. For this objective, we will look at several factors such as the number of trips which are assigned a different composition compared to the original plan and off-balances in the rolling stock inventories of stations.

A trade-off exists between the various objectives. For instance, utilizing large compositions for each train service ensures that a large number of passengers is able to secure a seat. Of course, this greatly increases the operational costs. We will study the importance of the different objectives and weigh them against each other.

Several assumptions are necessary to model our problem. Firstly, we assume that a rolling stock schedule only has to be made for the passenger trains. This could be passenger trains that are part of the timetable or empty rolling stock units that are transported to shunting yards or other stations. We assume that other requirements, such as maintenance and regular cleaning of rolling stock units, are planned in advance for all rolling stock units. The time and location at which such activities take place are fixed and a feasible rolling stock planning must therefore satisfy all such requirements.

We also assume that shunting movements can be performed during the night. In case the rolling stock inventories at the end of the day do not match the planned starting inventories at the start of the next day, we assume that this is restored during the night outside of the considered time horizon. Furthermore, we make the assumption that enough shunting drivers are present to perform all shunting movements. Other aspects of shunting crew scheduling, such as the creation of a feasible duty for each shunting driver which includes a meal break and sufficient time to travel from the ending location of one task to the starting location of the next task, are outside the scope of this thesis.

Our starting point is that the original timetable we use is conflict-free. We define a conflict as any violation of constraints, which in the case of the headway constraints, typically consists of two or more trains that do not satisfy the minimum headway time. The safety headway that must be present between two different trains is around three minutes, but this number varies depending on the exact location and trains that are involved. All headway times which are present in the original timetable that we use are viable and lead to no conflicts. The shunting movements that take place need to be planned around the planned trains in the timetable.

3 Literature review

Several papers have investigated methods for adjusting and rescheduling of rolling stock and the associated shunting movements in the case of disruptions. For this purpose, the explored solution methods range from solving Mixed Integer Linear Programming (MILP) formulations to developing neighborhood search algorithms and other heuristic approaches.

Cacchiani et al. (2014) give an overview of the disruption management process, which is typically split into several operational problems. They summarize the most important models and solution methods that have been developed in recent years for the rescheduling of the train timetable, the rolling stock and the crew. Furthermore, they describe various algorithms that integrate different rescheduling phases. Such methods are expected to find better solutions, at the cost of more complex models.

Fioole et al. (2006) provide the basis for much of the research done regarding the rescheduling of rolling stock. They address the rolling stock circulation problem, which, given the arrival and departure times of trains and the expected number of passengers on each train, assigns the available rolling stock to the different train services. Their composition model allows for the combining and splitting of trains by coupling or uncoupling rolling stock units before the train in question departs from its origin station, when it is dwelling at an intermediate station or after it has arrived at its destination. By minimizing an objective function that includes both operational costs and the service quality for passengers, the model is able to provide solutions within a few minutes of computation time that comply with market and quality requirements and that have been used for the creation of the weekly rolling stock schedule of NS.

Nielsen (2011) formalizes the Rolling Stock Rescheduling Problem (RSRP) by proposing a solution framework that separately generates circulations and duties for the rolling stock units. The circulation generation is performed with both the composition model of Fioole et al. (2006) and a task model, which is based on a constrained multi-commodity flow of the tasks of rolling stock units. The duties are created for the rescheduled circulation with a Mixed Integer Program (MIP) that aims to minimize changes from the planned duties.

Wagenaar et al. (2017) build on the RSRP by including the possibility of using dead-heading trips, which corresponds to moving empty rolling stock units during and after a disruption. Furthermore, their model incorporates information about the passenger flows to account for changes in passenger demands after a disruption. Their computational results show that using dead-heading trips reduces the number of cancelled train services and the number of seat-shortages compared to a model without dead-heading trips, and that larger train compositions are assigned to trips after the end of a disruption to account for the larger influx of passengers based on the adjusted passenger demands, resulting in fewer seat-shortages as well.

Hoogervorst (2021) describes the integrated Rolling Stock and Shunting Driver Rescheduling Problem (RSSDRP), which assigns compositions of rolling stock units to trips in such a way that all restrictions on the availability and the movement of rolling stock units for composition changes are satisfied and in such a way that feasible duties can be found for shunting drivers that perform these composition changes.

They tackle the problem in two ways. Firstly, they propose a Benders decomposition approach which generates cuts in case the rolling stock schedule does not allow for a feasible allocation of the shunting tasks among the shunting drivers. The master problem in this decomposition is based on the composition model proposed by Fioole et al. (2006) for the rescheduling of rolling stock, while the subproblem is a set partitioning formulation for the rescheduling of the shunting drivers. The second approach includes an MILP formulation that presents an arc-based model for the shunting driver rescheduling problem. Both methods are tested on instances of NS and both methods find optimal solutions for the integrated problem within reasonable running times whilst preventing infeasibilities that occur in case the two problems are solved sequentially.

Van Veen (2021) develops heuristic methods for the RSRP which combine the computational flexibility of the composition model formulation by Nielsen (2011) with a variable neighborhood descent that consists of two neighborhoods, namely the optimization of the deployment of one vehicle type in the composition model and the two-opt duty neighborhood of Hoogervorst et al. (2021) which switches duties between different vehicle types. Their computational results show that for four disruption instances on the Dutch railway network, the developed heuristic methods successfully decrease the running times compared to the exact composition model. Additionally, they show that it is beneficial for the comfort of passengers to incorporate dynamic passenger flows in the models and that a trade-off is present between passenger comfort and the stick-to-the-plan mentality of NS.

The previously mentioned papers have only looked at the possible compositions for each train service, without directly taking into consideration the logistics of the shunting yards themselves. Haahr et al. (2017) inspect the Train Unit Shunting Problem (TUSP), which assigns rolling stock from shunting yards to scheduled train services. This is done by matching rolling stock units to arriving and departing train services and allocating the created matchings to the shunting yard tracks and station platforms. The available infrastructure of the shunting yards is taken into account by imposing restrictions related to the length of the available tracks and platforms and the ordering and types of the residing trains. The problem is modelled as a constraint program and solved using column generation and a randomized greedy constructive heuristic. They additionally introduce a two-stage method, which solves the matching problem of train trips and the parking problem of the found matchings on the shunting tracks sequentially. Together with the randomized greedy constructive heuristic, these two methods provide feasible solutions in short computation times. Haahr and Lusby (2017) further integrate the shunting of rolling stock units in the rolling stock scheduling and propose two branch-and-cut procedures for solving this problem, which differ with respect to the formulation they use for the rolling stock scheduling. Similarly to Hoogervorst (2021), their results show that integrating the two phases leads to a more constrained problem which proves to be more beneficial than solving the two problems sequentially, as this can lead to infeasibilities, and only slightly increases the objective value compared to only solving the rolling stock scheduling problem.

Nielsen et al. (2012) implement a rolling horizon approach for the real-time disruption management of rolling stock: the rolling stock is rescheduled within a rolling horizon of limited length and the found plan is periodically updated with new information as time progresses. This approach is a heuristic problem decomposition which guides the rolling horizon solution to the desired rolling stock distribution by penalizing off-balances of rolling stock units at the end of the planning day. Due to the smaller size of the individual rolling stock scheduling problems, feasible solutions can be found for both the rolling stock and shunting plans in short computation times.

Van den Broek et al. (2022) split up the TUSP into four different aspects: the matching of incoming and outgoing units, the scheduling of the service tasks, the assignment of units to parking tracks and the routing of the units over the service site. They propose a local search model for the TUSP based on simulated annealing, which evaluates all four aspects simultaneously to create a shunting plan. This is done by creating a precedence graph which models the activities that take place on the service site. Their local search approach iteratively makes local changes to this train activity graph to improve the shunting plan. Comparing their heuristic to a tool built by NS for this purpose, called OPG, they find that their algorithm outperforms the OPG by being able to plan more trains and by being able to handle large real-world scenarios.

Our work primarily builds upon the composition model by Fioole et al. (2006) and the RSRP formulation by Nielsen (2011), which we use for the creation of the initial rolling stock schedule and the adjusted rolling stock schedule during disruptions, respectively. We create an iterative rolling stock rescheduling algorithm in which we combine the latter model with constraints of the Train Timetable Adjustment Problem (TTAP) formulation by Van Aken et al. (2017). This problem looks at the task of adjusting the timetable for planned maintenance instances whilst minimizing the deviations from the original timetable in terms of the number of cancellations and the incurred delays. We implement a tailored version of this model which detects, given the exact infrastructure of specific stations, which shunting movements can take place. The alternation between these two models provides the basis for our approach and will be elaborated on in the sections to come.

4 Input generation

As described in the problem definition, the aim of this thesis is to create a rolling stock rescheduling algorithm which performs shunting movements that are feasible with respect to infrastructure requirements that hold in practice, such as a minimum headway time between two trains that use the same track or crossing. This section describes all inputs that are necessary for this purpose and all inputs that we generate ourselves. First, we will describe all inputs that are fixed in Section 4.1. We will then describe the generation of the initial rolling stock schedule in Section 4.2 and the modelling of the railway infrastructure in Section 4.3.

4.1 Fixed inputs

For our algorithm, we will use the following fixed inputs:

- The original timetable. This is the timetable for an entire day without disruptions, which contains all passenger and freight train services that run for every hour of the day, described by all infra points that each train service visits and the times at which all arrival and departure events at these infra points take place. Each service in the timetable corresponds to a sequence of *trips*. A trip is a journey between two infra points, with a departure time at the departure track of the respective infra point and an arrival time at the arrival track of the respective infra point. The trips are connected by *transitions*, which describe which trips are performed in succession.

- Information on the different rolling stock types, the allowed compositions and composition changes. Each rolling stock type is unique with respect to its length and its seating capacity and only a specific number of rolling stock units is available in the network. In this thesis, we assume that (un)coupling of rolling stock units is only possible at the front and at the back of the composition. Furthermore, the compositions we work with consist of rolling stock units of the same type.

- Disruption instances. In the context of rolling stock rescheduling, any disruption that leads to deviations in the rolling stock schedule can be researched and resolved. This ranges from a composition that is too short or too long for a specific trip to complete blockages between two stations which result in all trips that make use of this infra line being cancelled. In this thesis, we will focus on disruptions of the latter kind. We assume that for the disruption instances, the adjusted timetable is given, which means that all affected trips are taken out of it. Since train services can partially still be completed up until the blockage, the transitions are adjusted as well. For instance, a train service can get stranded at a station adjacent to the disruption. A possible new transition is then to perform the service in reverse by connecting the last performed trip to a trip in opposite direction. This could be the first trip of the train service in opposite direction that is also cut off due to the disruption. We will look at disruptions which take place during rush hours, since this is when the largest number of trains are running and therefore makes for the most interesting problem instances. We will work with disruptions that last for two hours, as this is generally a realistic disruption duration and because this covers the entire rush hour. The developed rolling stock rescheduling algorithm can, however, be applied to any disruption.

Next, we will describe the generation of the initial rolling stock schedule.

4.2 Initial rolling stock schedule

One of the goals of rolling stock rescheduling is to create a schedule that deviates as little as possible from the original schedule. For this purpose, it is necessary to create an initial rolling stock schedule. We choose to generate an initial rolling stock schedule ourselves, as a real-life schedule does not satisfy the assumptions that we specify in this thesis. For instance, in practice more composition changes may be possible than exclusively (un)coupling rolling units at the front and at the back of compositions. Additionally, a real-life rolling stock schedule may not be optimal. This means that the quality of our algorithm, which generates the optimal adjusted schedule in case of the disruption, cannot objectively be determined.

We will look a formulation of the composition model by Fioole et al. (2006). The model allocates a composition of rolling stock units to every trip in a set of trips in the timetable $\mathcal{T}$. Each trip $t \in \mathcal{T}$ has an associated departure time $\tau_d(t)$, arrival time $\tau_a(t)$, departure station $s_d(t)$ and arrival station $s_a(t)$. The successor trip of trip t is denoted by $v(t)$; the composition that performs trip t performs trip $v(t)$ right after finishing trip t. The sets $\mathcal{T}_0$ and $\mathcal{T}_1$ denote the sets of trips which have no predecessor and successor trip, respectively. Furthermore, the set $\mathcal{C}$ contains all possible rolling stock compositions for each trip. Only certain rolling stock units can be combined to form compositions; often, only rolling stock of the same type is coupled to each other. A composition specifies which rolling stock type makes up the composition, as well as the position of each individual rolling stock unit in the composition. The order of the individual units is relevant for the shunting possibilities at stations; some stations only allow for (un)coupling rolling stock units at the front or rear of the composition, depending on the station infrastructure. It is also possible that a train service turns around, leading to a reversal of the composition without any shunting movements. Let C_t denote the set of compositions that are allowed for trip t. This set requires that the right rolling stock units are used for each trip, that enough seating capacity is available to provide a certain passenger service level and that the length of the composition does not exceed a certain upper bound. The set Γ_t denotes the possible transitions for trip t, which consist of pairs of compositions (c, c') such that $c \in C_t$ and $c' \in C_{v(t)}$ and such that the composition change from c to c' is allowed after trip t. The set of stations and/or infra points at which rolling stock units are available and can be (un)coupled is denoted by $\mathcal{S}$. The time it takes for an uncoupled rolling stock unit to be available for coupling at station s is denoted by $\rho(s)$. All rolling stock units types m are contained in the set M. Furthermore, n_m denotes the number of type m rolling stock units that are available on the railway network. Moreover, $\alpha^m_{c,c'}$ and $\beta^m_{c,c'}$ denote the number of uncoupled and coupled units, respectively, of type m during transition (c, c'). We also denote by $carr_c$ the number of carriages in composition c and by $seat^t_c$ the expected number of seat-shortages for trip t when composition c is used. The length of trip t in kilometers is represented by l_t and the duration of trip t in hours is represented by d_t.

