A Bayesian Rolling Horizon Approach for Rolling Stock Rotation Planning with Predictive Maintenance

Felix Prause[1]
Zuse Institute Berlin, Germany

Ralf Borndörfer
Zuse Institute Berlin, Germany

Abstract

We consider the rolling stock rotation planning problem with predictive maintenance (RSRP-PdM), where a timetable given by a set of trips must be operated by a fleet of vehicles. Here, the health states of the vehicles are assumed to be random variables, and their maintenance schedule should be planned based on their predicted failure probabilities. Utilizing the Bayesian update step of the Kalman filter, we develop a rolling horizon approach for RSRP-PdM, in which the predicted health state distributions are updated as new data become available. This approach reduces the uncertainty of the health states and thus improves the decision-making basis for maintenance planning. To solve the instances, we employ a local neighborhood search, which is a modification of a heuristic for RSRP-PdM, and demonstrate its effectiveness. Using this solution algorithm, the presented approach is compared with the results of common maintenance strategies on test instances derived from real-world timetables. The obtained results show the benefits of the rolling horizon approach.

2012 ACM Subject Classification Applied computing → Transportation; Mathematics of computing → Bayesian computation; Mathematics of computing → Mathematical optimization

Keywords and phrases Rolling stock rotation planning, Predictive maintenance, Rolling horizon approach, Bayesian inference, Local neighborhood search

Digital Object Identifier 10.4230/OASIcs.ATMOS.2024.13

Funding This work was supported by the innovation funding program ProFIT (grant no. 10186617), funded by the State of Berlin and co-funded by the European Union.

1 Introduction

Rail transport is one of the most positive modes of transport concerning environmental friendliness and sustainability. Its volume is likely to increase further in the future. This leads to an increased complexity in the planning of vehicle rotations and results in more challenging scenarios, particularly with respect to maintenance scheduling.

A maintenance strategy that has become increasingly important in recent years is predictive maintenance (PdM). One reason is the availability of sensors and the ability to analyze the data they generate using machine learning. In addition, PdM has obvious advantages in terms of economic and ecological factors. These advantages are based on the fact that the costs for spare parts are lower if the currently installed components are used until the end of their service life. As this also minimizes the number of spare parts used, the environmental impact is likewise reduced.

To combine the advantages of rail transport with those of PdM, it is necessary to develop approaches that integrate predictive maintenance planning into the optimization of rolling stock rotations.

[1] Corresponding author.

24th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2024).
Editors: Paul C. Bouman and Spyros C. Kontogiannis; Article No. 13; pp. 13:1–13:19
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One problem that arises in the application of PdM strategies is that the considered health states are usually unobservable quantities that have to be derived from measurements using some model. There are various steps in this process that introduce uncertainty into the obtained states. For example, measurement errors may occur when recording the observable quantities, or the employed model may be imprecise or only approximate. Since the states are subsequently projected into the future to serve as a basis for maintenance decisions, and the exact operating conditions cannot be known at that point in time, the uncertainty of the health states increases the further they are projected into the future. This uncertainty must be taken into account and addressed when optimizing the vehicle rotations.

Contribution

We consider the rolling stock rotation planning problem with predictive maintenance (RSRP-PdM), as presented in [31], and propose a rolling horizon approach that reduces the uncertainty of the health states and their future predictions. For this purpose, the RSRP-PdM is solved and the determined vehicle rotations are partially operated until measurements of the health states become available, for example, by analyzing sensor data collected during operation when the vehicles are parked overnight. Using these observations of the health conditions, the current predictions of the states are then updated by Bayesian inference. This reduces the variance of the health states and results in a subsequent RSRP-PdM instance.

Furthermore, we introduce a local neighborhood search, which is a modification of the heuristic presented in our earlier paper [29], to solve the occurring RSRP-PdM instances.

Outline

The article is structured as follows: First, in Section 2, we review and discuss the literature on the two arising tasks, i.e., the rolling stock rotation planning problem (RSRP) and predictive maintenance (PdM). Next, the problem formulation of RSRP-PdM, as stated in [31], is reproduced in Section 3. In Section 4, we describe the utilized Bayesian inference procedure. We provide a description of the rolling horizon approach in Section 5 and present a heuristic that extends the algorithm proposed in [29]. Section 6 then introduces the considered maintenance strategies, which are compared with each other in the subsequent computational experiments. Finally, we draw a conclusion on the obtained results in Section 7.

2 Related Work

Both aforementioned topics, i.e., PdM as well as RSRP, have already been thoroughly described and studied in the literature. In the following, we provide a brief overview of articles dealing with these two subjects.

The Rolling Stock Rotation Planning Problem (RSRP)

In RSRP, we are given a fleet of vehicles and a timetable whose trips must be operated. Furthermore, maintenance requirements are defined that the vehicles have to fulfill. The task is then to determine rotations for the rolling stock that operate all trips, satisfy the specified maintenance conditions, and have minimum costs. The articles addressing RSRP can be categorized according to the following three aspects:

- Regarding the model used to represent the vehicle rotations.
- According to the applied solution approach.
- Based on the employed maintenance strategy.

For a survey of the RSRP literature, we refer to [32]. An extensive comparison of different articles with an emphasis on additionally considered constraints can be found in [31].

The RSRP is usually modeled by space-time graphs, where the nodes correspond to events specified by a location and a time point, and the arcs represent the different actions of the vehicles, e.g., [8, 25, 38]. Then, there is the sequence model, in which the nodes depict the trips that need to be operated and the arcs indicate whether two tasks can be conducted in succession, see [6]. Next, we have the hypergraph model utilized by [4, 18, 32]. This model is a generalization of the space-time graph in which the hyperarcs represent vehicle compositions and their orientation during operation. Finally, there is the state-expanded event-graph, see [29, 31]. This is a space-time graph that is extended by additional dimensions to implicitly track the resource flow that enforces the maintenance constraints. In contrast to the previously mentioned approaches, this model allows for non-linear degradation functions, which is of particular relevance as the wear of mechanical components often exhibits such a behavior.

In all these models, the solutions to RSRP are given by flows that cover the trip arcs or nodes sufficiently often and satisfy the additional maintenance or capacity constraints. These induced flow problems are then solved by using branch and bound [8], the direct application of mixed-integer linear programs [17, 20, 21, 34, 39], column generation [2, 4, 32], branch and price [14, 25, 32], or the utilization of heuristics [5, 6, 18, 38].

The literature primarily employs preventive maintenance strategies. These include time-based maintenance [2, 4, 8, 14, 32] and distance-based maintenance [2, 4, 17, 21, 25, 32]. Recently, also condition-based and predictive maintenance regimes have been introduced for rail transport. Here, the maintenance decisions are either based on the vehicle states [5, 20], a classification of the vehicle conditions into degradation stages [39], or the remaining useful life (RUL) of the vehicles [34]. Finally, there are solution approaches that schedule the maintenance based on the predicted failure probability of the rolling stock [29, 31].

Predictive Maintenance (PdM)

The literature regarding PdM generally focuses on the prediction of the RUL or indices representing the future health conditions of the vehicles. We refer to [22] for different concepts of health indices. These indices usually do not consider the vehicles themselves, but rather the mechanical, electrical, or hydraulic components that are installed in them. The employed approaches are often distinguished into data-driven and model- or physics-based ones, see [1]. In addition, there are hybrid methods that combine both. A comprehensive literature review on PdM can be found in [13].

