Erasmus University Rotterdam
Erasmus School of Economics
In collaboration with Netherlands Railways

Master Thesis Operations Research and Quantitative Logistics

Rolling Stock Rescheduling with Dynamic Passenger Flows in case of a Disruption

Author: Ires van Veen (546152)
Supervisor: Dennis Huisman
Second Assessor: Rolf van Lieshout

September 20, 2021

Abstract

This thesis introduces a new heuristic for the Rolling Stock Rescheduling Problem (RSRP) and provides insight on the impact of stick-to-the-plan oriented rescheduling on passenger comfort and operational costs. The RSRP considers the assignment of vehicle units to trips in case of a disruption and stick-to-the-plan implies that it is strongly undesirable to deviate from the original rolling stock schedule. Moreover, we consider and study the change in passenger flows after a disruption.

The developed heuristic combines the computational advantage of the Composition Model of Nielsen (2011) under limited circumstances with the flexibility of a heuristic. The approach consists of a start heuristic and two neighborhoods in a Variable Neighborhood Descent, namely the Composition Model Neighborhood that optimizes the deployment of one vehicle type and the Two-Opt Duty Neighborhood of Hoogervorst et al. (2021) which swaps duties between the vehicle types.

In computational experiments on four disruptions on the network of Netherlands Railways (NS), the developed heuristic gives a significant decrease in running time in comparison to the exact Composition Model while the increase in the total objective value is reasonably small. Moreover, we show that passenger comfort benefits from rescheduling with dynamic instead of static passenger flows and that a trade-off exists between passenger comfort and stick-to-the-plan. Finally, computational experiments show that focusing on stick-to-the-plan comes at the cost of operational objectives.

Concluding, first, the decrease in computation time and the flexibility of the developed heuristic could be attractive for practical applications and, second, passengers comfort and operational costs can benefit greatly from taking into account dynamic passenger flows in rescheduling.

KeyWords: rolling stock rescheduling, disruption management, passenger railway operations, dynamic passenger flows, Composition Model, heuristic, Variable Neighborhood Descent

Preface

This thesis researches rolling stock rescheduling with dynamic passenger flows in case of a disruption. The main contributions of this research are the introduction of a new heuristic and the insight that stick-to-the-plan rescheduling negatively impacts passenger comfort and operational costs. This thesis was my graduation project of the MSc Operations Research and Quantitative Logistics at the Erasmus University Rotterdam. This research is performed in collaboration with Netherlands Railways (NS) where I was an intern in the Performance Management & Innovation (PI) department.

I would like to thank Dennis Huisman, who was both my supervisor from the Erasmus University and from NS, for all his input and feedback on this thesis. In addition, I would like to thank everyone from team BOS and department PI who contributed to my thesis. In particular, I would like to thank Pieter-Jan, Simone and Nicole for their input during the brainstorm sessions and beyond.

Finally, I would like to thank my friends and family for their support during my thesis. In special, I would like to thank Dirk and Edo for their critical feedback, my roommates, Britt, Sanne, Ellen, Atty, and Anouk, for listening to my endless stories on the functioning (or not functioning) of my models, and Martijn for both.

Ires van Veen
Rotterdam, September 20, 2021

Contents

Introduction

Currently, rolling stock rescheduling in passenger railway operations still involves a great deal of manual adjustments in case of a disruption. This means that in practice, employees manually search for a new, feasible assignment of vehicle units to trips. In addition, the current approach of rolling stock rescheduling is aimed at returning to the original plan as soon as possible, rather than finding a solution that takes into account passenger comfort by minimizing standing hours. Given that on average 16 disruptions occur every day on the Dutch passenger railway network[1], one can understand that the application of rolling stock rescheduling support can be of great value for the seating probability and, therefore, customer satisfaction.

We conduct research into a rolling stock rescheduling algorithm that takes into account the change of passenger flows after a disruption. In this thesis, a disruption is considered to be a complete blockage of a track. In case of a disruption, the timetable is adjusted, as trips are canceled due to the inability to drive on blocked tracks. Subsequently, the passenger demand forecast is adapted to the new situation.

The rolling stock rescheduling model is aimed to assign rolling stock to the trips in the (adapted) timetable using the adapted passenger forecast. This rescheduling model takes into account operational objectives, passenger comfort, and the amount of deviations from the original rolling stock plan. Operational objectives include reducing cancellations, shunting and (un)coupling, additional conductors, type changes, carriage kilometers, and off-balances. Passenger comfort is determined by the total expected amount of hours that passengers are standing while they are on the train.

We looked into literature on disruption management, passenger flows, and rolling stock management. For disruption management, Dollevoet et al. (2017) show that first rescheduling the timetable, then rolling stock, and finally crew is feasible in practice. Therefore, it makes sense to look at the rolling stock rescheduling problem separately. To study the impact of disruptions on passenger flows, most studies employ graph-based models, such as Kroon et al. (2015).

Methods for rolling stock rescheduling can generally be divided into three groups: exact flow-based methods (for example the Composition Model of Nielsen (2011)), exact path-based methods (for example in Lusby et al. (2017)), and recently introduced heuristic methods (Hoogervorst et al. (2021)). Among the exact methods, the Composition Model is computationally most attractive. The recently introduced heuristic is computationally less attractive than the Composition Model under limited circumstances. However, this heuristic has shown to be able to give comparable computational results under circumstances where the feasible solution space is larger due to allowing flexible turning. Flexible turning implies that we can change the way trips succeed one another. The heuristic of Hoogervorst et al. (2021) also has the advantage of providing intermediate results and can more easily incorporate real-life practicalities. This thesis explores the

[1] www.rijdendetreinen.nl/statistieken/2019

possibilities of combining the computational attractiveness of the Composition Model under limited circumstances with the advantages of a heuristic.

The main research question of this thesis is:

How can we model rolling stock rescheduling with dynamic passenger flows in case of a disruption and how does it perform in terms of passenger comfort, computation time, operational costs, and deviations from the original plan?

In order to answer the main question, we split the problem into four sub research questions:

1. *How do the passenger flows change after a disruption?*
2. *How can we reschedule rolling stock to account for timetable changes and changing passenger flows?*
3. *What is the influence of taking into account dynamic passenger flows in rolling stock rescheduling on passenger comfort, operational costs, and deviations from the original schedule?*
4. *How does the developed model perform in terms of passenger comfort, computation time, operational costs, and deviations from the original plan in comparison to methods from literature?*

This research is performed by answering the sub research questions sequentially. We first look into a method to take changing passenger flows into account. Then we look into current models for rolling stock rescheduling and develop a new model. For question 3, we perform two experiments: on one hand, we study the impact of current, stick-to-the-plan based rescheduling. On the other hand, we reschedule rolling stock using static passenger flows and dynamic passenger flows. In both experiments, we study the impact on the passenger comfort, operational objectives and deviation from the original rolling stock schedule. To answer question 4, we will perform a comparison between the developed method and the Composition Model.

This thesis consists of seven chapters, the bibliography, and appendices. The problem description in Chapter 1 consists of three parts: first, we explain all terms and concepts that are used in this thesis. Second, we elaborate on the objectives of rolling stock rescheduling. Third, we provide a generic problem definition. In Chapter 2, we discuss literature on disruption management, passenger flows, and rolling stock rescheduling and give a conclusion about which methods seem most suitable for the problem at hand. In Chapter 3 we elaborate on the models that are used to generate the input of our rescheduling models. In Chapter 4 we present the two methods that are used for rolling stock rescheduling. In Chapter 5, we perform computational experiments on the models and provide a discussion on the results. Then, in Chapter 6, we summarize these results in a conclusion and answer our research questions. Finally, we give suggestions for future research in Chapter 7.

1 Problem description

In this chapter, we provide a description of the problem to be addressed. First, in Section 1.1, we discuss all terms and concepts that are relevant in this thesis. Secondly, we discuss the rescheduling objectives in Section 1.2. Finally, we give a detailed problem definition in Section 1.3.

1.1 Terms and concepts

The terms and concepts are sorted by topic. Firstly, we discuss the disruptions that are considered in this thesis in detail in Section 1.1.1. Secondly, the timetable and its related concepts are discussed in Section 1.1.2. Thirdly, we look into rolling stock in Section 1.1.3. Fourthly, we take a look at the railway infrastructure that we are considering in Section 1.1.4. Finally, in Section 1.1.5, we discuss passengers and passenger demand. Throughout this section, we introduce some notation. The corresponding sets are given in Table 1.2 for later reference.

1.1.1 Disruptions

A *disruption* is considered a complete blockage of a track between two stations. This means that no trips on this trajectory can proceed. The decision rules that define the new timetable after such a blockage are given. Due to such a blockage, there is a change in the passenger demand at surrounding routes or on trips occurring at that route after the disruption.

1.1.2 Timetable

The train network consists of train lines $l \in L$ on which trains drive from a begin to an end station. The timetable consists of a set of trips T on those lines. A trip $t \in T$ is a journey from a departure station at a departure time to an arrival station at a arrival time. The timetable for each day is given.

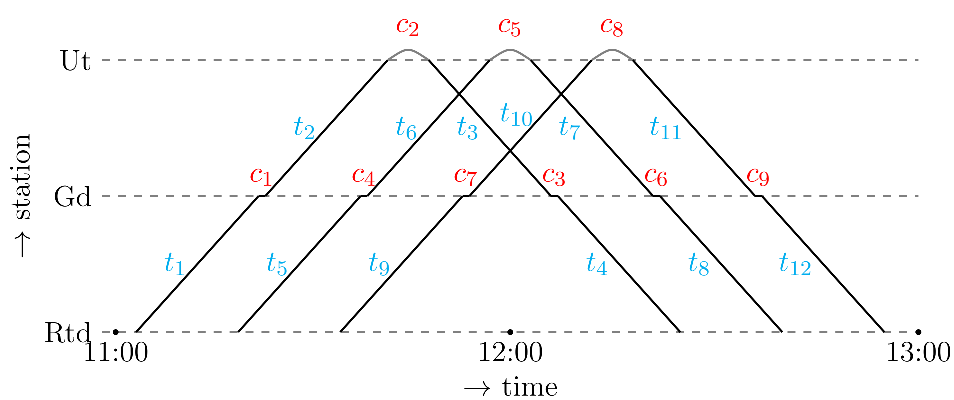

Figure 1.1: Circulation of three trains between Rotterdam (Rtd), Gouda (Gd) and Utrecht (Ut). The trips with trip numbers are given in cyan, the transition with transition numbers are given in red

Trips are connected through transitions, describing which trips are sequentially performed by the same train. Each transition $c \in C$ has a set of incoming trips T_c^- and outgoing trips T_c^+. The set of all transitions in the considered time window is C. An example of a line with trips and transitions is given in Figure 1.1. Usually, the size of set T_c^- and T_c^+ is one, but in rare cases, multiple incoming trips combine at a station and continue their journey together in one trip (or the other way around). In this thesis, however, we limit ourselves to the case where each transition has a maximum of one ingoing and one outgoing trip. In Figure 1.1 we see nine transitions where T_c^- and T_c^+ have size one. For example, $T_{c_1}^- = \{t_1\}$ and $T_{c_1}^+ = \{t_2\}$.

The transitions are given, but in case of *flexible turning*, it is possible to swap the outgoing trips of two transitions. Flexible turning is not considered in the methods in this thesis. Nevertheless, we briefly explain the concept because the term does appear in this thesis. An example of *flexible turning* is given in Figure 1.2. We here see that $T_{c_2}^+$ changes from $\{t_3\}$ to $\{t_7\}$ and $T_{c_5}^+$ changes from $\{t_7\}$ to $\{t_3\}$. The new turning pattern could be beneficial if, for example, t_7 and t_8 have a higher passenger demand than t_3 and t_4 and the composition of t_1 and t_2 has higher passenger capacity than that of t_5 and t_6.

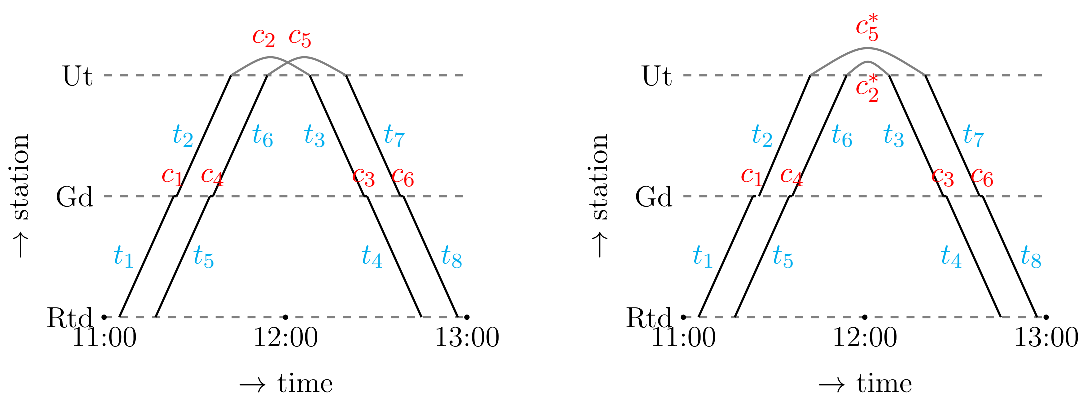

Figure 1.2: Circulation of three trains between Rotterdam and Utrecht. On the left, we see the original turning pattern at terminal station Utrecht. On the right, we see an example of an adjusted turning that is possible with the use of flexible turning

The timetable consists of trips and transitions. This timetable is given and is adjusted to the disrupted situation in case of a disruption. This means that no more trips on a route can take place if a track between two stations on this route is blocked. Moreover, trips that precede or succeed the trip on the blocked track may be canceled. As some transitions are no longer possible due to the cancellation of trips, the transitions are also adjusted. For example, the trip that takes place just before the canceled trip and strands at a station can be transitioned to a trip that departs later on from this same station. How the trips and transitions are adjusted after a disruption is given and is not part of this research.

1.1.3 Rolling stock

Rolling stock is the collective name for vehicles that move on tracks. The rolling stock types considered in this thesis are the Intercity (IC) trains in the Netherlands. Six types of train units are considered in this thesis: namely ICM-3, ICM-4, VIRM-4, VIRM-6, DDZ-4, and DDZ-6. The characters before the dash indicate the general train unit type, the number after the dash indicates the number of train carriages on the unit. The set M consists of all train unit types. Each train unit type has a certain seating capacity, length, and a number of carriages, which are given in Table 1.1.

Table 1.1: Considered InterCity train types, with type name, length and seating capacity

Train type	**Length (m)**	**Seating capacity (seats)**	**Number of carriages**
ICM-3	80.6	228	3
ICM-4	107.1	299	4
VIRM-4	108.6	405	4
VIRM-6	162.1	597	6
DDZ-4	101.8	373	4
DDZ-6	154	607	6

All considered train types are self-propelled and have a driver seat at both ends of the train. Trains units can be combined to form longer trains that have a higher passenger capacity. This sequence of trains is named the train *composition* and all possible train compositions are in set P. The set of train compositions that can be executed on trip t are denoted by P_t. In general, it is possible to combine two train unit types of the same general type (ICM, VIRM, or DDZ) as long as the composition does not exceed the maximum number of train carriages of trip t. The latter is limited by the platform lengths of the stations on the line of trip t.

The order of trains within the composition is relevant. For example, if the train would consist of a train unit of type a and b, we can have train composition ab and ba where in the former b and in the latter a is at the front of the train. This order of train unit types is relevant because only certain *composition changes* are allowed at stations. These possible composition changes at transition c are denoted by the set Q_c, which consists of all possible combinations of incoming compositions and outgoing compositions. At turning stations, stations where trains arrive from the same direction in which they depart, the order of carriages in a train composition reverses. This reversal is included in the possible composition changes. In order to perform composition changes, *(un)coupling* may be required. Coupling is the act of combining two trains into one, whereas uncoupling is splitting one train into two trains. An example of why the order of the train units in the composition is relevant is that at many stations it is only possible to couple a train unit to the front of the train or to uncouple a train unit from the back of the train.

The goal of rolling stock (re)scheduling is to assign rolling stock compositions to all trips in set T, which is called *rolling stock assignment*. This results in a *duty* for each available train unit. A duty is a sequence of trips and the position of the unit on each trip for a train unit. The set of all duties is set D.

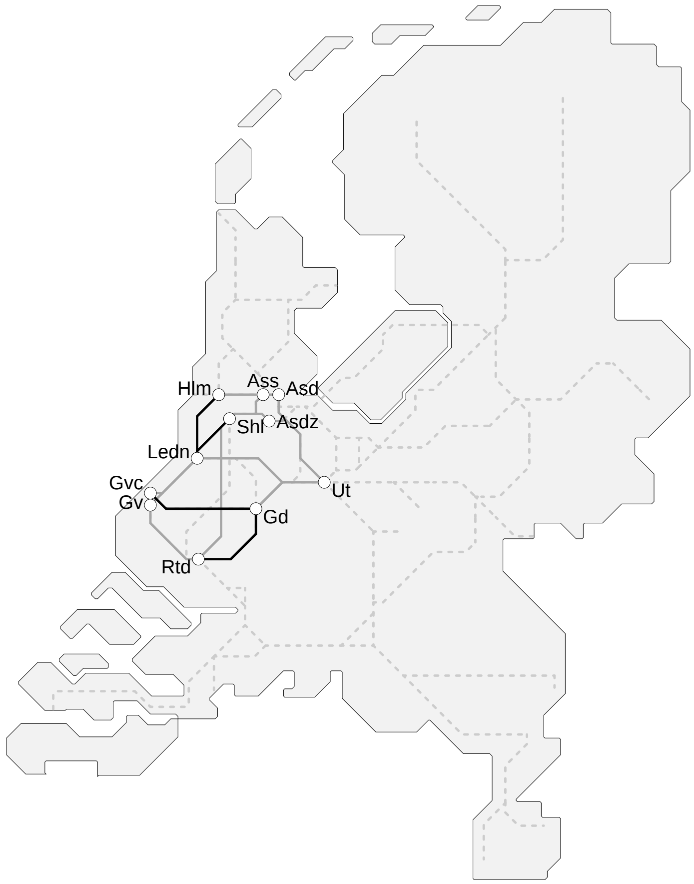

Figure 1.3: Network and subnetwork for change of passenger flows. Grey and black solid lines: network on which we study the change of passenger flows. Black solid lines: routes on which we generate disruptions. Dotted lines: NS network as a whole

1.1.4 Railway infrastructure

The model for dynamic passenger flows that is developed in this research is applied to a subnetwork of Netherlands Railways (NS). NS is the biggest railway operator in the Netherlands that operates on the tracks of the infrastructure manager ProRail. The network that we consider is presented in Figure 1.3. The gray and black solid lines represent the network on which we study the change of passenger flows. The black lines indicate the routes on which we generate disruptions. This subnetwork is chosen for two reasons. First of all, it is the part of the network with the highest passenger demand. Secondly, the network contains detour routes in case a disruption occurs. This makes the subnetwork interesting for this study as it allows us to respond to passenger demand changes on these detour routes. The dotted lines represent the network of NS as a whole. For rolling stock rescheduling, we consider the whole network of NS on which InterCity trains operate. This includes the dotted lines in Figure 1.3.

We make a distinction between *depot stations* and *passenger stations*. At depot stations, an *inventory* is kept and the trains can change their *composition*. At passenger stations, passengers can enter or leave the train but composition changes can not occur. The set S contains all stations in our network. In the original rolling stock schedule, each station has a given inventory at the begin and end of the planning period. If a disruption occurs and rolling stock rescheduling is performed, the actual end inventory could deviate from the planned end inventory. If this is the case, we speak of an *off-balance*.

