A Case Study on Large Vehicles Scheduling for Railway Infrastructure Maintenance: Modelling and Sensitivity Analysis

Research Paper

Jannes Glaubitz[1], Thomas Wolff[1], Henry Gräser[1], Philipp Sommerfeldt[1], Julian Reisch[2], David Rößler-von Saß[1], and Natalia Kliewer[1]

[1] Freie Universität Berlin, Department of Business Information Systems, Berlin, Germany
jannes.glaubitz@fu-berlin.com
[2] Synoptics GmbH, Dresden, Germany
julian.reisch@synoptics.de

Abstract. Railway infrastructure maintenance is a critical operational task that must balance efficiency, cost, and minimal disruption to train services. This paper presents an optimization-driven approach for scheduling large maintenance vehicles within a preventive maintenance framework.
Our computational study, based on real-world data from Deutsche Bahn, evaluates key factors influencing maintenance scheduling efficiency, including fleet size, scheduling flexibility, and depot distance constraints. To that end, a greedy heuristic and a Mixed Integer Programming (MIP) model are applied to solve the complex underlying optimization problem. Sensitivity analyses reveal that container reusability and adjustments in machine mobility have a more significant impact on demand fulfillment than simply increasing fleet size.
By leveraging heuristic optimization, our study offers actionable insights for strategic decision-making in railway maintenance planning. The results emphasize the importance of adaptive scheduling policies, improved operational management, and optimized resource utilization to enhance maintenance efficiency while minimizing operational disruptions.
Keywords: Railway Track Maintenance Planning, Maintenance Track Possession Problem, Operations Research, Mixed Integer Programming

1 Introduction

Focus of recent public debate in Germany is the maintenance, restoration, and expansion of the existing railway infrastructure. The potentially resulting track maintenance disruptions significantly impact the efficiency and reliability of rail transport for both passenger and freight services. Ongoing renewal and maintenance activities often result in delays and capacity restrictions, reducing the overall attractiveness of the railway system. Recognizing this challenge, the Scientific Advisory Board of the Ministry of Transport has emphasized the need for stability and reliability in rail infrastructure planning (Eckstein et al. 2024).

To mitigate these issues, Deutsche Bahn has radically changed their policy for planning railway construction activities by introducing a timed blockade system aimed

20th International Conference on Wirtschaftsinformatik
September 2025, Münster, Germany

at stabilizing the railway network (DB InfraGO 2024). This system utilizes predefined time windows, known as containers, to allocate dedicated track closures for maintenance tasks. By clustering maintenance activities within these standardized slots, disruptions are minimized, and operational efficiency is enhanced. A key advantage of this approach is that these time windows are integrated into the annual timetable, enabling short-notice maintenance without requiring extensive manual re-planning. In our work, we deal with two types of containers. The first type, the so-called *track containers*, is predetermined and used to schedule ad hoc maintenance during the year. Preventive maintenance, which is done by a limited number of heavy machinery, has to be scheduled inside *preventive containers*. Since those types of machines cannot be routed easily through the network, one has to explicitly predetermine the schedule in which preventive maintenance will be carried out.

This study examines the development and utilization of preventive containers, which are specifically designated for heavy machinery performing preventive maintenance such as milling, grinding and tamping. Unlike standard track containers (Reisch et al. 2024), which are planned independently, preventive containers are strategically aligned with existing maintenance schedules. This coordination allows for systematic task allocation, promoting a more efficient and proactive approach to infrastructure upkeep. The problem of scheduling preventive containers is introduced in Reisch et al. (2025).

Railway operation is characterized by highly integrated multi-stakeholder, multi-stage planning and scheduling problems, such as the maintenance scheduling problem in this work. Solving such problems requires the integration of vast, various data sources and the application of sophisticated information systems. However, neither in practice nor in theory do we find the required level of technological and organizational planning. In reality, finding suitable containers and their use in conducting maintenance activities are treated as separate planning problems, solved in different units. By conducting a sensitivity analysis of the maintenance scheduling problem, and deriving recommendations for previous planning strategies, this work attempts to bridge the resulting gap, providing a blueprint for the required integration of these planning stages.

1.1 Related work

Based on the classification of Lidén (2015) and the corresponding taxonomy by Sedghi et al. (2021) the described optimization problem incorporates elements from multiple planning problems within the field of railway track maintenance planning and scheduling (RTMP&S), not only from tactical but also from operational planning.