Data-driven models usually rely on machine learning techniques like classical neural networks (NNs), recurrent neural networks such as long short-term memory (LSTM) networks [10], support vector machines (SVMs) and decision trees [23], or deep learning [11]. For a survey of data-driven approaches applied in the railroad sector, see [9].

Model-based approaches, on the other hand, are based on some assumptions regarding the degradation process and often utilize Bayesian updating procedures [7, 16, 27].

Finally, there are hybrid approaches that combine model-based methods with machine learning. These include, for example, relevance vector machines, i.e., Bayesian variants of SVMs [35], Bayesian inference applied to the output of NNs [15, 26], or the assumption of a piecewise linear degradation behavior, followed by the subsequent combination of the outputs of multiple NNs using a Kalman filter [24]. In addition, a concept for extending deep learning models to Bayesian NNs is presented in [28], which enables the NNs to determine probability distributions for the predicted RUL.

Note that the hybrid approaches and the Bayesian methods have the advantage of obtaining a probability distribution for the RUL that captures the uncertainty of the prediction.

3 Problem Description

We consider the rolling stock rotation planning problem with predictive maintenance (RSRP-PdM), as presented in [31], and recall its description in the following. Suppose we are given a train timetable $\mathcal{T}$ consisting of various trips that need to be operated, and we have a homogeneous fleet of vehicles $\mathcal{V}$ at our disposal to conduct them. Each trip $t \in \mathcal{T}$ features a departure and an arrival location, i.e., $l_t^d, l_t^a \in \mathcal{L}$, as well as a departure and an arrival time, i.e., $k_t^d, k_t^a \in \mathcal{K}$. Here, $\mathcal{L}$ is the set of locations and $\mathcal{K}$ is the time horizon, which consists of a finite set of time points. Furthermore, we associate an integer $n_t \in \mathbb{Z}_{>0}$ with each trip, which indicates how many vehicles are required to operate t.

The task of RSRP-PdM is not only to find a feasible sequence of trips for each vehicle, i.e., a set of trips in which each pair of time-consecutive trips can be operated in succession, but also to schedule the maintenance of the vehicles. The maintenance actions can be carried out at the maintenance locations $\mathcal{L}_M \subseteq \mathcal{L}$ and should be based on the predicted health states of the vehicles. These health states are considered to be random variables since they cannot be measured directly and are thus prone to measurement, determination, and prediction errors. In addition, they are assumed to be normally distributed, i.e., $H_{v,k} \sim \mathcal{N}(\mu, \sigma^2)$ for some $\mu \in [0,1]$ and $\sigma^2 > 0$. Hence, all health states are distributed by the family of normal distributions with parameter space $\Theta = [0,1] \times \mathbb{R}_{>0}$ and can thus be characterized by their corresponding parameters $\theta \in \Theta$. In this sense, we assume that the initial state of each vehicle $v \in \mathcal{V}$ is given by an initial parameter $\theta_{v,0} \in \Theta$ that determines $H_{v,0}$. Note that a health value of one corresponds to a condition that is as good as new, while a value of zero or less signifies that the vehicle has a breakdown. The failure probability of a vehicle at a certain time is therefore given by

$$\mathbb{P}_f(v,k) := \mathbb{P}[H_{v,k} \leq 0] = \int_{-\infty}^{0} \frac{\exp\left(-\frac{\mu_{v,k}^2}{2\sigma_{v,k}^2}\right)}{\sqrt{2\pi\sigma_{v,k}^2}}\, dx = \frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{-\mu_{v,k}}{\sqrt{2\sigma_{v,k}^2}}\right)\right),$$

where erf is the Gauss error function.

The deterioration of the vehicles and their maintenance is then described by modifying the parameters that represent their health states. Therefore, we associate a degradation function $\Delta_t : \Theta \to \Theta$ with each of the trips $t \in \mathcal{T}$, whose application describes the wear that occurs during the operation of t. Let $\tau_t := k_t^a - k_t^d$ be the duration of t, then the parameters of v after conducting t are determined by $\theta_{v,k+\tau_t} = \Delta_t(\theta_{v,k})$. To obtain a reasonable deterioration behavior, we further demand $\mu_{v,k+\tau_t} \leq \mu_{v,k}$ and $\sigma_{v,k+\tau_t}^2 \geq \sigma_{v,k}^2$ for $(\mu_{v,k+\tau_t}, \sigma_{v,k+\tau_t}^2) = \Delta_t(\mu_{v,k}, \sigma_{v,k}^2)$, i.e., the mean of the health state decreases, while the uncertainty about the condition grows. Similarly, we can associate degradation functions with the other activities of the vehicles such as waiting or deadheading, which might depend on the duration or the mileage of the corresponding task. Maintenance is also described by a wear function, which does not cause a deterioration of the condition, but resets the parameters of the health state to a certain value $\theta_m \in \Theta$.

We further assume that we occasionally receive measurements $y_{v,k}$ of the true health value of v at time k that exhibit noise originating from the measuring process. This noise is supposed to be normally distributed around zero with known variance $\sigma_y^2 > 0$.

Combining these notions, the task of RSRP-PdM is to determine a feasible assignment of the trips to the vehicles such that each trip is operated by the required number of vehicles. The objective here is to find an assignment of minimum total cost that takes the potential failure costs into account, i.e., the product of the predicted failure probability during the operation of the trips and the costs associated with the breakdown of a vehicle. Finally, we require that the rotations must be balanced, i.e., the number of vehicles located at each destination at the beginning and at the end of the time horizon have to coincide. This constraint is important as it gives rise to schedules that can be repeated periodically.

4 Bayesian Inference

In RSRP-PdM, as described in Section 3, the health states of the vehicles and their predictions are assumed to be random variables to reflect their uncertainty. Since these random variables are distributed by members of a family of probability distributions, their associated probability density functions (PDFs) can be characterized by their parameters, and the degradation functions describe how these parameters change. The deterioration process can therefore be understood as a dynamical system in which the parameters of the vehicle conditions represent the system states, and the degradation functions define the state transitions.

One problem that arises in practice is that the degradation functions are also subject to uncertainty since they can only be derived from historical data, and the actual operating conditions are unknown at the time of planning. If this uncertainty is factored in, the variance, and thus the uncertainty, of the predicted health states increases the further into the future they are projected. However, these predicted conditions form the basis for maintenance planning. Therefore, it should be attempted to reduce their variance to obtain more accurate estimates of the actual health states. This can be achieved by updating the predicted states with measurements, which is a filtering problem, see [36]. Using standard terminology, we utilize the *rule of Bayes* for these updates, see, for example, [36].

▶ **Theorem 1** (Rule of Bayes [36]). *Let θ and y be random variables representing a parameter estimate and a measurement, then it holds*

$$\mathbb{P}[\theta \mid y] = \frac{\mathbb{P}[y \mid \theta] \cdot \mathbb{P}[\theta]}{\mathbb{P}[y]} \propto \mathbb{P}[y \mid \theta] \cdot \mathbb{P}[\theta].$$

Here, $\mathbb{P}[\theta]$ is the *prior belief* about θ before obtaining measurement y, $\mathbb{P}[y \mid \theta]$ describes the *likelihood*, i.e., the relationship between the true state and y, which represents the measurement error. $\mathbb{P}[\theta \mid y]$ is the belief about θ after the information about y is incorporated, i.e., the *posterior belief*, and $\mathbb{P}[y]$ is the *marginal probability*, which can be interpreted as a normalization constant ensuring that $\mathbb{P}[\theta \mid y]$ has an integral of one. Furthermore, $f(x) \propto g(x)$ signifies that $f(x)$ is proportional to $g(x)$, i.e., there exists a constant $c \in \mathbb{R}$ such that $f(x) = c \cdot g(x)$ for all $x \in \mathbb{R}$. For further details, we refer to [36].