1.1.5 Passengers

In a train network, we encounter passenger demand. This passenger demand can be expressed in two ways: first, we can consider trip demand, which is the expected amount of passengers on each trip t. Note that each trip t is unique, so no two trips exist with the same departure and arrival stations and/or times. Second, we encounter origin-destination demand. Origin-destination demand is given as the expected amount of passengers that want to travel from station $a \in S$ to station $b \in S$ for each combination of stations in time interval $n \in N$.

Table 1.2: Used sets in the problem

Sets	**Description**
L	train lines
T	trips
C	transitions
D	duties
M	train unit types
P	all feasible train compositions
S	stations
N	time intervals of the origin destination demand
P_t	feasible train compositions for trip t
T_c^-	incoming trips in transition c
T_c^+	outgoing trips in transition c
Q_c	allowed composition changes at transition c

1.2 Rescheduling objectives

In this thesis, a model for rolling stock rescheduling is created. In this model, several objectives are taken into account, namely:

- Computation time: Rescheduling must happen fast, as this allows the train operator to quickly communicate and execute the decisions. Especially in complicated railway organizations, this computation time is essential because of the collaboration between different railway actors (see Section 2.1 on disruption management in railway operations).
- Cancellations: The cancellation of trains in the rolling stock rescheduling phase is undesirable, as it increases the passengers' travel time and affects crew rescheduling.
- Stick-to-the-plan: It is desirable that the adjusted schedule is similar to the original plan, as both passengers, crew, and controllers are familiar with that plan and may prefer the rolling stock to be executed as originally intended. This similarity is measured by keeping track of the amount of trips that is performed by a different composition than planned.

- Operational objectives: From an operational point of view, it is beneficial to minimize the following operational objectives in rolling stock rescheduling:
  - New shunting movements: If the new schedule requires additional coupling or uncoupling of vehicles at a transition, this means that the local planner needs to plan new shunting movements which might be difficult or impossible to plan.
  - Off-balances: We want to have as few off-balances as possible, as having ample reserve vehicles at each station allows us to absorb disturbances and irregularities and because of planned services (maintenance or washing).
  - Additional required conductors: After rolling stock rescheduling, crew rescheduling is performed. It is undesirable that a train requires additional conductors, as this makes crew rescheduling more difficult.
  - Different vehicle type: A different vehicle type might make crew rescheduling more complicated, as not all train drivers are allowed to drive all types of vehicles.
  - Carriage kilometers: Driving a train costs money and therefore a train operator wants to minimize the number of kilometers that are driven by trains
- Passenger standing time: It is desirable that each passenger that enters a train is able to find a seat. However, this is not always the case and passengers need to stand while on the train. The total amount of time that passengers are standing on the train must be minimized. To determine the passenger standing time, we must make use of the adjusted passenger demand due to a disruption or irregularity. This objective is also referred to as seating shortage.

A trade-off exists between the objectives. For example, driving each trip with a composition with high passenger capacity will of course have a significant positive impact on the passenger standing time. This is, however, not desirable as it will greatly increase the operational costs. A comparable example can be made up for stick-to-the-plan and passenger standing time in case the passenger demand on a trip changes. It is, therefore, interesting to study the trade-off between the objectives.

1.3 Problem definition

The goal of this research is to develop a model for rolling stock rescheduling after a disruption while taking into account the dynamic passenger flows. A disruption is a complete blockage of a track. The passenger flows are considered dynamic, as the model adapts the passenger demand forecast to the new situation. In other words, if the timetable changes due to a track blockage, what is the best way to reassign train units to trips?

The rolling stock rescheduling model minimizes the objectives in Section 1.2. The formal definition of the rolling stock rescheduling problem (RSRP) is to assign a composition $p \in P$ to each trip $t \in T$ and thus generate a set of duties D. It must hold that the composition on each trip $t \in T$ is in P_t and that the corresponding composition changes on each transition $c \in C$ is in Q_c.

The inputs of the model are the duration and location of the disruption, the timetable changes at disruption, the corresponding sets as in Table 1.2, and information on the original timetable, network layout, passenger demand, rolling stock, stations, and original rolling stock schedule.

The output of the model is a new rolling stock schedule. This schedule contains both the composition for each trip as well as the composition change at each transition. Also, we can access the new passenger demand forecast and investigate the value of all objectives of Section 1.2.

2 | Literature review

This chapter discusses literature on topics that are related to this thesis. General papers on railway (re)scheduling are those of Caprara et al. (2007) and Cacchiani et al. (2014). Caprara et al. (2007) elaborates on passenger railway optimization in general. Cacchiani et al. (2014) gives a general review on railway systems and a review on literature and models on timetable, rolling stock, and crew rescheduling.

In this thesis, we perform a literature review on disruption management (Section 2.1), passenger flows (Section 2.2), and rolling stock rescheduling (Section 2.3), as these topics form the core of this research.

2.1 Disruption management

Disruption management is defined as follows by Jespersen-Groth et al. (2009): *"The joint approach of the involved organizations to deal with the impact of disruptions in order to ensure the best possible service for the passengers."* They mention that the infrastructure manager has the responsibility of using the railway network as efficiently as possible. The train operator, however, mainly focuses on customer service. These conflicting interests may be counterproductive in the decision-making process after a disruption. The decisions made by the operator and the infrastructure manager must be mutually approved, so decision time is essential. The involved organizations, in our case, are ProRail (infrastructure manager) and NS (train operator).

Moreover, Jespersen-Groth et al. (2009) state that the disruption management process sequentially solves timetable, rolling stock, and crew adjustment. Dollevoet et al. (2017) applies these three steps as an iterative framework. When it is not possible to find either a feasible rolling stock or crew schedule in the second and third step, they re-adjust the timetable heuristically such that they can reschedule successfully. They show that such a framework for real-time railway rescheduling can find an overall feasible solution after a disruption. This research relies on the assumption that the duration of the disruption is known at the beginning of the rescheduling procedure.

In order to handle cases where the duration of the disruption is not yet known at the beginning, Nielsen et al. (2012) developed a rolling horizon approach. At the time of the disruption, the approach only considers rolling stock decisions within a certain amount of time (horizon) after the disruption. When more information on the situation gradually becomes available or a certain amount of time has passed, the decision horizon is shifted and rolling stock rescheduling is again performed for the new time horizon. The rolling horizon approach is proven effective with the Composition model by Nielsen et al. (2012), but can in principle be applied using any rolling stock rescheduling method. Moreover, Nielsen et al. (2012) show a trade-off between solution quality and computation time and minor effects on the shunting plans. The operational costs significantly increase in comparison to the situation where the disruption duration is known at the beginning, but Nielsen et al. (2012) state that *"the values themselves are quite appealing in practice"*.

2.2 Passenger flow

In this section, we discuss papers that are related to passenger flow. It is a logical consequence that the passenger flow on routes that surround the blocked route change because (a part of the) passengers still want(s) to reach their final destination. Therefore, it seems sensible to consider passenger flows no longer as static but as dynamic input. We first discuss papers on passenger flows in the railway context. Secondly, we look into papers on passenger flows in other industries.

Kroon et al. (2015) developed an iterative heuristic for dynamic passenger flows. Using a feedback loop, they iteratively determine the passenger flows and the rescheduled rolling stock schedule, as both affect each other. Passengers are divided into passenger groups, where a group contains a number of passengers (n_p) that want to travel at a certain time from an origin station to a destination station. Passengers are assumed to have a total travel time deadline, to take the travel itinerary with the earliest arrival, to be fully informed of the timetable, and to not enter a train when it is full (meaning that there exists an interaction between groups when trains near capacity). For each passenger group, there must exist a flow with value n_p from their origin to destination station in a graph where each node represents an arrival or departure of a trip and where each arc connects an arrival with the corresponding departure. The new passenger flows, which are influenced by the former rolling stock schedule, are then used to construct a new rolling stock schedule, and so forth.

Cadarso et al. (2013) also considers the flow of passengers by defining passenger groups. Passengers in group g with size n_g can travel from their origin to destination through paths P_g, where the probability of group g choosing path p is $P(p|w)$. $P(p|w)$ is determined using a multinomial logit model and is based on the traveling, transfer, and waiting time. It is assumed that passengers within a passenger group can choose different paths, passengers can always enter a train, passengers have a certain deadline after which they leave the system, and that there exists no dynamic interaction between passenger demand and rolling stock capacity. Using $P(p|w)$, n_g and the origin and destination of group g, the number of passengers on each trip is determined. Cadarso et al. (2013) show that a single iteration, in which the passengers' path only depends on the timetable, already predicts passengers' behavior well. A second iteration, which takes the rescheduled rolling stock assignment into account, barely changes the passenger flow.

The final paper on railway passenger flows that we consider, is that of Wagenaar et al. (2017b). They use a different approach in which they incorporate the change in passenger flows in the rolling stock rescheduling phase. This relies on the assumptions that passengers can not leave the railway system, the passenger demand does not change, and passengers do not take a detour when a disruption occurs. Their paper does take into account that after a canceled trip, the demand for the next train will increase. These assumptions do not correspond to the situation we are looking at and therefore this model is not suitable for this study.

For metro passenger flow prediction, there exist high-level prediction tools for the number of passengers, namely those of Wei & Chen (2012) and Liu et al. (2019). Wei & Chen (2012) use neural networks and Liu et al. (2019) use deep learning to predict passenger flows in the short term. A difference between metro and railway passenger flows is the big influence of other transport modes in a metro network. Although this influence also exists in our network, this falls beyond the scope of our research. Also, an initial, high-level prediction of passenger demand is given as an input for our research, which makes the above methods superfluous.

In the aforementioned papers, the exact passenger demand is not known and passengers are free to use the itinerary of their choice on the day of travel. However, other industries, such as the airline industry, operate with a booking system and are thus able to control their passenger flows. In Dumas & Soumis (2008), a model is introduced to account for dynamic passenger flows in the airline industry during the booking period. This model is again based on a graph with origins and destinations as nodes and travel itineraries as arcs that must supply demand. The model assumes that when a flight is full, the passenger behavior (either not to book a flight or to book a certain other flight) is known. Dumas et al. (2009) iteratively uses this model to re-optimize their fleet assignment. This iterative model formed the basis for the model of Kroon et al. (2015).

2.3 Rolling stock rescheduling

For a long time, the Rolling Stock Rescheduling Problem (RSRP) was not often addressed in literature, as opposed to the Rolling Stock Scheduling Problem (RSSP). In recent years more attention has been paid to the RSRP and we discuss these papers in this section.

The majority of methods from literature can be divided into two categories: flow-based and path-based methods. In general, a flow-based method determines the flow of rolling stock units through consecutive trips, where the arcs represent the trips and the nodes represent an arrival, departure, or pass-through at a certain time at a station. In a path-based method, it is determined which paths are executed by rolling stock units, where a path is a feasible sequence of trips that can be covered by a certain unit. We can logically observe that the amount of paths in such a path-based model grows exponentially in the number of trips.

The first research focused on the RSRP is that of Nielsen (2011). He extends the flow-based Composition model of Fioole et al. (2006), which is a model that gives the assignment of compositions to trips by solving a MIP with commercial optimization software. He shows that this model can be used not only for scheduling but also for rescheduling rolling stock and concludes that it is the best model to solve the RSRP. Moreover, he uses the Duty Path Model to transform the composition assignment to the assignment of duties to rolling stock as real-life operation requires a duty for each rolling stock unit. The combination of the extended Composition model and the Duty Path Model is capable of finding a good solution in a sufficient amount of time and is therefore often used as a well-performing benchmark, for example in Haahr et al. (2016).

The Composition model is often extended in order to add more factors that can be of value in practice, for example in Wagenaar et al. (2017a) and Wagenaar et al. (2017b). Wagenaar et al. (2017a) show that maintenance constraints can be added to the Composition model. Wagenaar et al. (2017b) perform rolling stock rescheduling with dead-heading trips and adjusted passenger demand. Although they mention that dead-heading trips reduce the number of canceled trips, we will not add dead-heading trips to our model because this is uncommon and not preferred at NS. What is also interesting about this paper is that the adjustment of passenger flows is incorporated in the MIP of the Composition model. This extension does, however, rely on assumptions that do not fit our problem setting (see Section 2.2).

Next, we discuss path-based methods in literature. Both Lusby et al. (2017) and Haahr et al. (2014) deal with the large number of path variables in their strongly comparable MIP by using a Branch-and-Price algorithm. The Branch-and-Price algorithm solves the subproblems using a Shortest Path Problem in an acyclic time-space network that consists of trips (arcs) and events (nodes), such as arrivals, departures, and pass-throughs.

The main advantage of such a path-based approach is that units can be tracked individually, so it is easier to extend the model to handle constraints on individual units, for example for maintenance. However, the two mentioned path-based methods rely on assumptions that are not realistic in practice, such as that the composition within the train and the feasibility of the composition changes are not taken into account. The methods of Lusby et al. (2017) and Haahr et al. (2014) can quickly find a feasible solution, but require a longer computation time than the Composition model to reach optimality.

In the literature, several papers apply and compare a flow- to a path-based method. Haahr et al. (2016) compare the Composition model of Nielsen (2011) with an extension of Haahr et al. (2014) in which they take the composition of the trains into account, which is solved using a column generation approach. The Composition model is again more attractive computationally but has the disadvantage that it is not possible to track individual units.

Hoogervorst et al. (2020) also performs this comparison on rolling stock rescheduling in case of delays. They use an adapted version of the Composition model and a path-based model and again show that the flow-based Composition model computationally performs better.

As mentioned earlier, Wagenaar et al. (2017a) showed that the flow-based Composition model can also be extended to be able to handle maintenance constraints, although it is less evident than in path-based models. Although handling maintenance constraints falls beyond the scope of this research, this does make the path-based models of Lusby et al. (2017) and Haahr et al. (2014) less attractive to use.

The aforementioned methods required the use of a commercial solver. Hoogervorst et al. (2021) however, uses a different approach than previous papers. This new approach is neither flow- nor path-based and does not require a commercial solver. They perform rescheduling with Variable Neighborhood Search (VNS) using three neighborhoods. The first neighborhood, the Two-Opt Duty Neighborhood, swaps (a part of) a duty between two train units. Secondly, the Adjusted Path Neighborhood improves the assignment for one rolling stock type, using an augmenting path approach. The third neighborhood, the Composition Change Neighborhood, swaps the composition change between n transitions. The first two neighborhoods are used in the Local Search (LS) method by applying Variable Neighborhood Descent (VND). The final two neighborhoods are used for shaking the rolling stock schedule. Hoogervorst et al. (2021) mention that the Adjusted Path Neighborhood, as well as the Two-Opt Duty Neighborhood, only changes a few duties at a time. Also, an additional fourth neighborhood is introduced in both the LS and in shaking. This neighborhood allows for flexible turnings, meaning that the assignment of arrivals to departures at a terminal station is changed.

The heuristic of Hoogervorst et al. (2021) is computationally less attractive than the Composition model when flexible turning is not considered in both models. However, the heuristic can outperform the exact method in some cases when flexible turning is included. Moreover, the heuristic has the advantage that intermediate results are usable and has the advantage that it provides more flexibility to extend the model with real-life practicalities, which makes it well applicable in practice. Finally, they mention that the heuristic can be improved with the use of additional neighborhoods.

2.4 Conclusion

In this research, we develop a method for rolling stock rescheduling while taking passenger flow into account.

We focus on developing a rolling stock rescheduling model because an iterative approach in which rolling stock rescheduling is seen as a separate problem had proven to give feasible solutions. Moreover, we will not implement a rolling horizon framework, as this rolling horizon approach can be applied using any rolling stock rescheduling model and unknown disruption duration falls beyond the scope of this thesis.

To account for dynamic passenger flows in case of a disruption or irregularity, we will develop a graph-based model, such as the model of Kroon et al. (2015). This approach is the best fit as it allows us to include many possible decisions and constraints of passengers and translate origin-destination demand to trip demand. Moreover, we apply our dynamic passenger flow model before rolling stock rescheduling only, as Cadarso et al. (2013) shows that the rolling stock assignment barely affects passenger flow.

This thesis explores the possibility of combining the heuristic of Hoogervorst et al. (2021) with the Composition model, by applying the Composition model as a neighborhood in the heuristic. This Composition model Neighborhood has the same main idea as the Adjusted Path Neighborhood, namely to optimize the deployment of one vehicle type at a time. The advantage of this new neighborhood is that, by applying the Composition model to one vehicle type at a time, it optimizes the entire deployment of this vehicle type and not just finds one path change. By combining the heuristic with the Composition model, we hope to maintain the advantages of the Composition model, namely its computational attractiveness under limited circumstances, and the practical advantages of the heuristic, namely its flexibility and intermediate results. Also, the field of applying Local Search methods to the RSRP is relatively unexplored and thus the introduction of additional neighborhoods can do much to improve the field's knowledge on and the quality of methods for the RSRP.

3 Input Generation

As stated in the problem definition, the goal of this research is to develop a model for rolling stock rescheduling after a disruption while taking into account dynamic passenger flows. We can think of this as if the disruption, the corresponding adjusted timetable, and the new passenger flow are, along with the original rolling stock schedule, the input to our rolling stock algorithm. This chapter discusses all input and generated input that is used for rolling stock rescheduling. The information flow is presented in Figure 3.1.

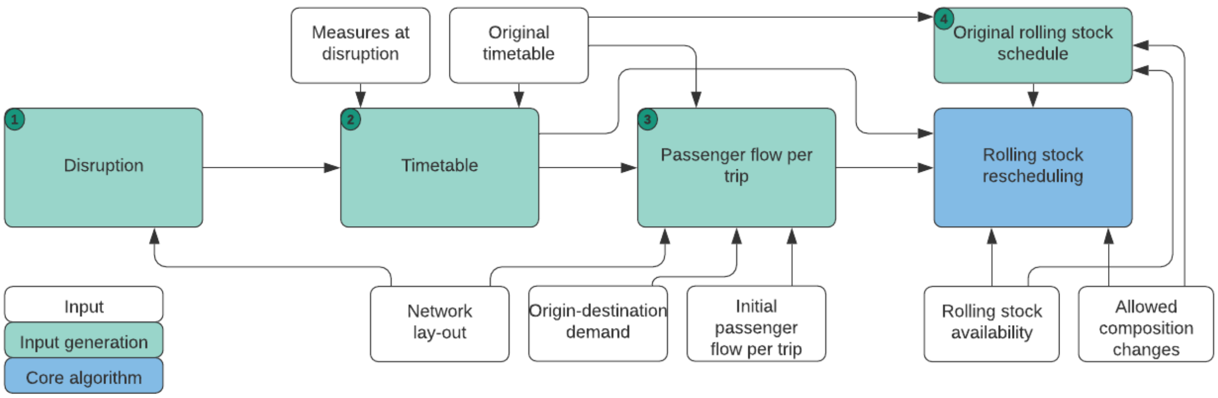

Figure 3.1: Input and input generation overview. In white we see all given input. The numbered green boxes represent input that is generated ourselves. In blue we see the core algorithm: rolling stock rescheduling

The white boxes in Figure 3.1 are given input. We provide a brief overview of what this input entails:

- Original timetable: The original timetable contains the set of trips T, transitions C, and corresponding parameters on the arrival and departure station and time.
- Measures at disruption: For the considered disruptions, it is given which trips $t \in T$ do not proceed and how the set of transitions C changes.
- Initial passenger flow per trip: This data set contains the expected number of passengers on each trip, generated by a department of NS that is specialized in demand forecasting.
- Origin-destination demand: For each pair of considered stations, the estimated origin-destination demand is given.
- Network lay-out: This is the network layout as given in Figure 1.3. Furthermore, it provides us with information on which stations (un)coupling of vehicle units is possible.
- Rolling stock availability: This concerns the amount of rolling stock unit types that are available in the network.