In a preceding work by Reisch et al. (2024), the authors aim at optimally planning the track containers with regard to minimizing train interference. Reisch et al. (2025) follow this research line by tactically scheduling the machines for preventive containers, so that a maximum of demands can be fulfilled. They classify the problem as a Track Maintenance Possession Problem (TMPP) and solve it as a Maximum Satisfiability problem. The paper at hand tackles the very same problem, but with emphasis on the exploration of organizational aspects of the use case.

Early work on the TMPP can be found in Budai et al. (2006). Overall, TMPP focuses on identifying appropriate time windows for track possessions on specific track segments, emphasizing the synergistic effects between timetabling capacity and the execution of maintenance activities. To be more precise regarding its correct positioning within the research, our work does not include or analyze the effects on timetabling, nor does it search for the optimal schedule of possessions in the rail network. Instead, it focuses on long-term vehicle routing with specific rules for the creation of possessions using the container concept. A related line of research addresses the long-term crew scheduling in railway maintenance known as the Production Gang Scheduling Problem (PGSP) or Curfew Planning Problem (CPP). While sharing domain-specific aspects such as restrictions on simultaneous maintenance on high-traffic railway lines, our setting differs fundamentally by focusing on machines rather than crews (Peng & Ouyang 2012, Nemani et al. 2010).

1.2 Contribution

While Reisch et al. (2025) present an exhaustive study on computability aspects and algorithms on this particular problem, the paper at hand seeks to answer the following research questions aiming at the tactical and strategic aspects as proposed in Eckstein et al. (2024).

RQ1: How well are predetermined containers and available rolling stock suited to fulfill preventive maintenance demands on the railway infrastructure?

RQ2: How and to what extent can demand fulfillment for preventive railway maintenance be increased by adjusting key limiting factors such as fleet size, scheduling flexibility, and depot distance constraints?

RQ3: How should the general technological and organizational conditions be adapted to better reflect actual preventive maintenance demands and practical requirements in order to increase operational efficiency in the maintenance scheduling process?

1.3 Structure

The remainder of this paper is structured as follows: Section 2 introduces the planning process and its optimization problems in the context of preventive railway maintenance. It also defines the specific planning problem addressed in this study. Section 3 outlines the solution approach applied in the case study as well as the design of the sensitivity experiments. Section 4 details the experimental setup using real-world data and presents the results of both the baseline scenario and the sensitivity analysis. Section 5 discusses the results in relation to the research questions and derives practical insights that can be fed back to practitioners and into the planning processes. Finally, Section 6 summarizes the key findings and outlines potential directions for future research.

2 Preventive railway maintenance

Railway maintenance planning is typically divided into strategic, tactical, and operational levels, depending on how far in advance the planning occurs (Lidén 2015). Strategic planning happens years ahead, tactical planning up to one year before execution, and operational planning covers the short-term phase from hours to a few weeks. Since the planning problem here takes place before the annual timetable is drawn up, the planning of preventive maintenance activities falls into the scope of tactical planning. The planning of preventive maintenance is embedded in the scheduling of track containers and other operational aspects, including work timing and concrete resource usage. Therefore, solving and investigating the long-term maintenance vehicle routing problem and specifically the underlying organizational structures, like the preventive containers, can result in the identification of certain synergies between the planning stages.

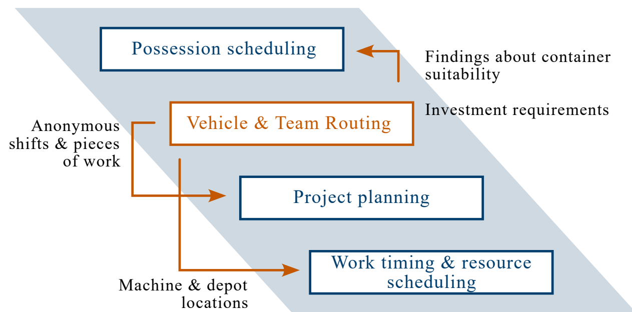

Figure 1. Railway maintenance planning stages addressed in this paper, based on planning problems identified in Lidén (2015)

2.1 Problem description

In this problem, the total length of completed maintenance demands N in a railway network is maximized by scheduling a set of machines V to perform maintenance tasks over the course of a fixed planning horizon. To achieve this goal, several complex characteristics have to be respected: Exclusive use of track sections for every machine, a heterogeneous fleet of vehicles, restrictions regarding the relocation between different depots and accommodating regular railway operations.