We now transfer the notions of Theorem 1 to the situation in RSRP-PdM. Recall that we assumed in Section 3 that the health states are normally distributed with a mean between zero and one. Suppose we are given a vehicle v at time k_0 with health state $H_{v,k_0} \sim \mathcal{N}(\mu_{k_0}, \sigma^2_{k_0})$ that operates some services $\mathcal{S} = \{s_1, \dots, s_n\}$, i.e., a set consisting of trips, waiting times, deadhead trips, and maintenance actions in their chronological order. Then, a prediction of the health state of v after the operation of $\mathcal{S}$ can be determined by applying the degradation functions of the services in $\mathcal{S}$ to the parameters characterizing H_{v,k_0}. If we set $\Delta_{\mathcal{S}} \coloneqq \Delta_{s_n} \circ \cdots \circ \Delta_{s_1}$ and $k \coloneqq k_0 + \tau_{s_1} + \cdots + \tau_{s_n}$, the predicted state is therefore $\hat{H}_{v,k} \sim \mathcal{N}(\hat{\mu}_k, \hat{\sigma}^2_k)$ with $(\hat{\mu}_k, \hat{\sigma}^2_k) = \Delta_{\mathcal{S}}(\mu_{k_0}, \sigma^2_{k_0})$.

Here, $\hat{H}_{v,k}$ represents the current belief about the true health state $H_{v,k}$. If we now obtain a measurement y of the true health state with a measuring error whose variance is specified by $\sigma_y^2 > 0$, then we have $y \mid \hat{H}_{v,k} \sim \mathcal{N}(\hat{H}_{v,k}, \sigma_y^2)$. Thus, we can apply the rule of Bayes to obtain $H_{v,k} = \hat{H}_{v,k} \mid y$. Since the considered random variables are all normally distributed, this inference corresponds precisely to the update step of the Kalman filter, as described in [36].

▶ **Definition 2** (Kalman Filter Update Step). *Let $\hat{\mu} \in \mathbb{R}^n$ be the predicted mean of the estimated state and $\hat{P} \in \mathbb{R}^{n\times n}$ the corresponding predicted covariance matrix. Furthermore, let $H \in \mathbb{R}^{m\times n}$ be the measurement model, $y \in \mathbb{R}^m$ a measurement of the true state and $R \in \mathbb{R}^{m\times m}$ the corresponding covariance matrix of the measurement. Then, the* Kalman filter update step *is defined as follows:*

$$K = \hat{P}H^T \left(H\hat{P}H^T + R\right)^{-1}$$
$$\mu = \hat{\mu} + K\left(y - H\hat{\mu}\right)$$
$$P = \left(I_n - KH\right)\hat{P},$$

where $K \in \mathbb{R}^{n\times m}$ is the Kalman gain*, $\mu \in \mathbb{R}^n$ is the updated mean of the estimated state, $P \in \mathbb{R}^{n\times n}$ is the corresponding updated covariance matrix, and $I_n \in \mathbb{R}^{n\times n}$ is the identity matrix of size n.*

Considering Definition 2 for the one-dimensional case and assuming that the measurement is a direct observation of the true state, i.e., $H = 1$, we obtain the following corollary:

▶ **Corollary 3.** *Let $\mu_\theta \in \mathbb{R}$ be the predicted mean of the estimated state and $\sigma_\theta^2 > 0$ the corresponding predicted variance. Furthermore, let $\mu_y \in \mathbb{R}$ be a direct measurement of the true state and $\sigma_y^2 > 0$ the corresponding variance of the measurement noise. Then, applying the Kalman filter update step yields the updated state estimate*

$$X \sim \mathcal{N}\left(\frac{\mu_\theta\sigma_y^2 + \mu_y\sigma_\theta^2}{\sigma_\theta^2 + \sigma_y^2}, \frac{\sigma_\theta^2\sigma_y^2}{\sigma_\theta^2 + \sigma_y^2}\right).$$

The inferred health states can therefore be directly determined by applying Corollary 3. Moreover, they are also normally distributed and thus members of the considered family of probability distributions.

An example illustrating the impact of inference is given in Figure 1. The two graphs show the development of a vehicle's predicted/updated health state over time. The corresponding PDFs are given from right to left and show the state after the operation of ten trips in each case, starting with the initial health $H_{v,0}$ on the far right. It can be observed that the variance in the scenario without inference continues to increase, while it remains in the same order of magnitude if the states are updated with measurements after every ten trips. As a result, $H_{v,50}$ has a high probability of taking values between 0.25 and 0.45 in the second case, whereas in the first case, it contains essentially no information since its variance is too large.

Application to Non-normally Distributed Health States

However, the health states do not necessarily have to be represented by normally distributed random variables. Other possible models could, for example, be based on families of probability distributions that are used in reliability theory, such as the family of Weibull or gamma distributions, see, e.g., [12].

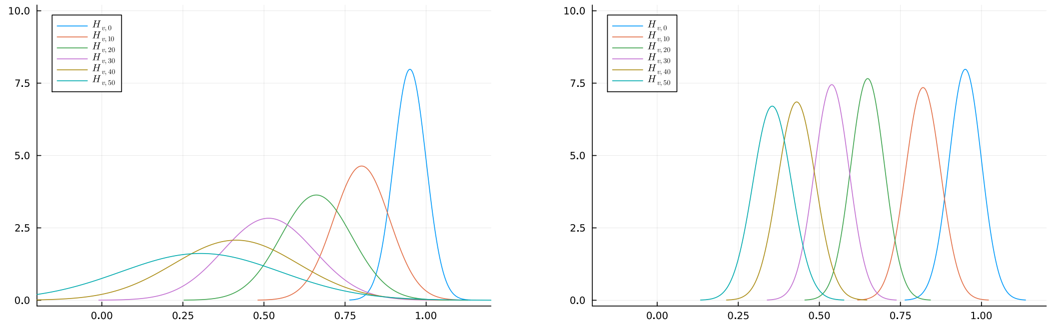

(a) PDF of the predicted health state over time, without state updates.

(b) PDF of the health state over time, using inference after every tenth trip.

Figure 1 Comparison of the health state PDF development over time, with and without inference.

Furthermore, it would also be conceivable that the states are distributed by discrete distributions or PDFs that do not belong to any common family of distributions. In these cases, we have to choose a family of distributions whose members fit the given data as well as possible since the employed approach is based on the assumption that the PDFs of the random variables representing the health states all belong to a parametric family.

Suppose now that we have determined such a family of distributions for modeling the health states and are able to obtain measurements of the true states, which are either given by point estimates with a certain measuring error or described by a PDF. If we then apply the rule of Bayes to infer the posterior of the health state after incorporating the measurement, as described above, the resulting PDF does not necessarily have to be a member of the selected family. It may even be that it does not belong to any common family of distributions.