- Allowed composition changes: In this thesis, it is assumed that it is allowed to either couple a vehicle unit at the front or uncouple a vehicle unit at the back of an existing composition.

We now discuss disruption generation and timetable adjustment, green box (1) and (2) in Figure 3.1. The severity, location, duration, and time of the disruption must be determined. First of all, as discussed in Section 1.1.1, we only consider a complete blockage of a track. We consider such a blockage on one route at a time. Furthermore, we consider disruptions with a duration of 2 hours because this is a realistic duration[1] and this duration allows the system to develop a steady changed passenger flow over the course of the disruption. What we mean by a steady changed passenger flow is that the duration is long enough for the passengers to choose alternative routes and not to wait for the first train after the disruption. This causes a change in the passenger flow on routes that surround the blocked track and allows us to study the impact of this change in passenger flow on rolling stock rescheduling. Moreover, the simulated disruption occur during rush hour, as this is when most passengers are present in the network. In this report, we limit ourselves to this kind of disruption to answer our research questions. However, the rolling stock rescheduling model can be applied to any disrupted timetable.

We now discuss timetable adjustment, which is a straightforward procedure as the measures at a disruption are given. Using these measures, we can adjust the trips $t \in T$ and transitions $c \in C$ and use this in the subsequent steps of our Input Generation procedure.

The following sections discuss the other input generation steps. First, we discuss how the adapted passenger demand per trip is generated in Section 3.1 using a graph-based model. Second, we discuss how and why we create an initial rolling stock schedule using a Mixed Integer Program (MIP) in Section 3.2.

3.1 Passenger demand after a disruption

This section discusses in detail how we determine the new passenger demand per trip after a disruption. We determine this new demand by applying a graph-based model which was inspired by Kroon et al. (2015). Inputs for this model are the new and old timetable, initial passenger flow forecast per trip, origin-destination demand, and the network layout. The output is a new passenger flow forecast per trip.

The general concept of the passenger flow model is as follows: first, we determine the amount of passengers on each trip before the disruption. Then, we remove the trips from the model that are canceled due to the disruptions and again determine the flow on each trip. This allows us to calculate the percentage increase or decrease in passenger flow on each trip. The given initial passenger flow forecast is now adjusted using this change in passenger flow. This initial forecast is an elaborate forecast that takes into account more factors than taken into account in the passenger flow model of this thesis. Therefore, applying the changes to the initial forecast results in a more accurate forecast of the passenger demand than the passenger flows that result directly from the passenger flow model.

[1] www.rijdendetreinen.nl/statistieken

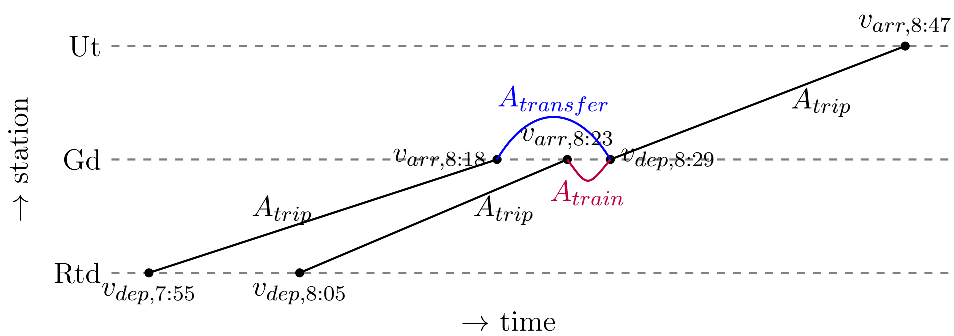

Figure 3.2: Timetable graph for three trips between Rotterdam, Gouda and Utrecht. The train that departs at 8:05AM at Rotterdam continues its journey to Utrecht at 8:29AM. We see the trip arcs in black, transfer arcs in blue and train arcs in purple. Furthermore, each node is labeled with either departure or arrival and the corresponding time

Timetable graph

We now elaborate on how we determine the passenger flow in the system both after and before the disruption. This is done by determining the flow over the *timetable graph* $G_{timetable} = (V, A)$ that represents possible sequences of trips (journeys) in the timetable. The nodes $V = V_{arr} \cup V_{dep}$ represent either the arrival (V_{arr}) or departure (V_{dep}) of a trip. Each node $v \in V$ has an arrival or departure station s_v, arrival or departure time t_v^V and a train number n_t that remains the same from the begin station to the terminal station of the line.

Furthermore, the arcs $A = A_{trip} \cup A_{train} \cup A_{transfer}$ connect the nodes in V by either representing a trip (A_{trip}) or a possible connection between two trips (A_{train} or $A_{transfer}$) by either staying in the same train or by transferring to another train. In Figure 3.2 we see an example of a passenger graph that consists of three trips between Rotterdam, Gouda, and Utrecht.

The arcs A_{trip} connect the departure node of a trip to the arrival node of that same trip. A_{trip} is formally defined as the arcs

$$A_{trip} = \{\big((s_{v_1}, t_{v_1}^V), (s_{v_2}, t_{v_2}^V)\big) \quad \forall \quad (\ (v_1, v_2) \in (V_{dep} \times V_{arr}) \mid v_1 \text{ and } v_2 \text{ belong to the same trip }) \ \}$$

The length of the arc $t_a^A \quad \forall \quad a \in A_{trip}$ represents the duration of the trip. For example, if there exists a trip that departs from Rotterdam at 9:05AM and arrives in Gouda at 9:23AM, we encounter an arc from the departure node in Rotterdam at 9:05AM to the arrival node in Gouda at 9:23AM with a length of 18 minutes.

The arcs A_{train} connect the arrival of a trip to the departure of a trip that corresponds to the same train. A_{train} is formally defined as:

$$A_{train} = \{\big((s_{v_1}, t_{v_1}^V), (s_{v_2}, t_{v_2}^V)\big) \quad \forall \quad (\ (v_1, v_2) \in (V_{arr} \times V_{dep}) \mid (n_{v_1} = n_{v_2} \ \& \ s_{v_1} = s_{v_2})\}$$

Furthermore, the transitions $A_{transfer}$ connect all possible transitions between arrival nodes and departure nodes that are not on the same train. A transition is possible if the difference between the arrival and departure time is larger than the *transfer time* t^{tf}. Arc $A_{transfer}$ is formally defined as:

$$A_{transfer} = \{((s_{v_1}, t^V_{v_1}), (s_{v_2}, t^V_{v_2})) \quad \forall \quad ((v_1, v_2) \in (V_{arr} \times V_{dep}) \mid (s_{v_1} = s_{v_2} \ \& \ t^V_{v_1} + t^{\text{tf}} \leq t^V_{v_2} \ \& \ n_{v_1} \neq n_{v_2}))\}$$

The length of the arcs $a \in (A_{transfer} \cup A_{train})$ correspond to the duration of the transfer from one trip to the next. Considering the previous example, we would add an arc from the arrival node in Gouda at 9:23AM to each departure node that we can transfer to. There exist arcs to both the departure node of the same train ($a \in A_{train}$) and to departure nodes of other train lines ($a \in A_{transfer}$).

Passenger flow

The origin-destination passenger demand is added to the graph by adding passenger demand to the graph under Assumption 1.

Assumption 1 *Passengers arrive at their origin station at uniformly distributed times*

The uniformly arriving passengers are aggregated into passenger groups g with a given origin station o_g, destination station d_g and arrival time at the origin station t^G_g. G is the set of all passenger groups. The arrival times of the passenger groups correspond to the arrival of a train at the origin station. For example, if at a station s a train departs every 10 minutes, all uniformly arriving passengers at that origin station with the same destination station in the time interval of 10 minutes are aggregated in one passenger group which has a total size of n_g. The total passenger group size is determined by summing over the origin-destination demand of the time intervals that fall within the disruption time interval. This total passenger group size is then added as demand in the graph under Assumption 2.

Assumption 2 *The passenger demand is constant over the considered time window*

The demand on each trip corresponds to the flow on each trip arc after adding the origin-destination passenger demand of each passenger group to the graph by determining the shortest path for each passenger group through the graph. This is done by adding a source node v_g for each passenger group that is connected to the timetable graph $G_{timetable}$ by adding time arcs if it is possible to board the departing train after the arrival of the passenger group. Meaning that we add an arc $\forall g \in G$ from v_g to $v_{trip} \in V_{dep}$ if $s_{v_g} = s_{v_{trip}}$ and $t^G_g \leq t^V_{v_{trip}}$. Moreover, we also restrict that $t^V_{v_{trip}} - t^G_g \leq 60$ *minutes* for two reasons: first, each train departs on each line at least once an hour so a passenger will never wait for a train for longer than an hour. Second, this ensures that the graph does not become unnecessarily complex by not adding redundant arcs.

We also add a dummy node $v_{dummy,s}$ for each station $s \in S$ that we consider in the passenger flow graph. We connect each $v_{trip} \in V_{arr}$ to $v_{dummy,s}$ if $s = s_{v_{trip}}$. The dummy station nodes are the sink nodes that correspond to the destination of each passenger group. Finally, the shortest path of each passenger group is determined. The size of this group is added to the flow of the trips that are on the shortest path. This implies Assumption 3.

Assumption 3 *Passengers can always travel by the trips that occur in the determined shortest path*

In practice, Assumption 3 implies that we do not consider the capacity of rolling stock that is, or will be, assigned to a trip.

Flow after disruption

We now discuss how we adapt the passenger graph to the situation after a disruption. This is done by recalculating the passenger flow on the altered timetable graph $G_{timetable}$ by removing the arcs and nodes from the original graph that correspond to canceled trips. We know in advance exactly which arcs and nodes must be removed, because we rely on Assumption 4:

Assumption 4 *The duration of the disruption is known in advance and is known by the passengers*

In order to not redetermine the shortest paths for all passenger groups (including the passengers that are not affected by the disruption), we only redetermine the travel journey for passenger groups that traveled by canceled trips in their original journey. We can use this approach as the measures at disruptions only cancel trips and do not create new trips. This implies that we only remove nodes and arcs from the graph and do not create new nodes and arcs. It is trivial that by solely removing arcs, the shortest path of an unaffected passenger group does not change.

Concerning the travel path of the affected passengers, we fix the trips in the travel journey of passengers that depart before the start of the disruption. For example, starting at 8AM we simulate a disruption between Rotterdam and Gouda and consider a passenger group that enters the system at 7:45AM and wants to travel from Utrecht to Rotterdam. This passenger group, in the original passenger graph, first travels from Utrecht to Gouda from 7:48AM until 8:06AM and then travels from Gouda to Rotterdam at 8:08AM until 8:25AM. In the disrupted scenario, the second train is disrupted and therefore, we redetermine the shortest path for this passenger group from Gouda to Rotterdam at 8:06AM.

After we redetermine the shortest path for each affected passenger group, we can recalculate the passenger demand on the trips and hence determine the expected increase and decrease in passenger demand on each trip. As mentioned earlier, we apply these increases or decreases on the (high-level) passenger forecast per trip.

Passenger behavior

In essence, we determine the shortest path through the timetable graph for each passenger group by applying Dijkstra's algorithm both before and after disruption. We make the following assumptions on passenger behavior to determine the routes that passengers are likely to take:

- A passenger prefers to have as few transfers as possible on their journey. Therefore, it prefers a journey with the same length of time with fewer transfers over a journey with more transfers. This is taken into account by increasing the length of the time arcs $A_{transfer}$ with a *transfer penalty*.
- If a passenger group can choose from multiple equally long journeys (including the transfer penalty), the passengers are distributed over the journeys uniformly.
- If a passenger group can choose a second journey which is α% longer, β% of the passengers chooses the second shortest route. α is referred to as the *alternative path tolerance* and β is referred to as the *alternative path fraction*.
- Passengers prefer to travel in an Intercity train over a Sprinter train. Therefore, we add a *Sprinter penalty* to the arcs that correspond to trips driven by Sprinters.
- If the journey takes γ minutes longer after disruption, δ% of the passengers would leave the system. We name γ our *delay tolerance* and name δ our *delay leave fraction*. We also state that passengers that were already on their way, as explained in the previous example, do not leave the system after a disruption.

3.2 Initial rolling stock schedule

The Rolling Stock Rescheduling Problem (RSRP) requires knowledge of the original rolling stock problem, as stick-to-the plan is one of the objectives in our RSRP. Therefore, it is necessary to determine an initial rolling stock schedule. This schedule is determined using the Composition Model developed by Fioole et al. (2006). The notation is based on the adjusted Composition Model of Nielsen (2011) as we use this adjusted Composition Model for rescheduling in Section 4.1.

There are two main reasons that we do not use a rolling stock schedule from reality. First of all, the real-life schedule does not obey the assumptions that we restrict the rolling stock (re)schedule to in this thesis, which means that this real-life schedule is not directly usable. For example, we assume that we can only couple one train unit at the front or uncouple one train unit at the back, but in reality, more composition changes might be allowed. Furthermore, the real-life schedule might not be optimal. If we then reschedule rolling stock to optimality using this real-life schedule, we cannot objectively assess the quality of our rescheduling model.

We first discuss the sets of the Composition Model, followed by the variables, parameters, and the Mixed Integer Problem (MIP). We use the previously discussed set of trips T, set of transitions C, set of vehicle unit types M, and set of stations S. Furthermore, we introduce the set $\mu(t)$, containing the allowed compositions on trip $t \in T$, which is, for example, constrained by the platform lengths. Furthermore, we have a set of allowed composition changes Q and the set $\rho(c) \in Q$, which contains the feasible composition changes for transition $c \in C$. This set applies the restriction that at stations where it is allowed to change compositions, we can decide to couple a vehicle unit at the front and decouple a vehicle from the back of a composition.

We now discuss the used parameters of the Composition Model. First of all, we have $\sigma_t \in C$ and $\pi_t \in C$ which denote, respectively, the successor and predecessor transitions of trip $t \in T$.

Moreover, $p_{q,t}$ is the composition of the ingoing trip t and $p'_{q,t}$ is the composition of the outgoing trip t of composition change q. Furthermore, we use the parameters $\alpha_{q,m}$ and $\beta_{q,m}$ which, respectively represent how many units of type $m \in M$ are uncoupled and coupled at composition change $q \in Q$. Also, we denote as τ_c^- and τ_c^+ as the time at which a vehicle must be available to perform respectively coupling or uncoupling at transition c. Finally, km_t is the amount of kilometers of trip t, dur_t is the duration of trip t, d_t is the demand on trip t, l_t is the maximum number of carriages on trip t, $carr_p$ is the amount of carriages in composition p and cap_p is the total seating capacity of composition p. Finally, i_m is the total number of vehicle units of type m that is available in the system.

Table 3.1: Decision variables of Composition Model

Variable	**Domain**	**Description**
$X_{t,p}$	$\mathbb{B}$	equals 1 if trip $t \in T$ is performed with composition $p \in P$ and 0 otherwise
$Z_{c,q}$	$\mathbb{B}$	equals 1 if transition $c \in C$ is performed with composition change $q \in \rho(c)$ and 0 otherwise
$I_{c,m}$	$\mathbb{Z}^+$	the number of units of type $m \in M$ in inventory at station $s(c)$ after transition c
$I^0_{s,m}$	$\mathbb{Z}^+$	the number of available units of type $m \in M$ at station $s \in S$ at the beginning of the day
$I^{end}_{s,m}$	$\mathbb{Z}^+$	the number of available units of type $m \in M$ at station $s \in S$ at the end of the day
$C_{c,m}$	$\mathbb{Z}^+$	the number of coupled vehicle units of type $m \in M$ at transition c
$U_{c,m}$	$\mathbb{Z}^+$	the number of uncoupled vehicle units of type $m \in M$ at transition c
$I^{\text{diff}}_{s,m}$	$\mathbb{Z}^+$	the absolute difference between $I^0_{s,m}$ and $I^{end}_{s,m}$ at station $s \in S$ for type $m \in M$

The variables of the model are given in Table 3.1. The objective of the MIP is given in Equation (3.1) with weights w_1 until w_5. The constraints are given in Equation (3.2) - (3.16).

$$\begin{aligned}\mathbf{min} \; & \sum_{t\in T}\sum_{p\in\eta(t)} w_1 \cdot \text{carr}_p \cdot \text{km}_t \cdot X_{t,p} + \\ & \sum_{t\in T}\sum_{p\in\eta(t)} w_2 \cdot \max\left\{d_t - \text{cap}_p, 0\right\} \cdot \text{dur}_t \cdot X_{t,p} + \\ & \sum_{s\in S}\sum_{m\in M} w_3 \cdot I^{\text{diff}}_{s,m} + \\ & \sum_{c\in C}\sum_{m\in M} w_4 \cdot C_{c,m} + \sum_{c\in C}\sum_{m\in M} w_5 \cdot U_{c,m}\end{aligned} \tag{3.1}$$

Subject to

$$\sum_{p\in\eta(t)} X_{t,p} = 1 \qquad \forall t \in T \tag{3.2}$$

$$X_{t,p} \cdot carr_p \leq l_t \qquad \forall t \in T, p \in P \tag{3.3}$$

$$X_{t,p} = \sum_{q\in\rho(\sigma_t):p_{q,t}=p} Z_{\sigma_t,q} \qquad \forall t \in T, p \in \eta(t) \tag{3.4}$$

$$X_{t,p} = \sum_{q \in \rho(\pi_t): p'_{q,t} = p} Z_{\pi_t,q} \qquad \forall t \in T, p \in \eta(t) \tag{3.5}$$

$$C_{c,m} = \sum_{q \in \rho(c)} \beta_{q,m} \cdot Z_{c,q} \qquad \forall c \in C, m \in M \tag{3.6}$$

$$U_{c,m} = \sum_{q \in \rho(c)} \alpha_{q,m} \cdot Z_{c,q} \qquad \forall c \in C, m \in M \tag{3.7}$$

$$I_{c,m} = I^0_{s(c),m} - \sum_{\substack{c' \in \mathcal{C}: s(c') = s(c), \\ \tau^+_{c'} \leq \tau^+_c}} C_{c',m} + \sum_{\substack{c' \in \mathcal{C}: s(c') = s(c), \\ \tau^-_{c'} \leq \tau^+_c}} U_{c',m} \qquad \forall c \in C, m \in M \tag{3.8}$$

$$I^{end}_{s,m} = I^0_{s,m} - \sum_{c \in C: s(c) = s} C_{c,m} + \sum_{c \in C: s(c) = s} U_{c,m} \qquad \forall s \in S, m \in M \tag{3.9}$$

$$I^{\text{diff}}_{s,m} \geq I^{end}_{s,m} - I^0_{s,m} \qquad \forall s \in S, m \in M \tag{3.10}$$

$$I^{\text{diff}}_{s,m} \geq I^0_{s,m} - I^{end}_{s,m} \qquad \forall s \in S, m \in M \tag{3.11}$$

$$\sum_{s \in S} I^0_{s,m} = i_m \qquad \forall m \in M \tag{3.12}$$

$$X_{t,p} \in \{0, 1\} \qquad \forall t \in T, p \in \eta(t) \tag{3.13}$$

$$C_{c,m}, U_{c,m}, I_{c,m} \in \mathbb{R}^+ \qquad \forall c \in C, m \in M \tag{3.14}$$

$$I^0_{s,m}, I^{end}_{s,m}, I^{\text{diff}}_{s,m} \in \mathbb{R}^+ \qquad \forall s \in S, m \in M \tag{3.15}$$

$$Z_{c,q} \in \mathbb{R}^+ \qquad \forall c \in C, q \in \rho(c) \tag{3.16}$$

We now discuss the objective and constraints. The objective consists of six parts: first, we encounter the minimization of the number of carriage kilometers. Secondly, we minimize the under capacity of the number of seats in the train measured over time. Moreover, it is penalized if the final inventory of the shunting yard does not equal the initial inventory, in order to make the schedule feasible in the long run. Finally, we penalize the number of couplings and uncouplings.