Containers C are abstract spatio-temporal units that represent a mono-directional blockage of sections on a railway track at certain time intervals. While data for the fixed track containers is given, we can take the spatial information of each container and plan new ones for preventive maintenance. The use of these preventive containers is subject to extensive restrictions. Containers can overlap such that they share a set of track sections. For each container $c \in C$, the set O_c denotes all containers the overlap with c. As soon as a maintenance task is performed in a container, it automatically blocks all containers that overlap with the corresponding section of the track for other machines during the blocking period. The blocking times within a container are generally eight hours per night, corresponding to the maximum shift length h, in order to minimize the impact on regular train services. Additionally, to ensure the continuous operability of the rail network, containers are not allowed to block too many routes of central connections. These central connections are organized in a set of so-called corridors K. We define K^{rst} as a set of diversion sets and for each of those sets, at least one corridor must remain free in order to maintain ongoing train operations.

The planning horizon is subdivided into periods P. Within a period, a container can only be blocked a fixed number of times in order to avoid overloading the rail network. A key difficulty is the exact allocation of maintenance requirements to suitable machines, as each requirement has specific parameters in terms of maintenance depth and maintenance type. In this planning problem, a fleet of vehicles with varying capacities and technical specifications must be scheduled to perform different types of tasks. These tasks differ in their general nature, such as milling, grinding or tamping. Additionally, tasks categorized as grinding vary further based on the required depth of material removal. Each type of grinding machine has a predefined maximum grinding depth, which directly influences the speed $vel_{v,n}$ at which maintenance tasks can be completed. Central elements of maintenance planning are the depots S. In each period, every machine has to choose a starting depot $dep_{v,p}$, from which it begins its long-term route through the railway network, with respect to two distance thresholds: A weekly distance r for every depot change as well as the distance a machine can move away from its starting depot during one period R. While Reisch et al. (2025) constructed shifts with valid movements through the network for 8-hour machine shifts, this study abstracts this process by allowing a machine to cover all demands within a container in a single night. Trips from and to depots (pull-out, pull-in trips) are ignored. This approach does not account for real spatial constraints within the container or the distance a machine must travel between multiple maintenance demands. Additionally, a machine is permitted to perform different types of maintenance activities within a single shift.

2.2 MILP formulation

A mixed-integer programming (MIP) model was developed, as presented in Constraints (2.1) – (2.26). The objective function (2.1) maximizes the length of the maintained railway infrastructure. Constraints (2.2) – (2.4) ensure that each demand assigned within the planning horizon is executed by a technically suitable vehicle on a single day, while respecting vehicle capacities and the daily maximum working time.

Table 1. Sets, Parameters, and Decision Variables

Sets & Functions	**Description**
V	Set of vehicles
N	Set of preventive maintenance demands
N_c	Subset of preventive maintenance demands located in container $c \in C$
C	Set of preventive containers
C_k	Subset of containers located on corridor $k \in K$
O_c	Set of containers overlapping with container $c \in C$
K	Set of corridors
K^{rst}	Set of corridor diversion sets
S	Set of depots
$S_{s,dist}$	Set of depots within $dist$ kilometers of depot $s \in S$
S_c	Subset of depots that are reachable for container $c \in C$
P	Set of time periods
W	Set of weeks within the planning horizon
W_p	Set of weeks in period $p \in P$
D	Set of days in the planing horizon
D_p	Subset of days within period $p \in P$
D_c^{rst}	Subset of days when container $c \in C$ is restricted

Name	Description	Range/Unit
$cw(d)$	Function that returns the calendar week of day d	
Parameters	**Description**	**Range/Unit**
l_n	Length of demand $n \in N$	[m]
$vel_{v,n}$	Speed for $v \in V, n \in N$	$[\frac{m}{h}]$
$dep_{v,p}$	Starting depot for all vehicles $v \in V$ in period $p \in P$	
ω_{0_p}	Starting week of period $p \in P$	[week]
r	Weekly travel range to a new depot	[m]
R	Max. allowed distance from the start depot per period	[m]
h	Max. operational duration of a machine per day	[h]
$\delta_{n,v}$	1, if vehicle $v \in V$ fits the technical requirement of demand $n \in N$; else 0	{0,1}
Decision Variables	**Description**	**Range/Unit**
$x_{n,v,d}$	Demand $n \in N$ is fulfilled by vehicle $v \in V$ on day $d \in D$	$\{0,1\}$
y_n	Demand $n \in N$ is fulfilled	$\{0,1\}$
$z_{c,n}$	Container $c \in C$ is reserved for demand $n \in N_c$	$\{0,1\}$
$u_{c,d}$	Container $c \in C$ is used on day $d \in D$	$\{0,1\}$
$w_{v,c,d}$	Vehicle $v \in V$ uses container $c \in C$ on day $d \in D$	$\{0,1\}$
$t_{k,d}$	Corridor $k \in K$ is used on day $d \in D$	$\{0,1\}$
$r_{v,s,\omega}$	Vehicle $v \in V$ uses depot $s \in S$ in week $\omega \in W$	$\{0,1\}$