To provide a mathematical formulation for the problem, we define the binary variable x^t_c which takes on the value 1 if composition $c \in \mathcal{C}$ is chosen for trip $t \in \mathcal{T}$, and 0 otherwise. Additionally, the binary variable $z^t_{c,c'}$ has value 1 if trip $t \in \mathcal{T}$ has composition $c \in \mathcal{C}$ and trip $v(t)$ has composition $c' \in \mathcal{C}$, and 0 otherwise. The variables c^t_m and u^t_m denote the number of rolling stock units m that are coupled to the composition right before trip t and uncoupled from the composition right after trip t, respectively. Additionally, the variable I^t_m expresses the number of rolling stock units of type m that are in inventory at station $s_d(t)$ right after the departure of trip t. The variables $I^s_{m,0}$, $I^s_{m,\text{end}}$ and $I^s_{m,\text{diff}}$ denote the initial inventory, the inventory at the end of the day and the absolute difference between the initial and ending inventory of rolling stock type m units at station s. For the objective function, we define five weights w_1, w_2, w_3, w_4 and w_5. A Mixed Integer Program (MIP) formulation is then as follows:

$$\begin{aligned}
\min \quad & \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}_t} w_1 \cdot carr_c \cdot l_t \cdot x^t_c + \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}_t} w_2 \cdot seat^t_c \cdot d_t \cdot x^t_c \\
& + \sum_{t \in \mathcal{T}} \sum_{m \in M} w_3 \cdot c^t_m + \sum_{t \in \mathcal{T}} \sum_{m \in M} w_4 \cdot u^t_m + \sum_{s \in \mathcal{S}} \sum_{m \in M} w_5 \cdot I^s_{m,\text{diff}} && (1) \\
\text{s.t.} \quad & \sum_{c \in \mathcal{C}_t} x^t_c = 1 && \forall t \in \mathcal{T} \quad (2) \\
& x^t_c = \sum_{c' \in \mathcal{C}_{v(t)} : (c,c') \in \Gamma_t} z^t_{c,c'} && \forall t \in \mathcal{T} \setminus \mathcal{T}_1,\ c \in \mathcal{C}_t \quad (3) \\
& x^{v(t)}_{c'} = \sum_{c \in \mathcal{C}_t : (c,c') \in \Gamma_t} z^t_{c,c'} && \forall t \in \mathcal{T} \setminus \mathcal{T}_1,\ c' \in \mathcal{C}_{v(t)} \quad (4) \\
& c^{v(t)}_m = \sum_{(c,c') \in \Gamma_t} \beta^m_{c,c'} z^t_{c,c'} && \forall t \in \mathcal{T} \setminus \mathcal{T}_1,\ m \in M \quad (5) \\
& u^t_m = \sum_{(c,c') \in \Gamma_t} \alpha^m_{c,c'} z^t_{c,c'} && \forall t \in \mathcal{T} \setminus \mathcal{T}_1,\ m \in M \quad (6) \\
& \sum_{s \in \mathcal{S}} I^s_{m,0} = n_m && \forall m \in M \quad (7) \\
& I^t_m = I^{s_d(t)}_{m,0} - \sum_{\substack{t' \in \mathcal{T} : s_d(t') = s_d(t), \\ \tau_d(t') \le \tau_d(t)}} c^{t'}_m + \sum_{\substack{t' \in \mathcal{T} : s_a(t') = s_d(t), \\ \tau_a(t') \le \tau_d(t) - \rho(s_d(t))}} u^{t'}_m && \forall t \in \mathcal{T},\ m \in M \quad (8) \\
& I^s_{m,\text{end}} = I^s_{m,0} - \sum_{t \in \mathcal{T} : s_d(t) = s} c^t_m + \sum_{t \in \mathcal{T} : s_a(t) = s} u^t_m && \forall s \in \mathcal{S},\ m \in M \quad (9) \\
& I^s_{m,\text{diff}} \ge I^s_{m,\text{end}} - I^s_{m,0} && \forall s \in \mathcal{S},\ m \in M \quad (10) \\
& I^s_{m,\text{diff}} \ge I^s_{m,0} - I^s_{m,\text{end}} && \forall s \in \mathcal{S},\ m \in M \quad (11) \\
& c^t_m, u^t_m, I^t_m \in \mathbb{R}^+ && \forall t \in \mathcal{T},\ m \in M \quad (12) \\
& I^s_{m,0}, I^s_{m,\text{end}}, I^s_{m,\text{diff}} \in \mathbb{R}^+ && \forall s \in \mathcal{S},\ m \in M \quad (13) \\
& x^t_c \in \{0,1\} && \forall t \in \mathcal{T},\ c \in \mathcal{C}_t \quad (14) \\
& z^t_{c,c'} \in \{0,1\} && \forall t \in \mathcal{T},\ (c,c') \in \Gamma_t \quad (15)
\end{aligned}$$

The objective function as given by (1) minimizes the number of carriage kilometers driven, the number of seat-shortages measured over time, the number of shunting movements and off-balances in the inventory of rolling stock units at the end of the day. Constraints (2) ensure that exactly one rolling stock composition is chosen for each trip t. Constraints (3) link the binary variables by guaranteeing that a composition c is chosen for trip t if and only if an allowed composition change is chosen such that composition c' is used for the successor trip $v(t)$ of trip t. Constraints (4) take on a similar role for the predecessor trips. Constraints (5) and (6) model the number of coupled and uncoupled rolling stock units for each composition change, respectively. Constraints (7) specify the initial rolling stock inventories at each station. Constraints (8) describe the inventory of rolling stock units after each trip. This is done by taking the initial inventory at the departure station of trip t, subtracting the number of rolling stock units that have been coupled to departing trains and adding the number of rolling stock units that have been uncoupled from arriving trains up until the departure of trip t. Constraints (9) model the ending inventories of each rolling stock unit type at each station, by adding all coupled units to and subtracting all uncoupled units from the initial inventory of the considered station. Constraints (10) and (11) ensure that the variable $I^s_{m,\text{diff}}$ takes on the absolute value of the difference between the ending and initial inventory of rolling stock unit type m at station s. Constraints (12) - (15) model the domains of the variables.

This model provides a composition allocation for each train trip, as well as the composition changes that take place after each trip. These composition changes specify the times and locations of the (un)coupling actions that are to take place. However, this model does not take into account the exact infrastructure that is used for each train movement; it is possible that the generated shunting movements cause conflicts with the planned trains. In the next section, we will describe our framework for modelling the infrastructure of stations and their shunting yards.

4.3 Railway infrastructure

To formulate the headway and station capacity constraints, it is necessary to model the exact infrastructure of the inspected stations. The main station that we examine is Utrecht Centraal.

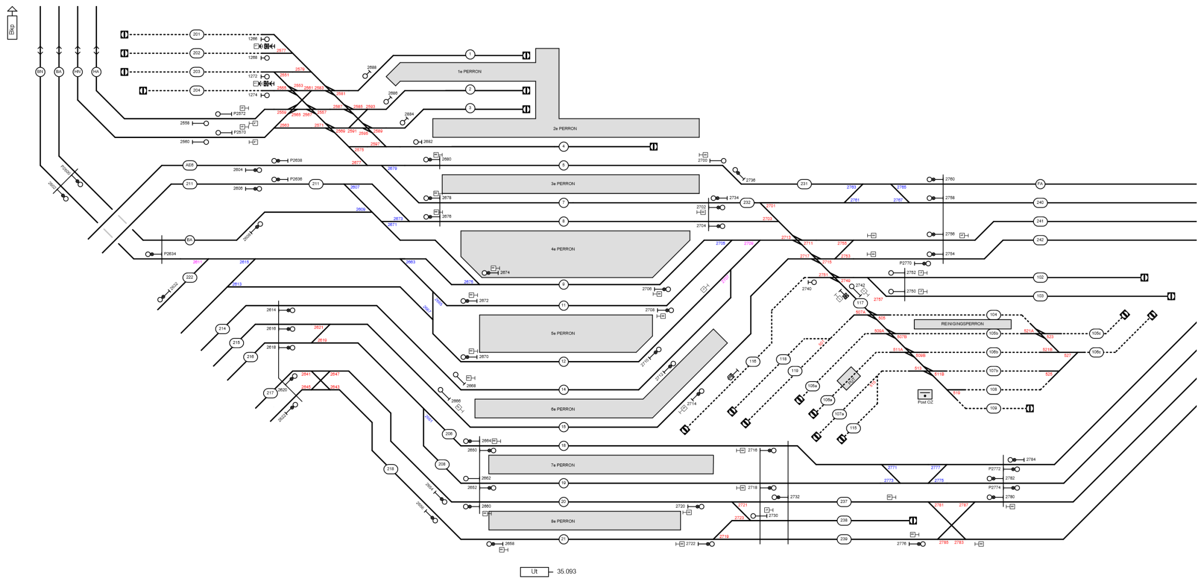

Figure 1: The exact infrastructure of station Utrecht Centraal.

Figure 1 shows the exact railway infrastructure of station Utrecht Centraal, the largest station at the center of the Netherlands. The station contains sixteen platform tracks and three shunting yards. The level of detail with which we look at a station in this thesis consists of taking into account all platforms, tracks, switches, crossings and shunting tracks. Each of these structures is modelled as an individual point on which the headway and capacity constraints must hold between each pair of trains that traverses this point.

The next step is to model the exact route that each planned train trip and shunting movement traverses. Since the timetable that we take as input contains the platform and/or track that each train trip stops at or passes through, the route of planned trains can easily be extracted by following the connected structures in the direction of the next station of the train service.

A similar procedure holds in place for the shunting movements. We make the assumption that for all platforms, a unique fixed route is followed for all shunting movements that originate from the same platform. Other papers, such as the one by Haahr et al. (2017), leaves open the exact track that rolling stock units reside on and solves a feasibility problem which assigns each rolling stock unit to an available track. However, we make the assumption that during the day, most shunting tracks will be empty and that not enough shunting movements will take place, even after the occurrence of a disruption, that would hinder the availability of the rolling stock units parked there. Especially for a station as big as Utrecht Centraal, which has three shunting yards consisting of 45 shunting tracks total, this seems like a reasonable assumption. Furthermore, the main issue that is researched in this thesis is the feasibility of the shunting movements with respect to the planned trains; we leave the exact setups on the shunting yards outside the scope of this thesis.

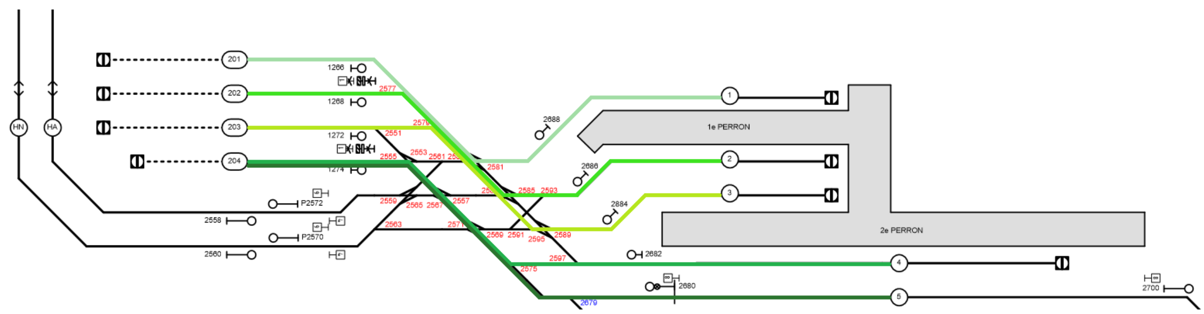

Figure 2: The shunting movement routes for platform tracks 1 to 5 of Utrecht Centraal are shown in green.

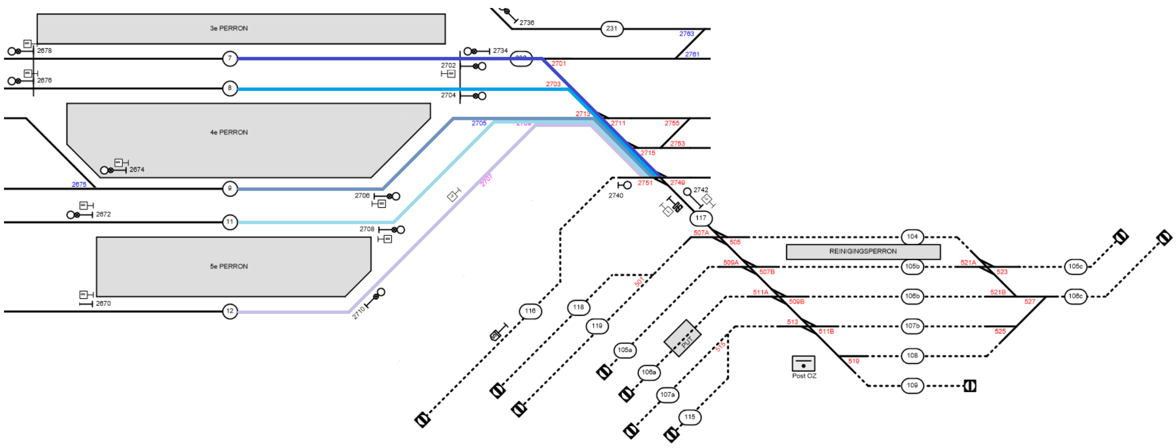

Figure 3: The shunting movement routes for platform tracks 7 to 12 of Utrecht Centraal are shown in blue.

Figures 2 and 3 show the shunting movement routes for platform tracks 1 to 5 and 7 to 12, respectively, of Utrecht Centraal. The only shunting yard available from tracks 1 to 5 is located on the west side of the platforms and contains four shunting tracks. Tracing the line from track 1, the route to the shunting yard is as follows: Ut1, Ut2581+Ut2583, Ut2579, Ut2577, Ut201. A similar route can be seen for the remaining tracks. The shunting yard to the east of tracks 7 to 12 are all exclusively accessible via the crossing Ut2749+Ut2751. This is where the route therefore ends in case a shunting movement transports a rolling stock unit to this shunting yard; all shunting movements use this crossing at some point on their path and the exact shunting tracks are assumed to be utilized in a conflict-free manner. The remaining tracks of Utrecht Centraal are modelled similarly. Note that for any train service that visits Utrecht Centraal, the exact routes into and out of Utrecht Centraal are also modelled in this way.

Using the routes of all train services and shunting movements, we model all tracks, switches and crossings as points on which a minimum headway time h must be present between each pair of trains. We work with a timetable which is conflict-free, which means that without the insertion of the shunting movements, the minimum headway time required is satisfied for all pairs of trains with overlapping routes. To check the feasibility of the shunting movements, we impose the headway restrictions between each pair of shunting movement and planned train on each individual point on the shunting route. Since we allow for the shifting of planned trains, the headway constraints need to be present between each pair of planned trains as well. In the next section, we will elaborate on the exact formulation of these constraints and the fitting procedure of the shunting movements.