We now look into the constraints of the Composition Model. Constraints (3.2) ensure that a composition is assigned to each trip. Note that in the initial schedule each planned trip must be executed. Hence, an empty composition is not part of the set $\mu(t) \forall t \in T$. Constraints (3.3) ensure that a composition can not be longer than the maximum number of carriages on that trip. Constraints (3.4) and (3.5) make sure that the compositions and composition changes match each other. Equations (3.4) do this by stating that when a composition p is chosen for trip t, in the successor transition σ_t a composition change must be chosen that has composition p as incoming composition. Equations (3.5) do the same for the predecessor transition. Note that if the set of incoming trips of a transition c is empty, we limit $\sigma(c)$ such that the incoming composition is empty. For transitions c with an empty set of outgoing trips, we ensure that the outgoing composition is empty.

Constraints (3.6) ensure that the number of vehicle units of type m that is coupled at transition c is in correspondence with the chosen composition change at the transition. Constraints (3.7) do the same for the number of uncoupled units.

Constraints (3.8) determine the number of units of type $m \in M$ in inventory at station $s(c)$ after transition c, by summing over the number of uncoupled and coupled vehicle units that occurred earlier in time than transition c at station $s(c)$. Equations (3.9) determine the final inventory, by adding all the couplings and subtracting all the decouplings of the considered time window to the initial inventory.

Equations (3.10) and (3.11) ensure that $I^{\text{diff}}_{s,m}$ is the absolute difference between the initial and final inventory. Furthermore, (3.12) make sure that the total initial inventory of each type equals the total number of available units of that type. Finally, (3.13) - (3.16) specify the domain of the variables. All variables, except for $X_{t,p}$ can be chosen continuous as they all stem from the discrete, binary variable $X_{t,p}$.

In case we want to limit ourselves to only one vehicle type per line, so allowing only DDZ, VIRM, or ICM vehicles on a line, we add the variables q_t^{trip} and q_l^{line}. The variable q_t^{trip} depicts the type of trip $t \in T$ and q_l^{line} the type of line $l \in L$. q_t^{trip} and q_l^{line} equal 1 if the type is DDZ, 2 if the type is VIRM, and 3 if the type is ICM. Furthermore, we introduce the set T_l which contains all trips of line $l \in L$ and the parameter $type_p$ which denotes the type (DDZ, VIRM or ICM expressed as an integer) of composition $p \in P$.

$$q_t^{trip} = \sum_{p \in P} X_{t,p} \cdot type_p \qquad \forall t \in T \qquad (3.17)$$

$$q_t^{trip} = q_l^{line} \qquad \forall t \in T_l, \forall l \in L \qquad (3.18)$$

$$q_t^{trip} \in \{1, 2, 3\} \qquad \forall t \in T \qquad (3.19)$$

$$q_l^{line} \in \{1, 2, 3\} \qquad \forall l \in L \qquad (3.20)$$

4 Rolling Stock Rescheduling

We apply two methods to perform rolling stock rescheduling. First of all, we use an adjusted version of the Composition Model, which was already explained in Section 3.2. Secondly, we apply a heuristic that consists of a start heuristic, followed by a Local Search method that explores two neighborhoods. This chapter first discusses the adjustments to the Composition Model in Section 4.1. Secondly, we discuss the heuristic in Section 4.2.

4.1 Composition Model

This section explains how the model of Section 3.2 is adjusted to make the model suitable for rescheduling. In essence, the model remains the same. However, the sets of trips T and transitions C are updated by applying the given set of measures and the objective is updated by adding stick-to-the-plan and additional operational objectives. The latter causes the introduction of new parameters, variables, and constraints, which we discuss, in this order, in this section. Finally, we state the new objective.

First, we introduce parameters that are related to the original rolling stock schedule: p_t^0 denotes the original composition of trip $t \in T$, oh_t denotes the original number of planned conductors for trip $t \in T$, oC_c (oU_c) denotes the number of coupled (uncoupled) units at transition $c \in C$ in the original schedule, and $i_{s,m}^0$ ($i_{s,m}^{end}$) denote the initial (final) inventory of unit type $m \in M$ at station $s \in S$. Moreover, nh_p is the number of conductors required for composition $p \in P$, dep_t is the departure time of trip $t \in T$, and t_{disr} is the starting time of the disruption. Also, we use the parameter $M_{p_1,p_2}^{\text{diff}}$ which equals 1 if composition $p_1 \in P$ and $p_2 \in P$ are of a different rolling stock type and equals 0 otherwise.

Six new variables are introduced, which are given in Table 4.1. All variables from Table 3.1 are still required for the rescheduling model, except for $I_{s,m}^{\text{diff}}$.

Table 4.1: Decision variables of Composition Model for rolling stock rescheduling

Variable	**Domain**	**Description**
h_t	$\mathbb{Z}^+$	the number of additional conductors required for trip $t \in T$
m_t	$\mathbb{B}$	equals 1 if trip $t \in T$ is executed by another vehicle unit type than in the original schedule and zero otherwise
p_t^{diff}	$\mathbb{B}$	equals 1 if trip $t \in T$ is rescheduled to a different composition and zero otherwise
nC_c	$\mathbb{Z}^+$	the number of additional coupled vehicle units at transition $c \in C$
nU_c	$\mathbb{Z}^+$	the number of additional uncoupled vehicle units at transition $c \in C$
$I_{s,m}^{\text{off}}$	$\mathbb{Z}^+$	the off-balance at station $s \in S$ for type $m \in M$

The new objective is given in Equation (4.1) with weights w_5 until w_{12}. Constraints (3.2) - (3.9) and Constraints (3.13) - (3.16) are still required for the rescheduling model. Furthermore, we add Constraints (4.2) - (4.13).

$$
\begin{aligned}
\mathbf{min} \; & \sum_{t \in T} \sum_{p \in \eta(t)} w_5 \cdot \mathrm{carr}_p \cdot \mathrm{km}_t \cdot X_{t,p} + \sum_{t \in T} w_6 \cdot X_{t,0} + \\
& \sum_{t \in T} \sum_{p \in \eta(t)} w_7 \cdot \max\left\{d_t - \mathrm{cap}_p, 0\right\} \cdot \mathrm{dur}_t \cdot X_{t,p} + \sum_{s \in S} \sum_{m \in M} w_8 \cdot I^{\mathrm{off}}_{s,m} + \\
& \sum_{t \in T} w_9 \cdot h_t + \sum_{t \in T} w_{10} \cdot m_t + \sum_{c \in C} w_{11} \cdot nC_c + \sum_{c \in C} w_{12} \cdot nU_c + \sum_{t \in T} w_{13} \cdot p^{\mathrm{diff}}_t
\end{aligned}
\tag{4.1}
$$

The objective in equation (4.1) penalizes all undesired actions, namely, respectively, carriage kilometers, cancellations, passenger standing time, off-balances, additional conductors, change of vehicle unit type, additional coupled vehicle units, additional uncoupled vehicle units, and composition changes.

$$h_t \geq \sum_{p \in P} (nh_p \cdot X_{t,p}) - oh_t \qquad \forall t \in T \tag{4.2}$$

$$m_t = \sum_{p \in P} \left(M_{p^0_t,p} \cdot X_{t,p} \right) \qquad \forall t \in T \tag{4.3}$$

$$p^{\mathrm{diff}}_t \geq \sum_{p \in P \,\mid\, p \neq p^0_t} X_{t,p} \qquad \forall t \in T \tag{4.4}$$

$$nC_c \geq \sum_{m \in M} C_{c,m} - oC_c \qquad \forall c \in C \tag{4.5}$$

$$nU_c \geq \sum_{m \in M} U_{c,m} - oU_c \qquad \forall c \in C \tag{4.6}$$

$$I^{\mathrm{off}}_{s,m} \geq I^{end}_{s,m} - i^{end}_{s,m} \qquad \forall s \in S, m \in M \tag{4.7}$$

$$I^{\mathrm{off}}_{s,m} \geq i^{end}_{s,m} - I^{end}_{s,m} \qquad \forall s \in S, m \in M \tag{4.8}$$

$$I^0_{s,m} = i^0_{s,m} \qquad \forall s \in S, m \in M \tag{4.9}$$

$$X_{t,p^0_t} = 1 \qquad \forall t \in T \quad \mid \quad dep_t \leq t_{disr} \tag{4.10}$$

$$h_t \in \mathbb{Z}^+ \qquad \forall t \in T \tag{4.11}$$

$$m_t \in \mathbb{B} \qquad \forall t \in T \tag{4.12}$$

$$nC_c, nU_c \in \mathbb{Z}^+ \qquad \forall c \in C \tag{4.13}$$

Constraints (4.2) ensure the value of the additional number of conductors that are required, by subtracting the original number of required conductors from the number of conductors in the rescheduling solution. Be aware that (4.11) states that h_t must be greater than zero, whereas the elimination of a required conductor is not taken into account in the objective. Furthermore, the values of m_t are ensured by Constraints (4.3), by looking at whether the former and the new chosen vehicle unit types are equal.

Constraints (4.4) ensure that p^{diff}_t equals one if a composition is chosen that does not equal the original composition. Constraints (4.5) and (4.6) establish the number of additional (un)coupled vehicle units. Again, note that nC_c and nU_c are greater than zero (see Constraints (4.13)) and thus only the additional number of (un)coupled vehicle units is stored.

Furthermore, Constraints (4.7) and (4.8) make sure that $I^{\text{off}}_{s,m}$ is the absolute difference between the scheduled final inventory and the rescheduled final inventory. Also, (4.9) secures the value of the initial inventory level, which equals the planned initial inventory level. Constraints (4.10) secure the chosen composition until the disruption has happened, which is the moment we can start rescheduling. Finally, Equations (4.11) - (4.13) state the domain of the variables.

4.2 Heuristic

This section describes the heuristic for rolling stock rescheduling in detail. The heuristic consists of three parts: a start heuristic, the Two-Opt Duty Neighborhood, and the Composition Model Neighborhood. The two neighborhoods are applied as Local Search neighborhoods in a Variable Neighborhood Descent (VND).

The start heuristic obtains a feasible rolling stock schedule after a disruption. This start heuristic takes into account the objectives as defined in Section 1.2 as this provides a good starting point for the rest of the heuristic. After the start heuristic is executed, the Two-Opt Duty and the Composition Model Neighborhood are iteratively applied as VND. The main idea of the Two-Opt Duty Neighborhood is to swap the remaining parts of two duties if the corresponding trains are at the same station at the same time. The swap that results in the best objective value is applied. The main idea of the Composition Model Neighborhood is to optimize the deployment of rolling stock for one vehicle unit at a time. In this way, we make use of the computational benefits of the Composition Model under limited circumstances while maintaining the flexibility of a heuristic. Moreover, the two neighborhoods complement each other because the Composition Model Neighborhood takes care of optimization within one vehicle unit type, while the Two-Opt Duty Neighborhood looks for an optimal exchange of units between the different types.

The heuristic is formalized in Algorithm 1. As we see, we start the heuristic by determining the set of duties D. We repeat that a duty is a list of trips and the positions on the trips that are executed by one train unit during the planning horizon. We determine the initial set D by first establishing the fundamental duties of the initial rolling stock schedule. A fundamental duty is a sequence of trips performed by one unit from coupling to decoupling. Then, using the Duty Flow Model from Nielsen (2011), we determine which vehicles perform which fundamental duties, resulting in a duty for each vehicle unit. Multiple combinations of fundamental duties are possible, meaning that from one initial schedule determined by the Composition Model, multiple different sets of duties can be derived. This does not hold the other way around: for a given set of duties, it is always fixed which trip is performed using which composition. Moreover, we can determine the objective value associated with a given set of duties D by means of the function $f(D)$. This function takes into account the same objectives as the MIP in Section 4.1.

After determining the original set of duties, we transform this to a set of duties after disruption by applying the start heuristic. Then we continue to the VND: in each iteration, we first try to find an improvement of the objective value using the Two-Opt Duty Neighborhood as this neighborhood is least computationally heavy. If the Two-Opt Duty Neighborhood can no longer find an improvement, we move on to the Composition Model Neighborhood. If the Compositions Model Neighborhoods succeeds in finding an improvement, we go back to the Two-Opt Duty Neighborhood. If the Composition Model Neighborhood also does not find an improvement, the heuristic terminates.

Algorithm 1: Rolling stock rescheduling heuristic

```
Data: See Figure 3.1
Result: Rolling stock schedule after disruption
D ← transformScheduleToDuties();
D ← startHeuristic(disruption);
improvement ← true;
while improvement do
    improvement ← false;
    D' ← findImprovement(D, N_{two-opt});
    if f(D') < f(D) then
        D ← D';
        improvement ← true;
    else
        D' ← findImprovement(D, N_{comp});
        if f(D') < f(D) then
            D ← D';
            improvement ← true;
        end
    end
end
```

We now discuss the other heuristic steps in detail. We discuss the start heuristic in Section 4.2.1, the Two-Opt Duty Neighborhood in Section 4.2.2 and finally discuss the Composition Model Neighborhood in Section 4.2.3.

4.2.1 Start heuristic

The start heuristic works as follows: first, we cut off an original duty whenever the duty encounters a trip that is canceled due to the disruption. The remaining trips of a cut off duty are stored and are referred to as 'empty duty'. Recall that in Subsection 1.1.2 we explained that due to the disruption new transitions are added to the set of transitions. In the second step of the start heuristic, we connect the two duties that correspond to such a new transition. We do this by merging an empty duty and a cut off duty if the last trip of the cut off duty is connected to the first trip of the empty duty by such a new transition. These first two steps are visualized in Figure 4.1. In this figure, we encounter a blockage between Rotterdam and Gouda that results in the cancellation of trips (the black dotted lines). We see that t_{10}, t_{11} and t_{12} occur in a duty after a canceled trip and hence form an empty duty. Also, t_{14}, t_{15} and t_{16} form an empty duty. This empty duty, however, is connected by a new transition to the duty of trip t_1, t_2 and t_3.

Thirdly, we try to find a vehicle unit from inventory to perform a remaining empty duty. For the situation in Figure 4.1, this implies that we want to assign an available unit at Gouda to t_{10}, t_{11} and t_{12}. In this thesis, we consider three ways to prioritize the assignment of available vehicle units to empty duties, which we compare in Chapter 5. In the first method, we prioritize the empty

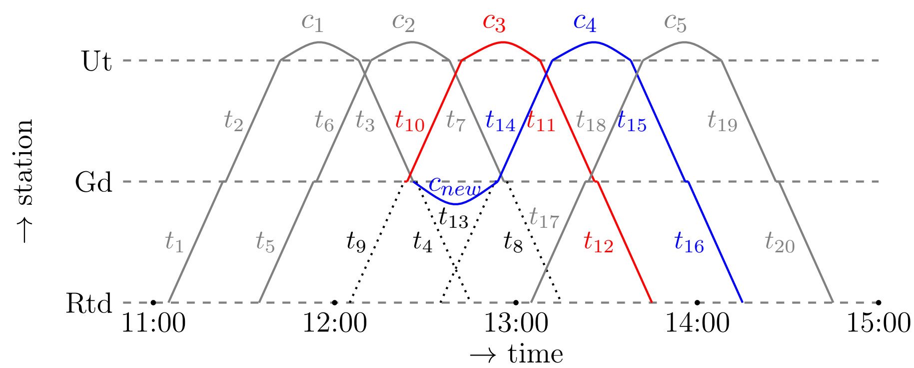

Figure 4.1: Circulation of five trains between Rotterdam and Utrecht, where 5 trips (black, dotted) between Rotterdam and Gouda are cancelled due to a disruption. In blue, we see an empty duty that is filled by a new transition. In red we see an empty duty that can not be filled by a new transition.

duties chronologically, so we first try to find an available vehicle unit at the begin station of the empty duty that starts earliest in time. The second method prioritizes the empty duties from the longest empty duty to the shortest. Finally, the third method prioritizes the empty duties by the number of currently canceled trips in the empty duties. A trip is currently canceled if there is no vehicle unit assigned to the trip.

Note that in all of the above, we only append an empty duty to another duty if the general type is the same. For example, if the cut off duty type is DDZ-4, we can add an empty duty that is of type DDZ-4 or DDZ-6 as this ensures that when later on, other vehicle units are added to the composition, the composition is likely to remain feasible. Moreover, for the same reason, we make sure that the new number of vehicle units that is assigned to a trip cannot exceed the original number of planned vehicle units. Finally, if it were to happen that the start heuristic causes an infeasible composition, for example due to its length, we cut off the duty from the point where the composition becomes infeasible.

The start heuristic is formalized in Algorithm 2. We first loop through all duties and check if these duties contain canceled trips. We do this by looping through the set of trips, which are ordered by time of occurrence. We start with index 0 in line 4. In case we find a canceled trip, we add a new duty to the set of duties after disruption that consists of the trips from the start of the duty (index 0) to the trip right before the cancellation (index n - 1). We then store the index of the trip after the canceled trip and continue to loop through the trips. In case we, again, encounter a canceled trip, line 11 until 15 add a duty that consists of the trips from the last cancellation until the trip right before the new cancellation to the set of empty duties. This continues until all trips of a duty are checked. In case we do not find (another) canceled trip, we must add the final sequence of trips to the set of duties. We do this in line 18 until 23. In case we did not find a canceled trip at all, we add the unaltered duty to the duties after disruption. After checking for canceled trips, we sort the empty duties by one of the earlier mentioned methods. Then, we try to fill the empty duties.

The first for loop of the algorithm, line 3 until 24, is of $\mathcal{O}(|D| \cdot |T|)$ and the second part, line

26 until 31, is of order $\mathcal{O}(|D|^2)$. As the number of trips is larger than the number of duties, we can conclude that the start heuristic is of order $\mathcal{O}(|D| \cdot |T|)$.

Algorithm 2: Start heuristic

Data: See Figure 3.1: T_c = set of cancelled trips, D = set of duties, C_c = set of new transitions
Result: Starting solution for rolling stock schedule after disruption

1 $D^{\text{empty}} \leftarrow \{\}$;
2 $D^{\text{disruption}} \leftarrow \{\}$;
3 **foreach** $d \in D$ **do**
4   addFromIndex $\leftarrow 0$;
5   **foreach** *trip* $t_n \in d$ **do**
6   **if** $t_n \in T_c$ **then**
7   **if** $\mathit{addFromIndex} = 0$ **then**
8   newDuty $\leftarrow t_0 ... t_{n-1}$;
9                 $D^{\text{disruption}} \leftarrow D^{\text{disruption}} \cup \{\text{newDuty}\}$;
10   addFromIndex $\leftarrow n + 1$;
11   **else**
12   newDuty $\leftarrow t_{\mathit{addFromIndex}} ... t_{n-1}$; // in case the duty contains multiple canceled trips
13                 $D^{\text{empty}} \leftarrow D^{\text{empty}} \cup \{\text{newDuty}\}$;
14   addFromIndex $\leftarrow n + 1$;
15   **end**
16   **end**
17   **end**
18   **if** $\mathit{addFromIndex} = 0$ **then**
19         $D^{\text{disruption}} \leftarrow D^{\text{disruption}} \cup \{\text{d}\}$;
20   **else**
21   newDuty $\leftarrow t_{\mathit{addFromIndex}} ... t_N$;
22         $D^{\text{empty}} \leftarrow D^{\text{empty}} \cup \{\text{newDuty}\}$;
23   **end**
24 **end**
25 $D^{\text{empty, sorted}} \leftarrow \mathit{sort}(D^{\text{empty}})$;
26 **foreach** $d \in D^{\mathit{empty,\ sorted}}$ **do**
27     $D^{\text{disruption}} \leftarrow \mathit{connectByNewTransition}(C_c, d, D^{\mathit{disruption}})$;
28   **if** *not connected by new transition* **then**
29         $D^{\text{disruption}} \leftarrow \mathit{findAvailableVehicle}(d, D^{\mathit{disruption}})$;
30   **end**
31 **end**

4.2.2 Two-Opt Duty Neighborhood

In the Two-Opt Duty Neighborhood, we try to find a swap of the remaining part of two duties that decreases the objective value of the rolling stock schedule after disruption. This swap occurs between different rolling stock types, meaning that this neighborhood improves the assignment of rolling stock types to trips. An addition to the method of Hoogervorst et al. (2021) is that we also add the empty duties that result from the start heuristic to this neighborhood.