In these cases, the posterior of the state would have to be approximated by a member of the considered family of distributions. This would, for example, be possible by employing a Markov chain Monte Carlo algorithm to sample from the posterior distribution of the health state. Subsequently, we determine the PDF of the family that best fits the obtained data w.r.t. some statistical distance. Another option would be the utilization of a variational Bayesian method, see, e.g., [37].

5 Solution Approach

In this section, we introduce the rolling horizon approach. This algorithm determines a solution for RSRP-PdM by sequentially generating and solving sub-instances. Here, the instances occurring in each iteration are solved by a heuristic, which is also described below.

5.1 The Rolling Horizon Approach

The idea of the rolling horizon approach is the following: Suppose we have a solution for RSRP-PdM where the expected deterioration caused by the trips is only known approximately. Then, the maintenance services of the vehicles are planned entirely based on predictions of the health states, whose uncertainty increases the further they are projected into the future. However, if it is possible to occasionally obtain measurements of the health states and incorporate them into the planning of the rotations, these can be adjusted such that the vehicles are assigned to trips that better match their conditions. In addition, maintenance needs can be better estimated, and the vehicles can be serviced at more appropriate times.

This can be accomplished by first solving the initial RSRP-PdM instance (`solveInstance`) and operating the resulting vehicle rotations until measurements y of the health states become available at a certain time k_y (`operateRotations`). Subsequently, a new problem instance is created by restricting the timetable of the considered instance to trips whose departure time is greater than or equal to k_y (`restrictTrips`). The positions of the vehicles are then updated by assigning them the location at which they are at time k_y or at which they terminate the operation they are conducting at that time (`updatePositions`). In addition, the predicted health states of the vehicles are updated with y using Bayesian inference, as described in Section 4 (`inferStates`). This procedure is outlined in Algorithm 1 and results in a new RSRP-PdM instance, which is then considered in the next iteration of the approach. The algorithm stops when all trips of the original instance have been completed.

Algorithm 1 One iteration of the rolling horizon approach.

Input : RSRP-PdM instance $\mathcal{I}$, measurements y of the health states at time k_y
Output : Updated RSRP-PdM instance for the remaining time

$s \leftarrow$ `solveInstance`$(\mathcal{I})$
`operateRotations`(s, k_y)
$\mathcal{I}' \leftarrow$ `restrictTrips`$(\mathcal{I}, k_y)$
$\mathcal{I}' \leftarrow$ `updatePositions`$(\mathcal{I}', s, k_y)$
$\mathcal{I}' \leftarrow$ `inferStates`$(\mathcal{I}', y)$
return $\mathcal{I}'$

5.2 A Local Neighborhood Search That Considers Transition Costs

To solve the occurring RSRP-PdM instances, we utilize a modified version of the multi-swap heuristic proposed in [29]. This algorithm initially solves the underlying RSRP of the instance using an integer linear program that neglects the maintenance constraints. Afterwards, maintenance is scheduled in a second step by determining a shortest path in the state-expanded event-graph (SEEG). Then, a local neighborhood search is employed to improve the rotations of the current solution, where maintenance planning is again done by searching for a shortest path in the SEEG.

This local neighborhood search works as follows: Given two vehicle rotations, the first step is to determine the possible swap positions, i.e., the times at which both vehicles can reach and operate the next trip of the other. This groups the trips of the rotations into sets that can be exchanged without violating the feasibility of the corresponding vehicle schedules. Subsequently, the generated subsets of trips are randomly swapped to obtain new rotations. An example of this procedure is shown in Figure 2a. Here, the rotations of vehicles v_1 and v_2 (top) are first grouped into subsets of interchangeable trips (middle). Afterwards, some of these related subsets are swapped to create two new vehicle schedules (bottom).

In [29], the swapping decisions are sampled from a discrete uniform distribution $\mathcal{U}\{0, 1\}$, i.e., the swaps are performed with a probability of one-half. However, some swaps are more beneficial than others in terms of costs, and their selection may accelerate the process of finding good solutions. Therefore, the transition costs of the vehicles should be included when determining the exchange probabilities. These costs are referred to as $c_{ij} \in \mathbb{R}_{\geq 0}$, for $i, j \in \{1, 2\}$, and are associated with the waiting times and deadhead trips that are necessary for vehicle v_i to reach the next trip of vehicle v_j. For example, consider the scenario depicted in Figure 2b with costs $c_{11} = c_{22} = 1$ and $c_{12} = c_{21} = 10$. If we disregard maintenance

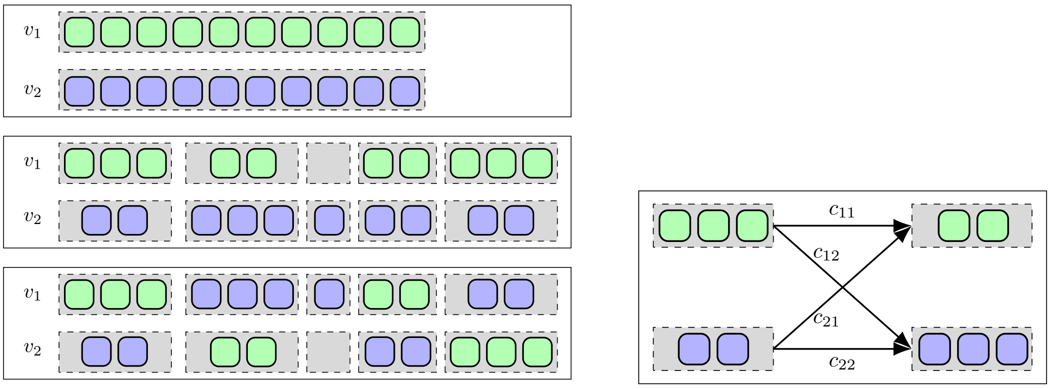

(a) The various stages of the local neighborhood search.

(b) The costs used in the weighted variant of the swapping procedure.

Figure 2 Visualization of the stages of the employed heuristic (a) and of the transition costs of the vehicles between trip subsets of the second stage (b).

decisions, a swap of the second pair of trip subsets would lead to an increase in transition costs. Therefore, it would not be advisable to swap with a probability of one-half at this position. For this reason, we use a probability of

$$\mathbb{P}_s = \begin{cases} \frac{c_{11}+c_{22}}{c_{11}+c_{12}+c_{21}+c_{22}} & \text{if } c_{11}+c_{12}+c_{21}+c_{22} \neq 0 \\ \frac{1}{2} & \text{else} \end{cases}$$

for sampling the swapping decisions. However, to ensure a certain degree of exploration, the $\mathbb{P}_s$ were finally rounded to values in $[0.05, 0.95]$. The swapping decisions for each pair of trip subsets are then made in chronological order with the associated probability $\mathbb{P}_s$. Here, we set $\mathbb{P}_s := 1 - \mathbb{P}_s$ if the previous decision was a swap to account for the resulting exchange of c_{11} and c_{12}, as well as of c_{22} and c_{21}.

A comparison of the solutions obtained with the multi-swap heuristic using this swapping procedure with those of the algorithm from [29] can be found in Appendix A. The results demonstrate the effectiveness of the modified swapping approach.

6 Computational Results

In this section, we examine the results of the Bayesian rolling horizon approach proposed in Section 5.1. For this purpose, we compare its solutions with those of two other maintenance strategies, namely predictive maintenance without Bayesian inference and preventive maintenance. To solve the occurring RSRP and RSRP-PdM instances, we apply the *weighted multi-swap heuristic* presented in Section 5.2 and use scenarios derived from the instances given in [30] for testing.