5 An iterative framework for rolling stock rescheduling

In this section, we will discuss the methods we develop for rolling stock rescheduling with infrastructure constraints. Section 5.1 adjusts the composition model in such a way that a new rolling stock schedule can be made after the occurrence of a disruption. Section 5.2 introduces the railway availability constraints which ensure that all train trips and shunting movements can be performed without any conflicts. Section 5.3 discusses the interaction between the different models and proposes a solution method which finds a schedule that satisfies the constraints of every model.

5.1 The composition model for rolling stock rescheduling

The Rolling Stock Rescheduling Problem (RSRP) as described by Nielsen (2011) looks at the real-time rescheduling of rolling stock for maintenance and after disruptions lead to cancellation of services in the timetable. Given the modified timetable that includes the adjusted and cancelled train trips as a result of the disruptions, the available rolling stock and the shunting options, this problem aims to find an adjusted rolling stock plan that specifies which rolling stock compositions are assigned to which trips.

The model used for the purpose of rolling stock rescheduling is largely the same as the composition model that we introduced previously for the generation of the initial rolling stock schedule. To account for the disruption, the set of trips $\mathcal{T}$ and the set of allowed composition changes Γ_t for each trip t are adjusted. Furthermore, we add operational and stick-to-the-plan objectives to the model, which require the introduction of additional parameters, variables and constraints.

We denote by p_0^t the original composition of trip t in the original schedule, by oc^t and ou^t the original number of coupled units right before trip t and the original number of uncoupled units right after trip t, respectively, and by $i_{m,0}^s$ and $i_{m,\text{end}}^s$ respectively the original initial and ending inventory of rolling stock unit type m at station s. Additionally, τ_{disr} denotes the starting time of the disruption. We introduce the binary variable p_{diff}^t, which takes on the value 1 if trip t is executed with a different composition than the originally planned composition p_0^t and 0 otherwise. The variables nc_t and nu_t represent the number of additional coupled units right before trip t and the number of additional uncoupled units right after trip t, respectively. The variable $I_{m,\text{off}}^s$ denotes the off-balance in the ending inventory of rolling stock unit type m of station s. Furthermore, we define seven new weights w_6 - w_{12} for the different factors of the objective function. The model then looks as follows:

$$\begin{aligned}\min \sum_{t\in\mathcal{T}}\sum_{c\in\mathcal{C}_t} w_6 \cdot carr_c \cdot l_t \cdot x_c^t + \sum_{t\in\mathcal{T}}\sum_{c\in\mathcal{C}_t} w_7 \cdot seat_c^t \cdot d_t \cdot x_c^t + \sum_{t\in\mathcal{T}} w_8 \cdot x_0^t \\ + \sum_{s\in\mathcal{S}}\sum_{m\in M} w_9 \cdot I_{m,\text{off}}^s + \sum_{t\in\mathcal{T}} w_{10} \cdot nc_t + \sum_{t\in\mathcal{T}} w_{11} \cdot nu_t + \sum_{t\in\mathcal{T}} w_{12} \cdot p_{\text{diff}}^t \end{aligned} \tag{16}$$

s.t. (2) – (9), (12) – (15)

$$p_{\text{diff}}^t \geq \sum_{c\in\mathcal{C}:c\neq p_0^t} x_c^t \qquad \forall t \in \mathcal{T} \tag{17}$$

$$nc_t \geq \sum_{m\in M} c_m^t - oc_t \qquad \forall t \in \mathcal{T} \tag{18}$$

$$nu_t \geq \sum_{m \in M} u_m^t - ou_t \qquad \forall t \in \mathcal{T} \quad (19)$$

$$I_{m,\text{off}}^s \geq I_{m,\text{end}}^s - i_{m,\text{end}}^s \qquad \forall s \in \mathcal{S},\ m \in M \quad (20)$$

$$I_{m,\text{off}}^s \geq i_{m,\text{end}}^s - I_{m,\text{end}}^s \qquad \forall s \in \mathcal{S},\ m \in M \quad (21)$$

$$I_{m,0}^s = i_{m,0}^s \qquad \forall s \in \mathcal{S},\ m \in M \quad (22)$$

$$x_{p_0^t}^t = 1 \qquad \forall t \in \mathcal{T} : \tau_d(t) \leq \tau_{\text{disr}} \quad (23)$$

$$p_{\text{diff}}^t \in \{0, 1\} \qquad \forall t \in \mathcal{T} \quad (24)$$

$$nc_t, nu_t \in \mathbb{Z}^+ \qquad \forall t \in \mathcal{T} \quad (25)$$

$$I_{m,\text{off}}^s \in \mathbb{Z}^+ \qquad \forall s \in S, m \in M \quad (26)$$

The objective function (16) aims to minimize the number of driven carriage kilometers, the number of seat-shortage hours, the number of cancellations, off-balances in the ending inventories, the number of additional (un)coupled units and the number of different compositions compared to the initial rolling stock schedule. Constraints (17) ensure that the binary variable p_{diff}^t takes on the value 1 if a different composition is chosen for trip t. Constraints (18) and (19) model the additional coupled and uncoupled units, respectively. Since the variables nc_t and nu_t are non-negative, only additional shunting movements are taken into account. Constraints (20) and (21) ensure that the variables $I_{m,\text{off}}^s$ are equal to the absolute value of the off-balance between the ending inventory of the original plan and the ending inventory of the adjusted plan. Constraints (22) specify that the initial inventories are equal to the planned initial inventories. Constraints (23) fix the chosen compositions for all trips that take place before the disruption occurs. Constraints (24) - (26) specify the domains of the decision variables. The rescheduling model then consists of Constraints (2) - (9) and (12) - (26).

5.2 The infrastructure model for railway availability

The purpose of the constraints that are covered in this section is to ensure that all shunting movements cause no headway or infrastructure availability conflicts with each other and with other planned trains. We will utilise constraints similar to those in the Train Timetable Adjustment Problem (TTAP) as described by Van Aken et al. (2017), which aims at finding an alternative timetable that minimizes the deviation from the original timetable in terms of the incurred delays and cancellations of train services. This problem models all trips in the timetable as train events and ensures that the minimum headway time is kept between all trains and also that the station and track capacities are satisfied. It is important to note that the TTAP works with a cyclic timetable, whereas a rolling stock plan is made for the entire day. Therefore, all TTAP constraints are rewritten in such a way that the cyclicity is removed.

All constraints are formulated for one infra point $s \in \mathcal{S}$ only; the model can be extended to include multiple, if not all, infra points by modelling additional infrastructure and creating the associated constraints. The disruption takes place adjacent to s, such that rescheduling is required at and around s. The set M contains all train lines M_l and all shunting movements M_r that take place at station s, with the latter consisting of trips from a specific track of the infra point s to a shunting track at (one of) its shunting yard(s), or vice versa. We define the set of tracks $\mathbb{T}$ and the set of crossings C at station s. Station tracks and shunting tracks are often connected through crossings, on which the minimum headway time must be present as well. The headway constraints are in place for all train services and shunting movements that make use of the same track and/or crossing.

We create an event-activity network which is characterized by a directed graph $G = (E, A)$, where E represents the set of vertices, or events, and A represents the set of arcs, or activities. The set E contains all events regarding infra point s, like arrivals ($E_{arrival}$) and departures ($E_{departure}$). The set A contains the activities that connect these events, like running (A_{run}), dwelling (A_{dwell}) and turning (A_{turn}). We also define activities that are necessary to implement the constraints, like headway activities for trains running on the same track ($A_{headway}$) and headway activities for crossings ($A_{crossing}$). Furthermore, the set A_{shunt} denotes the shunting activities which for instance link the arrival of a passenger train at a station to an uncouple movement, i.e. the departure of a rolling stock unit which has been uncoupled from the incoming train and which is transported to the shunting yard. Each activity is bounded by a minimum and maximum amount of time that the activity can take. Additionally, we only consider pairs of events for the headway and crossing constraints that take place within thirty minutes of each other, since the headway constraints only need to be imposed for events that take place somewhat close to each other in time. We introduce our working area as $[0, T]$, with T indicating the end of the day in minutes. The formulation is then as follows:

$$\min \sum_{j \in E} w_j^{\text{deviation}} \cdot p_j + \sum_{m \in M} w_m^{\text{cancel}} \cdot X_m \tag{27}$$

$$\text{s.t. } 0 \leq v_j \leq T(1 - X_m) \qquad \forall j \in E_m, \forall m \in M \tag{28}$$

$$l_{i,j}(1 - X_m) \leq v_j - v_i + q_{i,j}T \leq u_{i,j}(1 - X_m) \qquad \forall (i,j) \in A^m_{\text{run}} \cup A^m_{\text{dwell}} \cup A^m_{\text{turn}}, \forall m \in M \tag{29}$$

$$l_{i,j}(1 - X_m - X_n) \leq v_j - v_i + q_{i,j}T \leq u_{i,j}(1 - X_m - X_n) + (T-1)(X_m + X_n) \qquad \forall (i,j) \in A_{\text{shunt}},\ \forall m, n \in M \tag{30}$$

$$h_{i,j}(1 - X_m - X_n) \leq v_j - v_i + q_{i,j}T \leq (T - h_{i,j})(1 - X_m - X_n) + (T-1)(X_m + X_n) \qquad \forall (i,j) \in A^t_{\text{headway}} \cup A^c_{\text{crossing}},\ \forall t \in \mathbb{T},\ \forall c \in C,\ \forall m, n \in M \tag{31}$$

$$d_j = v_j - \pi_j(1 - X_m) \qquad \forall j \in E_m,\ \forall m \in M \tag{32}$$

$$-d^+_{\max} \leq d_j \leq d^+_{\max} \qquad \forall j \in E \tag{33}$$

$$p_j \geq d_j \qquad \forall j \in E \tag{34}$$

$$p_j \geq -d_j \qquad \forall j \in E \tag{35}$$

$$v_j, p_j, d_j \in \mathbb{R}^+ \qquad \forall j \in E \tag{36}$$

$$q_{i,j} \in \{0, 1\} \qquad \forall (i,j) \in A \tag{37}$$

$$X_m \in \{0, 1\} \qquad \forall m \in M \tag{38}$$

The objective function (27) minimizes a weighted sum of the incurred deviations from the original timetable and the cancellations of train services. The variable p_j represents the absolute value of the deviation of event j compared to the original timetable. The binary variable X_m is equal to 1 if train line m is cancelled and 0 otherwise. The parameters $w_j^{\text{deviation}}$ and w_m^{cancel} represent the weights of the deviations and the cancellations, respectively. Constraints (28) limit the event time v_j of every event j to range between 0 and T. If the train line m which event j is part of is cancelled, v_j is equal 0. Constraints (29) set the time of every running, dwelling and turning activity to range between the activity-specific predefined lower bound $l_{i,j}$ and upper bound $u_{i,j}$. Both the left-hand side and right-hand side are set to 0 in case the train line is cancelled to establish the feasibility of the event times which are set to 0 by Constraints (28). To keep track of the order of events, a binary variable $q_{i,j}$ is introduced for every arc (i, j). The variable $q_{i,j}$ is equal to 0 if event i takes place chronologically before event j and 1 otherwise. Constraints (30) ensure that the minimum (un)couple time $l_{i,j}$ is present between the arrival of a composition and the uncouple movements of its unit(s) and between the arrival of its coupled unit(s) and the departure of the new composition. In the former case, event j of route m corresponds to the uncouple movement and event i of route n corresponds to the arrival of the old composition. In the latter case, event i of route n corresponds to the couple movement and event j of route m corresponds to the departure of the new composition. Additionally, the parameter $u_{i,j}$ represents the maximum amount of time by which an uncouple movement can take place after the arrival of the old composition and the maximum amount of time by which a couple movement can take place before the departure of the new composition. To model the headway and crossing requirements, Constraints (31) create a safety distance between trains m and n for all headway and crossing activities. The difference between the two event times must be at least the specified headway time $h_{i,j}$. If at least one of the considered train lines is cancelled, the headway constraint no longer has to hold. The variables X_m and X_n are incorporated to ensure this.

All previous constraints contain the variable X_m to allow for the cancellation of trains. Another set of constraints is needed to allow for deviations of the event times. These are described by Constraints (32)-(35). The deviation is modelled by (32). The variable d_j measures the deviation in the event time v_j in minutes. The parameter π_j indicates the event time of event j in the original timetable. The deviation is limited by Constraints (33). The parameter d^+_{max} indicates the maximum allowed deviation. Since we want to penalize deviations irrespective of their sign, we use the variable p_j which is equal to the absolute value of the deviation. This is done by Constraints (34) and (35). Constraints (36) - (38) indicate the domains of the decision variables.

Note that in this formulation of the problem, it is possible to cancel train services. For our application, the cancellation of train services is included to conclude that a shunting movement cannot be performed. We choose to give priority to all planned trains and to only perform shunting movements in case they can be fit in between all planned trains without drastically shifting the existing services, which means that a shunting movement is deemed to be impossible in case performing this movement leads to cancellations of other train services. However, the cancellation of planned trains will not occur in our experiments as we will assign a much larger penalty to this than to the cancellation of a shunting movement.

5.3 An iterative rolling stock rescheduling algorithm

The two previous sections provide an in-depth description of the two main models that we use in our rolling stock rescheduling algorithm. In this section, we will characterize the interaction between the two models and describe the feedback steps that are necessary to provide a feasible rolling stock schedule. Algorithm 1 gives an overview of our approach.

Algorithm 1 The iterative rolling stock rescheduling algorithm.

```
INPUT: Initial timetable, initial rolling stock schedule and a disruption
OUTPUT: An adjusted rolling stock schedule for the disruption
Initialization
Commence disruption, update trips and transition set
while Schedule is not feasible do
    Create interim rolling stock schedule by solving the composition model
    Extract new shunting movements
    if No new shunting movements then
        Schedule is feasible
        break
    end
    Fit shunting movements by solving the infrastructure model
    Cancel impossible shunting movements
    Shift affected trains
end
```

Note that the initial timetable that we work with is conflict-free, which means that without including any shunting movements, all minimum headway time constraints and capacity constraints are satisfied. The manner in which we ensure this is by performing one iteration of the infrastructure model (27) - (38) on the initial timetable, in which we do not allow the cancellation of any trains. After processing the disruption information in the inputs, the algorithm revolves around iteratively solving the composition model and the infrastructure model whilst incorporating the relevant feedback for the next iteration. In the next sections, we will elaborate on each individual phase.