MILP model

$$max \quad \sum_{n \in N} y_n \cdot l_n \tag{2.1}$$

s.t.

$$y_n \le \sum_{v \in V} \sum_{d \in D_p} x_{n,v,d}, \quad \forall n \in N, p \in P \tag{2.2}$$

$$\sum_{n \in N} \frac{l_n}{vel_{v,n}} \cdot x_{n,v,d} \le h, \quad \forall v \in V, p \in P, d \in D_p \tag{2.3}$$

$$x_{n,v,d} \le \delta_{n,v}, \quad \forall n \in N, v \in V, p \in P, d \in D_p \tag{2.4}$$

$$\sum_{c \in C} z_{c,n} = y_n, \quad \forall n \in N \tag{2.5}$$

$$w_{v,c,d} \ge x_{n,v,d} + z_{c,n} - 1, \quad \forall n \in N, v \in V, c \in C, p \in P, d \in D_p \tag{2.6}$$

$$\sum_{d \in D_p} u_{c,d} \le 1, \quad \forall c \in C, p \in P \tag{2.7}$$

$$\sum_{v \in V} w_{v,c,d} \le u_{c,d}, \quad \forall c \in C, p \in P, d \in D_p \tag{2.8}$$

$$\sum_{c' \in O_c} u_{c',d} \le 0, \quad \forall c \in C, d \in D_c^{rst} \tag{2.9}$$

$$\sum_{c \in C} w_{v,c,d} \le 1, \quad \forall v \in V, p \in P, d \in D_p \tag{2.10}$$

$$\sum_{\substack{v' \in V \\ v' \neq v}} \sum_{c' \in O_c} w_{v',c',d} \le 1 - w_{v,c,d}, \quad \forall v \in V, c \in C, p \in P, d \in D_p \tag{2.11}$$

$$|C_k| \cdot t_{k,d} \ge \sum_{c \in C_k} u_{c,d}, \quad \forall k \in K, p \in P, d \in D_p \tag{2.12}$$

$$\sum_{k \in div} t_{k,d} \le |div| - 1, \quad \forall p \in P, d \in D_p, div \in K^{rst} \tag{2.13}$$

$$\sum_{v \in V} r_{v,s,\omega} \le 1, \quad \forall s \in S, p \in P, \omega \in W_p \tag{2.14}$$

$$r_{v,\omega_{0_p},dep_{v,p}} = 1, \quad \forall v \in V, p \in P \tag{2.15}$$

$$\sum_{s \in S} r_{v,s,\omega} = 1, \quad \forall v \in V, p \in P, \omega \in W_p \tag{2.16}$$

$$\sum_{s \in S_c} r_{v,s,cw(d)} \ge w_{v,c,d}, \quad \forall v \in V, c \in C, p \in P, d \in D_p \tag{2.17}$$

$$r_{v,s,\omega} \le 0, \quad \forall v \in V, p \in P, \omega \in W_p \setminus \omega_{0_p}, s \notin S_{dep_{v,p},\mathrm{R}} \tag{2.18}$$

$$r_{v,s,\omega} + r_{v,s',\omega-1} \le 1, \quad \forall v \in V, p \in P, \omega \in W_p \setminus \omega_{0_p}, s \in S, s' \notin S_{s,r} \tag{2.19}$$