Computational Setup

The data structures and algorithms were implemented in Julia v1.9.4 [3], and Gurobi v10.0.2 [19] was employed to solve the integer programs that are used to find initial rotations for the heuristics. The computations were conducted on a computer with Intel(R) Xeon(R) Gold 6342 @ 2.80GHz CPUs, eight cores, and 64GB of RAM. Finally, all approaches had a time limit of seven hours, with the rolling horizon approach using one hour to solve the RSRP-PdM instance of each day of the given week.

6.1 Test Instances

The considered test scenarios, used for the evaluation and comparison of the maintenance regimes, are based on the instances T1 – T6 constructed in [30]. Each of them has an individual timetable, which needs to be completed within one week, and contains the information about the available fleet, i.e., the operating costs of the vehicles and their initial positions. In addition, the distances between the considered locations and the costs associated with trips, waiting times, deadheading, and maintenance are indicated. The costs for vehicle breakdowns are also defined.

The components that specify the conditions of the vehicles are their doors, and it is assumed that they can undergo 1,500 cycles before failing. This assumption is based on the real-world data used in [33] and influences the number of required maintenance services. Furthermore, the expected deterioration caused by the operation of each trip, i.e., the number of expected cycles, is given by the mean and variance of a normal distribution derived from the passenger volume at the served stations, see [30]. Note that the vehicle conditions in the conducted computations are assumed to deteriorate only during the operation of the trips. In the original instances, a non-linear degradation behavior is considered, whereas we assume a linear one, as this improves the comparability of the obtained solutions.

From each of these instances, we derive ten scenarios by sampling the initial states of the vehicles, i.e., the number of already performed cycles, from the discrete uniform distribution $\mathcal{U}\{0, 1200\}$. In addition, the number of cycles that actually arise during each trip is sampled from the corresponding normal distribution. These values are later used to determine whether vehicle failures occur during the operation of the obtained schedules and to derive the measurements utilized in the Bayesian approach. This gives rise to the various scenarios T1-01 – T6-10, which are used for the computations below.

Furthermore, we assume that, if the states are considered as random variables, the initial states of the vehicles, the conditions after maintenance, and the measurements each have a variance of 25, i.e., the corresponding values are accurate to within ten cycles with a probability of 90%. Finally, for the predictive approaches, we apply a transformation that converts the number of cycles into a health value between zero and one. Here, zero cycles correspond to a new condition, i.e., a value of one, while 1,500 cycles correspond to a vehicle failure and thus a value of zero. The variances were also adjusted accordingly.

6.2 Compared Maintenance Strategies

Now, we describe the maintenance strategies considered in the computational experiments. Since we assume a linear degradation behavior, one-dimensional quantities simply have to be added. However, if the states and occurring cycles of the trips are considered as random variables, the convolution of their PDFs must be determined. Nevertheless, since both are presumed to be normally distributed, this convolution can be calculated by summing the mean values and the variances of the corresponding PDFs.

Preventive Maintenance

The first considered maintenance strategy is preventive maintenance (PM), which represents the status quo of maintenance planning. In this approach, one-dimensional quantities such as the traveled distance or the elapsed time are usually considered, and maintenance is performed before the values exceed a certain threshold.

Applied to the considered scenarios, we assume that maintenance decisions are based on the number of operated cycles, but only the mean values of the distributions characterizing the deterioration due to the trips are used. Since the vehicles are assumed to complete 1,500 cycles before failing, we must first define a suitable threshold to avoid vehicle breakdowns. For this purpose, we examine the preventive strategy with safety margins of 5% and 10%.

Based on the 3σ *rule*, i.e., the fact that approximately 99.7% of the values in a normal distribution lie within an interval of three standard deviations around the mean, we require that $\mu + 3\sigma \leq 1,500$ holds for all parameters that the predicted health state of a vehicle could possess before maintenance. Let v be a vehicle in new condition, i.e., with parameters $(\mu, \sigma^2) = (0, 25)$. Then, we apply the degradation functions of randomly selected trips to v until the mean value of its predicted state is 1,350 and 1,425, respectively. With 50,000 repetitions, this results in a maximum variance of 750 or 800. Thus, a margin of 5% does not fulfill the property required above and might therefore lead to failures. Hence, we use a threshold of 1,350 cycles, i.e., a safety margin of 10%, for PM.

Predictive Maintenance

The next strategy used is the direct application of predictive maintenance (PdM), as described in [29, 31]. Here, maintenance planning is based on the predicted failure probabilities of the vehicles. However, the health states are not updated. Since the degradation behavior is assumed to be linear, the deterioration of a vehicle v due to a trip t is expressed by the convolution of the PDF of the random variable representing the health state of v with the PDF of the random variable describing the expected number of cycles occurring during the operation of t. As previously mentioned, the initial vehicle states and the conditions after maintenance are assumed to have a variance of 25.

Predictive Maintenance with Bayesian Inference

Finally, we consider the maintenance strategy developed in this article, namely predictive maintenance with Bayesian inference (PdM-B). For this, we utilize the predictive maintenance regime PdM described above but additionally assume that it is possible to receive measurements of the health states when the vehicles are parked overnight. As explained in Section 4, we assume that these measurements are normally distributed and have a fixed variance $\sigma_y^2 = 25$, which reflects the measuring error. After obtaining the measurements of the states, we apply Algorithm 1, using Corollary 3 to update the predicted vehicle states and reduce their variance. This yields a new RSRP-PdM instance, which is subsequently solved.

6.3 Results

The results of the computations are summarized in Table 1 and show the number of maintenance services and the total costs of the solutions after averaging the scenarios derived from each instance. The best results for each instance are marked in bold. A more detailed itemization of the costs by type can be found in Tables 3–5 in Appendix B. Here, only the actually incurred costs are taken into account, i.e., the expected failure costs, which are considered in the predictive maintenance strategies for maintenance planning, are ignored. In addition, all determined vehicle rotations were compared with the number of cycles that actually occurred, i.e., the number of cycles that were sampled for each trip when the scenarios were created. However, none of the generated solutions resulted in a vehicle failure.

Table 1 Number of maintenance services and total costs of the solutions generated by the considered maintenance strategies after averaging the scenarios of each instance.

Instance	Maintenance services			Total costs		
	PM	PdM	PdM-B	PM	PdM	PdM-B
T1	6.3	**5.8**	**5.8**	272,822	**271,781**	271,937
T2	**4.9**	**4.9**	**4.9**	462,219	466,564	**458,174**
T3	14.9	13.3	**13.2**	1,421,806	**1,411,300**	1,412,457
T4	6.0	5.5	**5.1**	212,015	212,069	**211,831**
T5	7.0	6.8	**6.2**	346,513	344,891	**342,976**
T6	21.1	18.8	**18.7**	**2,369,415**	2,377,673	2,375,013
Σ	60.2	55.1	**53.9**	5,084,790	5,084,278	**5,072,388**

A comparison of PM with PdM shows that both approaches achieve better results in terms of costs than the other for three of the six instances. Overall, however, they have almost the same total costs. The reason for this is that PdM reduces the number of maintenance services but does so at the expense of higher deadhead costs, compare Tables 3 and 4. Nevertheless, excluding instance T2, PdM performed fewer maintenance services for each instance and was able to reduce the total number of service actions by 5.1.

