

The Maintenance Location Choice Problem for railway rolling stock

Zomer, Jordi ; Bešinović, Nikola; de Weerd, Mathijs M.; Goverde, Rob M.P.

DOI

[10.1016/j.jrtpm.2021.100268](https://doi.org/10.1016/j.jrtpm.2021.100268)

Publication date

2021

Document Version

Final published version

Published in

Journal of Rail Transport Planning & Management

Citation (APA)

Zomer, J., Bešinović, N., de Weerd, M. M., & Goverde, R. M. P. (2021). The Maintenance Location Choice Problem for railway rolling stock. *Journal of Rail Transport Planning & Management*, 20, 1-14. Article 100268. <https://doi.org/10.1016/j.jrtpm.2021.100268>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Rail Transport Planning & Management

journal homepage: www.elsevier.com/locate/jrtpm

The Maintenance Location Choice Problem for railway rolling stock

Jordi Zomer, Nikola Bešinović*, Mathijs M. de Weerd, Rob M.P. Goverde

Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Keywords:

Railway planning
Maintenance
Rolling stock
Location choice
Optimization

ABSTRACT

Due to increasing railway use, the capacity at railway yards and maintenance locations is becoming limiting to accommodate existing rolling stock. To reduce capacity issues at maintenance locations during nighttime, railway undertakings consider performing more daytime maintenance, but the choice at which locations personnel needs to be stationed for daytime maintenance is not straightforward. Among other things, it depends on the planned rolling stock circulation and the maintenance activities that need to be performed. This paper presents the Maintenance Location Choice Problem (MLCP) and provides a Mixed Integer Linear Programming model for this problem. The model demonstrates that for a representative rolling stock circulation from The Netherlands Railways a substantial amount of maintenance activities can be performed during daytime. Also, it is shown that the location choice delivered by the model is robust under various time horizons and rolling stock circulations. Moreover, the running time for optimizing the model is considered acceptable for planning purposes.

1. Introduction

In many countries, the demand for rail transport is increasing. For instance, in The Netherlands the total number of passenger kilometers increased with over 30% in the last 20 years. To accommodate the increasing travel demand, the utilization of the available rolling stock units may increase and more rolling stock units may be added to the network.

In order for rail transport to function properly, the rolling stock that operates on the railway network needs to receive maintenance on a regular basis. Rolling stock includes all units that move over railway tracks, such as locomotives, passenger wagons and freight wagons. Rolling stock units are fixed compositions that consist of various carriages. The aim of maintenance is to ensure that the rolling stock remains available, safe and comfortable for passengers (Dinmohammadi et al., 2016). Maintenance activities can be divided into two categories: (1) regular maintenance, corresponding to maintenance activities with higher frequencies (days or weeks) and shorter duration (at most several hours), and (2) heavy maintenance, corresponding to maintenance types with lower frequencies (every several months or less) and longer duration (up to multiple days) (Andrés et al., 2015). The current work focuses on regular maintenance. Examples of regular maintenance activities that apply at NS are interior cleaning, performed every 24 h during approximately an hour, and technical checks, performed every 48 h during approximately 10 min.

In general, the maximum interval between consecutive maintenance activities is governed by strict rules that are imposed by railway authorities. Maintenance activities are carried out at so-called maintenance locations, which are railway yards with maintenance facilities, spread over the network. For a maintenance location to be operational, it is necessary that personnel is stationed at that maintenance location. The number of personnel stationed at each location is a railway undertaking's decision and determines, together with the maintenance location design, the availability of maintenance locations. An operator can choose to open a facility at certain moments of the day or night by stationing personnel at this location. In particular, a distinction can be made

* Corresponding author.

E-mail address: n.besinovic@tudelft.nl (N. Bešinović).<https://doi.org/10.1016/j.jrtpm.2021.100268>

Received 22 December 2020; Received in revised form 6 July 2021; Accepted 26 July 2021

Available online 4 October 2021

2210-9706/© 2021 The Authors.

Published by Elsevier Ltd.

This is an open access article under the CC BY license

<http://creativecommons.org/licenses/by/4.0/>.

between daytime operations (which means that a location is opened during the day) and nighttime operations (which means that a location is opened during the night). This distinction is clearly visible in The Netherlands where maintenance is usually carried out during nighttime because the rolling stock is mostly in use during daytime.

Increased railway use has several effects, two of the most important being (1) more rolling stock units on the network and (2) higher utilization leading to less time to actually do maintenance. Both affect the pressure on maintenance locations directly.

Due to increasing use of the railway network, two developments are observed. First, more rolling stock units are purchased and added to the network. Second, existing rolling stock units are utilized more, leaving less time for maintenance. As a result of these developments, the capacity of maintenance locations during nighttime is under pressure and reaching its limits. As a result, the Netherlands Railways (Nederlandse Spoorwegen, abbreviated NS) are considering to perform more maintenance activities during daytime. This raises the question at which existing maintenance locations teams need to be stationed to perform daytime maintenance, referred to as the maintenance location choice. These maintenance location choice decisions can be made several months before operations, when the rolling stock circulation, which describes all planned rolling stock movements for a given period, is usually fixed for the next period. The challenge is to choose maintenance locations such that more maintenance activities can be performed during daytime instead of during nighttime, without modifying the rolling stock circulation. Also, due to the fact that there are fixed costs associated to opening a maintenance location for daytime maintenance, the number of maintenance locations that can be opened for daytime maintenance is assumed to be restricted and given. This reflects the fact that these locations are scarce and that the railway operator decides explicitly upon the number of locations that it wants to open for maintenance during daytime.