MILP model continued

$$x_{n,v,d} \in \mathbb{Z}_2 \qquad \forall n \in N, v \in V, d \in D \tag{2.20}$$

$$y_n \in \mathbb{Z}_2 \qquad \forall n \in N \tag{2.21}$$

$$z_{c,n} \in \mathbb{Z}_2 \qquad \forall c \in C, n \in N_c \tag{2.22}$$

$$u_{c,d} \in \mathbb{Z}_2 \qquad \forall c \in C, d \in D \tag{2.23}$$

$$w_{v,c,d} \in \mathbb{Z}_2 \qquad \forall v \in V, c \in C, d \in D \tag{2.24}$$

$$t_{k,d} \in \mathbb{Z}_2 \qquad \forall k \in K, d \in D \tag{2.25}$$

$$r_{v,s,\omega} \in \mathbb{Z}_2 \qquad \forall v \in V, s \in S, \omega \in W \tag{2.26}$$

Constraints (2.5) – (2.11) formalize the interactions between containers, vehicles, and demands by ensuring consistent assignment of demands to exactly one container and one vehicle (2.5, 2.6), limiting container usage to once per period (2.7), prohibiting usage on days with overlapping track possessions (2.9), restricting each vehicle to one container per day (2.6), and enforcing mutual exclusion of overlapping containers across vehicles (2.11). Constraints (2.12) and (2.13) ensure that at least one transit corridor remains open at all times. Constraints (2.14) – (2.17) govern depot usage by limiting each depot to one machine (2.14), enforcing unique depot assignments (2.15, 2.16), and restricting machine movement by a weekly range r and a period-wise maximum distance R (2.17 – 2.19).

3 Methods: Solution approach

The problem is approached using a greedy heuristic that efficiently assigns maintenance machines to tasks while adhering to operational constraints. For further improvement, an exact solver solves the MILP model for additional refinement, using the start depots determined by the greedy heuristic.

The greedy heuristic operates iteratively over multiple planning periods, assigning maintenance machines to tasks based on immediate optimization criteria (Algorithm 1). Initially, machines are assigned to depots and for each day within the planning period, available machines select maintenance demands in a prioritized manner. A knapsack-based approach is employed to determine the most efficient allocation of resources, ensuring that available machine capacity is utilized effectively. Containers are selected based on the length of demands a vehicle can fulfill within them, and depot assignments are updated weekly to reflect machine movement constraints. While the greedy algorithm provides a fast and scalable solution, it inherently risks converging to a local optimum due to its myopic decision-making process.

Gurobi acting as the exact solver complements the heuristic by using the calculated solution as a warm-start. It systematically searches for better solutions while ensuring compliance with all constraints. Given the complexity of the problem, the solver is allocated a 24-hour runtime to refine the schedule within practical limits.

Algorithm 1 Greedy Solver Algorithm

Initialize: Load data and define sets (see Table 2.2)
for each period $p \in P$ **do**
    Assign start depots $dep_{v,p} \in S$ for each vehicle $v \in V$
    for each day $d \in \mathrm{D}_p$ **do**
        for each vehicle $v \in V$ **do**
            if a new week starts **then**
                Update depot assignment for each vehicle $v \in V$
            end if
            Select the best container $c^* \in C$
            Assign demands $n \in N_{c^*}$ to vehicle v using a knapsack approach
            Update schedule and resource usage
        end for
    end for
end for

4 Case study: Scheduling for preventive maintenance at Deutsche Bahn

4.1 Experimental setup

The case study is based on a real-world dataset from DB InfraGO AG with initial processing by Synoptics GmbH with 5,574 preventive maintenance demands with a total length of 6,878 km, 2,720 containers, 14 machines of eight different types and 210 depots. Before integration into the MILP model, the data was cleaned for structural and semantic consistency. Oversized demands were split, and small, similar ones merged to reduce complexity. Direction mismatches were corrected based on track orientation. Mapping showed that some demands had no matching container and some containers were unused, with 13 percent of demands unmatched and 46 percent of containers unused. The final dataset included 6,852 demands and 5,651 km of schedulable track and is summarized in Table 2. All experiments used Python with Gurobi (time limit: 24 hours) and were run on the high-performance cluster of Freie Universität Berlin (Bennett et al. 2020) with 256 GB RAM and an Intel Xeon CPU.