PdM-B, on the other hand, is able to reduce the number of conducted maintenance operations even further. To be precise, an average of 6.3 fewer maintenance operations are required compared to PM. Moreover, it can compensate for the disadvantage of PdM and reduces the deadhead costs to such an extent that it generates the vehicle schedules with the lowest total costs. It was thus able to achieve lower costs than PM for all instances except T6 and reduced the number of maintenance services for all instances except for T2. In particular, the number of service actions for instances with many vehicles, i.e., for T3 and T6, is reduced. Looking at the entire network, i.e., all instances combined, PdM-B was able to decrease the number of maintenance actions by 10.5%, while the costs are 0.24% lower. If the fixed costs, i.e., the costs required to operate the trips, are neglected, the overall cost advantage increases to 0.97%.

The cost differences, expressed as percentages, of PdM-B compared to PM and PdM are shown in Table 6. These differences range between -1.02% and 0.24% and between -1.8% and 0.08%, respectively. If the fixed costs are again excluded, the cost advantages for the individual instances increase to up to 5.72% and 8.05%, respectively.

The main cost benefits of PdM-B in comparison to PM are less maintenance (T1), lower deadhead costs (T2), or the combination of both (T3 and T5), see Tables 3 and 5. For T4, the costs for maintenance and deadhead trips are also decreased, but PdM-B deploys additional vehicles that may be necessary to absorb a more severe deterioration identified by the incorporation of the measurements. In addition, some maintenance can be spared by using these vehicles, which also reduces the costs. In the case of T6, PdM-B achieved solutions with higher costs than PM, which is due to the increased deadhead costs, although the number of maintenance tasks could be lowered.

Compared to PdM, PdM-B achieves lower costs for four of the six instances and lower costs overall, specifically 0.23% less. When the fixed costs are neglected, the benefit is 0.93%. In addition, the number of performed maintenance actions was reduced by a total of 1.2. The cost savings are again due to lower deadhead costs (T2 and T6) and to a combination of lower deadhead and maintenance costs (T4 and T5). For T1 and T3, the solutions generated by PdM-B have slightly higher costs than those generated by PdM. In the first case, this is due to the utilization of additional vehicles, which reduces deadhead costs, while in the second case, slightly higher deadhead costs are accepted to assign the vehicles to trips that better match their conditions.

7 Conclusion

In this article, we presented a rolling horizon approach for RSRP-PdM that incorporates health state measurements using Bayesian inference. We also extended the local neighborhood search from [29] to include transition costs when determining the swapping probabilities.

For this purpose, we first gave a literature review of the two topics RSRP and PdM. Then, we recalled the problem formulation of RSRP-PdM and introduced the notions that arise in the context of Bayesian inference. Subsequently, the Bayesian rolling horizon approach was introduced and we described the modified local neighborhood search. Finally, we conducted computational experiments with three different maintenance strategies, namely preventive and predictive maintenance, as well as predictive maintenance with Bayesian inference, and compared their solutions.

The results show that the iterative scheduling of the Bayesian approach is advantageous over both preventive and predictive maintenance without updating. Not only can the number of maintenance actions be reduced by 10.5% compared to the conventional strategy of preventive maintenance, but also the total costs of all instances combined are decreased by 0.24%, or by 0.97% if the fixed costs are excluded. In comparison to predictive maintenance without updating, the costs can likewise be reduced by 0.23% and 0.93% respectively due to fewer deadhead trips, while the number of maintenance operations is also slightly decreased. This demonstrates the effectiveness and benefits of the Bayesian rolling horizon approach.

In addition, we have shown that taking transition costs into account improves the performance of the multi-swap heuristic, both in terms of the solution value after a short and a long computation time.

References

1. Dawn An, Nam H. Kim, and Joo-Ho Choi. Practical options for selecting data-driven or physics-based prognostics algorithms with reviews. *Reliability Engineering & System Safety*, 133:223–236, 2015. `doi:10.1016/j.ress.2014.09.014`.
2. Javier Andrés, Luis Cadarso, and Ángel Marín. Maintenance scheduling in rolling stock circulations in rapid transit networks. *Transportation Research Procedia*, 10:524–533, 2015. `doi:10.1016/j.trpro.2015.09.006`.
3. Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. `doi:10.1137/141000671`.
4. Ralf Borndörfer, Markus Reuther, Thomas Schlechte, Kerstin Waas, and Steffen Weider. Integrated optimization of rolling stock rotations for intercity railways. *Transportation Science*, 50(3):863–877, 2016. `doi:10.1287/trsc.2015.0633`.
5. Omar Bougacha, Christophe Varnier, and Noureddine Zerhouni. Impact of decision horizon on post-prognostics maintenance and missions scheduling: a railways case study. *International Journal of Rail Transportation*, 10(4):516–546, 2022. `doi:10.1080/23248378.2021.1940329`.
6. Valentina Cacchiani, Alberto Caprara, and Paolo Toth. A fast heuristic algorithm for the train unit assignment problem. In *12th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012. `doi:10.4230/OASIcs.ATMOS.2012.1`.

7 Nan Chen and Kwok Leung Tsui. Condition monitoring and remaining useful life prediction using degradation signals: Revisited. *IiE Transactions*, 45(9):939–952, 2013. doi:10.1080/0740817X.2012.706376.

8 Jean-François Cordeau, François Soumis, and Jacques Desrosiers. Simultaneous assignment of locomotives and cars to passenger trains. *Operations research*, 49(4):531–548, 2001. doi:10.1287/opre.49.4.531.11226.

9 Narjes Davari, Bruno Veloso, Gustavo de Assis Costa, Pedro Mota Pereira, Rita P. Ribeiro, and João Gama. A survey on data-driven predictive maintenance for the railway industry. *Sensors*, 21(17):5739, 2021. doi:10.3390/s21175739.

10 Luigi de Simone, Enzo Caputo, Marcello Cinque, Antonio Galli, Vincenzo Moscato, Stefano Russo, Guido Cesaro, Vincenzo Criscuolo, and Giuseppe Giannini. Lstm-based failure prediction for railway rolling stock equipment. *Expert Systems with Applications*, 222:119767, 2023. doi:10.1016/j.eswa.2023.119767.

11 André Listou Ellefsen, Emil Bjørlykhaug, Vilmar Æsøy, Sergey Ushakov, and Houxiang Zhang. Remaining useful life predictions for turbofan engine degradation using semi-supervised deep architecture. *Reliability Engineering & System Safety*, 183:240–251, 2019. doi:10.1016/j.ress.2018.11.027.

12 Frank Emmert-Streib and Matthias Dehmer. Introduction to survival analysis in practice. *Machine Learning and Knowledge Extraction*, 1(3):1013–1038, 2019. doi:10.3390/make1030058.

13 Aurora Esteban, Amelia Zafra, and Sebastian Ventura. Data mining in predictive maintenance systems: A taxonomy and systematic review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(5):e1471, 2022. doi:10.1002/widm.1471.

14 Yuan Gao, Jun Xia, Andrea D'Ariano, and Lixing Yang. Weekly rolling stock planning in chinese high-speed rail networks. *Transportation Research Part B: Methodological*, 158:295–322, 2022. doi:10.1016/j.trb.2022.02.005.

15 Nagi Z. Gebraeel and Mark A. Lawley. A neural network degradation model for computing and updating residual life distributions. *IEEE Transactions on Automation Science and Engineering*, 5(1):154–163, 2008. doi:10.1109/TASE.2007.910302.