An example of such a swap is shown in Figure 4.2. The red vehicle type was supposed to be coupled at transition c_1 and the blue vehicle type at transition c_2 and was available to couple to transition c_1. The remaining duties of these two vehicle types are from this point on swapped. This can be beneficial if the red type has a higher seating capacity and trips t_5 - t_8 have a higher

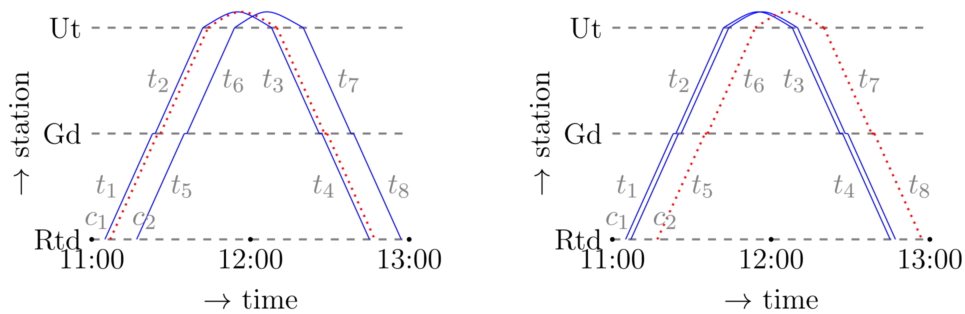

Figure 4.2: Circulation of three trains between Rotterdam and Utrecht with one type in red and one type in blue. On the left, we see the initial rolling stock schedule. On the right, we see one swap of the Two-Opt Duty Neighborhood between the red and blue vehicle type. Superfluous labels of trips and transitions have been omitted for clarity

passenger demand than trips t_1 - t_4.

A swap between d_1 and d_2 at transition c is allowed if the vehicle unit corresponding to either d_1 or d_2 (or both) is coupled at transition c, if it is allowed to shunt at the corresponding station and if both vehicle unit corresponding to duties d_1 and d_2 are present at the transition station. In order to efficiently find all possible swaps, we add, for each transition $c \in C$, all pairs of duties that can be swapped at transition c to the set $S(c)$.

Algorithm 3: Two-Opt Duty Neighborhood

Data: Duty List (D)
Result: Best Two-Opt Duty swap

```
1  S ← getPossibleDutySwaps(D, S_previous);
2  objectiveDecrease ← 0;
3  bestSwap ← {} ;
4  foreach c ∈ C do
5      foreach {d1, d2} ∈ S(c) do
6          D' ← swap(d1, d2);
7          objectiveDecrease' = getObjectiveDecrease(D, D');
8          if objectiveDecrease' > objectiveDecrease then
9              objectiveDecrease ← objectiveDecrease';
10             bestSwap ← {d1, d2};
11         end
12     end
13 end
```

We now find the best possible swap by implementing Algorithm 3. This algorithm, among others, performs *getPossibleDutySwaps(D)* in line 1 and *getObjectiveDecrease(D, D')* in line 7. *GetPossibleDutySwaps(D)* finds, for each transition, the set $S(c)$ by looping through all duties. This means that *GetPossibleDutySwaps(D)* is of order $\mathcal{O}(|D| \cdot |C|)$ The sets $S(c)$ are created once from scratch and later, after each iteration, adjusted to the new situation for computational efficiency. The function *getObjectiveDecrease(D, D')* determines the objective decrease in a smart manner, by only evaluating the new and old objective of the adjusted trips. This means that *getObjectiveDecrease(D, D')* is of order $\mathcal{O}(|T|)$. The overall Two-Opt Duty Neighborhood algorithm is of order $\mathcal{O}(|D|^2 \cdot |C| \cdot |T|)$, because it requires for each transition for each combination of possible duties an evaluation of the objective value. Moreover, *GetPossibleDutySwaps(D)* is of smaller order than the for loop in lines 4 until 13 of Algorithm 3.

4.2.3 Composition Model Neighborhood

The main idea of the Composition Model Neighborhood is to optimize the deployment of one vehicle unit type at a time. The Composition Model Neighborhood combines the Composition Model for rolling stock rescheduling of Nielsen (2011) with the main idea of the Adjusted Path Neighborhood in Hoogervorst et al. (2021). The main idea of the Adjusted Path Neighborhood is to find, among all unit types, a change in the path of one vehicle unit of that type that decreases the objective value most. Hoogervorst et al. (2021) mentions that this Adjusted Path Neighborhood, as well as the Two-Opt Duty Neighborhood, only changes a few duties at a time. The Composition Model Neighborhood, however, optimizes the entire deployment of one vehicle unit type at a time instead of finding one path change. Moreover, as discussed in Section 2.3, it was shown that currently the Composition Model is computationally most attractive under limited circumstances but has the disadvantage of being less flexible. By applying the Composition Model as a neighborhood in our heuristic, we can make use of the computational advantage but maintain the flexibility of a heuristic.

The MIP of the Composition Model Neighborhood is in essence the same as for the regular Composition Model for Rolling Stock Rescheduling as seen in Section 4.1. Therefore, the MIP of the Composition Model Neighborhood consists of the objective of Equation (4.1) with Constraints (3.2) - (3.9), Constraints (3.13) - (3.16), and Constraints (4.2) - (4.13). However, we reduce the sizes of the sets in the model.

Consider that we optimize the rolling stock schedule for type m'. This type can be combined with type m'', as our six rolling stock types can be combined in pairs of two. Now, we change the set of trips T, transitions C types M, compositions P, allowed compositions $\mu(t)$ and allowed composition changes $\rho(c)$ by removing all elements of these sets that do not correspond to type m' and type m''. We change the set T and C by only considering the trips that were deployed by type m' and m'' and the corresponding transitions. Moreover, we can restrict $\mu(t)$ and $\rho(c)$ as the number of units of type m'' on each trip is fixed.

As the allowed composition $\mu(t)$ of each trip is limited to the right amount of m'' on that trip, the variables $U_{c,m''}$ and $C_{c,m''}$ are also restricted. This means that $I_{c,m''}$ and $I^{\infty}_{s,m''}$ are also restricted. Concluding, no additional constraints are necessary in order to remain the rolling stock schedule of m'' the same. However, in order to speed up the model, the known values of $U_{c,m''}$, $C_{c,m''}$, $I_{c,m''}$ and $I^{\infty}_{s,m''}$ are inserted into the program.

We now discuss the structure of the Composition Model Neighborhood. We could perform the MIP of the Composition Model Neighborhood for each type in each iteration of the neighborhood and apply only the best improvement. However, since the different general types operate completely separately from each other, we can make better use of the optimized MIP's. As mentioned earlier, our six rolling stock types can be combined in pairs of two. In Section 1.1.3, we state that we consider the following vehicle unit types: DDZ-4, DDZ-6, VIRM-4, VIRM-6, ICM-3, and

ICM-4. The general types are DDZ, VIRM, and ICM. When we enter the Composition Model Neighborhood, the assignment of general types to trips is fixed. In other words, it is fixed which trips and transitions are executed by either DDZ, VIRM, or ICM. Therefore, we can, for each general type, adjust the deployment of one of the sub types without changing anything about the deployment of the other general types. For example, if we change the deployment of DDZ-4, this only changes the compositions of trips that are assigned to the DDZ type and not of the VIRM and ICM type. As solving the Composition Model requires significant computation time, we implement the improvement of the deployment of one of the sub types of each general type.

Algorithm 4: Composition Model Neighborhood

Data: Duty List (D), current objective value
Result: New duty list D

1 **foreach** $\{m_1, m_2\} \in \{\{DDZ\text{-}4, DDZ\text{-}6\}, \{VIRM\text{-}4, VIRM\text{-}6\}, \{ICM\text{-}3, ICM\text{-}4\}\}$ **do**
2   continue $\leftarrow$ checkContinue(m_1, m_2);
3   **if** *continue* **then**
4         $m' \leftarrow m_1$;
5         $m'' \leftarrow m_2$;
6         $\text{obj}_{m_1} \leftarrow$ performCompositionModel(m', m'');
7         $m' \leftarrow m_2$;
8         $m'' \leftarrow m_1$;
9         $\text{obj}_{m_2} \leftarrow$ performCompositionModel(m', m'');
10   **if** $obj_{m_1} \leq obj_{m_2}$ ***and*** $obj_{m_1} \leq obj_{current}$ **then**
11             $D \leftarrow removeDutiesOfType(m_1)$;
12             $D \leftarrow D \;\cup$ improved duties of type m_1;
13   **else if** $obj_{m_2} \leq obj_{m_1}$ ***and*** $obj_{m_2} \leq obj_{current}$ **then**
14             $D \leftarrow removeDutiesOfType(m_2)$;
15             $D \leftarrow D \;\cup$ improved duties of type m_2;
16   **end**
17   **end**
18 **end**

The Algorithm for one iteration of the Composition Model Neighborhood thus looks as in Algorithm 4. We explain the idea of this algorithm by walking through one iteration of the Composition Model Neighborhood with an example: suppose we consider general types a and b, with sub types a_1, a_2, b_1 and b_2. We start by considering general type a: first, in line 2, the function *checkContinue(*m_1, m_2*)* checks whether the previous time the Two-Opt Duty or Composition Model Neighborhood was performed an improvement was found to type a_1 or type a_2. If this was the case, we first re-optimize type a_1 while fixing a_2 and store the obtained objective value as obj_{a_1}. Second, we re-optimize type a_2 while fixing a_1 and obtain obj_{a_2}. If we have found an improved objective value among obj_{a_1} and obj_{a_2}, we implement the change of deployment of the type a_1 or a_2 that corresponds to the lowest objective value. Then we do the same for general type b: we re-optimize b_1 while fixing b_2 and re-optimize b_2 while fixing b_1 and implement the best improvement in objective value (if an improvement was found). Note that the changes that have already been made to general type a do not affect general type b. Now, we proceed to the next iteration of the VND in Algorithm 1.

If we would, in contrast to the example above, only apply the best improvement among general types a and b, we could encounter the following example: we find the best improvement in type a_1 and also find an improvement in type b_2. We would apply the (best) improvement of type a_1 and would not use the improvement in type b_2 in this iteration. This would be a waste as the two general types operate separately within one iteration of the Composition Model Neighborhood, the Composition Model has a non-negligible computation time, and we did find an improvement in general type b.

In one iteration of the Composition Model neighborhood we solve a MIP in CPlex multiple times. It is interesting to study whether allowing an optimality gap could result in a further reduction of the computation time of the rolling stock rescheduling heuristic. By tolerating such a gap, CPlex terminates the branch-and-bound procedure when an integer solution is found that has a certain relative gap with the objective of the best remaining node in the branch-and-bound tree. It is important that this tolerance does not significantly reduce the final objective value of the heuristic.

5 | Results & Discussion

In this chapter, we conduct computational experiments in order to answer our research questions and provide insight into the performance of the models as described in Chapter 3 and Chapter 4.

We begin by defining the exact setting of our experiments in Section 5.1. Secondly, we show how the passenger flows change when a disruption is simulated in Section 5.2. Thirdly, we describe how the Composition Model that was used for our initial rolling stock schedule performs in Section 5.3.

Then we look into rolling stock rescheduling: First, we study the influence of taking into account passenger- and operational objectives in rolling stock rescheduling by comparing the rescheduling results of seven different parameter sets of the Composition Model in Section 5.4. Then, we look into the difference in passenger-related performance when we perform rescheduling on the original passenger flow prediction in Section 5.5

Finally, we study the potential of a heuristic method for rolling stock rescheduling. We look into the influence of heuristic choices in Section 5.6. Also, we study the potential of the heuristic by comparing the performance of this method to the Composition Model in Section 5.7.

5.1 Experimental setup

Two computers were used for the computational results. The initial rolling stock schedule of Section 5.3 is computed on an Intel(R) Xeon(R) E5-1650CPU @3.60GHz with 9.2GB of available computer RAM. The other experiments are, due to circumstances, performed on an Intel(R) Core(TM) i7-4710MQ CPU @2.50GHz with 2.82GB available computer RAM.

We simulated four complete blockages, with locations as given by the black solid lines in Figure 1.3. These are the Intercity routes Den Haag - Gouda, Leiden - Haarlem, Leiden-Schiphol, Gouda - Rotterdam, which can from now on be referred to as respectively disruption 1, 2, 3, and 4. We simulate these disruptions from 8 AM until 10AM on a Tuesday.

5.2 Passenger flow

In this section, the changes in passenger flow after disruption are presented and discussed. The results are visually presented in Figure 5.1.

The model requires parameters on passenger behavior. Based on conversations with employees of NS, a *transfer penalty* of 20 minutes is chosen. Furthermore, we applied an *alternative path tolerance* of 20% and an *alternative path fraction* of 20%. Moreover, we added a *Sprinter penalty* of 20 minutes. The *delay tolerance* is 30 minutes and the *delay leave fraction* is 30%. Finally, we chose a *transfer time* of 4 minutes.

The running time of the passenger flow model after a disruption is between 60 and 80 seconds. This running time could be drastically decreased by storing the already obtained shortest paths more efficiently but this is beyond the scope of this thesis.

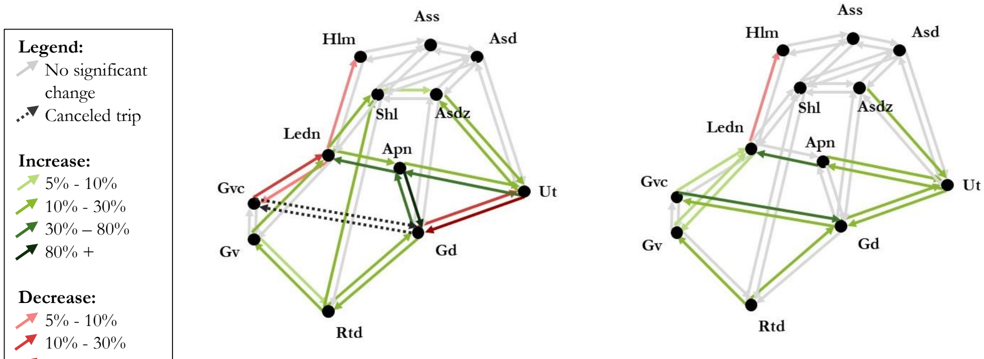

(a) Disruption 1: Den Haag Centraal – Gouda

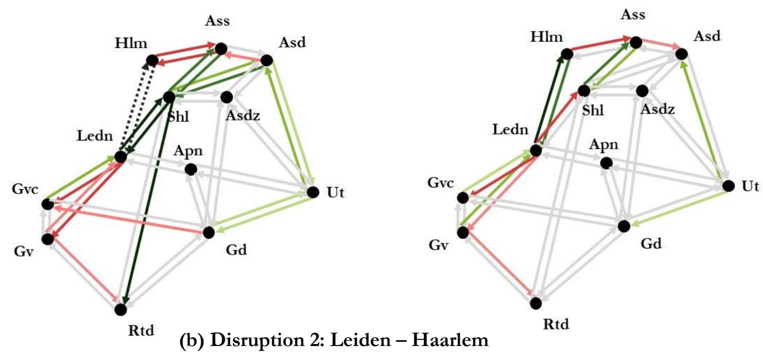

(b) Disruption 2: Leiden – Haarlem

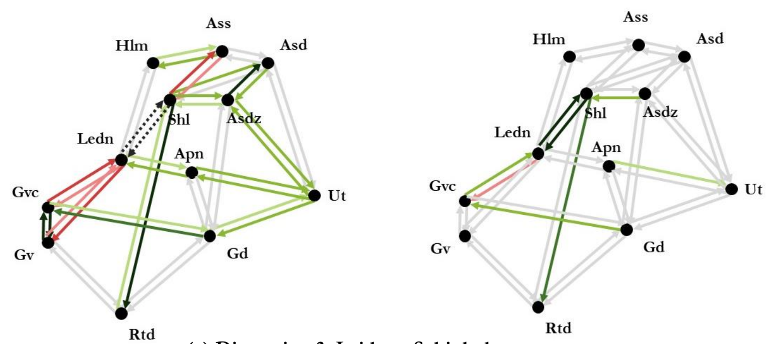

(c) Disruption 3: Leiden - Schiphol

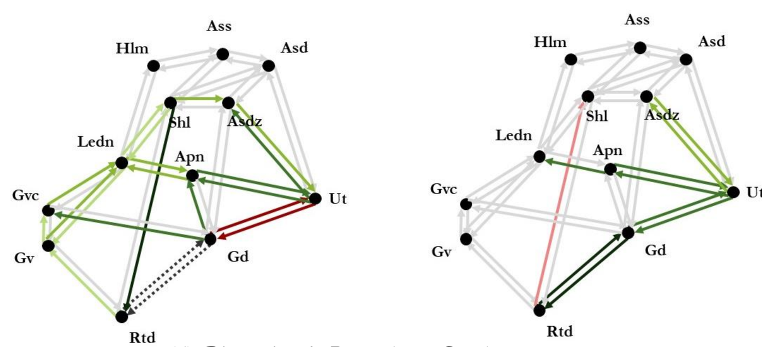

(d) Disruption 4: Rotterdam - Gouda

Figure 5.1: Passenger flow results. Left image: Average decrease or increase in passenger flow on that route during the disruption. Right image: Average decrease of increase in passenger flow in the half hour after the disruption. No significant change means that the increase or decrease is smaller than 5%

Figure 5.1 shows the changes in passenger flows in case of a disruption. The color depicts the change in passenger flow over a directed arc. The alert reader observes that station Apn is added to the graph. This station is added as on routes around this station a lot of passenger flow increases or decreases occur after a disruption, which makes it interesting to observe from a passenger point of view. However, on the route Apn - Gd there is no InterCity train, which means that the passenger flow increase on this route is not used in our rescheduling model.

A foreseeable observation is that on paths that can be used as alternatives to the disrupted route the passenger flow increases during a disruption. For example, in Figure 5.1.a, we see that if trips from Gd - Gvc are canceled, the passenger flow at routes Gd - Rtd - Gv and Ut - Apn - Ledn increases. Also, we see that in the half hour following the disruption, the disrupted path encounters an increase in passengers due to passengers that have waited until the disruption was over.