5.3.1 Initialization

In the initialization step of the algorithm, we process the disruption in the given inputs. Our disruptions consist of complete blockages of railway tracks between two adjacent stations. The set of trips $\mathcal{T}$ is adjusted by cancelling all trips that are planned to traverse the blocked railway tracks during the duration of the disruption. This means that all trips that follow the cancelled trip on the same train line are cancelled as well due to the inability to reach the remaining stations after the blockage. The sets of transitions Γ_t for the affected trips t are adjusted as well. For the type of disruptions that we consider, we fix the transitions by introducing *short-turnings*. This measure assigns services that are approaching a blockage to the cancelled services going in opposite direction. An example of a short-turning is displayed in Figure 4 for a complete blockage on the tracks linking the stations Gouda (Gd) and Rotterdam Centraal (Rtd) between 12:00 and 13:00 in the afternoon. The trips are denoted by $t_n, n = 1 \dots 20$ and the transitions at the end of the train services are denoted by $c_n, n = 1 \dots 5$.

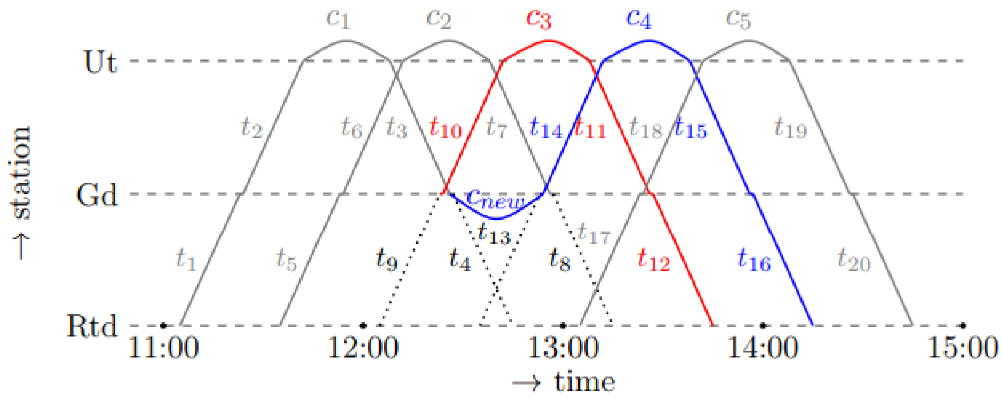

Figure 4: A time-space diagram of five train services between Utrecht Centraal (Ut) and Rotterdam Centraal (Rtd). The dotted lines represent trips that are cancelled due to a disruption. The blue line represents a train service that can be completed due to the introduction of a new transition. The red line represents a train service that cannot be salvaged by a new transition.

Trips t_4, t_8 and the train services that correspond to trips t_9 - t_{12} and t_{13} - t_{16} cannot be started due to the occurrence of the disruption. We can now connect trips t_3 and t_{14}, as the train composition that is used for trip t_3 is stuck at Gouda and sufficient time is available to perform the short-turning into trip t_{14}. We create all such possible transitions for the trips that are affected by the disruptions and add them to the respective transition sets.

5.3.2 Shunting movement extraction and selection

After solving the composition model as described in Section 5.1, a rolling stock schedule that contains the compositions for each trip and the composition changes for each transition is obtained. Based on this information, the necessary shunting movements can be transformed into trips and added to the set of train lines M. We introduce the parameters s^- and s^+ which represent the minimum and maximum time in minutes, respectively, that must and may be present between the arrival of a train composition at its track and the uncoupling movements that involve rolling stock unit(s) belonging to the considered composition. For couple movements, the coupled unit needs to be retrieved from the shunting yard and must be present at its desired track at least s^- and at most s^+ before the departure of the train it is coupled to. Furthermore, we introduce the parameter r which represents the time it takes in seconds to travel between each track, crossing and/or switch within a station.

Several selections of shunting movements are now possible and will be explored in the iterative framework. The models as formulated have the ability to find a rolling stock schedule which is feasible with respect to all infrastructure constraints. However, in practice a rolling stock dispatcher may not have the power to delay passenger trains and the mathematically optimal solution may therefore hold little value. For this reason, we will define several parameters which limit the severity of the suggested measures. The parameter s_n represents the number of shunting movements that are fed forward to the infrastructure model in iteration n of the iterative algorithm. In case not all shunting movements are chosen to be fit, several criteria could be used to select the order of the shunting movements:

- Select the shunting movements that are necessary to solve the disruption first. This particularly pertains to the shunting movements that are performed for the train services that are directly affected by the disruption; for instance, it might be possible that as the result of a disruption, not enough rolling stock units are available to perform a specific service. A rolling stock unit may then, if possible, be extracted from the shunting yard. This selection procedure would give such shunting movements priority.
- Select the shunting movements that take place on the least busy tracks first. These shunting movements have the highest probability of being feasible without having to shuffle any planned trains and can therefore form a good basis for trying to fit as many shunting movements as possible.
- Select the shunting movements in the way that they appear chronologically. As a disruption lasts for a finite amount of time, working through the shunting movements chronologically will eventually lead to a feasible schedule.

We furthermore define t_m as the maximum number of trains that can be shifted simultaneously in iteration m of the infrastructure model, and d_m as the maximum total deviation incurred for all planned trains in iteration m of the infrastructure model. The goal is to find a distribution of the parameters such that the given solutions do not impose measures that are too drastic in practice.

5.3.3 Feedback incorporation

All shunting movements are added to the set of train services and the infrastructure model (27) - (38) is then solved. The problem then outputs which trips incur deviations in time and which shunting movements cannot be performed. This information is then given back to the composition model. In principle, this is done by updating the time deviations for all train trips. It is possible that shunting movements are cancelled; this information can be incorporated by removing the option of changing the composition at this transition. The currently used composition must therefore be kept and used for the trip directly after the transition. If more than one shunting movement at a time is cancelled, the previously described procedure is kept in place if the paths of the shunting movements do not cross and/or the shunting movements do not take place close to each other in time. Else, it is possible that performing one shunting movement prevents the possibility of performing the other(s). In this case, the addition of each shunting movement is performed individually and separately and the best one(s) as measured by the improvement in the objective value is (are) kept for the next iteration.

It is possible that more adjustments are suggested than are feasible in practice. For this purpose, we define a threshold objective value v_m for iteration m of the infrastructure model as follows: if including the chosen selection of shunting movements in the infrastructure model leads to an objective value of at most v_m, the solution is accepted and all information is incorporated in the next iteration of the composition model. Else, we resolve the model with fewer shunting movements until the objective value falls below the threshold.

6 Computational results

In this section, we present computational results for our main methods evaluated on five different disruption instances. We start with a description of the five disruption instances in Section 6.1, after which we discuss the initial parameter settings in Section 6.2. Sections 6.3 and 6.4 discuss the evaluation of the solutions of the initial rolling stock schedule generation and the iterative rolling stock rescheduling algorithm, respectively. We then analyze the computational results and investigate the influence of the weights and parameters by performing a sensitivity analysis on our methods in Section 6.5.

Experiments are executed on a PC with a 3.0 GHZ Intel Core i5-9500 processor and 16.0 GB RAM. The solution methods are implemented in Java using the IDE Eclipse 2022-03 (version 14.23) and the MILP model is solved using the commercial solver IBM ILOG CPLEX 20.1 with default settings. We impose a time limit of one hour on the creation of the initial rolling stock schedule, a time limit of ten minutes on each iteration of the composition model and a time limit of ten minutes on each iteration of the infrastructure model.

6.1 Disruption instances

We implement five disruptions instances which take place at and around Utrecht Centraal (Ut). The five disruption instances are shown in Figure 5 and 6.

Figure 5: Disruption instances 1, 2 and 4.

Figure 6: Disruption instances 3 and 5.

The first disruption consists of a full blockage between Utrecht Centraal and Den Haag, between which lies Gouda. This is indicated by the green and yellow line in Figure 5; the infra lines Utrecht Centraal-Gouda and Gouda-Den Haag are inaccessible. The second disruption takes place between Utrecht Centraal and Rotterdam Centraal, between which also lies Gouda. This disruption is shown by the green and orange lines. The third disruption takes place between 's-Hertogenbosch and Eindhoven, as shown by the light blue line in Figure 6. Furthermore, the fourth disruption takes place in the north of the country between Alkmaar and Zaandam and is shown by the pink line in Figure 5. Finally, the fifth disruption is indicated by the purple line in Figure 6 and takes place between Ede-Wageningen and Arnhem Centraal. For all disruption instances, we will model the infrastructure of Utrecht Centraal as shown in Section 4.3 and the shunting movements that take place there.

6.2 Parameter settings

In this section, we will introduce the base values of the parameters that we use for our solution methods and we will discuss possible options for varying these values.

The composition model has been extensively researched in the past and is also utilized by NS in several pieces of their software. The values we assign to the many parameters in the composition model for both creating an initial rolling stock schedule and for rescheduling after a disruption will therefore largely be equal to the values as chosen by NS. Table 2 shows the values of the objective value parameters for both versions of the composition model as introduced in Sections 4.2 and 5.1.

Table 2: Parameter values for the composition model.

Description	Parameter	Value
Carriage kilometers	w_1	0.13
Seat-shortage hours	w_2	60
Couple movements	w_3	50
Uncouple movements	w_4	50
Inventory differences	w_5	50
Carriage kilometers	w_6	0.13
Seat-shortage hours	w_7	60
Cancellations	w_8	1,000,000
Inventory off-balances	w_9	1,000
New couple movements	w_{10}	100
New uncouple movements	w_{11}	100
Different compositions	w_{12}	100,000

The values of the carriage kilometers and seat-shortage hours weights correspond to the values used by NS. We choose to assign values of equal scale to the other operational objectives. After running small experiments, we found that these values provide a good balance between the different components of the objective. For the rescheduling model, we maintain the same values for the carriage kilometers and seat-shortage hours. Additionally, we heavily penalize cancellations as a result of rolling stock rescheduling as this is the least desirable outcome for passengers. We also assign a large weight to different compositions compared to the initial schedule to lower the number of shunting movements that take place. Furthermore, inventory off-balances are assigned a higher penalty than the other operational objectives as off-balances may elicit the movement of empty rolling stock units at the end of the day.

Table 3: Parameter values for the infrastructure model.

Description	Parameter	Value
Planned train deviations	$w_j^{deviation}$	60
Planned train cancellations	w_m^{cancel}	1,000,000
Maximum deviation of planned trains	d_{max}^+	1
Shunting movement deviations	$w_j^{deviation}$	0.1
Shunting movement cancellations	w_m^{cancel}	10,000
Maximum deviation of shunting movements	d_{max}^+	15
Headway time	$h_{i,j}$	3
Time horizon	T	1,440
Minimum (un)couple time	s^-	5
Maximum (un)couple time	s^+	10
Travel time between tracks, crossings and/or switches	r	0.1

Table 3 shows the base parameter values that are used for the infrastructure model. We choose to heavily penalize any deviations and cancellations for planned trains, as they have priority over performing the shunting movements. We only allow deviations of up to one minute in our model, as in practice shifting the times of trips in the timetable is not a preferred method of increasing the likelihood of being able to perform shunting movements. However, incurring a delay of at most a minute could open up the necessary windows without drastically inconveniencing the passengers. Note that the model allows for both positive and negative deviations in time. In practice, incoming trains can only be delayed and not brought forward in time. For this reason, we only allow positive deviations for planned trains and both positive and negative delays for the shunting movements as these are not contained in the original timetable.

Furthermore, for the lower bounds $l_{i,j}$, we take the values of the running and dwelling activities as they are in the timetable. As a start, we do not allow the time between two events to be shorter than the current run and dwell times to ensure that events are not brought forward in time. The upper bound $u_{i,j}$ is set to three minutes above the lower bound. In case the pairs of events come from the set A_{shunt}, we set the lower bound s^-, which corresponds to $l_{i,j}$ in Constraints (30) in Section 5.2, to five minutes, since this is the amount of time we assume to be necessary for a coupled unit to be present on the track before the new composition departs and for all passengers to get in and out of the train before a unit can be uncoupled. We set the upper bound s^+ to ten minutes, as shunting movements must not take place much earlier or later than the departure or arrival of the considered composition. We also assume that it takes a fixed time of 0.1 minute, or six seconds, to drive between each pair of tracks, crossings and/or switches. In reality, shunting movements do not take place at a fixed speed since some acceleration time is needed at the start of the movement. However, we assume that this time is caught by the five minute buffer that is contained in s^- and that once the rolling stock unit has reached a constant speed, a fixed time of six seconds serves as a close representation of reality.

Table 4 shows the initial values for the parameters that are present in the iterative framework. We try to implement measures that could realistically be performed, like exclusively trying to shift at most two trains by at most a minute each and by allowing for a maximum total deviation of at most three minutes amongst all planned trains per iteration of the infrastructure model. Initially, the shunting movements are fit one by one, following the order in which they appear in the timetable. A shunting movement is deemed to be impossible in case the objective value of the model is larger than 10,000; this corresponds to the case where either a train service or a shunting movement is cancelled, as then the cancellation penalty of 1,000,000 or 10,000 is incurred.

Table 4: Parameter values for the iterative framework.

Description	Parameter	Value
Number of shunting movements	s_n	1
Maximum number of shifted trains	t_m	2
Maximum total deviation	d_m	3
Maximum objective value	v_m	10,000

We will perform a sensitivity analysis on various of the parameters in Section 6.5. The results in Sections 6.3 and 6.4 are obtained with the parameter settings as described above.

6.3 Initial rolling stock schedule generation

The generation of the initial rolling stock schedule is performed in the manner described in Section 4.2 with the base parameter settings as described in Section 6.2. The stations at which we initially allow compositions to change are Eindhoven Centraal, Roosendaal, Deventer, Zwolle and Alkmaar. Table 5 shows the objective value, the running time and the values of the objective indicators for the generation of the initial rolling stock schedule.

Table 5: Resulting statistics of the generation of the initial rolling stock schedule.

Indicator	Value
Objective value	487,358.58
Running time (s)	2670.88
Carriage kilometers	1,200,766.0
Seat-shortage hours	5,247
Couple movements	757
Uncouple movements	757
Inventory differences	100

The model succeeds in finding a feasible initial rolling stock schedule in less than 45 minutes. The remaining values will be used as a benchmark with which we compare the performance of the iterative algorithm. Note that in the initial schedule, no trips are cancelled and each trip is therefore assigned a composition which consists of at least one rolling stock unit.