Table 2. Configuration of the baseline experiment

Configuration: Baseline experiment	
Periods: 2	Weeks: 32
Max. Working Time: 8 h	Num. Demands: 6,852
Num. Containers: 1,463	Num. Depots: 210
Depot Distance (r): 200 km	Depot Distance (R): 350 km
Container Reusability per Period: 1	Machines by Type: 1:1, 2:3, 3:1, 4:2, 5:1, 6:2, 7:3, 8:1

4.2 Results

The optimization aimed to maximize the maintained track length while adhering to operational constraints. The Mixed Integer Programming (MIP) approach, initialized with a warm-start solution from the greedy procedure, achieved a maintained track length of 5,005,038 meters. In comparison, the Greedy method alone covered 4,785,888 meters, corresponding to an improvement of approximately 4%. The MIP model, formulated with 9.3 million decision variables and 46.5 million constraints, required up to 256 GB of RAM. Despite these computational resources, the solver was unable to reach an optimal solution within the 24-hour time limit. After 6,024 seconds (around 1 hour and 40 minutes), no further improvements were observed, leaving an 11% MIP gap, indicating that further optimization was computationally prohibitive within the given time. In contrast, the Greedy method produced a solution in just 280 seconds but, due to its heuristic nature, was unable to evaluate explicit bounds, leading to a suboptimal result. The MIP approach identified an upper bound of 5,556,977 meters, significantly exceeding the best-found solution, suggesting that the bound calculation was not particularly tight or informative. As shown in Figure 2, which displays the results of the Greedy algorithm, the vehicles tend to work less as time progresses, but when the period change occurs, machine activity increases again. This pattern reflects the scheduling constraints and the influence of period resets on resource utilization. Consequently, this explains the overall vehicle utilization observed in the final optimized schedule, where only 44.8% of available workdays were used, and machines operated for an average of 5.8 hours per workday, but when considering the entire planning horizon, the actual working time per vehicle averaged just 2.6 hours per day. This suggests that simply increasing the number of machines would likely not improve results, as the limiting factor lies in scheduling efficiency rather than fleet size. To further investigate these findings, the sensitivity analysis will explore the impact of various parameters on the optimization results and assess whether alternative adjustments could lead to improvements.

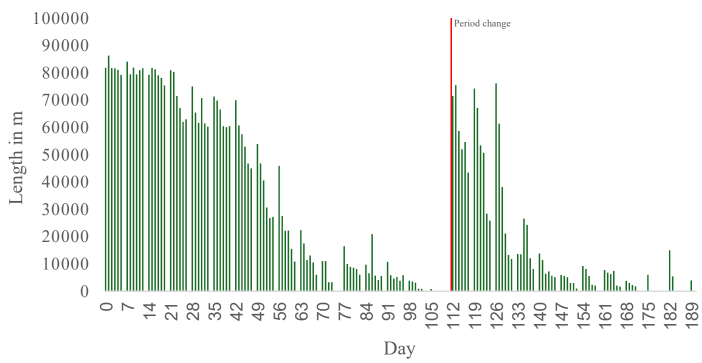

Figure 2. Greedy: Machine deployment schedule for the baseline scenario.

4.3 Sensitivity analysis

This section examines the impact of key operational parameters on maintenance scheduling efficiency through sensitivity experiments on fleet size, scheduling flexibility, and depot distance constraints (RQ2). The analysis relies on a greedy heuristic, selected for its computational feasibility and ability to generate sufficiently accurate results. In repeated random selections with a reduced number of demands and a shortened planning horizon, the greedy approach achieved objective values on average 3.2% worse than the exact solution, with a standard deviation of 1.3%. In this study, the MIP solver failed to find an optimal solution within 24 hours, yielding a bound too loose for meaningful insights – a challenge also reported by Reisch et al. (2025) when addressing the same problem. Given that the greedy heuristic produces results within minutes while achieving nearly the same maintained track length as MIP, it serves as a practical and reliable foundation for the sensitivity analysis. An overview of the results is provided in Table 4.3.

Fleet size expansion: The impact of fleet size expansion was assessed by adding one machine of various types to evaluate changes in demand coverage. For each demand type, only the most efficient machine types were selected. The results indicate only marginal improvements – and in some cases even a deterioration of the objective value. One possible explanation is the nature of the greedy algorithm: with each additional machine, containers tend to be booked out more quickly, reducing accessibility in subsequent iterations. The MIP solver might be able to sort this out and produce a result superior to baseline, however the observed values potentially indicate the limited effect of a fleet size expansion.

Scheduling flexibility: Two key aspects of scheduling flexibility were analyzed, namely container reusability and the planning horizon. Allowing each container to be booked multiple times per period significantly improved demand coverage, thereby supporting the hypothesis that container availability – rather than fleet size – constitutes the primary bottleneck. Allowing each container to be booked twice per period increases the satisfiability rate to 93.5%, whereas permitting four bookings per period raises it further to 98.2%, almost full coverage. Similarly adjusting the planning horizon across multiple periods also led to substantial improvements, since containers can be booked once each period and vehicles are allowed to start from different depots in each period.