Another, less obvious, observation is that the passenger flow decreases at the routes adjacent to the begin- and end station of the canceled route. This can be explained by the fact that trips on the blocked route are used not only to reach stations directly adjacent to the blocked route, but also to reach stations adjacent to the ends of the blocked routes. Alternative routes will now be used for this purpose, which may reduce the passenger demand on routes adjacent to the canceled route. In some cases, this effect expands to more routes. This is visible in Figure 5.1.b: We here observe that a decrease occurs on the route Gd - Gvc. This can, among others, be explained by the fact that passengers from Hlm, who normally travel to Gd via Gvc, travel via Ut after the disruption, as reflected in the increase in passenger demand around Utrecht.

Another interesting phenomenon is that on a route the increase or decrease in demand is not always the same in both directions. This is likely due to the timing of trips, causing a different shortest route in one direction than in the other.

Finally, we look into changes in passenger demand that seem unpredictable at first glance. This might be the case for the passenger demand increase from Ut to Asd for disruption 2. This can be explained by the fact that passengers who started their journey during the disruption have this journey in their itinerary, for example, passengers that travel from Gvc or Gv to Asd. Other notable cases have in common that a decrease in the flow of one part of the passenger groups and an increase of another part occur simultaneously. This occurs between Ledn and Hlm for disruption 3: an increase occurs because this route is a detour route for the disruption but a decrease occurs as a part of the travelers with origin-destination Hlm and Shl (or vice versa) do no longer travel via Ledn. This is also the case for disruption 4 between Rtd and Shl: we encounter an increase because this is a detour route but also a decrease as passengers coming from Gd would normally travel via Rtd to reach Shl and adjacent stations. This also explains the decrease in travel demand between Rtd and Shl after the disruption.

5.3 Initial rolling stock schedule

This section elaborates on the Composition model for rolling stock scheduling as given in Section 3.2. In Table 5.1, we see the chosen parameters and a summary of the performance of the Composition Model for generating our initial rolling stock schedule. Moreover, we run the model under the assumption that only one vehicle type is allowed per line. The parameters for carriage kilometers and seating shortage have the value used at NS. It was chosen to make the other operational variables equally large. In small experiments, the chosen values showed a good balance between the amount of seating shortage, carriage kilometers and operational objectives. Moreover, we ran the model with the limit that only one vehicle type per line is allowed.

The optimization in CPlex is truncated after 5 hours, after which the model achieves an optimality gap of 0.55%. We can conclude that the running time of approximately 5 hours is high even when you do not run the model to optimality. However, reducing the running time of the initial rolling stock schedule is beyond the scope of this thesis.

We see in Table 5.1 that, obviously, the number of couplings and uncouplings is the same. Furthermore, we cannot yet draw any conclusions from these results but these numbers could be relevant as a benchmark in later sections. Finally, note that the initial schedule does not contain canceled trips because initially, the allowed set of compositions for each trip does not contain empty compositions.

Table 5.1: Overview of parameters and results of the initial rolling stock schedule. On the left, we see the value parameters, its weight parameter in Equation (3.1) and the value. On the right, we see the corresponding running time, optimality gap, carriage kilometers, number of canceled trips, seating shortage for the initial rolling stock schedule generated by the Composition Model with a limit of one vehicle unit type per line

Parameter	**Weight in objective**	**Value**
Carriage kilometers (km)	w_1 (unit km^{-1})	0.13
Seating shortage (hour)	w_3 (unit $hour^{-1}$)	60
Inventory difference	w_4	50
Coupling	w_5	50
Uncoupling	w_6	50

Measure	**Value**
Carriage kilometers	1,231,542.0 km
Seating shortage	5,189.61 hours
Inventory difference	100 units
Total number of couplings	757
Total number of uncouplings	757
Running time:	5 hours
Optimality gap:	0,55%

5.4 Rolling stock rescheduling parameters

The goal of this section is to study the impact of stick-to-the-plan controlled rescheduling, which currently is the main objective in rolling stock rescheduling. This section provides a use-case for considering passenger- and operational objectives in rolling stock rescheduling. We study this by performing computational experiments on the Composition Model of Section 4.1 for different parameter instances. We perform this study on the Composition Model as this solves the rescheduling problem to optimality.

As you can see in Table 5.2, we differentiate the weight that is put on penalizing performing a trip with a different composition than planned while keeping the other penalties constant. The constant penalties were determined as follows: the balance between carriage kilometers and seating shortage remains the same as in the initial schedule. Moreover, we penalize cancellations heavily, as cancellations are very undesirable. Finally, the values of the operational parameters are carefully determined by conducting small computational experiments, which will not be discussed further.

Table 5.2: Overview of the parameters for rolling stock rescheduling. On the left, we see the fixed parameters with description of the parameter and the weight in Equation (4.1). On the right, we see the stick-to-the-plan parameter value that is varied among the seven parameter instances.

Fixed parameters		
Description	**Index**	**Value**
Carriage kilometers (km)	w_5 (unit km^{-1})	0.13
Cancellation	w_6	10,000,000
Seating shortage (hour)	w_7 (unit $hour^{-1}$)	60
Off-balance	w_8	1,000
Extra HC	w_9	100
Different Type	w_{10}	100
New Coupling	w_{11}	100
New Uncoupling	w_{12}	100

Different composition: w_{13}	
Parameter Instance	**Value**
1	0.01
2	50
3	100
4	300
5	600
6	1,200
7	100,000

However, we note that the penalty on an off-balance is set higher than the other operational objectives because an off-balance might cause moving empty vehicles, which is strongly undesirable.

Parameter instances 1 and 7 in Table 5.2 examine the two extremes of rolling stock rescheduling. Parameter instance 1 does not consider stick-to-the-plan at all, as it does not penalize executing a trip with a different composition as planned. Parameter instance 7 focuses on stick-to-the-plan controlled rescheduling as it strongly penalizes deviating from the original composition, which outweighs passenger comfort and operational objectives. Parameter instance 7 can be seen as the current situation. However, the current determined rolling stock schedule in case of a disruption is likely to be far from the optimal solution of parameter instance 7 because rolling stock rescheduling is currently performed manually. Parameter instance 2 until 6 study the middle ground between instances 1 and 7.

We perform rolling stock rescheduling on the four disruptions using the seven parameter instances. Table 5.3 gives information on the running time of the models and the number of canceled trips. The bar graphs in Figure 5.2 show the average value of the seating shortage, operational costs, and the number of different compositions over the four disruptions. In Figure 5.3, the relation between the three objectives is shown by plotting the average value over the disruptions for each parameter instances against each other. Note that in order to determine the operational costs, we use the weights of Table 5.2 but do not include canceled trips. Moreover, in the Appendix, the obtained values for each disruption and each parameter instance are presented in Table A.1 and the average value of all operational objectives are given in Figure A.1.

Table 5.3: Running time and number of trips for the six parameter instances

	Parameter instance						
Measure	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Run time (average over disruptions) (s)	405	242	233	189	168	137	110
Canceled trips disruption 1, 2 and 4	0	0	0	0	0	0	0
Canceled trips disruption 3	3	3	3	3	3	3	3

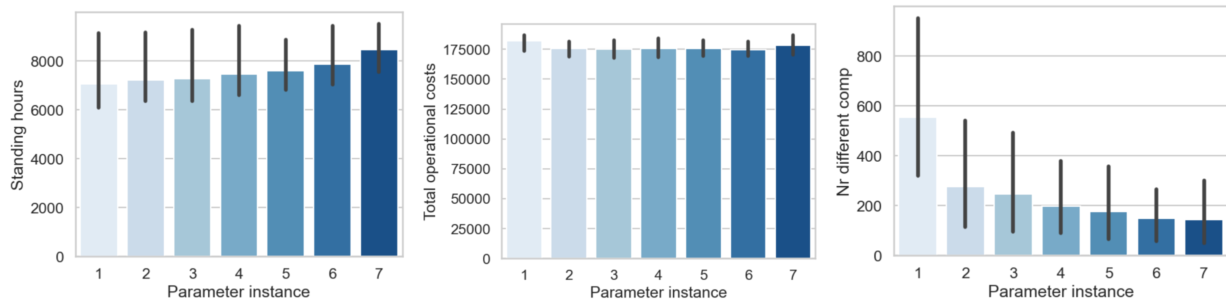

((a)) Total seating shortage ((b)) Operational costs ((c)) Composition changes

Figure 5.2: Bar graphs of the average number of standing hours, operational costs and composition changes of the seven parameter instances over the four disruptions

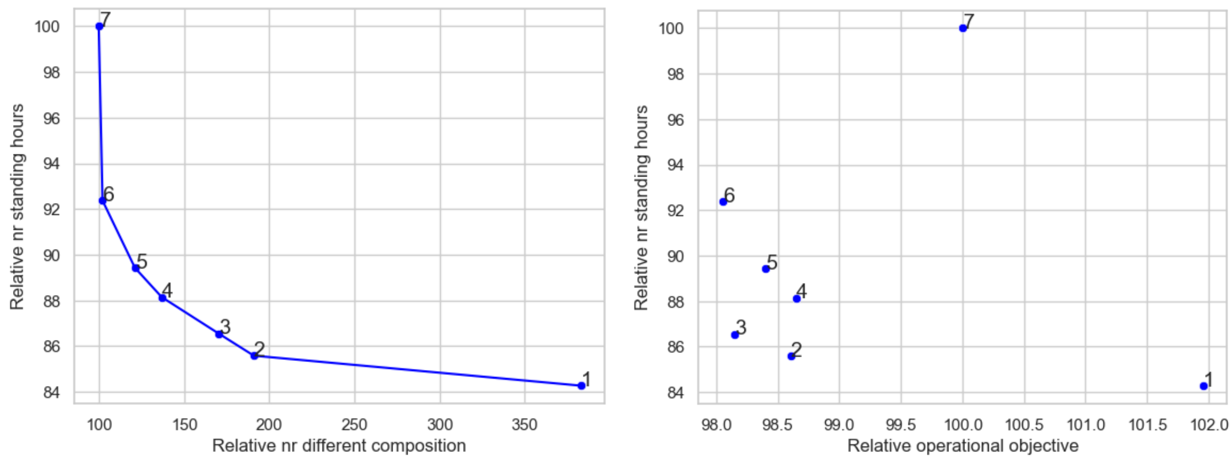

((a)) Relative seating shortage vs relative nr different compositions

((b)) Relative seating shortage vs relative operational costs

Figure 5.3: Plots of the relative value in comparison to parameter instance 7 of the seating shortage, number of different compositions and operational costs. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

We now discuss the obtained results on the seven parameter instances and begin by discussing the potential of considering passenger and operational objectives in rescheduling. As expected, we see in Figure 5.2 that when we put more weight on stick-to-the-plan objectives, the number of composition differences decreases and the seating shortage increases. We also observe this in the convex line in Figure 5.3(a). This figure clearly shows that stick-to-the-plan rescheduling comes at the expense of passenger comfort.

Moreover, we see an interesting relation between the penalty on stick-to-the-plan and operational costs. We observe in Figure 5.2(b) and Figure 5.3(b) that for the two extremes, namely parameter instances 1 and 7, the operational costs are significantly higher than for the other parameter instances. We take a look at the separate operational objectives in Figure A.1 in the Appendix to determine where these high operational costs come from. In this figure, we observe a decreasing trend in the number of type changes, additional conductors, and new (un)couplings and the decrease in carriage kilometers if we put more emphasis on stick-to-the-plan. This causes high operational costs for parameter instance 1. Moreover, we see an opposite trend for the number of off-balances, which causes the relatively high operational costs for parameter instance 7.

It is interesting to briefly study whether the operational costs in relation to the seating shortage would be (semi-)convex in case off-balances are not included in rescheduling. Figure A.2 in the Appendix shows that this is indeed the case. We can, therefore, conclude that the interesting relation between stick-to-the-plan and operational costs is due to off-balances.

We see that in terms of operational costs parameter instances 2 until 6 do not perform significantly different. Therefore, it is up to the operator to determine which point in the trade-off between standing hours and the number of different compositions is preferred. For example for parameter instance 4, we see that putting less emphasis on stick-to-the-plan results in a decrease in passenger standing hours of around 12% and a decrease in the operational costs of around 1.4%. We can conclude that by including both passenger comfort, operational objectives, and stick-to-the-plan in rescheduling, we can significantly reduce the seating shortage and the operational costs at a cost of more deviations from the original plan.

We observe in Table 5.3 that the running time of the Composition Model increases as we decrease the weight of the stick-to-the-plan objective. This makes sense because we are giving the MIP less guidance by reducing the penalty of deviating from the existing schedule. The running times are reasonable but could be reduced to make the model more suitable for practical purposes. A second observation from Table 5.3 that the model is not able to find a solution without canceled trips for disruption 3. This might be the case because of the restrictions we impose on the compositions and composition changes.

Another observation that we briefly discuss is that we see in Table A.1 that the total number of standing hours is much larger for disruption 2 than for the other disruptions. This is due to a very high passenger demand on some trips. On these trips, the demand is higher than the maximal capacity of any rolling stock composition. Moreover, for disruption 2 parameter instances 5 and 6 result in exactly the same schedule. It is likely the case that the feasible region in which no cancellations occur for disruption 2 is very limited.

5.5 Influence dynamic passenger flow rolling stock rescheduling

In this section, we investigate how the dynamic passenger flows, as generated by our passenger flow model, affect rolling stock rescheduling. This tells us whether we can avoid standing hours by not optimizing under the assumption that passenger demand remains constant after a disturbance. We do this by, first, rescheduling under static passenger demand and, second, studying the performance under the adapted passenger demand. We compare this performance to the performance of the same rescheduling model under dynamic passenger flows, as seen in Section 5.4. We do this by adding the results of the Composition Model under static passenger flow to Figure 5.3(a) and Figure 5.3(b) which resulted in Figure 5.4. Moreover, in the Appendix, we added the full results of the model with static passenger flows in Table B.1.

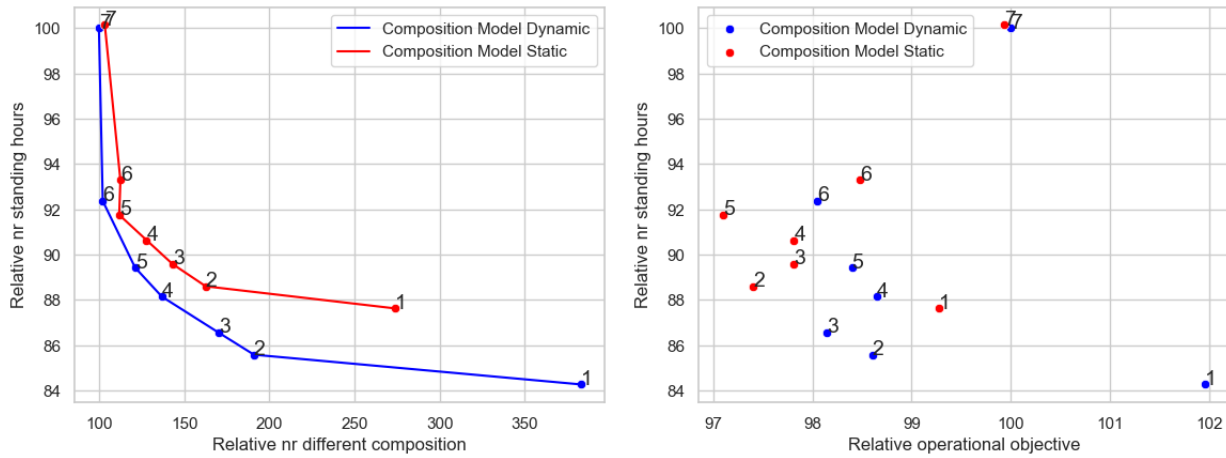

((a)) Relative seating shortage vs relative nr different compositions

((b)) Relative seating shortage vs relative operational costs

Figure 5.4: Comparison of the Composition Model when performed under static or dynamic passenger flow. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

First of all, we observe in Figure 5.4(a) that the line of optimizing under static passenger flow is always above that of dynamic passenger flow. We can therefore state that rescheduling with dynamic passenger flow is beneficial for passenger comfort. We see that the line of the static model is not strictly convex. This may seem mathematically incorrect, but this is not the case as the standing hours that are considered in the objective of the static model are not the same as the standing hours that are used to make Figure 5.4(a). In this figure, we namely determine the standing hours of the dynamic passenger flow of the schedule that is made with static passenger flows.

We observe in Figure 5.5(a) that rescheduling with static passenger flow leads to a lower number of trips driven with a different composition than planned for parameter instance 1 until 5. This makes sense since rescheduling is now performed with the same passenger demand as in the initial schedule. This means that the planned composition is likely to fit well with the passenger demand used in rescheduling. Therefore, fewer trips are performed with a different composition as planned.

It is difficult to draw conclusions about the operational costs in Figure 5.4(b) because the points are fairly close to each other. However, the operational costs of the static model seem slightly lower. Moreover, parameter instance 1 results in much lower operational costs for the static model than the dynamic model. This probably also has to do with the fact that the original schedule is likely to fit well with the passenger demand that is used in rescheduling with static passenger demand.

In some cases, the schedule that is generated with static or dynamic passenger flows is the same. First of all, we observe this for parameter instance 7 for disruption 1 and 4 and for parameter instance 6 for disruption 1. This can be explained by the high penalty on a different composition than planned for these parameter instances which makes the two models similar. This also explains that when we focus less on stick-to-the-plan the two lines are further apart. For disruption 2 and

3, the solutions of parameter instance 7 slightly deviate. We expect that if we had chosen value 0.0 instead of 0.01 for w_{13} in parameter instance 7, both models would have resulted in the same schedule. Second, we see that for disruptions 2 and parameter instances 5 and 6, the models result in the same schedule. As we already explained in Section 5.4, this is likely due to the small feasible region.

The differences in results of other disruptions and parameter instances show that the model responds to the changes in passenger demand in case we consider dynamic passenger flows. In practice, this would imply that, for example, a train on a route that has an increase in passenger demand is lengthened. Concluding, we have shown that rescheduling with dynamic passenger demand increases passenger comfort at a cost of deviations from the original schedule.

5.6 Influence heuristic choices

This section discusses the influence of heuristic choices. First of all, we look into the influence of in what order the inventory is allocated to empty duties in Subsection 5.6.1. Second, in Subsection 5.6.2, we look at it is the case that the two neighborhoods complement each other, as explained in Section 4.2. Third, we study whether allowing a MIP gap could improve the heuristic in Section 5.6.3

5.6.1 Start heuristic

Figure 5.5 and Table 5.4 provide results on the three different ways (chronologically, by length, or by the number of cancellations) of filling up the empty duties with available vehicle units from inventory. Moreover, we provide results on the heuristic if we would not fill up the empty duties at all. The graph in Figure 5.5 shows the results after running through the VND of Algorithm 1 with both the Two-Opt Duty and the Composition Model Neighborhood. Another result is that all methods are able to find the same number of cancellations as the Composition Model (see Table 5.3), except when we do not fill up the empty duties. In this case, we encounter 4 cancellations in the final solution for 6 out of 7 parameter instances for disruption 3. Finally, we note that we do not present the running times because the four methods result in similar running times.

First of all, we can conclude that filling the empty duties with any method has a great added value over not filling the empty duties. We can conclude this from Figure 5.5 because the magenta line lies far above the other lines, meaning that this method performs significantly worse than the

Table 5.4: Comparison of the initial number of canceled trips between filling empty duties chronologically, by length of by number of canceled trips, or not at all.