6.4 Iterative algorithm performance

This section elaborates on the performance of the iterative rolling stock rescheduling algorithm for the five instances. We will first perform rolling stock rescheduling without shunting at Utrecht Centraal, as described in Section 5.1. We will refer to this as the *basic reschedule*. Then, we add the station Utrecht Centraal to the set of stations at which composition changes are allowed during the duration of the disruption.

We will show different statistics regarding the shunting movements that take place at Utrecht Centraal, whilst we initially do not allow shunting there. We distinguish between four types of transitions and composition changes that take place at Utrecht Centraal:

- Regular turnings. These are turnings of compositions that take place at the end of train services, such that (some of) the same rolling stock units can be utilized to perform the next train service in opposite direction. The initial rolling stock schedule contains planned composition changes at such transitions, which can for example be the reversal of the composition in case the composition does not consist solely of rolling stock units of the same (sub)type, or the (un)coupling of units to account for an increase or decrease in passenger demand due to for instance the start or end of a rush hour. Since these turnings take place regardless of the occurrence of a disruption, we will assume that (un)couple movements at such turnings are always possible and we will not include such shunting movements in our statistics.
- Short-turnings. Such turnings are in our application introduced as a way of countering disruptions, as described in Section 5.3.1. They connect train services that reach stations adjacent to the disruption to train services going in the opposite direction that cannot be performed anymore due to the disruption. In the basic reschedule, we do not allow for rolling stock units to be (un)coupled during such transitions at Utrecht Centraal. To evaluate the performance of our iterative algorithm, the remaining experiments do include the option of (un)coupling rolling stock units during short-turnings at Utrecht Centraal.
- Stranded services. These transitions correspond to train services that reach a station adjacent to the disruption and are unable to complete the remaining route. No short-turning option is available for these services and the incoming composition needs to be removed from its platform. We assume that for such services, the entire composition needs to be parked and each unit needs to be moved to a shunting yard. Since these shunting movements must be executed to not hinder the rest of the timetable that is unaffected by the disruption, we will assume that they are always performed and include them in the number of new uncoupled units at Utrecht Centraal statistic. It is possible that our model determines that these shunting movements cannot be performed. In practice, it is often chosen to cancel other passenger trains that are limiting the availability of the shunting yard. We will report such occurrences separately.
- New composition changes. In our iterative algorithm, we allow for shunting at Utrecht Centraal during the duration of the disruption. As a result, it is possible that the composition model outputs a rolling stock schedule which suggests that train services which are unaffected by the disruption should change their composition and as a result (un)couple rolling stock units. Such composition changes do not occur in the setting of the basic reschedule, but will be included in the statistics of our remaining experiments.

6.4.1 Instance 1: Utrecht Centraal - Den Haag Centraal

The first disruption instance consists of a full blockage between Utrecht Centraal and Den Haag Centraal. Three train lines, 1700, 2000 and 11700 are affected, of which line 2000 is fully cancelled and the lines 1700 and 11700 can still be operated up until reaching either Utrecht Centraal or Den Haag Centraal, at which the train services turn around.

To evaluate the performance of our algorithm, we first perform rolling stock rescheduling without allowing for composition changes at Utrecht Centraal, as described in Section 5.1. Table 6 shows for the basic reschedule the running time, the objective value and the values of the following statistics: the number of carriage kilometers driven, the number of seat-shortage hours for passengers, the number of seat-shortage hours during the duration of the disruption, the number of new coupled and uncoupled units compared to the original rolling stock schedule, the number of trips with different compositions compared to the initial schedule and the number of units by which the rolling stock inventories at the end of the day differ. We also show the number of new coupled and uncoupled units at Utrecht Centraal, as well as the number of trips that have a different composition that start or end at Utrecht Centraal. Furthermore, the inventory differences at Utrecht Centraal are also displayed.

Additionally, Table 6 shows the resulting statistics of performing the iterative rolling stock rescheduling algorithm. The table presents general statistics regarding the iterative framework which include the running time, the number of iterations, the number of suggested and cancelled shunting movements, the objective value of the composition model after the first iteration, the optimality gap between the final schedule as obtained from the composition model and the schedule after the first iteration and the percentage improvement of the final schedule over the basic reschedule. Note that the objective value of the basic reschedule forms an upper bound on the objective value of the final schedule, as the basic reschedule does not allow for shunting at Utrecht Centraal and our algorithm can identify feasible shunting movements which improve the obtained rolling stock schedule. Further note that the objective value after the first iteration of the composition model forms a lower bound, as this schedule assumes that all shunting movements at Utrecht Centraal are possible. For the composition model, we display the overall running time of all iterations, the final objective value and the values of the previously mentioned objectives of the basic reschedule. We furthermore show for the infrastructure model the overall running time, the number of shifted trains, the total incurred deviation in departure and arrival times over all planned train trips and the average incurred deviation over the affected trains. Note that the running times for each model also contain the preprocessing time of the data, the initialization of the model and the conversion of the output to feedback and feedforward.

Table 6: Resulting statistics of the rolling stock rescheduling algorithm for instance 1.

Approach	Basic reschedule	Iterative algorithm
Composition model		
Running time (s)	88.28	162.12
Objective value	11,667,191.10	9,067,208.88
Carriage kilometers	1,186,070	1,186,076
Seat-shortage hours	8,090	8,039
Seat-shortage hours during disruption	3,389	3,385
#New coupled units	15	19
#New uncoupled units	29	32
#New coupled units at Ut	0	3
#New uncoupled units at Ut	4	6
Different compositions	110	84
Different compositions at Ut during disruption	2	5
Inventory differences	12	14
Inventory differences at Ut	2	1
Infrastructure model		
Running time (s)	-	191.18
#Shifted trains	-	4
Total deviation (min)	-	3.00
Average deviation per shifted train (min)	-	0.75
General		
Running time (s)	-	353.30
#Iterations	-	2
#Shunting movements suggested	-	7
#Shunting movements cancelled	-	1
Objective value after first iteration	-	9,067,208.88
Optimality gap (%)	-	0.00
Improvement over basic reschedule (%)	-	28.67

Table 6 shows that the basic reschedule terminates within less than two minutes and generates a rolling stock schedule in which two stranded train services are uncoupled, which makes up the four new uncoupled units at Utrecht Centraal. We find a feasible solution using the iterative algorithm in less than six minutes, with the composition model and the infrastructure model taking around three minutes each. The algorithm suggests to perform seven shunting movements in total at Utrecht Centraal, of which six can be performed whilst incurring a total deviation of three minutes across four trains. The remaining shunting movement that cannot be performed corresponds to an incoming train of line 2000, which cannot continue its journey due to the disruption and is therefore stranded. This trip does not have a successor trip and the composition model outputs that after the completion of this trip, all rolling stock units that make up the incoming composition are to be brought to a shunting yard. The infrastructure model finds a clash with a sprinter of line 7300 to Veenendaal Centrum and our algorithm therefore concludes that the shunting movement cannot be performed. In practice, the solution to such disruptions would be to cancel the sprinter train and to move the stranded train off the platform tracks. The composition model as described in Sections 4.2 and 5.1 does not detect such occurrences. We therefore find the same objective value as after the first iteration with an optimality gap of 0.00%.

Our algorithm could potentially be adjusted by reporting to the composition model that the uncoupling movement takes place as planned and that instead, the clashing sprinter is assigned the empty composition and therefore cancelled. This way, we more closely resemble the measures that are taken in real-life for rolling stock rescheduling. Our current implementation does not, however, directly return the train service that causes a clash and also does not recognize the scenarios in which the clashing train service should be cancelled. Further fine-tuning is therefore required to incorporate this aspect and will be considered out of the scope of this thesis. Note that since we assume that such shunting movements always take place, the objective value after the first iteration equals the objective value of the final schedule; with our current implementation, a manual penalty would need to be added to incorporate this occurrence in the obtained objective value.

Further comparing the two performances, the number of seat-shortage hours and the number of different compositions have decreased compared to the rolling stock schedule that is created without the ability to shunt at Utrecht Centraal. The number of new couplings and uncouplings has increased, which can partially be attributed to the new shunting movements performed at Utrecht Centraal. Additionally, the number of driven carriage kilometers has slightly increased. There is an additional inventory difference of two units over the entire network, but the inventory difference at Utrecht Centraal has decreased by one unit.

Overall, our iterative algorithm succeeds in improving the rolling stock schedule for this instance. We find a decrease in the number of seat-shortage hours, at the cost of delaying four planned trains by on average a little under a minute each. With a running time of less than six minutes for the first instance, our algorithm could find uses for real-time applications.

6.4.2 Instance 2: Utrecht Centraal - Rotterdam Centraal

The second disruption instance takes place between Utrecht Centraal and Rotterdam Centraal. The train lines 500, 600, 2800, 4000 are affected and can be operated until reaching either Utrecht Centraal or Rotterdam Centraal, at which they turn around, and the line 7700 is fully cancelled. Table 7 shows the results of the basic reschedule and of the iterative algorithm in the same way as in the previous section.

Table 7: Resulting statistics of the rolling stock rescheduling algorithm for instance 2.

Approach	Basic reschedule	Iterative algorithm
Composition model		
Running time (s)	70.84	369.11
Objective value	5,591,447.62	5,591,447.62
Carriage kilometers	1,193,614	1,193,614
Seat-shortage hours	7,036	7,036
Seat-shortage hours during disruption	3,297	3,297
#New coupled units	11	11
#New uncoupled units	15	15
#New coupled units at Ut	0	0
#New uncoupled units at Ut	2	2
Different compositions	50	50
Different compositions at Ut during disruption	0	0
Inventory differences	2	2
Inventory differences at Ut	1	1
Infrastructure model		
Running time (s)	-	82.74
#Shifted trains	-	0
Total deviation (min)	-	0
Average deviation per shifted train (min)	-	-
General		
Running time (s)	-	451.86
#Iterations	-	4
#Shunting movements suggested	-	4
#Shunting movements cancelled	-	4
Objective value after first iteration	-	5,392,770.66
Optimality gap (%)	-	3.68
Improvement over basic reschedule (%)	-	0.00

For the second instance, our algorithm terminates in a running time of less than eight minutes, with the composition model taking a little over six minutes and the infrastructure model taking over one minute. This instance checks four shunting movements, of which none are executable. One of these shunting movements corresponds to completely uncoupling a stranded train service at Utrecht Centraal. This shunting movement is included in the basic reschedule and leads to two new uncoupled units at Utrecht Centraal. As was the case for instance 1, this shunting movement is blocked by a sprinter of line 7300 to Veenendaal Centrum. This sprinter would in practice be cancelled to perform this uncoupling movement. Two of the other suggested shunting movements correspond to short-turning transitions and the last one corresponds to a new composition change of an unaffected train service. Since none of these shunting movements can be performed, the objective value and the values of the operational objectives of the iterative algorithm are identical to those of the basic reschedule and no improvement is found. Note that the objective value after the first iteration of the composition model corresponds to the rolling stock schedule where the stranded train service can be uncoupled and where the short-turning shunting movements and the suggested new composition change can be performed. This objective value is lower than the one obtained with our algorithm and it shows the potential improvement in case the suggested shunting movements at Utrecht Centraal could have been performed.

6.4.3 Instance 3: ’s-Hertogenbosch - Eindhoven Centraal

The third disruption instance takes place between ’s-Hertogenbosch and Eindhoven Centraal. The train lines 800, 3500 and 3900 are affected and can still be operated up until reaching either ’s-Hertogenbosch or Eindhoven Centraal, at which they turn around. Table 8 shows the results of the basic reschedule and of the iterative algorithm in the same way as the previous sections.

Table 8: Resulting statistics of the rolling stock rescheduling algorithm for instance 3.

Approach	Basic reschedule	Iterative algorithm
Composition model		
Running time (s)	59.86	307.37
Objective value	15,558,125.80	13,670,947.30
Carriage kilometers	1,195,060	1,190,610
Seat-shortage hours	6,319	6,493
Seat-shortage hours during disruption	2,054	1,906
#New coupled units	12	15
#New uncoupled units	19	23
#New coupled units at Ut	0	4
#New uncoupled units at Ut	0	3
Different compositions	150	131
Different compositions at Ut during disruption	10	10
Inventory differences	6	10
Inventory differences at Ut	0	1
Infrastructure model		
Running time (s)	-	202.75
#Shifted trains	-	5
Total deviation (min)	-	4.50
Average deviation per shifted train (min)	-	0.90
General		
Running time (s)	-	510.13
#Iterations	-	2
#Shunting movements suggested	-	9
#Shunting movements cancelled	-	1
Objective value after first iteration	-	13,355,961.14
Optimality gap (%)	-	2.36
Improvement over basic reschedule (%)	-	14.02

Our algorithm performs two iterations whilst solving the third instance. During the first iteration, nine shunting movements are suggested of which one is cancelled, after which the last iteration checks the feasibility of the rolling stock schedule with the possible shunting movements. The solution to this instance provides the greatest improvement compared to the initial rolling stock schedule so far in terms of the operational and passenger service level objectives, as the number of driven carriage kilometers decreases by over four thousand and the number of seat-shortage hours during the disruption decreases by over a hundred. Additionally, nineteen fewer different compositions are chosen compared to the initial rolling stock schedule, at the cost of in total three extra new couplings and four extra new uncouplings. Note that there are four new coupled units and three new uncoupled units at Utrecht Centraal which come from seven shunting movements total, whilst there are eight shunting movements which have been approved. The cancelled shunting movement corresponds to the uncoupling of a train unit and the shunting movement which is approved but not performed in the final schedule corresponds to a coupling movement; due to the unavailability of this specific rolling stock unit after the cancellation of the uncouple movement, the latter movement also ceases to exist. Furthermore, the number of inventory differences has increased by four, of which one is at Utrecht Centraal, and the number of seat-shortage hours overall has also increased by nearly two hundred. Overall, the main contribution of our algorithm for this problem instance is found in reducing the number of carriage kilometers driven and sticking closer to the original plan by lowering the number of different compositions, at the cost of more seat-shortage hours overall, more inventory differences and more (un)couplings. The running time for this algorithm is close to nine minutes, which can still prove to be beneficial for real-time disruption management.