16 Nagi Z. Gebraeel, Mark A. Lawley, Rong Li, and Jennifer K. Ryan. Residual-life distributions from component degradation signals: A bayesian approach. *IiE Transactions*, 37(6):543–557, 2005. doi:10.1080/07408170590929018.

17 Giovanni Luca Giacco, Andrea D'Ariano, and Dario Pacciarelli. Rolling stock rostering optimization under maintenance constraints. *Journal of Intelligent Transportation Systems*, 18(1):95–105, 2014. doi:10.1080/15472450.2013.801712.

18 Boris Grimm, Ralf Borndörfer, Markus Reuther, and Thomas Schlechte. A cut separation approach for the rolling stock rotation problem with vehicle maintenance. In *19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/OASIcs.ATMOS.2019.1.

19 Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, Version 10.0.2. https://www.gurobi.com, 2024.

20 Nathalie Herr, Jean-Marc Nicod, Christophe Varnier, Noureddine Zerhouni, and Pierre Dersin. Predictive maintenance of moving systems. In *2017 Prognostics and System Health Management Conference (PHM-Harbin)*, pages 1–6. IEEE, 2017. doi:10.1109/PHM.2017.8079111.

21 Satoshi Kato, Naoto Fukumura, Susumu Morito, Koichi Goto, and Narumi Nakamura. A mixed integer linear programming approach to a rolling stock rostering problem with splitting and combining. In *RailNorrköping 2019. 8th International Conference on Railway Operations Modelling and Analysis (ICROMA), Norrköping, Sweden, June 17th–20th, 2019*, pages 548–564. Linköping University Electronic Press, 2019.

22 Yaguo Lei, Naipeng Li, Liang Guo, Ningbo Li, Tao Yan, and Jing Lin. Machinery health prognostics: A systematic review from data acquisition to rul prediction. *Mechanical systems and signal processing*, 104:799–834, 2018. doi:10.1016/j.ymssp.2017.11.016.

23 Hongfei Li, Dhaivat Parikh, Qing He, Buyue Qian, Zhiguo Li, Dongping Fang, and Arun Hampapur. Improving rail network velocity: A machine learning approach to predictive maintenance. *Transportation Research Part C: Emerging Technologies*, 45:17–26, 2014. doi:10.1016/j.trc.2014.04.013.

24 Pin Lim, Chi Keong Goh, Kay Chen Tan, and Partha Dutta. Estimation of remaining useful life based on switching kalman filter neural network ensemble. In *Annual Conference of the PHM Society*, volume 6(1), 2014. doi:10.36001/phmconf.2014.v6i1.2348.

25 Richard M. Lusby, Jørgen Thorlund Haahr, Jesper Larsen, and David Pisinger. A branch-and-price algorithm for railway rolling stock rescheduling. *Transportation Research Part B: Methodological*, 99:228–250, 2017. doi:10.1016/j.trb.2017.03.003.

26 Kursat Rasim Mestav, Jaime Luengo-Rozas, and Lang Tong. State estimation for unobservable distribution systems via deep neural networks. In *2018 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2018. doi:10.1109/PESGM.2018.8586649.

27 Roberto Nappi, Gianluca Cutrera, Antonio Vigliotti, and Giuseppe Franzè. A predictive-based maintenance approach for rolling stocks vehicles. In *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, volume 1, pages 793–798. IEEE, 2020. doi:10.1109/ETFA46521.2020.9212183.

28 Weiwen Peng, Zhi-Sheng Ye, and Nan Chen. Bayesian deep-learning-based health prognostics toward prognostics uncertainty. *IEEE Transactions on Industrial Electronics*, 67(3):2283–2293, 2019. doi:10.1109/TIE.2019.2907440.

29 Felix Prause. A multi-swap heuristic for rolling stock rotation planning with predictive maintenance. In *Proceedings of the 11th International Network Optimization Conference (INOC), Dublin, Ireland, March 11-23, 2024*, pages 58–63, 2024. doi:10.48786/inoc.2024.11.

30 Felix Prause and Ralf Borndörfer. Construction of a test library for the rolling stock rotation problem with predictive maintenance. Technical Report 23-20, ZIB, Takustr. 7, 14195 Berlin, 2023.

31 Felix Prause, Ralf Borndörfer, Boris Grimm, and Alexander Tesch. Approximating rolling stock rotations with integrated predictive maintenance. *Journal of Rail Transport Planning & Management*, 30:100434, 2024. doi:10.1016/j.jrtpm.2024.100434.

32 Markus Reuther. *Mathematical optimization of rolling stock rotations*. PhD thesis, Technische Universität Berlin (Germany), 2017. doi:10.14279/depositonce-5865.

33 Rita P. Ribeiro, Pedro Pereira, and João Gama. Sequential anomalies: a study in the railway industry. *Machine Learning*, 105:127–153, 2016. doi:10.1007/s10994-016-5584-6.

34 Pegah Rokhforoz and Olga Fink. Hierarchical multi-agent predictive maintenance scheduling for trains using price-based approach. *Computers & Industrial Engineering*, 159:107475, 2021. doi:10.1016/j.cie.2021.107475.

35 Bhaskar Saha, Kai Goebel, Scott Poll, and Jon Christophersen. Prognostics methods for battery health monitoring using a bayesian framework. *IEEE Transactions on instrumentation and measurement*, 58(2):291–296, 2008. doi:10.1109/TIM.2008.2005965.

36 Simo Särkkä and Lennart Svensson. *Bayesian filtering and smoothing*, volume 17. Cambridge university press, 2023.

37 Shiliang Sun. A review of deterministic approximate inference techniques for bayesian machine learning. *Neural Computing and Applications*, 23:2039–2050, 2013. doi:10.1007/s00521-013-1445-4.

38 Per Thorlacius, Jesper Larsen, and Marco Laumanns. An integrated rolling stock planning model for the copenhagen suburban passenger railway. *Journal of Rail Transport Planning & Management*, 5(4):240–262, 2015. doi:10.1016/j.jrtpm.2015.11.001.

39 Meng-Ju Wu and Yung-Cheng Lai. Train-set assignment optimization with predictive maintenance. In *RailNorrköping 2019. 8th International Conference on Railway Operations Modelling and Analysis (ICROMA), Norrköping, Sweden, June 17th–20th, 2019*, pages 1131–1139. Linköping University Electronic Press, 2019.

A Computational Evaluation of the Weighted Swapping Procedure

In the following, we evaluate the effectiveness of the multi-swap heuristic utilizing the weighted swapping procedure presented in Section 5.2. We will refer to this approach as *weighted multi-swap* (WMS) and compare its results with those obtained by the algorithm from [29]. This uses a swapping probability of one-half and is therefore denoted as *equal probability multi-swap* (EPMS) in the following.

For testing, we use the instances of the preventive maintenance strategy described in Section 6. The detailed results are listed in Appendix A and show the values of the solutions obtained by EPMS and WMS after 360 seconds and after one hour of computation time. The best results for each instance are marked in bold.