	Initial number of cancellations for each start heuristic			
Disruption	**Chronological**	**Length**	**Nr cancellations**	**Do not fill**
1	100	110	88	331
2	13	14	14	330
3	89	69	69	271
4	18	4	0	218
Average	55	49.25	42.75	287.5

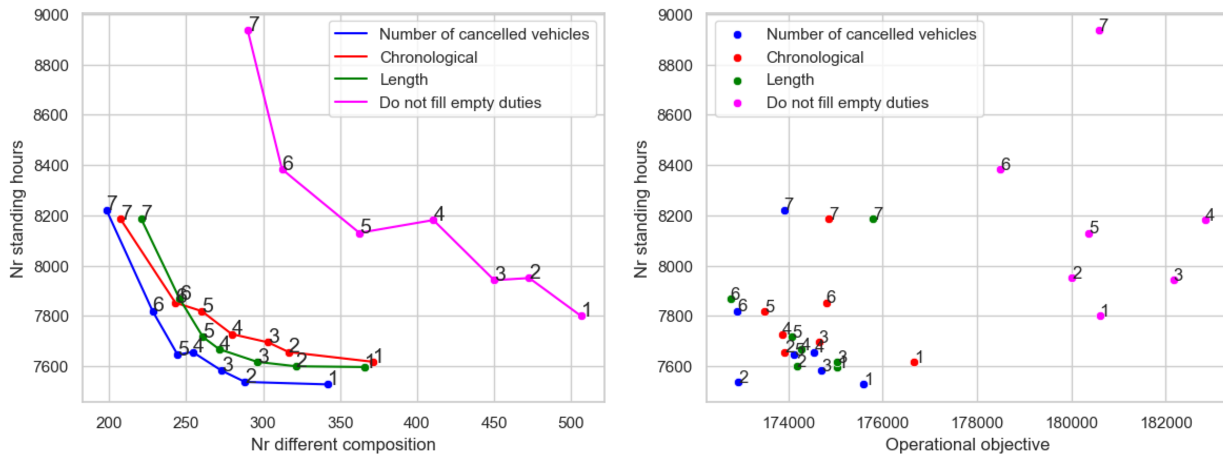

((a)) Seating shortage vs nr different compositions

((b)) Seating shortage vs operational costs

Figure 5.5: Comparison of the four methods of filling the empty duties by plotting the nr of standing hours against the number of different compositions and the nr of standing hours against. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

others. Moreover, we see in Table 5.4 that the initial number of cancellations is much higher if we do not fill the empty duties. A final argument that this method performs worse is that, in some cases, the number of cancellations in the final solution is higher than in the exact solution (see Table 5.3).

Secondly, we explain why we choose to continue filling empty duties based on the number of cancellations. We see in Figure 5.5(a) that this method provides the best results because the line of this method lies below the lines of the other methods. In Figure 5.5(b), we see that the operational costs seem similar for the three sorting methods. Finally, we see that the initial number of cancellations is the lowest for filling up using the number of cancellations. It is also interesting to note that for disruption 4, the start heuristic finds a starting solution without any canceled trips.

We end this section with a discussion on the start heuristic methods as a whole. We see in Figure 5.5(a) that the methods have a shape that is similar to a convex line but is not fully convex. This is caused by the fact that we apply a VND to find our final solution. A VND, without shaking, is likely to get stuck in a local optimum. For each disruption and each parameter instance, this can turn out just differently, causing the final solution in some cases to be close and sometimes a little further from the global optimum. This causes a wiggly and non-convex line. Moreover, it is interesting that the operational costs have similar characteristics as in the Composition Model. We namely see that the points of parameter instances 1 and 7 generally lie outside the cluster of points of the other parameter instances and show a worse relationship between standing hours and operational costs. We discuss the operational costs of the heuristic in more depth in Section 5.7.

5.6.2 Combination of neighborhoods

This section shows that it is the case that the two neighborhoods together result in the best performance because they complement each other.

Table 5.5: Number of cancellations after start heuristic and after running the Two-Opt Duty Neighborhood until no more improvement is found

Disruption	**Initial nr canceled trips**	**Nr canceled trips after** $N_{two-opt}$
1	88	86
2	14	2
3	69	45
4	0	0

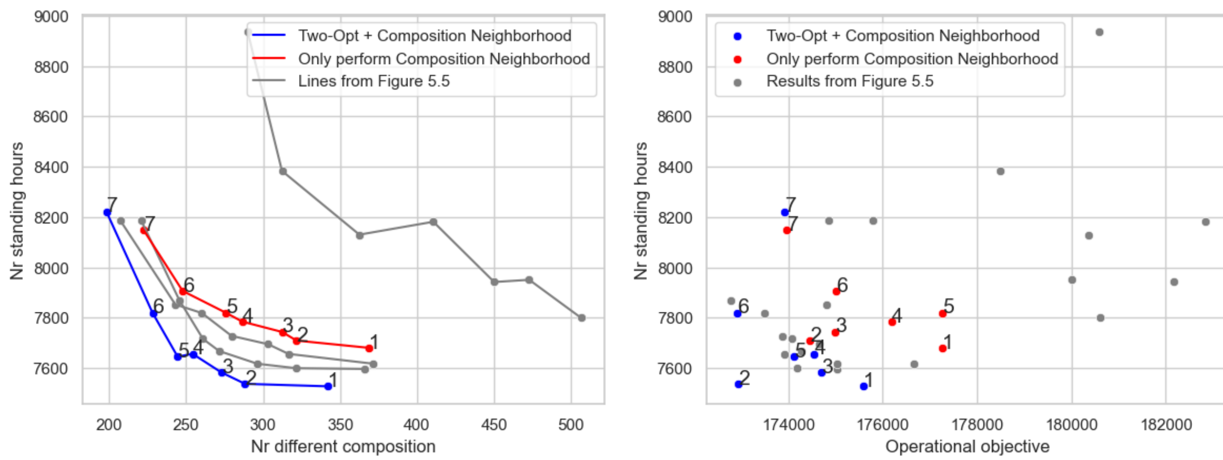

((a)) Seating shortage vs nr different compositions

((b)) Seating shortage vs operational costs

Figure 5.6: Comparison of the heuristic with only the Composition Neighborhood (red) or both neighborhoods (blue). The grey lines are a copy of the magenta, red and green line in Figure 5.5 to clarify the relative differences. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

First of all, Table 5.5 shows that executing the Two-Opt Duty Neighborhood until no more improvement is found succeeds in reducing the number of canceled trips but still results in a highly undesirable amount of cancellations for disruption 1 and 3. Table 5.5 only presents the number of cancellations, because the high amount of cancellations makes it difficult to compare the other objectives. Disruption 4 is a different case, as its initial schedule already has 0 cancellations. However, the output of the program of disruption 4 shows that the Two-Opt Duty Neighborhood terminates after a small number of iterations and does not cause a significant improvement in the objective value. We can conclude that the Two-Opt Duty on itself does not perform well.

Second, we look into the performance of the Composition Model Neighborhood on itself. A big difference compared to the Two-Opt Duty Neighborhood is that the Composition Model Neighborhood is able to find the same number of cancellations for each disruption and for each parameter instance as in Table 5.3. Therefore, we take a look at the value of the rescheduling objectives. In Figure 5.6, we see that the combination of the start heuristic and Composition Model Neighborhood performs worse than the model with the start heuristic and both neighborhoods. The red line is also above the gray lines that correspond to filling up the empty duties chronologically and by length (see Figure 5.5). We can therefore state that the Two-Opt Duty and Composition

Model Neighborhood complement each other because the combination of the two neighborhoods leads to a significant improvement in performance in comparison to the performance of the separate neighborhoods.

5.6.3 MIP gap tolerance

In this section, we study whether allowing an optimality gap in the MIP of the sub problem in the heuristic could help reduce the computation time without negatively affecting the objective value. We do this by running the heuristic without a MIP gap and with a MIP gap of 0.001, 0.01, 0.1, 0.25, and 0.5. Table 5.6 shows the change in run time and Figure 5.7 gives an overview of the objective value of the difference MIP gaps. Furthermore, we add that the different MIP gaps cause a very insignificant change in operational cost.

In Table 5.6, we see that allowing a MIP gap of 0.25 decreases the run time of the heuristic most. Moreover, we see in Figure 5.7 that all lines are very close together. This means that a MIP gap does not cause a significant change in the objective value. Concluding, allowing a MIP gap of 0.25 reduces the computation time without negatively affecting the objective value.

Table 5.6: Run times for several MIP gaps in the Composition Neighborhoods

	Run times (s) for MIP gap					
Parameter instance	**0**	**0.001**	**0.01**	**0.1**	**0.25**	**0.5**
1	59.8	51.7	55.9	53.4	53.0	63.0
2	58.4	54.6	57.6	54.9	54.5	66.1
3	54.9	50.3	51.6	52.8	48.8	59.8
4	56.2	53.6	56.1	53.6	49.4	54.2
5	62.0	60.2	58.5	59.3	53.1	54.6
6	64.2	61.0	58.6	59.7	52.0	53.6
7	92.3	91.6	83.5	89.5	78.1	76.6
Average	64.0	60.4	60.3	60.4	55.6	61.1
Reduction wrt no MIP gap		-6%	-6%	-6%	-13%	-4%

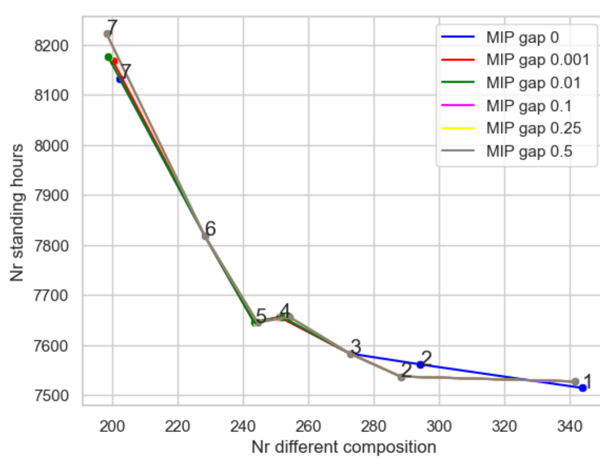

Figure 5.7: Seating shortage versus the number of different compositions for the 6 different MIP gaps. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

5.7 Heuristic performance

In this section, we compare the performance of the heuristic to the exact model. Section 5.6 has shown that the following settings for the heuristic lead to the best results: the Composition Model neighborhood has a MIP gap of 0.25 and the start heuristic fills the empty duties by the number of canceled trips. We compare the two methods by, again, plotting the average seating shortage against the number of different compositions and the operational costs for the seven parameter instances of Table 5.2. This results in Figure 5.8. Moreover, we compare the running time of the two approaches and present the optimality gap in Table 5.7. Finally, we provide all results of the heuristic in Table C.1 in the Appendix.

We can make several interesting observations from Figure 5.8 and Table 5.7. First of all, we observe in Figure 5.8(a) that the heuristic performs worse than the Composition Model in terms of standing hours versus the number of different compositions because the red line has a significant distance from the blue line. Moreover, we observe that the solutions for the different parameter instances are much closer to each other in the heuristic. This shows that the heuristic has trouble getting away from its starting solution and, again, supports the hypothesis that the heuristic gets stuck in a local minimum.

Second, we take a look at Figure 5.8(b). This graph shows that for the simulated disruptions, the operational costs for the heuristic are on average lower than in the Composition Model. We take a look at the average of each operational objective in Table C.1 such that we can draw general conclusions about the operational performance of the heuristic. We see that, on average, all operational objectives except for the number of additional required conductors and carriage kilometers improve. Moreover, we see that the heuristic never causes a type change. This can

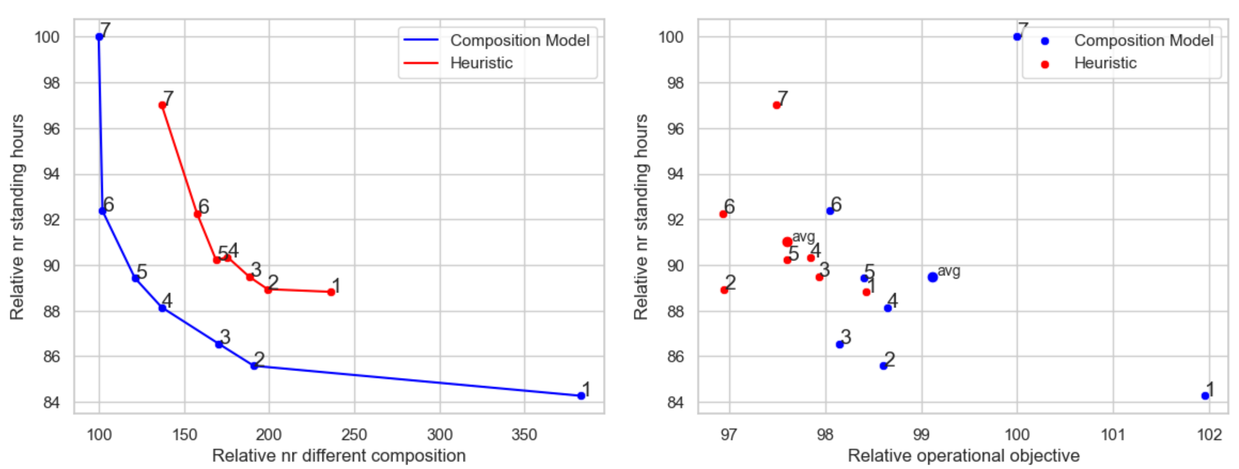

((a)) Relative seating shortage vs relative nr different compositions

((b)) Relative seating shortage vs relative operational costs and the average over the parameter instances

Figure 5.8: Comparison of the Composition Model and the heuristic. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

Table 5.7: Running time of the Composition Model and heuristic, the decrease in running time and the average optimality gap over the disruptions. The gap is the relative difference in the total objective value with the weights as in Table 5.2 for each parameter instance

	Parameter instance						
	1	**2**	**3**	**4**	**5**	**6**	**7**
Run time Composition Model	405	242	233	189	168	137	110
Run time Heuristic	53	55	49	49	53	52	78
Decrease in running time	87.0%	77.4%	79.0%	73.9%	68.4%	62.0%	28.9%
Average optimality gap	2.2%	2.1%	1.9%	2.0%	2.8%	3.3%	15.4%

be explained by the fact that the Two-Opt Duty neighborhood only performs one swap at a time and that it is difficult to find a single swap between different types that remains feasible at later couplings. Moreover, this proves that implementing an improvement of each general type in the Composition Model Neighborhood does not impact the final solution because the general types are never exchanged in the heuristic.

Third, we see a significant decrease in computation time for parameter instances 1 until 6. For these instances, the heuristic finds a solution within a minute. Moreover, we see that in the heuristic, the difference in running time between the parameter instances is much smaller. This might be due to the fixed structure of the heuristic which causes most iterations to take place regardless of which parameter instance is applied. We see that the running time of the heuristic for parameter instance 7 is higher than the other instances. This shows that the heuristic struggles in rescheduling with stick-to-the-plan as the main objective. This might be caused by the fact that the heuristic only makes a few changes at a time. Together with the fact that the Composition Model has the lowest running time for parameter instance 7, this results in a relatively low decrease in running time.

Finally, we compare the overall performance by looking at the optimality gap in Table 5.7. We observe an average gap of between 1.9% and 3.3% for parameter instances 1 until 6. This is a reasonable gap if we consider the high reduction in computational time for these parameter instances. Table 5.7 shows that the objective gap deviates per parameter instance and disruption. We see that for parameter instance 3, the gap is lowest and that the gap grows as we move further away from this instance. This could occur for (a combination of) two reasons: first, the start heuristic results in a start solution that has the best fit with the parameter instances that are in the middle and the VND struggles with moving away from this initial solution. Second, the neighborhoods have a tendency to result in solutions belonging to the middle parameter instances. These might also be the reasons for the high average optimality gap of parameter instance 7. However, it is likely that this effect is amplified by the high penalty on deviations from the original plan, which causes a small increase in the number of different compositions to result in a relatively large optimality gap.

We can conclude that the heuristic gives a significant reduction in running time at the cost of the total objective value. This expresses itself in an increase in the number of standing hours and more deviations from the original plan. The operational objectives, however, improve slightly for the simulated disruptions.

6 | Conclusion

In this conclusion, we first discuss the answers to the sub research questions as stated in the Introduction. The answers to these questions summarize and conclude the report and answer our main research question. Second, we discuss whether we achieved our goal as stated in Section 1.3, which is to develop a model for rolling stock rescheduling after a disruption while taking into account the dynamic passenger flows.

We answer the first research question, *"How do the passenger flows change after a disruption?"*, by modeling passenger flow on a directed graph. In this graph, the nodes represent either the departure or the arrival of a train. The arcs connect these nodes and represent either the course of a trip, a transfer from one train to another, or a train waiting at a station. This graph transfers the given origin-destination demand to trip demand by finding the shortest path through the graph under the assumptions on passenger behavior as provided in Section 3.1. From the results on the dynamic passenger flow model in Section 5.2, we can draw four main conclusions: first, we see that the passenger demand increases on detour routes. Second, we see that in the half hour after the disruption, the passenger demand on the disrupted route increases. Third, we see that the passenger flow decreases at the routes adjacent to the begin- and end station of the canceled route. Fourth, the passenger flow change can differentiate between the two directions of the same route.

The second research question is as follows: *"How can we reschedule rolling stock to account for timetable changes and changing passenger flows?"* We applied two methods for rolling stock rescheduling in this thesis. First, we applied an exact method from literature, namely the extended Composition Model as given in Nielsen (2011). Second, we developed a heuristic. Both models use the passenger demand as obtained from the passenger flow model. Moreover, the models make use of sets of trips and transitions that are adapted to the disruption.

The developed heuristic consists of a start heuristic and two neighborhoods: the Two-Opt Duty Neighborhood and the Composition Neighborhood. The main idea of the start heuristic is to find the best possible initial solution, by taking the following steps: first, the start heuristic cuts off duties that contain canceled trips. Next, the remaining parts of these canceled duties are, when possible, filled by available vehicle units. Section 5.6.1 shows that it is best to first fill the remaining part of the duties that contain most canceled trips. Furthermore, the Two-Opt Duty Neighborhood finds the best swap of the remaining parts of two duties if the corresponding trains are at the same station at the same time. The Composition Neighborhood optimizes the assignment of one vehicle unit type while keeping the assignment of the other types constant. The two neighborhoods complement each other, because the Composition Neighborhood takes care of optimization within one vehicle unit type, while the Two-Opt Duty Neighborhood looks for an optimal exchange of units between the different types. This is shown in a computational experiment in Section 5.6.2. The idea is that this heuristic makes use of the computational benefits of the Composition Model while maintaining the flexibility of a heuristic.

The Composition Model is used to answer the following research question: *What is the influence of taking into account dynamic passenger flows in rolling stock rescheduling on passenger comfort, operational costs, and deviations from the original schedule?"* We answer this question from two angles. Firstly, we have shown in Section 5.4 that there exists a trade-off between passenger comfort and stick-to-the-plan. Moreover, we showed that rescheduling with a large focus on either stick-to-the-plan or passenger comfort results in high operational costs. Secondly, Section 5.5 shows that rolling stock rescheduling with dynamic passenger flows instead of static passenger flows is beneficial for passenger comfort and mainly comes at a cost of more deviations from the original plan. Concluding, taking into account dynamic passenger flows along with stick-to-the-plan and operational costs results in a decrease in standing hours and operational costs at a cost of an increase in deviations from the original plan. Moreover, we have shown that the current practice of stick-to-the-plan rescheduling leads to as well a high amount of standing hours as an increase in operational costs.