6.4.4 Instance 4: Alkmaar - Zaandam

The fourth disruption instance takes place between Alkmaar and Zaandam. This disruption affects the train lines 800 and 3000, of which both can still be operated up until reaching either Alkmaar or Zaandam, at which they turn around. Table 9 shows the results of the basic reschedule and the iterative algorithm in the same way as the previous sections.

Table 9: Resulting statistics of the rolling stock rescheduling algorithm for instance 4.

Approach	**Basic reschedule**	**Iterative algorithm**
Composition model		
Running time (s)	72.25	179.88
Objective value	13,020,380.30	12,541,975.62
Carriage kilometers	1,195,502	1,196,474
Seat-shortage hours	5,761	6,064
Seat-shortage hours during disruption	1,654	2,046
#New coupled units	9	11
#New uncoupled units	14	20
#New coupled units at Ut	0	3
#New uncoupled units at Ut	0	2
Different compositions	125	120
Different compositions at Ut during disruption	4	6
Inventory differences	4	6
Inventory differences at Ut	0	1
Infrastructure model		
Running time (s)	-	109.53
#Shifted trains	-	3
Total deviation (min)	-	2.90
Average deviation per shifted train (min)	-	0.97
General		
Running time (s)	-	289.41
#Iterations	-	2
#Shunting movements suggested	-	5
#Shunting movements cancelled	-	0
Objective value after first iteration	-	12,541,975.62
Optimality gap (%)	-	0.00
Improvement over basic reschedule (%)	-	3.81

The iterative algorithm performs two iterations within a running time of less than five minutes and finds five shunting movements which can be performed at Utrecht Centraal, at the cost of delaying three planned trains by on average almost a minute each. All five of these shunting movements can be performed and are incorporated in the obtained rolling stock schedule, leading to an objective value equal to the objective value after the first iteration, which is lower than that of the basic reschedule. The final rolling stock schedule provides no operational benefits compared to the basic reschedule; the number of driven carriage kilometers, seat-shortage hours, new (un)coupled units and inventory differences have all increased to accomplish a decrease of five fewer different compositions compared to the original schedule. The performance of our iterative algorithm for this instance shows off the potential it has in improving a rolling stock schedule with respect to stick-to-the-plan objective; it remains to be seen in the sensitivity analysis whether other objectives can be enhanced as well by adjusting their weight in the objective function of the composition model.

6.4.5 Instance 5: Ede-Wageningen - Arnhem Centraal

The fifth disruption instance takes place between Ede-Wageningen and Arnhem Centraal. This disruption affects the train services 3000 and 3100, of which both can still be performed up until reaching either Ede-Wageningen or Arnhem Centraal, at which they turn around. Table 10 reports the results of the basic reschedule and the iterative algorithm in the same way as the previous sections.

Table 10: Resulting statistics of the rolling stock rescheduling algorithm for instance 5.

Approach	**Basic reschedule**	**Iterative algorithm**
Composition model		
Running time (s)	77.13	155.50
Objective value	8,035,207.26	5,921,914.46
Carriage kilometers	1,197,002	1,196,474
Seat-shortage hours	6,085	5,851
Seat-shortage hours during disruption	2,385	2,049
#New coupled units	9	11
#New uncoupled units	14	13
#New coupled units at Ut	0	3
#New uncoupled units at Ut	0	2
Different compositions	75	54
Different compositions at Ut during disruption	8	6
Inventory differences	2	4
Inventory differences at Ut	0	1
Infrastructure model		
Running time (s)	-	99.31
#Shifted trains	-	0
Total deviation (min)	-	0
Average deviation per shifted train (min)	-	-
General		
Running time (s)	-	254.81
#Iterations	-	2
#Shunting movements suggested	-	5
#Shunting movements cancelled	-	0
Objective value after first iteration	-	5,921,914.46
Optimality gap (%)	-	0.00
Improvement over basic reschedule (%)	-	35.69

Similar to the previous instance, our algorithm performs two iterations in a running time of less than five minutes and finds five shunting movements which can be performed. These five shunting movement all fit in the timetable without shifting any other trains. Unlike the fourth instance, our algorithm succeeds in finding a rolling stock schedule which provides improvement in several objectives; we find a decrease in the number of carriage kilometers driven and the number of seat-shortage hours. Additionally, we find a decrease in the total number of new uncoupled units, whilst there are three more new coupled units and two more uncoupled units at Utrecht Centraal compared to the basic reschedule. Overall, we find the largest percentual improvement over the basic reschedule for this instance.

6.4.6 Summary

We will now briefly summarize and evaluate the performance of our iterative algorithm on the five instances. Table 11 summarizes the most important statistics of our experiments so far, which include the objective values of the basic reschedule, the final schedule and after the first iteration, the improvement of the final schedule over the basic schedule, the optimality gap of the objective value of the final schedule compared to the objective value after the first iteration and the number of new shunting movements introduced at Utrecht Centraal. As mentioned previously, the objective value of the basic reschedule forms an upper bound on the objective value of the final schedule, as the former is computed in a setting in which shunting at Utrecht Centraal is prohibited. Likewise, the objective value after the first iteration of the composition model forms a lower bound, as the generated rolling stock schedule assumes that all composition changes and shunting movements at Utrecht Centraal can be performed.

Table 11: Summarizing statistics for all five instances.

Instance	1	2	3	4	5
Objective value of basic reschedule	11,667,191.10	5,591,447.62	15,558,125.80	13,020,380.30	8,035,207.26
Objective value of final schedule	9,067,208.88	5,591,447.62	13,670,947.30	12,541,975.62	5,921,914.46
Objective value after first iteration	9,067,208.88	5,392,770.66	13,355,961.14	12,541,975.62	5,921,914.46
Improvement over basic reschedule (%)	28.67	0.00	14.02	3.81	35.69
Optimality gap (%)	0.00	3.68	2.36	0.00	0.00
#New shunting movements at Ut	6	0	8	5	5

For the first, fourth and fifth instance, we obtain a rolling stock schedule which attains the optimal objective value as outputted by the first iteration in which all suggested shunting movements are preliminarily deemed to be possible. We find six, five and five shunting movements for these three instances respectively which can be performed at Utrecht Centraal. All shunting movements which are suggested for the second instance do not fit in between the other planned trains, and for the third instance we find an improvement of around fourteen percent by introducing eight new shunting movements. We conclude that for disruptions that take place directly adjacent to Utrecht Centraal and disruptions that take place somewhere else in the country, our approach is able to identify shunting movements at Utrecht Centraal which offer improvements for the adjusted rolling stock schedule with respect to operational, passenger service level and stick-to-the-plan objectives, if they can be performed.

6.5 Sensitivity analysis

As our algorithm and each of its components are driven by the values of the different weights and parameters, we explore the extent to which the results are influenced by the parameter choices. We perform a sensitivity analysis on the weight of choosing a different composition compared to the initial schedule, the maximum allowed deviation and the travel time between pairs of tracks, crossings and/or switches.

6.5.1 Weight of choosing a different composition compared to the initial schedule

For our previous experiments, we used a weight of 100,000 for the parameter w_{12}, which represents the penalty that is incurred for choosing a composition that is different than originally planned in the initial schedule for a single trip. In the objective value of the composition model, the stick-to-the-plan objective therefore dominated the other operational and passenger service level objectives. In this section, we will vary the value of this parameter as shown in Table 12 for each instance. The remaining parameter settings as discussed in Section 6.2 will remain the same.

Table 12: The values of the weight of different compositions compared to the initial schedule, w_{12}.

Parameter	**Values**			
w_{12}	0.01	100	1,000	100,000

Since the weight w_{12} resembles the penalty of choosing a different composition compared to the composition that was originally planned, this weight partially represents the importance of the stick-to-the-plan objective compared to the other operational and passenger service level objectives. We choose to experiment with values of w_{12} that are below the weights of the other objectives (0.01), that are at around the same level as the weights of the other objectives (100) and that are above the weights of the other objectives (1,000 and 100,000). We perform the basic reschedule and use our iterative algorithm for all instances; the results are shown in Appendix A.1. Table 13 shows the averages over all instances of the same statistics that were reported for the previous experiments for the basic reschedule.

Table 13: Average statistics over all instances for the basic reschedule for different weights of choosing a different composition.

Value of w_{12}	0.01	100	1,000	100,000
Running time (s)	228.40	188.45	99.92	73.67
Objective value	535,888.34	556,147.68	662,291.44	10,774,470.40
Carriage kilometers	1,198,948.4	1,199,253.2	1,195,447.0	1,193,450.0
Seat-shortage hours	5,912.2	5,972.4	6,253.4	6,658.2
Seat-shortage during disruption	2,481.6	2,499.6	2,532.60	2,555.80
#New coupled units	22.6	18.4	12.8	11.2
#New uncoupled units	28.4	25.8	19.6	18.2
#New coupled units at Ut	0.0	0.2	0.0	0.0
#New uncoupled units at Ut	1.2	1.4	1.4	1.2
Different compositions	346.8	153.8	110.2	102.0
Different compositions at Ut	7.6	6.8	5.2	4.8
Inventory differences	1.6	2.4	5.6	5.2
Inventory differences at Ut	0.4	0.4	0.4	0.6

Several trends can be identified in the average results of the basic reschedule. The running time decreases as the weight of choosing a different composition increases; since we more heavily penalize deviating from the initial rolling stock schedule, we guide the model more towards sticking to the original plan and it is therefore able to find a solution more quickly. Next, a clear trade-off is present between the different objectives. As we increase the weight of choosing a different composition and therefore put more emphasis on the stick-to-the-plan objective, the average number of (un)coupled units and different compositions decrease. Conversely, the number of seat-shortage hours and inventory differences increase, which means that we indeed find a decline in the operational objectives and passenger service level objectives. Furthermore, note that the average number of (un)coupled units and the inventory differences at Utrecht Centraal remain stable across the different weights. We initially do not allow composition changes at Utrecht Centraal, besides the special cases mentioned in Section 6.4, which causes these statistics to remain at their low values. However, the number of different compositions at Utrecht Centraal does decrease as the weight of choosing a different composition increases. This statistic contains all trips that either start or end at Utrecht Centraal and it is therefore likely that fewer different compositions are chosen earlier in the service at the stations at which composition changes are allowed before visiting Utrecht Centraal.

Table 14: Average statistics over all instances for the iterative algorithm for different weights of choosing a different composition.

Value of w_{12}	0.01	100	1,000	100,000
Composition model				
Running time (s)	733.72	499.11	298.17	234.80
Objective value	524,634.52	541,871.31	637,246.74	9,358,698.78
Carriage kilometers	1,196,388.0	1,196,437.4	1,194,875.2	1,192,649.6
Seat-shortage hours	5,766.0	5,857.0	6,048.8	6,696.6
Seat-shortage hours during disruption	2,350.2	2,377.0	2,459.2	2,536.6
#New coupled units	24.0	18.6	15.2	13.4
#New uncoupled units	31.8	25.6	21.8	20.6
#New coupled units at Ut	3.0	3.0	2.8	2.6
#New uncoupled units at Ut	4.2	4.2	3.2	3.0
Different compositions	310.2	136.0	94.8	87.8
Different compositions at Ut	10.2	9.8	6.6	5.4
Inventory differences	1.2	2.0	5.6	7.2
Inventory differences at Ut	0.4	0.4	1.0	1.0
Infrastructure model				
Running time (s)	154.35	142.32	124.49	137.10
#Shifted trains	3.0	3.0	2.2	2.4
Total deviation (min)	2.72	2.72	1.88	2.08
Average deviation per shifted train (min)	0.90	0.93	0.85	0.87
General				
Running time (s)	888.11	641.60	422.66	371.90
#Iterations	2.6	2.4	2.4	2.4
#Shunting movements suggested	8.0	7.4	6.2	6.0
#Shunting movements cancelled	1.4	1.2	0.8	1.2
Objective value after first iteration	524,310.18	541,478.70	636,360.37	9,255,966.15
Optimality gap (%)	0.06	0.07	0.14	1.11
Improvement over basic reschedule (%)	2.14	2.63	3.93	15.13

Table 14 shows the average statistics over all instances for the iterative algorithm. The main trends that we discussed for the basic reschedule are also present in the results of the iterative algorithm.

This includes a decrease in the average running times as the weight of choosing a different composition increases and the trade-off between the stick-to-the-plan objective and the operational and passenger service level objectives. Additionally, these experiments show changes in the statistics regarding the shunting movements. Compared to the base value of 100,000 that was used for the results in Section 6.4, a value of 0.01 leads to on average two more suggested shunting movements. The number of iterations also increases by on average 0.2, but all problem instances still do not require more than three iterations total. Furthermore, since we now allow shunting at Utrecht Centraal, there is a decreasing trend in the number of (un)coupled units at Utrecht Centraal and in the number of trips which start or end at Utrecht Centraal with a different composition as the weight of choosing a different composition increases. Moreover, as more shunting movements are suggested for lower weights of choosing a different composition, the number of shifted trains and the total deviation increase. The number of shunting movements that are cancelled and the average deviation per shifted train remain steady across the different weights; this suggests that these statistics are instead more dependent on the exact shunting movements that are suggested, which are largely the same for the different weights. The optimality gap and percentage improvement over the basic reschedule depend on the scale of the weights in the objective function. The smallest weight of choosing a different composition provides the smallest average optimality gap and also the smallest improvement over the basic reschedule.

To summarize, our algorithm has the ability to perform rolling stock rescheduling with emphasis on the objective of choice. Decreasing the penalty for choosing a different composition than the one that was originally planned improves the statistics that pertain to the operational and passenger service level objectives by introducing additional composition changes and shunting movements, at the cost of larger running times and more planned train deviations. With average running times of less than fifteen minutes depending on the chosen weight, our algorithm provides improvement over the basic reschedule within practical running times.

6.5.2 Maximum deviation per train event

In this section, we will solve the instances with different values of d_{max}, the maximum deviation per train event. We previously assumed that a deviation of one minute is acceptable. Table 15 shows the values of d_{max} that we experiment with.

Table 15: The values of the maximum allowed deviation per train event, d_{max}.