We can observe that EPMS finds the best result after 360 seconds only eight times and that it was able to determine the best solution after one hour for only eight scenarios. The majority of these cases occur for instances T3 and T4, both of which are medium-sized scenarios in terms of the number of trips. In addition, there is a tie for the best solution for six and seven instances, respectively. WMS, on the other hand, could find the best solution after 360 seconds for 46 of the scenarios and achieved a better result after one hour for 45 of the 60 instances. These results show that the cost-oriented swapping used in WMS offers an advantage over the procedure employed in EPMS, both in terms of finding good solutions quickly and finding solutions with low costs.

Table 2 Results of EPMS and WMS after 360 seconds and after one hour of computation time.

Instance	EPMS after 360 s	WMS after 360 s	EPMS	WMS
T1-01	274,016	**272,090**	272,514	**272,090**
T1-02	276,858	**276,120**	276,544	**276,120**
T1-03	276,684	**275,607**	275,836	**275,607**
T1-04	272,514	**272,090**	272,514	**272,090**
T1-05	272,514	**272,090**	**272,090**	**272,090**
T1-06	**270,075**	**270,075**	**270,075**	**270,075**
T1-07	277,825	**276,544**	**276,544**	**276,544**
T1-08	272,514	**272,090**	**272,090**	**272,090**
T1-09	**272,090**	**272,090**	**272,090**	**272,090**
T1-10	**270,075**	**270,075**	**270,075**	**270,075**
T2-01	507,420	**494,036**	484,718	**468,615**
T2-02	511,781	**488,175**	496,993	**474,073**
T2-03	503,256	**492,784**	493,873	**471,772**
T2-04	486,411	**476,156**	480,048	**468,335**
T2-05	498,193	**466,304**	485,160	**459,522**
T2-06	500,491	**482,167**	480,559	**472,028**
T2-07	506,559	**485,151**	489,451	**474,629**
T2-08	496,731	**473,915**	484,462	**461,136**
T2-09	499,435	**475,127**	486,415	**461,121**
T2-10	500,269	**484,520**	481,735	**468,810**

Continued on next page

Table 2 – continued from previous page

Instance	EPMS after 360 s	WMS after 360 s	EPMS	WMS
T3-01	1,443,584	**1,417,645**	1,429,586	**1,415,815**
T3-02	1,466,915	**1,436,870**	1,459,579	**1,435,336**
T3-03	**1,471,162**	1,486,642	**1,446,601**	1,458,286
T3-04	1,465,143	**1,442,264**	1,448,476	**1,433,629**
T3-05	1,510,924	**1,441,830**	1,469,280	**1,432,776**
T3-06	**1,456,730**	1,475,282	**1,437,910**	1,456,083
T3-07	1,470,481	**1,443,980**	1,447,427	**1,441,852**
T3-08	1,430,814	**1,424,607**	**1,422,649**	1,424,329
T3-09	1,433,021	**1,419,417**	1,427,896	**1,418,864**
T3-10	1,449,332	**1,412,973**	1,427,362	**1,411,817**
T4-01	**216,340**	216,664	**214,181**	215,506
T4-02	**219,954**	221,504	218,414	**216,350**
T4-03	222,980	**218,718**	220,409	**216,814**
T4-04	**219,139**	222,581	**217,756**	218,094
T4-05	**220,460**	221,212	**219,281**	219,683
T4-06	219,149	**219,112**	216,181	**215,830**
T4-07	219,666	**218,058**	**216,285**	216,839
T4-08	**216,103**	218,168	**215,807**	215,936
T4-09	223,116	**220,051**	221,855	**215,975**
T4-10	219,098	**217,057**	216,856	**214,176**
T5-01	347,529	**346,924**	347,356	**346,111**
T5-02	**349,658**	**349,658**	**349,658**	**349,658**
T5-03	349,642	**348,547**	349,359	**347,590**
T5-04	**345,630**	**345,630**	345,630	**345,175**
T5-05	357,457	**350,891**	356,291	**350,891**
T5-06	343,027	**342,714**	343,027	**342,052**
T5-07	358,069	**352,783**	354,852	**351,553**
T5-08	347,051	**345,996**	347,002	**345,996**
T5-09	351,079	**350,766**	351,079	**350,417**
T5-10	**345,039**	**345,039**	345,039	**343,742**
T6-01	2,394,624	**2,384,517**	2,379,655	**2,370,161**
T6-02	2,450,322	**2,430,110**	2,432,481	**2,392,949**
T6-03	2,400,259	**2,390,905**	2,389,128	**2,380,847**
T6-04	2,390,995	**2,390,943**	2,390,995	**2,370,129**
T6-05	2,397,184	**2,384,236**	2,387,226	**2,372,693**
T6-06	2,442,339	**2,426,068**	2,416,792	**2,390,244**
T6-07	**2,393,514**	2,404,469	2,389,038	**2,382,052**
T6-08	2,393,010	**2,389,088**	2,384,886	**2,373,894**
T6-09	2,428,272	**2,422,976**	2,415,699	**2,389,409**
T6-10	2,386,473	**2,382,799**	2,377,996	**2,362,586**

B Tables of Computational Results

Table 3 Costs by type after averaging the scenarios of each instance when the preventive maintenance strategy (PM) is applied.

Instance	Operating costs	Deadhead costs	Maintenance costs	Trip costs	Total costs
T1	11,507	1,342	12,600	247,373	272,822
T2	64,438	25,595	9,800	362,386	462,219
T3	308,384	58,388	29,800	1,025,235	1,421,806
T4	20,137	14,248	12,000	165,630	212,015
T5	41,425	6,459	14,000	284,629	346,513
T6	536,986	63,994	42,200	1,726,235	2,369,415
Σ	982,877	170,025	120,400	3,811,487	5,084,790

Table 4 Costs by type after averaging the scenarios of each instance when the predictive maintenance strategy (PdM) is applied.

Instance	Operating costs	Deadhead costs	Maintenance costs	Trip costs	Total costs
T1	11,507	1,301	11,600	247,373	271,781
T2	64,438	29,939	9,800	362,386	466,564
T3	308,384	51,082	26,600	1,025,235	1,411,300
T4	20,137	15,303	11,000	165,630	212,069
T5	41,425	5,237	13,600	284,629	344,891
T6	536,986	76,852	37,600	1,726,235	2,377,673
Σ	982,877	179,714	110,200	3,811,487	5,084,278

Table 5 Costs by type after averaging the scenarios of each instance when the preventive maintenance strategy with Bayesian inference (PdM-B) is applied.

Instance	Operating costs	Deadhead costs	Maintenance costs	Trip costs	Total costs
T1	11,737	1,227	11,600	247,373	271,937
T2	64,438	21,550	9,800	362,386	458,174
T3	308,384	52,439	26,400	1,025,235	1,412,457
T4	22,151	13,850	10,200	165,630	211,831
T5	41,425	4,523	12,400	284,629	342,976
T6	536,986	74,392	37,400	1,726,235	2,375,013
Σ	985,121	167,981	107,800	3,811,487	5,072,388

Table 6 The cost differences of PdM-B compared to PM and PdM in percent after averaging the scenarios of each instance.

Instance	Cost diff. in %		Cost diff. without fixed costs in %	
	PM	PdM	PM	PdM
T1	−0.32	+0.06	−3.48	+0.64
T2	−0.86	−1.80	−4.05	−8.05
T3	−0.66	+0.08	−2.36	+0.30
T4	−0.09	−0.11	−0.40	−0.51
T5	−1.02	−0.56	−5.72	−3.18
T6	+0.24	−0.11	+0.86	−0.41
Combined	−0.24	−0.23	−0.97	−0.93