The final sub research question is *"How does the developed model perform in terms of passenger comfort, computation time, operational costs, and deviations from the original plan in comparison to methods from literature?"* As seen in Section 5.7, the biggest improvement of the heuristics compared to the Composition Model from literature is in the reduction of the computational time. The computation time decreases by 62% to 87% for the parameter instances that do not solely focus on stick-to-the-plan rescheduling. This reduction in running time is impressive, as the running time of the heuristic from Hoogervorst et al. (2021) without flexible turning was not lower than the Composition Model. The reduction in running time for these instances comes at the cost of the total objective value with an optimality gap of between 1.9% and 3.3%. This increase in total objective seems reasonable when we consider the high decrease in computation time. As we see in Figure 5.8(a) the increase in total objective value expresses itself in an increase in the number of standing hours and more deviations from the original plan. The operational objectives, however, improve slightly for the simulated disruptions.

Finally, we conclude that we have reached the goal that was set in Section 1.3 because the developed heuristic reschedules rolling stock while taking into account the objectives of Section 1.2 and the dynamic passenger flows. The main contribution of the developed heuristic is the significant decrease in computation time in comparison to the method from literature. Moreover, the advantages of the heuristic in comparison to the exact Composition Model of Nielsen (2011) are the increased flexibility to incorporate real-life practicalities and the possibility of providing intermediate results. The disadvantages are the decrease in passenger comfort and the increase in deviations from the original schedule. Finally, we have shown in computational experiments on the exact Composition Model that the current practice of stick-to-the-plan rescheduling leads to a high amount of standing hours as well as an increase in operational costs.

7 | Recommendations

This section provides suggestions for future research. Firstly, we look into general suggestions that concern the problem setting and the computational experiments. Secondly, we look into possible extensions of the passenger flow model. Thirdly, we provide suggestions on rolling stock rescheduling. This final section provides ideas to relax the assumptions of both rescheduling models and provides suggestions to improve the heuristic.

7.1 Problem setting and computational experiments

The computational experiments were limited to considering only one disruption at a time with a known duration. It is interesting for future research to look into the impact of multiple disruptions at the same time on the passenger flow and quality of rolling stock rescheduling. The applied models are suitable for such problem settings, as they are able to reschedule rolling stock for any adapted timetable. Moreover, to address disruptions with unknown duration, one could apply the rolling stock rescheduling methods of this thesis in a rolling horizon framework. Such a framework was introduced by Nielsen et al. (2012).

In addition, it would be valuable to conduct computational experiments on more than four disruptions. We suggest generating disruptions on more routes, at more different times, and with different durations. This provides more insight into the average performance of the two rolling stock rescheduling methods. For example, it could determine whether the reduction in operational costs in the heuristic is coincidental to the simulated disruptions or generally true for the developed heuristic. Finally, a future study could apply the rolling stock rescheduling methods to all vehicle types.

7.2 Passenger flow

The passenger flow model in this thesis relies on a list of assumptions on passenger behavior. In order to include more factors on passenger behavior, one could apply a multinomial passenger flow model that takes into account historical data on passenger behavior.

Moreover, Assumption 3 implies that the capacity of rolling stock is not considered in determining the passenger flows. We could take this capacity into account by imposing a limit on the flow on each trip and redetermine the shortest path for (a part of) the passengers that wanted to travel on a trip that has reached its passengers' capacity. In rescheduling, taking into account the rolling stock capacity in determining the passenger flows requires iterating between the rolling stock rescheduling and the passenger flow model

In addition, Assumption 4 states that the passengers know the duration of the disruption in advance. In reality, this is often not the case. This could be taken into account in the model by applying the following steps: first, we determine the passenger flow in the model before the disruption. Second, we insert the disruption from the start time of the disruption to the end of the planning horizon and determine the shortest path for each passenger group. Third, we insert the end time of the disruption and recalculate the shortest path for each passenger group from their location at the end time of the disruption to their final destination.

Finally, the passenger flow model can be extended by applying it to the whole network. In this case, the origin-destination demand from and to all stations is taken into account, resulting in a more accurate representation of the passenger flow.

7.3 Rolling stock rescheduling

First of all, we look into the assumptions of the rolling stock models. Both models were restricted to only be able to couple one vehicle unit to the front of the train and decouple one vehicle unit from the back. In reality, more composition changes are allowed, which could be taken into account in future research. It is interesting to study the influence of allowing more composition changes on the computation time of the rolling stock rescheduling methods. We expect that the influence of the computation time of the Composition Model is greater than on the heuristic. This is expected because the problem size of the MIP of the Composition Model grows, but the problem size of the start heuristic and Two-Opt Neighborhood in the heuristic does not, as it depends on the number of trips, duties, and transitions. The problem size of the Composition Neighborhood also increases, but because it is performed iteratively with the Two-Opt Duty Neighborhood and solves a smaller problem, we expect the influence of the added composition changes on the computational time of the heuristic to be smaller. It would be interesting to study this hypothesis in future research because this would also confirm the argument that the heuristic is more flexible. Another argument for allowing more composition changes is that it would allow us to apply the model to the initial real-life rolling stock schedule. This would ensure that the computational experiments are closer to reality.

Furthermore, we provide suggestions to improve the heuristic. First of all, the computation time could be further decreased by running the Composition Model Neighborhood on the three general types in parallel on different cores of a computer. Furthermore, future studies could look into ideas for escaping the local minimum. A well-known method to escape a local optimum is shaking. Shaking implies that we would apply changes to the solution that in the short term decrease the objective value but in the long term allow us to find a better objective value as we escape from the local minimum. For example, the Two-Opt Duty Neighborhood could be applied for shaking. Another suggestion to escape the local minimum is to look into new neighborhoods for both local search and shaking. A suggestion for a neighborhood to be studied is Flexible Turning. It was already shown by Nielsen (2011) that adding flexible turning can decrease the objective value of the Composition Model but leads to significantly higher computation times. It is interesting to study how flexible turning can contribute to lowering the objective of our heuristic and how it affects computation time.

8 | Bibliography

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

Cadarso, L., Marín, Á., & Maróti, G. (2013). Recovery of disruptions in rapid transit networks. *Transportation Research Part E: Logistics and Transportation Review*, *53*, 15–33.

Caprara, A., Kroon, L., Monaci, M., Peeters, M., & Toth, P. (2007). Passenger railway optimization. *Handbooks in operations research and management science*, *14*, 129–187.

Dollevoet, T., Huisman, D., Kroon, L. G., Veelenturf, L. P., & Wagenaar, J. C. (2017). Application of an iterative framework for real-time railway rescheduling. *Computers & Operations Research*, *78*, 203–217.

Dumas, J., Aithnard, F., & Soumis, F. (2009). Improving the objective function of the fleet assignment problem. *Transportation Research Part B: Methodological*, *43*(4), 466–475.

Dumas, J., & Soumis, F. (2008). Passenger flow model for airline networks. *Transportation Science*, *42*(2), 197–207.

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

Haahr, J. T., Lusby, R. M., Larsen, J., & Pisinger, D. (2014). A branch-and-price framework for railway rolling stock rescheduling during disruptions. *DTU Management Engineering*.

Haahr, J. T., Wagenaar, J. C., Veelenturf, L. P., & Kroon, L. G. (2016). A comparison of two exact methods for passenger railway rolling stock (re) scheduling. *Transportation Research Part E: Logistics and Transportation Review*, *91*, 15–32.

Hoogervorst, R., Dollevoet, T., Maróti, G., & Huisman, D. (2020). Reducing passenger delays by rolling stock rescheduling. *Transportation Science*, *54*(3), 762–784.

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

Jespersen-Groth, J., Potthoff, D., Clausen, J., Huisman, D., Kroon, L., Maróti, G., & Nielsen, M. N. (2009). Disruption management in passenger railway transportation. In *Robust and online large-scale optimization* (pp. 399–421). Springer.

Kroon, L., Maróti, G., & Nielsen, L. (2015). Rescheduling of railway rolling stock with dynamic passenger flows. *Transportation Science*, *49*(2), 165–184.

Liu, Y., Liu, Z., & Jia, R. (2019). Deeppf: A deep learning based architecture for metro passenger flow prediction. *Transportation Research Part C: Emerging Technologies*, *101*, 18–34.

Lusby, R. M., Haahr, J. T., Larsen, J., & Pisinger, D. (2017). A branch-and-price algorithm for railway rolling stock rescheduling. *Transportation Research Part B: Methodological*, *99*, 228–250.

Nielsen, L. K. (2011). *Rolling stock rescheduling in passenger railways: Applications in short-term planning and in disruption management* . PhD thesis Erasmus University Rotterdam.

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

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

Wagenaar, J. C., Kroon, L. G., & Schmidt, M. (2017a). Maintenance appointments in railway rolling stock rescheduling. *Transportation Science*, *51*(4), 1138–1160.

Wei, Y., & Chen, M.-C. (2012). Forecasting the short-term metro passenger flow with empirical mode decomposition and neural networks. *Transportation Research Part C: Emerging Technologies*, *21*(1), 148–162.

Appendices

A | Parameters composition model

Table A.1: Rolling stock rescheduling composition model parameters results

Disruption	Parameter instance	Run time (s)	Total standing hours	New couplings	New uncouplings	Off-balances	Extra conductor	Different type	Carriage km	Canceled trips	Different comp
1	1	661	6664	40	56	6	66	28	1225636	0	553
1	2	245	6817	31	42	4	30	6	1221449	0	230
1	3	245	6824	31	45	4	55	6	1224884	0	263
1	4	248	6926	24	39	8	41	12	1221193	0	224
1	5	216	7098	24	37	10	15	12	1216776	0	180
1	6	147	7530	25	39	10	28	6	1218757	0	168
1	7	99	8266	21	34	18	24	6	1214577	0	161
2	1	269	9162	27	35	4	25	18	1225537	0	319
2	2	211	9299	19	25	4	11	0	1227674	0	114
2	3	198	9347	18	24	4	0	0	1226494	0	94
2	4	125	9357	17	23	4	8	0	1226230	0	90
2	5	112	9448	16	22	6	0	6	1224600	0	65
2	6	104	9448	16	22	6	0	6	1224600	0	70
2	7	123	9519	18	24	6	2	6	1224676	0	68
3	1	342	6280	42	51	4	72	35	1219385	3	952
3	2	200	6358	42	45	6	73	9	1220732	3	540
3	3	199	6351	42	50	4	97	6	1223224	3	492
3	4	174	6596	37	46	8	85	9	1220696	3	379
3	5	160	6769	35	40	8	78	9	1218240	3	357
3	6	166	6908	34	38	10	59	0	1217518	3	297
3	7	134	8558	28	36	18	44	6	1209572	3	301
4	1	353	6460	26	41	4	72	90	1230333	0	393
4	2	310	6537	20	31	6	29	56	1227916	0	222
4	3	291	6819	17	26	6	33	0	1228088	0	137
4	4	209	7001	18	26	6	26	7	1225332	0	100
4	5	185	7000	18	26	6	26	7	1225332	0	100
4	6	130	7426	11	20	8	0	6	1221190	0	57
4	7	83	7556	10	19	10	1	6	1220932	0	49
Average			7583	25.25	34.36	7.07	35.71	12.79	1222556	0.75	249.11

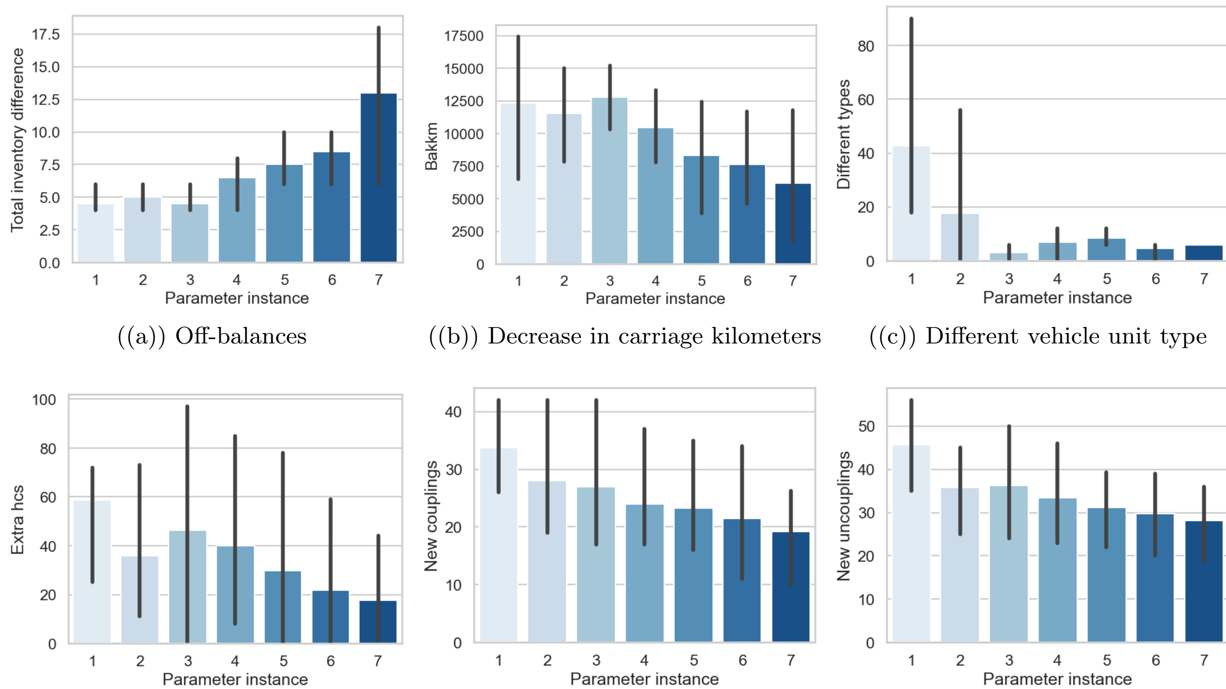

((a)) Off-balances ((b)) Decrease in carriage kilometers ((c)) Different vehicle unit type

((d)) Additional Conductors ((e)) New couplings ((f)) New uncouplings

Figure A.1: Bar graphs of the average values over the four disruptions for the seven parameter instances of the separate operational objectives: off-balances, decrease in carriage kilometers, number of different vehicle types, additional conductors, new couplings, and new uncouplings

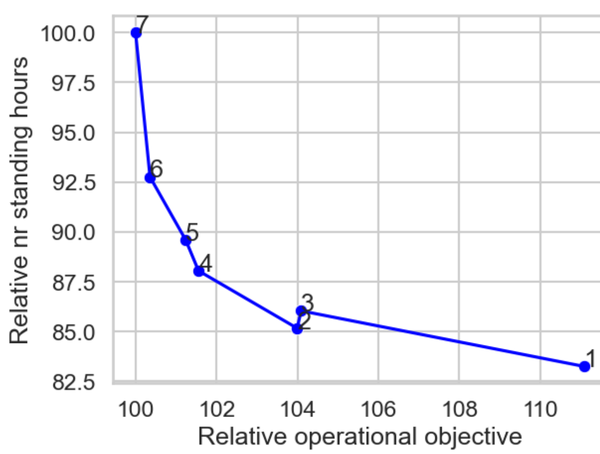

Figure A.2: Results of the Composition Model for rolling stock rescheduling if off-balances are not included in the objective. The figure shows the average value over the four disruption for each of the seven parameter instances as provided in Table 5.2

B | Rolling stock rescheduling with dynamic vs static passenger flow

Table B.1: Rolling stock rescheduling with static passenger flow results

Disruption	Parameter instance	Run time (s)	Total standing hours	New couplings	New uncouplings	Off-balances	Extra conductor	Different type	Carriage km	Canceled trips	Different comp
1	1	343	6891	36	46	4	57	40	1223869	0	432
1	2	257	7051	30	41	4	37	6	1220802	0	266
1	3	236	7141	26	39	4	43	6	1221083	0	214
1	4	227	7141	26	39	4	43	6	1221081	0	212
1	5	176	7280	26	38	6	0	6	1216829	0	146
1	6	97	7530	25	39	10	28	6	1218757	0	168
1	7	62	8266	21	34	18	24	6	1214577	0	161
2	1	277	9261	22	27	4	23	20	1227024	0	125
2	2	197	9321	18	23	4	0	0	1226224	0	0
2	3	119	9357	18	23	4	7	2	1226548	0	90
2	4	86	9448	17	22	6	2	8	1224918	0	73
2	5	73	9448	17	22	6	0	8	1224918	0	73
2	6	62	9448	16	22	6	0	8	1224766	0	71
2	7	73	9538	17	22	6	0	6	1224676	0	68
3	1	291	6428	32	41	6	46	49	1214718	3	846
3	2	181	6488	35	37	6	52	36	1217444	3	526
3	3	180	6593	39	45	6	84	30	1219146	3	408
3	4	153	6676	36	42	10	61	9	1216708	3	380
3	5	137	6916	34	38	10	51	9	1214758	3	354
3	6	139	7089	33	37	12	64	9	1215738	3	343
3	7	91	8587	29	35	18	44	6	1209572	3	301
4	1	1072	7123	19	29	4	15	42	1226399	0	184
4	2	296	7175	17	28	4	35	19	1225134	0	150
4	3	219	7278	13	22	6	23	0	1225192	0	117
4	4	201	7460	13	22	6	16	7	1222436	0	76
4	5	168	7460	13	22	6	15	7	1222436	0	75
4	6	85	7556	10	19	8	1	6	1220854	0	70
4	7	40	7556	10	19	10	1	6	1220932	0	69

C | Results of the heuristic for rolling stock rescheduling

Table C.1: Rolling stock rescheduling heuristic results

Disruption	Parameter instance	Run time (s)	Total standing hours	New couplings	New uncouplings	Off-balances	Extra conductor	Different type	Carriage KM	Canceled trips	Different comp
1	1	60	7305	37	52	6	68	0	1233233	0	440
1	2	65	7330	30	44	6	58	0	1226967	0	342
1	3	66	7432	30	46	8	71	0	1230536	0	322
1	4	59	7626	29	42	10	66	0	1230142	0	303
1	5	59	7707	27	38	8	61	0	1228456	0	289
1	6	67	7707	27	38	8	61	0	1228456	0	289
1	7	146	7677	26	36	10	35	0	1221551	0	259
2	1	44	9436	20	26	0	24	0	1227546	0	94
2	2	39	9457	19	25	0	17	0	1227162	0	80
2	3	39	9475	18	24	4	11	0	1227900	0	55
2	4	39	9486	17	23	4	11	0	1227636	0	51
2	5	40	9486	17	23	4	11	0	1227636	0	51
2	6	44	9468	17	23	0	11	0	1227470	0	40
2	7	43	9558	18	23	0	11	0	1227470	0	38
3	1	67	6486	30	36	10	100	0	1219166	3	717
3	2	77	6455	31	36	8	66	0	1219172	3	651
3	3	56	6519	26	35	8	72	0	1217354	3	634
3	4	63	6607	23	33	8	60	0	1215200	3	579
3	5	75	6482	26	31	8	70	0	1219076	3	554
3	6	65	7027	22	25	12	49	0	1217166	3	516
3	7	87	8285	25	31	16	73	0	1215362	3	436
4	1	41	6882	18	30	2	23	0	1227492	0	116
4	2	37	6907	16	24	2	20	0	1228164	0	80
4	3	34	6907	16	24	2	20	0	1228164	0	80
4	4	36	6907	15	25	2	26	0	1228580	0	84
4	5	38	6907	15	25	2	26	0	1228580	0	84
4	6	33	7074	14	22	2	18	0	1226756	0	68
4	7	37	7367	12	20	2	10	0	1225268	0	60
Average			7713	22	31	5	41	0	1225274	0.75	261
Percentual change wrt composition model			1.7%	-12.2%	-10.6%	-23.2%	14.9%	-100.0%	0.2%	0.0%	4.8%