Parameter	**Values**		
d_{max}	0	1	2

Allowing a maximum deviation of zero minutes represents a situation in which no planned trains can be shifted in time; shunting movements are therefore only performed if they already fit between all the other trains. Additionally, we investigate whether accepting an extra minute of deviation allows for more shunting movements to be performed. We furthermore choose to continue with the weight $w_{12} = 100$, as this increases the number of shunting movements that are suggested without drastically increasing the running times of the experiments. The results of the iterative algorithm are shown in Appendix A.2. The results of the basic reschedule correspond to the results in Table A1 in Appendix A.1 with $w_{12} = 100$. Table 16 contains the average statistics over all instances for the different values of d_{max}.

Table 16: Average statistics over all instances for the iterative algorithm for different maximum deviations.

Value of d_{max}	0	1	2
Composition model			
Running time (s)	644.95	499.11	471.33
Objective value	543,776.93	541,871.31	541,593.78
Carriage kilometers	1,197,411.0	1,196,437.4	1,196,586.4
Seat-shortage hours	5,884.6	5,857.0	5,874.8
Seat-shortage hours during disruption	2,368.8	2,377.0	2,381.0
#New coupled units	18.2	18.6	18.8
#New uncoupled units	25.8	25.6	25.8
#New coupled units at Ut	2.4	3.0	3.2
#New uncoupled units at Ut	3.6	4.2	4.6
Different compositions	132.0	136.0	134.0
Different compositions at Ut	18.0	9.8	10.2
Inventory differences	2.8	2.0	2.0
Inventory differences at Ut	0.4	0.4	0.4
Infrastructure model			
Running time (s)	90.71	142.32	154.62
#Shifted trains	0.0	3.0	3.8
Total deviation (min)	0.00	2.72	3.82
Average deviation per shifted train (min)	-	0.93	1.10
General			
Running time (s)	735.66	641.60	625.94
#Iterations	3.0	2.4	2.2
#Shunting movements suggested	8.2	7.4	7.4
#Shunting movements cancelled	3.0	1.2	0.6
Objective value after first iteration	541,478.70	541,478.70	541,478.70
Optimality gap (%)	0.42	0.07	0.02
Improvement over basic reschedule (%)	2.27	2.63	2.69

Our results show that the average overall running time increases as we decrease the maximum deviation per train event. The major increase in running time comes from the composition model, whilst we find a decrease in the running time of the infrastructure model. This can be explained by an increase in the number of iterations performed by our iterative algorithm for lower values of d_{max}; relatively many suggested shunting movements in the first iteration are cancelled as we do not provide any room for the shifting of planned trains, which leads our algorithm to perform additional iterations in which alternative shunting movements are suggested. The number of suggested and cancelled shunting movements therefore also increases as the maximum deviation per train event decreases. The infrastructure model can be solved significantly faster with a maximum deviation of zero, as the infrastructure model does not have the ability to shift any trains and therefore only recognizes the shunting movements which already fit. A large portion of the remaining running time corresponds to starting up the model and converting the output into usable feedback for the composition model.

Furthermore, we find a decrease in the optimality gap of the found solution compared to the objective value after the first iteration as the maximum deviation increases, and subsequently also an increase in the improvement over the basic reschedule. Since increasing the maximum allowed deviation provides the infrastructure model with more possibilities to fit the shunting movements, the resulting rolling stock schedule is always at least as good as that created with a lower maximum deviation.

However, note that the first iteration of the composition model provides the same rolling stock schedule and therefore the same objective value for all values of the maximum allowed deviation, as we use identical parameter settings for that part of the algorithm. Our current implementation of the composition model does not recognize which shunting movements can be performed without incurring any planned train delays; comparing the performance of our algorithm with a maximum allowed deviation of zero to the same objective value after the first iteration that is obtained with a positive maximum allowed deviation might therefore lead to a skewed image of the achieved potential, as the shunting movements which require the shifting of other trains simply cannot be performed. To counter this, it could for instance be determined prior to the start of the algorithm that the shunting movements which clash with at least one other train cannot be performed for $d_{max} = 0$. A preprocessing iteration of the infrastructure model which finds for each track in which time window a shunting movement to or from this track could be performed without shifting any planned trains could be included for this specific parameter setting. All transitions which take place outside the found time windows could then be prevented from containing a composition change. Adding such a preprocessing step would eliminate the impossible shunting movements before the start of the algorithm, which means that the rolling stock schedule after the first iteration would also be the best possible rolling stock schedule for this parameter setting. For positive maximum deviations, such a preprocessing step would be more difficult to implement as detecting a violated headway constraint does not necessarily rule out the possibility of performing a shunting movement in case we allow the shifting of other trains.

Further comparing the different objectives of the composition model, we find that there is no significant change for some of the resulting statistics as we vary the maximum allowed deviation. The average number of carriage kilometers, seat-shortage hours, different compositions and inventory differences are all close to each other for the three different parameter values. An increase in the number of (un)coupled units at Utrecht Centraal as the maximum allowed deviation increases follows from the extra shunting movements that fit and can be performed at Utrecht Centraal. Additionally, a larger maximum deviation leads to on average more shifted trains and more total deviation. The average deviation per shifted train for $d_{max} = 2$ is larger than one minute, which indicates that the extra minute of allowed delay provides the model with the necessary room to fit extra shunting movements which benefit the rolling stock schedule as created by the composition model. It depends on the preferences of the rolling stock dispatcher and the practical possibilities of delaying planned trains whether this extra minute is acceptable in practice and whether this rolling stock schedule is therefore an improvement over that with a lower maximum allowed deviation.

Note that we currently do not include a penalty for planned train deviations in the objective function of the composition model, because we assume that the specific value of d_{max} is the maximum deviation that is allowed in practice and that all shunting movements that obey this maximum deviation should be performed so as to improve the rolling stock schedule as outputted by the composition model as much as possible. The quality of the obtained rolling stock schedule can, however, not be assessed by purely looking at the objective value as reported by the composition model. Whilst the performance of extra shunting movements as suggested by the composition model does improve either the operational, passenger service level or stick-to-the-plan objective as specified by the individual weights, this could come at the cost of incurring planned train delays which therefore also needs to be taken into account. To more accurately compare the quality of the rolling stock schedules, the final objective value with which the optimality gap is calculated could therefore contain a separate penalty for the incurred planned train delays, which depends on the number of shifted trains and on the total incurred deviation, on top of the objective value of the final iteration of the composition model. Nevertheless, the severity of the penalty depends on what a rolling stock dispatcher deems to be acceptable and executable in practice.

In summary, the results in this section show that our algorithm successfully finds an adjusted rolling stock schedule for different maximum allowed deviations per train event in the infrastructure model. This includes a maximum deviation of zero minutes, which means that only shunting movements are allowed which do not require any other planned trains to be delayed. Increasing the maximum allowed deviation to two minutes allows more shunting movements to be fit with a lower running time compared to lower maximum deviations, at the cost of more incurred planned train deviations.

6.5.3 Travel time between tracks, crossings and/or switches

We have previously assumed that it takes a time of $r = 0.1$ minutes, or six seconds, to travel between pairs of tracks, crossings and/or switches on the exact infrastructure of the station that we consider. This holds for both the shunting movements and the planned trains that visit Utrecht Centraal somewhere in their journey. Additionally, we assume that a coupling movement must take place at least five minutes before the departure of the new composition and that an uncoupling movement must take place at least five minutes after the arrival of the old composition. These assumptions may deviate from what is possible in practice; trains need to accelerate when they have come to a complete stop and need to depart from a platform, which means that an increasing travel time for the first few pairs of structures may provide a more accurate representation of reality. We will instead experiment with larger values of r, as shown in Table 17.

Table 17: The values of the travel time between tracks, crossings and/or switches, r.

Parameter	**Values**		
r	0.1	0.2	0.3

Similar to our previous sensitivity analyses, we perform the basic reschedule and use our iterative algorithm to solve all instances for the aforementioned values of r. We again choose to continue with the different composition weight $w_{12} = 100$ and to allow a maximum deviation of $d_{max} = 1$ minute. The results of the iterative algorithm can be found in Appendix A.3. The results of the basic reschedule correspond to the results in Table A1 in Appendix A.1 with $w_{12} = 100$. Table 18 shows the average statistics over all instances for the different values of r.

Table 18: Average statistics over all instances for the iterative algorithm for different travel times between tracks, crossings and/or switches.

Value of r	0.1	0.2	0.3
Composition model			
Running time (s)	499.11	535.61	618.77
Objective value	541,871.31	541,879.40	545,815.20
Carriage kilometers	1,196,437.4	1,196,685.6	1,198,414.0
Seat-shortage hours	5,857.0	5,877.2	5,919.8
Seat-shortage hours during disruption	2,377.0	2,369.0	2,372.4
#New coupled units	18.6	18.6	17.8
#New uncoupled units	25.6	25.6	26.0
#New coupled units at Ut	3.0	3.0	2.2
#New uncoupled units at Ut	4.2	4.2	3.6
Different compositions	132.0	134.0	142.0
Different compositions at Ut	9.8	9.4	9.0
Inventory differences	2.0	2.0	2.4
Inventory differences at Ut	0.4	0.4	0.4
Infrastructure model			
Running time (s)	142.32	158.21	174.10
#Shifted trains	3.0	4.2	4.6
Total deviation (min)	2.72	3.20	3.30
Average deviation per shifted train (min)	0.93	0.80	0.68
General			
Running time (s)	641.60	693.82	792.87
#Iterations	2.4	2.4	2.8
#Shunting movements suggested	7.4	7.4	9.6
#Shunting movements cancelled	1.2	1.0	3.4
Objective value after first iteration	541,478.70	541,478.70	541,478.70
Optimality gap (%)	0.07	0.07	0.08
Improvement over basic reschedule (%)	2.63	2.63	1.89

In contrast to the conducted analyses of the previous sections, the parameter that we vary in this section does not affect the manner in which our iterative algorithm finds a solution. Changing the weight for choosing a different composition changes the importance of the stick-to-the-plan objective and changing the maximum allowed deviation changes the possibilities the infrastructure model has to fit shunting movements; varying the travel time between tracks, crossings and/or switches simply changes the exact arrival and departure times at each of the individual structures for both planned trains and shunting movements, which potentially influences the pairs of trains for which the minimum headway time requirements holds and does not hold anymore.

Nevertheless, increasing the travel time does affect the obtained solutions in several ways. The overall average running times increase as the travel time increases. Choosing a travel time of 0.3 minutes, or eighteen seconds, causes our algorithm to cancel on average two more shunting movements than for a travel time of six and twelve seconds, which leads to more iterations in which alternative shunting movements are suggested and therefore larger running times. This is also reflected in a decrease in the number of (un)coupled units at Utrecht Centraal, as fewer of the suggested shunting movements at Utrecht Centraal can be performed, and an increase in the overall number of different compositions and inventory differences.

Furthermore, note that the number of shifted trains and the total deviation increase as the travel time increases, but that the average deviation per shifted train decreases. The remaining operational and passenger service level statistics are all relatively close to each other for the different travel times. The best rolling stock schedule, as evaluated by the average objective value of the final schedule, can be found for a travel time of six seconds, with the lowest running times and the lowest overall deviations.

All in all, the goal of varying this parameter is to model reality as closely as possible. Increasing the what we assume to be a constant travel time may portray the real travel times more accurately, but this leads to larger running times and more cancelled shunting movements in the obtained rolling stock schedule. Our approach could potentially be extended by incorporating the acceleration and deceleration of rolling stock units and compositions in the travel times between the considered structures for dwelling passenger trains and for shunting movements.

7 Conclusions

In this section, we look back on the research that we presented in this thesis and discuss what can be done in future work. Section 7.1 presents a summary of our research and some concluding remarks. Section 7.2 discusses our recommendations for future research.

7.1 Summary and conclusions

In this thesis, we consider the problem of rolling stock rescheduling during disruptions whilst ensuring feasibility with respect to the available railway infrastructure for the performance of shunting movements. We introduce an iterative algorithm that alternates between solving the composition model, which outputs a rolling stock schedule that contains the rolling stock compositions that are used for each trip and the composition changes and shunting movements that take place between two consecutive trips, and the infrastructure model, which, given the disrupted timetable and the suggested shunting movements, finds a feasible time slot to perform each shunting movement, if one exists. The infrastructure model feeds the impossible shunting movements back to the composition model and the algorithm continues until a feasible rolling stock schedule is obtained.

We use five disruption instances on the Dutch railway network that contain full blockages at and around Utrecht Centraal to test our methods. We find that our algorithm succeeds in adjusting the rolling stock schedules in the face of disruptions within running times of around a few minutes. The infrastructure model is able to identify shunting movements that can take place at Utrecht Centraal and incorporates them into the adjusted rolling stock schedule, hence improving upon the existing rolling stock rescheduling methods which do not consider the possibility of shunting at a busy station like Utrecht Centraal. Our approach comes at the cost of delaying other passenger trains. However, our sensitivity analyses show that even in the situation where passenger train delays are not allowed, our approach still succeeds in finding and fitting several shunting movements and composition changes which would otherwise have been assumed to be impossible. Additionally, we identify the trade-off that is present between the different operational, passenger service level and stick-to-the-plan objectives and illustrate that our algorithm can function with a focus on any of these objectives.

7.2 Recommendations

We recommend future research to look at combining our existing framework with other previously explored problems that model different aspects of rolling stock rescheduling. One of these is the Train Unit Shunting Problem (TUSP), as described by Haahr et al. (2017). This problem creates matchings between incoming and outgoing rolling stock units at shunting yards and assigns them to the available shunting tracks such that all movements can be performed without any units being blocked in. In our research, we have assumed that enough space on the shunting tracks is available during the day. However, on smaller stations with smaller shunting yards, the storage of too many rolling stock units could impair the movement possibilities of rolling stock units parked near the back. We also assumed that all shunting movements follow a fixed route to one specific shunting yard and that one inventory of rolling stock units exists for each station, whilst a station can have multiple shunting yards and therefore multiple separate inventories. All of these aspects are contained in the TUSP and an iterative framework could be implemented to integrate these aspects into one solution method. Since the TUSP is a feasibility problem which takes the incoming and outgoing shunting movements as a given, an infeasible problem instance could be fed back to our framework with the conclusion that this combination of shunting movements cannot be performed due to shunting track limitations. A consideration then needs to be made regarding the exact shunting movements that are excluded from the solution.

Additionally, we recommend incorporating the Shunting Driver Scheduling Problem (SDSP), as described by Hoogervorst (2021). This problem aims at generating a set of feasible duties for all present shunting drivers. Hoogervorst (2021) solves an integrated version of the problem which simultaneously solves the Rolling Stock Scheduling Problem (RSSP), as is given by the composition model, and the SDSP with the means of a Benders decomposition. An interesting research direction could be the further incorporation of our infrastructure constraints, to ensure that each of the shunting tasks in the generated shunting driver duties is possible with respect to the available infrastructure.