Depot distance constraints: Another critical parameter is the limited operational radius around start depots. Our experiments show that increasing the maximum allowed distance for weekly depot changes (r) and the total distance per period (R) by 50% would allow for an additional 169 km of satisfied demands, corresponding to a 3% increase. This suggests that either allowing longer distances could further improve demand satisfaction.

Shift length adjustments: The influence of shift length on demand coverage was also analyzed. While longer shifts contributed to higher demand fulfillment, vehicle utilization rates varied significantly, indicating that some vehicles operated at much higher levels than others. This suggests that while extending shifts can enhance coverage, it may also lead to imbalanced fleet utilization.

5 Findings and discussion

Impact of container recycling: While machine availability does not appear to be a bottleneck in the planning process, the ability to recycle containers strongly impacts the coverage of maintenance demands. Regarding RQ2, our experiments indicate that

Table 3. Greedy: Maintained track length for different experiments. The highest results in each group are highlighted with bold font.

Experiment	Maintained Track Length	Total Demand Coverage
Baseline (Section 4.2)	4,785,888 m	84.68%
One Period	3,550,960 m	62.83%
Four Periods	**5,340,700 m**	**94.49%**
Container 2x per Period	5,282,022 m	93.46%
Container 4x per Period	**5,551,493 m**	**98.22%**
9h per shift	4,867,502 m	86.12%
10h per shift	4,952,969 m	87.63%
16h per shift	**5,243,418 m**	**92.8%**
Added Type 3 Machines	4,718,490 m	83.49%
Added Type 4 Machines	4,768,060 m	84.36%
Added Type 6 Machines	4,785,340 m	84.65%
Added Type 7 Machines	**4,822,090 m**	**85.34%**
Added Type 8 Machines	4,762,140 m	84.26%
Increased Depot Distance	4,955,360 m	87.70%

while other factors, such as increased flexibility in reselecting start depots, improve satisfiability, the ability to reuse certain containers has an even greater effect. Allowing each container to be booked four times per period increases the satisfiability rate to 98.2% – almost full coverage. From the perspective of infrastructure managers, it was clearly intended to only use each container once as a preventive container, because of the capacity reduction for regular railway traffic. Our analysis shows the additional container usage does not directly impact the capacity proportionally, as many containers remain unused. Furthermore, the evaluation of the container system for the first year reveals that many track containers remained unused (DB InfraGO AG 2025). As a result, numerous fixed track blockages – which neither allowed for preventive maintenance scheduling nor ultimately supported ad hoc maintenance – did not serve their intended purpose. This suggests that a prior assessment of which network sections should offer more capacity for preventive maintenance could be beneficial. In fact, in the meantime, DB InfraGO has changed their policy concerning container utilization: Containers can now be booked indiscriminately.

Container utilization and inefficiencies: As noted in Section 4.1, more than 1,000 km of maintenance demands were excluded during data preprocessing because they fell outside the coverage of defined track containers. Providing an answer to RQ1, our results indicate that certain containers are never booked, while others are underutilized. This suggests inefficiencies in planning, potentially due to excessive route lengths, which prevent machines from fully utilizing these containers due to time constraints. Since no other machine can enter a container after it has been booked in a shift, unfulfilled demands within a container remain unmet. We observe that the 500 containers with the highest demand concentration correspond to only 47 routes and have an average length of 103 km (69.8 km across all routes), while the operational range of the machines varies only between 1.8 km and 15 km. Pertaining to RQ3, this leads to the conclusion that existing containers should be split into smaller parts that offer more flexibility, for example, by allowing multiple machines to operate on consecutive sections.

Effect of operational radius: This suggests that allowing longer distances could further improve demand satisfaction. The underlying rationale for these constraints is that worker crews operate from fixed home bases or hotels and should not be required to travel excessive distances between weeks. Hence, greater radii do not seem to be practical for operations.
Impact of shift length: Our experiments show that shift length is a minor factor. Extending the maximum shift length by one or two hours increases demand coverage to 86.12% and 87.63%, respectively. Doubling the shift length improves coverage to 92.8%. However, in practice, this would incur substantial additional labor costs due to overtime payments or the need for additional workforce. Moreover, it would be ill-advised to extend the shift length, as this would counteract the goal to exploit low-traffic conditions at nighttime for maintenance work.