Furthermore, we recommend extending our models by incorporating the infrastructures of multiple stations. In our research, we have limited ourselves to investigating the shunting possibilities at Utrecht Centraal with one disruption at a time. We allow for shunting at the five stations that are used for the generation of the initial rolling stock schedule, at the stations that are adjacent to the disruption and at Utrecht Centraal. However, Utrecht Centraal is the only station at which we check the feasibility of the shunting movements. By allowing shunting at more stations, multiple disruption instances throughout the country can be solved at once and the shunting possibilities at those stations that are otherwise deemed as impossible can be explored.

Moreover, we recommend investigating disruptions consisting of multiple blockages which take place at different places in the country and which do not all start at the same time but do partially overlap in time. Our experiments have all contained only one blockage with one fixed duration, which means that all suggested shunting movements serve to counter this specific disruption. In case different shunting movements are suggested to counter different disruptions simultaneously, the number of shunting movements per iteration and the order in which they are tried may be of importance in the iterative algorithm.

Bibliography

Cacchiani, V., Huisman, D., Kidd, M., Kroon, L., Toth, P., Veelenturf, L., and Wagenaar, J. (2014). An overview of recovery models and algorithms for real-time railway rescheduling. *Transportation Research Part B: Methodological*, 63:15–37.

Fioole, P.-J., Kroon, L., Maróti, G., and Schrijver, A. (2006). A rolling stock circulation model for combining and splitting of passenger trains. *European Journal of Operational Research*, 174(2):1281–1297.

Freling, R., Lentink, R. M., Kroon, L. G., and Huisman, D. (2005). Shunting of passenger train units in a railway station. *Transportation Science*, 39(2):261–272.

Haahr, J. and Lusby, R. M. (2017). Integrating rolling stock scheduling with train unit shunting. *European Journal of Operational Research*, 259(2):452–468.

Haahr, J. T., Lusby, R. M., and Wagenaar, J. C. (2017). Optimization methods for the train unit shunting problem. *European Journal of Operational Research*, 262(3):981–995.

Hoogervorst, R. (2021). *Improving the Scheduling and Rescheduling of Rolling Stock: Solution Methods and Extensions*. Number 534. Erasmus Institute of Management (ERIM).

Hoogervorst, R., Dollevoet, T., Maróti, G., and Huisman, D. (2021). A variable neighborhood search heuristic for rolling stock rescheduling. *EURO Journal on transportation and logistics*, 10:100032.

Nielsen, L. K. (2011). *Rolling Stock Rescheduling in Passenger Railways: Applications in short-term planning and in disruption management*. Number EPS-2011-224-LIS.

Nielsen, L. K., Kroon, L., and Maróti, G. (2012). A rolling horizon approach for disruption management of railway rolling stock. *European Journal of Operational Research*, 220(2):496–509.

Van Aken, S., Bešinović, N., and Goverde, R. M. (2017). Solving large-scale train timetable adjustment problems under infrastructure maintenance possessions. *Journal of rail transport planning & management*, 7(3):141–156.

Van den Broek, R., Hoogeveen, H., Van den Akker, M., and Huisman, B. (2022). A local search algorithm for train unit shunting with service scheduling. *Transportation Science*, 56(1):141–161.

Van Veen, I. (2021). Rolling stock rescheduling with dynamic passenger flows in case of a disruption. Master's thesis.

Wagenaar, J., Kroon, L., and Fragkos, I. (2017). Rolling stock rescheduling in passenger railway transportation using dead-heading trips and adjusted passenger demand. *Transportation Research Part B: Methodological*, 101:140–161.

A Appendix

A.1 Different composition weight results

Table A1: Resulting statistics of the basic reschedule for different weights of choosing a different composition for all instances.

Disruption	Value of w_{12}	Running time (s)	Objective value	Carriage kilometers	Seat-shortage hours	Seat-shortage hours during disruption	#New coupled units	#New uncoupled units	#New coupled units at Utrecht	#New uncoupled units at Ut	Different compositions	Different compositions at Ut	Inventory differences	Inventory differences at Ut
1	0.1	200.38	581,130.29	1,199,560	6,618	3,228	31	41	0	4	349	5	2	1
1	100	188.84	603,491.94	1,199,438	6,742	3,253	29	40	1	5	169	4	4	1
1	1,000	119.56	727,453.96	1,193,892	6,921	3,251	21	33	0	5	126	4	12	1
2	0.1	231.96	569,310.64	1,197,552	6,568	3,215	19	23	0	2	207	2	0	0
2	100	183.96	581,058.34	1,196,558	6,689	3,230	13	18	0	2	82	2	2	0
2	1,000	94.02	640,402.42	1,195,174	6,762	3,303	14	17	0	2	58	0	6	1
3	0.1	246.76	520,117.34	1,199,578	5,546	2,007	27	32	0	0	617	13	2	0
3	100	205.20	550,105.90	1,201,430	5,371	2,023	21	29	0	0	232	13	2	0
3	1,000	100.06	700,847.10	1,194,870	6,005	2,054	12	20	0	0	162	10	6	0
4	0.1	215.11	485,118.88	1,199,606	5,073	1,615	17	23	0	0	310	7	2	0
4	100	186.30	507,631.34	1,199,718	5,191	1,634	15	22	0	0	164	4	2	0
4	1,000	104.22	632,546.48	1,196,296	5,494	1,665	8	14	0	0	130	4	2	0
5	0.1	247.77	523,764.55	1,198,446	5,756	2,343	19	23	0	0	251	11	2	1
5	100	177.97	538,450.86	1,199,122	5,869	2,358	14	20	0	0	122	8	2	1
5	1,000	81.75	610,207.26	1,197,002	6,085	2,390	9	14	0	0	75	8	2	0

Table A2: Resulting statistics of the iterative algorithm for different weights of choosing a different composition for all instances.

		General					Composition model													Infrastructure model			
Disruption	Value of w_{12}	Running time (s)	#Iterations	#Shunting movements suggested	#Shunting movements cancelled	Objective value after first iteration	Running time (s)	Objective value	Carriage kilometers	Seat-shortage hours	Seat-shortage hours during disruption	#New coupled units	#New uncoupled units	#New coupled units at Ut	#New uncoupled units at Ut	Different compositions	Different compositions at Ut	Inventory differences	Inventory differences at Ut	Running time (s)	#Shifted trains	Total deviation (min)	Average deviation per shifted train (min)
1	0.1	929,69	2	9	2	576,195.13	764.03	576,730.00	1,197,432	6,526	3,154	36	47	3	8	384	9	2	1	165.66	5	4.00	0.80
1	100	572.13	2	9	2	595,815.42	405.03	595,950.30	1,194,310	6,618	3,176	32	41	3	8	167	9	2	1	167.10	5	4.00	0.80
1	1,000	423.83	2	7	1	706,560.76	280.00	706,560.76	1,190,052	7,028	3,390	22	31	3	6	101	5	12	1	143.82	4	3.00	0.75
2	0.1	839.14	3	6	3	567,436.77	733.56	567,972.71	1,196,418	6,493	3,168	24	31	1	3	251	5	0	0	105.58	0	0.00	-
2	100	728.15	3	6	3	580,229.30	613.32	580,563.59	1,194,295	6,583	3,233	15	19	1	3	70	4	4	0	114.83	1	1.00	1.00
2	1,000	480.91	3	4	1	637,273.98	396.53	637,273.98	1,194.186	6,758	3,303	13	19	1	3	49	2	6	1	84.38	1	0.80	0.80
3	0.1	1,206.25	3	14	2	499,886.89	933.05	500,437.74	1,195,966	5,342	1,766	25	34	5	5	416	16	2	0	273.01	6	6.00	1.00
3	100	661.84	2	9	1	522,694.06	490.93	524,187.92	1,198,184	5,415	1,779	20	29	4	4	200	12	2	0	170.09	5	5.00	1.00
3	1,000	602.15	3	10	2	655,259.62	411.18	659,691.44	1,196,288	5,606	1,814	17	24	4	3	143	13	6	1	190.98	3	3.00	1.00
4	0.1	834.05	3	6	0	477,107.90	709.23	477,107.90	1,197,246	5,014	1,674	19	26	3	3	292	11	0	0	124.81	4	3.60	0.90
4	100	411.31	2	5	0	497,533.22	311.69	497,533.22	1,198,994	5,110	1,687	14	22	3	3	175	12	0	0	99.62	4	3.60	0.90
4	1,000	270.49	2	5	0	616,537.54	173.24	616,537.54	1,197,458	5,254	1,740	12	19	3	2	127	7	2	1	97.25	3	2.60	0.86
5	0.1	631.40	2	5	0	500,924.24	528.71	500,924.24	1,194,878	5,455	1,989	16	21	3	2	208	10	2	1	102.69	0	0.00	-
5	100	834.55	3	8	0	511,121.52	674.60	511,121.52	1,196,404	5,559	2,010	12	17	4	3	68	12	2	1	159.95	0	0.00	-
5	1,000	335.94	2	5	0	566,169.96	229.92	566,169.96	1,196,392	5,598	2,049	12	16	3	2	54	6	2	1	106.02	0.	0.00	-

A.2 Maximum deviation results

Table A3: Resulting statistics of the iterative algorithm for different maximum deviations for all instances.

		General					Composition model													Infrastructure model			
Disruption	Value of d_{max}	Running time (s)	#Iterations	#Shunting movements suggested	#Shunting movements cancelled	Objective value after first iteration	Running time (s)	Objective value	Carriage kilometers	Seat-shortage hours	Seat-shortage hours during disruption	#New coupled units	#New uncoupled units	#New coupled units at Ut	#New uncoupled units at Ut	Different compositions	Different compositions at Ut	Inventory differences	Inventory differences at Ut	Running time (s)	#Shifted trains	Total deviation (min)	Average deviation per shifted train (min)
1	0	786.55	3	11	7	595,815.42	681.65	602,239.70	1,200,190	6,664	3,178	31	44	1	6	183	55	4	1	104.90	0	0.00	-
1	2	617.74	2	9	1	595,815.42	433.00	595,815.42	1,194,534	6,622	3,165	32	42	3	9	167	10	2	1	184.74	7	6.40	0.91
2	0	768.72	3	6	4	580,229.30	706.80	580,565.26	1,195,542	6,645	3,225	15	19	1	3	72	3	4	0	61.92	0	0.00	-
2	2	618.33	2	5	2	580,229.30	515.52	580,565.26	1,195,542	6,671	3,230	15	19	1	3	72	3	4	0	102.80	1	1.50	1.50
3	0	741.78	3	10	2	522,694.06	625.27	524,642.36	1,198,109	5,381	1,778	21	29	4	4	198	12	2	0	116.51	0	0.00	-
3	2	621.83	2	9	0	522,694.06	421.77	522,933.46	1,197,462	5,412	1,778	21	29	5	5	188	14	2	0	200.06	7	7.60	1.09
4	0	612.31	3	6	2	497,533.22	544.29	500,315.82	1,196,814	5,174	1,618	12	20	2	2	139	8	2	0	68.02	0	0.00	-
4	2	419.24	2	6	0	497,533.22	316.07	497,533.22	1,198,994	5,110	1,687	14	22	3	3	175	12	0	0	103.20	4	3.60	0.90
5	0	768.95	3	8	0	511,121.52	666.76	511,121.52	1,196,400	5,559	2,045	12	17	4	3	68	12	2	1	102.18	0	0.00	-
5	2	852.58	3	8	0	511,121.52	670.28	511,121.52	1,196,400	5,559	2,045	12	17	4	3	68	12	2	1	182.29	0	0.00	-

A.3 Travel time between tracks, crossings and/or switches results

Table A4: Resulting statistics of the iterative algorithm for different travel times between tracks, crossings and/or switches for all instances.

		General					Composition model													Infrastructure model			
Disruption	Value of r	Running time (s)	#Iterations	#Shunting movements suggested	#Shunting movements cancelled	Objective value after first iteration	Running time (s)	Objective value	Carriage kilometers	Seat-shortage hours	Seat-shortage hours during disruption	#New coupled units	#New uncoupled units	#New coupled units at Ut	#New uncoupled units at Ut	Different compositions	Different compositions at Ut	Inventory differences	Inventory differences at Ut	Running time (s)	#Shifted trains	Total deviation (min)	Average deviation per shifted train (min)
1	0.2	563.78	2	9	2	595,815.42	386.59	595,950.30	1,194,304	6,641	3,142	32	41	3	8	167	9	2	1	177.19	3	3.00	1.00
1	0.3	661.52	2	9	3	595,815.42	470.56	601,672.56	1,200,062	6,652	3,121	30	44	1	7	185	8	4	1	190.96	3	2.30	0.77
2	0.2	854.94	3	6	2	580,229.30	729.45	580,516.06	1,195,542	6,674	3,232	15	19	1	3	74	4	4	0	125.49	1	1.00	1.00
2	0.3	626.21	2	5	2	580,229.30	508.64	580,528.26	1,195,542	6,669	3,197	15	19	1	3	74	2	4	0	117.57	2	0.80	0.40
3	0.2	654.92	2	9	1	522,694.06	460.83	524,175.92	1,198,184	5,402	1,774	20	29	4	4	186	10	2	0	194.09	7	6.00	0.86
3	0.3	1,099.58	4	13	4	522,694.06	943.18	526,924.80	1,199,840	5,489	1,744	18	28	4	4	192	13	2	0	156.40	6	5.60	0.93
4	0.2	450.61	2	5	0	497,533.22	339.29	497,533.22	1,198,994	5,110	1,687	14	22	3	3	175	12	0	0	111.32	6	4.20	0.70
4	0.3	422.20	2	5	1	497,533.22	317.66	497,677.00	1,197,980	5,111	1,652	13	21	2	2	178	10	0	0	104.55	5	3.40	0.68
5	0.2	944.86	3	8	0	511,121.52	761.90	511,121.52	1,196,404	5,559	2,010	12	17	4	3	68	12	2	1	182.97	4	1.80	0.45
5	0.3	1,154.83	4	16	7	511,121.52	853.81	522,273.38	1,198,646	5,678	2,148	13	18	3	2	81	12	2	1	301.02	7	4.40	0.63