6 Conclusions & outlook

The study underscores the need for rigorous process and problem understanding facilitating efficient planning. Moreover, it shows that the effect of changing organizational conditions far outweighs investments in auxiliary infrastructure – such as maintenance depots – or rolling stock. A more demand-conforming creation of maintenance containers, more flexibility pertaining to shift durations and machine mobility may significantly increase the coverage of preventive maintenance demands. Specifically, splitting overlong preventive containers into smaller parts might be beneficial to efficient planning and operation of preventive railway maintenance.

Future research should focus on mathematically decomposing the model, as this could enhance the efficiency of the solution process, both in terms of computational complexity and hardware requirements. Additionally, a sequential solving approach that splits the problem into multiple subproblems may help to overcome some of the limitations and modeling abstractions introduced in this work.

From a broader perspective, this study highlights the considerable efficiency gains that can be achieved through a more integrated approach to the various planning stages involved in railway infrastructure maintenance. Reflecting on our collaboration with DB Netze AG and insights gained from recent exchanges with Synoptics GmbH and DB InfraGO AG, it becomes evident that this domain offers fertile ground for contributions from the field of Information Systems research. Currently, the maintenance planning process remains costly and is characterized by multiple sources of inefficiency; ranging from fragmented data systems to the lack of integrated scheduling and coordination mechanisms. Addressing these challenges through intelligent information systems and advanced optimization models could substantially improve planning quality, reduce costs, and enhance the overall reliability of railway operations.

References

Bennett, L., Melchers, B. & Proppe, B. (2020), 'Curta: A general-purpose high-performance computer at Zedat, Freie Universität Berlin'.

Budai, G., Huisman, D. & Dekker, R. (2006), 'Scheduling preventive railway maintenance activities', *Journal of the Operational Research Society* **57**(9), 1035–1044.

DB InfraGO (2024), 'Einführung des getakteten Sperrzeitensystems zur strukturellen Systemstabilisierung und -beruhigung', `https://www.dbinfrago.com/web/aktuelles/kund-inneninformationen/kund-inneninformationen/2024-KW02-Sperrzeitensystem-12634666`. Accessed: 2025-03-06.

DB InfraGO AG (2025), 'Bauen im Takt – Modernisierung der Infrastruktur mit möglichst wenig Auswirkungen auf den Zugverkehr', `https://www.dbinfrago.com/web/schienennetz/fahren_und_bauen/bauen-im-takt-13070488`. Accessed: 2025-06-17.

Eckstein, L., Federrath, H., Fricke, H., Friedrich, M., Gühnemann, A., Hartmann, E., Jahn, C., Jipp, M., Kliewer, N., Knauff, M., Knorr, A., Martin, U., Mitusch, K., Oeter, S., Petzoldt, T., Siedentop, S., Sieg, G. & Vortisch, P. (2024), 'Gestaltung eines attraktiven kundenorientierten Schienenpersonenverkehrs', `https://bmdv.bund.de/SharedDocs/DE/Anlage/G/wissenschaftlicher-beirat-schienenpersonenverkehr.pdf?__blob=publicationFile`. Accessed: 2025-03-10.

Lidén, T. (2015), 'Railway infrastructure maintenance – a survey of planning problems and conducted research', *Transportation Research Procedia* **10**, 574–583.

Nemani, A. K., Bog, S. & Ahuja, R. K. (2010), 'Solving the Curfew Planning Problem', *Transportation Science* **44**(4), 506–523.

Peng, F. & Ouyang, Y. (2012), 'Track maintenance production team scheduling in railroad networks', *Transportation Research Part B: Methodological* **46**(10), 1474–1488.

Reisch, J., Glaubitz, J., Budweg, R. & Kliewer, N. (2024), Eine Potenzialabschätzung zur Reduktion der von Baustellen betroffenen Zugfahrten, *in* 'HEUREKA 2024', number FGSV 002/140 *in* 'HEUREKA', FGSV Verlag, Stuttgart.

Reisch, J., Großmann, P. & Weiß, R. (2025), 'A MaxSAT Model for Solving the Track Maintenance Possession Problem for the Railway Network in Germany'.

Sedghi, M., Kauppila, O., Bergquist, B., Vanhatalo, E. & Kulahci, M. (2021), 'A taxonomy of railway track maintenance planning and scheduling: A review and research trends', *Reliability Engineering & System Safety* **215**, 107827.