The problem of determining an optimal maintenance location choice that minimizes the number of maintenance activities during nighttime is complex, due to the symbiotic relationship between the maintenance location choice and the maintenance schedule. The maintenance location choice influences the maintenance schedule that can be made, which determines the number of activities performed during nighttime, which in turn is the main determinant for the optimal maintenance location choice. There are many possible location choices to evaluate. For each of these possible location choices, associated (optimal) maintenance schedules should also be taken into account, while the maintenance scheduling aspect is a complicated task as well due to the boundless amount of possibilities. The problem is complicated even further by the fact that two time windows apply and that priority should be given to daytime maintenance activities. Therefore, an automated procedure that optimizes the maintenance location choice by taking into account the scheduling aspect is required.

In planning, commonly three types of problems can be distinguished being strategical (long-term, typically multiple years before operations), tactical (medium-term, from week to one year before operations) and operational (short-term, within several days before operations). When a plan has to be adjusted during operations due to certain disturbances or disruptions, then it is referred to as (real-time) traffic management. In the problem under consideration, it is assumed that the rolling stock circulation is fixed. This assumption is justified since currently, the choice regarding which locations to open for daytime maintenance is made at a moment in time when the rolling stock circulation has already been determined. Moreover, the maintenance location choice needs to be determined several weeks in advance to be able to assign the right amounts of staff. As a result, the problem considered in the current research can be classified as a tactical planning problem.

This research introduces the Maintenance Location Choice Problem (MLCP), which formulates an optimal maintenance location choice to facilitate timely regular maintenance of all rolling stock, without changing their planned rolling stock circulation. The aim of the model is to reduce the amount of work that needs to be performed during nighttime. The model performance is demonstrated on realistic instances provided by NS.

The main contributions of the paper are the following:

- It proposes a new mathematical problem formulation for a maintenance location choice problem that does not alter the planned rolling stock circulation. In this way it addresses the location choice problem at a tactical level.
- It explicitly distinguishes between nighttime maintenance and daytime maintenance aiming to reduce the pressure on the capacity of maintenance locations during nighttime by performing more maintenance during daytime.
- It demonstrates the performance of the proposed model on real-life instances on the Dutch railway network, providing the main Dutch railway undertaking NS with suggestions on which maintenance locations would best be opened, with information on the amount of work that can be moved to daytime, and with an automated method to assign maintenance activities to maintenance opportunities.

Section 2 examines the scientific literature. Section 3 introduces the problem and Section 4 develops a model for solving it. The model's performance is investigated using various scenarios in Section 5. Section 6 gives the key takeaways and identifies some interesting directions for future research.

2. Literature review

This section discusses papers on rolling stock maintenance location choice. It covers both railway and aviation papers, given the systematic similarities of the two domains related to vehicle (rolling stock and aircraft) maintenance. The characteristics of papers are summarized and the differences are mapped out. The existing gaps are defined in the final part of the section.

Tönissen et al. (2019) aimed at building new maintenance locations in the railway network, which represents a decision problem at the strategic level due to the long-term nature of these decisions. Their focus was on determining optimal maintenance locations under unknown/uncertain train lines and fleet size (and mix). They proposed a two-stage stochastic mixed integer programming

model, in which the first stage is to open a facility, and in the second stage to minimize the routing cost for the first-stage location decision for each line plan scenario.

Tönissen and Arts (2018) built on Tönissen et al. (2019) by including recovery costs of maintenance location decisions, unplanned maintenance, multiple facility sizes and economies of scale (providing that a location twice as big is not twice as expensive). Since, as a result, the second-stage problem becomes NP-hard, an algorithm was provided with the aim to avoid having to solve the second stage for every scenario.

Tönissen and Arts (2020) further extended the ideas developed by Tönissen et al. (2019) and Tönissen and Arts (2018). They noted that in reality, the available resources and equipment influence the locations where specific rolling stock types can be maintained. To account for this, they included so-called allocation restrictions. The inclusion of these restrictions is important for the practical usability of the model, as is shown by the fact that they significantly influence the solutions found.

Canca and Barrena (2018) considered the simultaneous rolling stock allocation to lines and choice for depot locations in a rail-rapid transit context. They proposed a Mixed Integer Linear Programming (MILP) formulation which appears hard to solve. Therefore they proposed a three-step heuristic approach determining first the minimum number of vehicles needed for each line, subsequently the routes of rolling stock on each line, and lastly the circulation of rolling stock on lines over multiple days together with the depot choice.

Zomer (2019) considered the railway network in The Netherlands and designed a simulation model to estimate the expected effects of carrying out daytime maintenance at given maintenance locations using historical data. Although the use of historical rolling stock data has advantages, since it accurately describes reality (as opposed to rolling stock planning data), it cannot be used for situations in the future for which only planning data is available. Also, although simulation offers the opportunity to investigate effects for various scenarios of daytime maintenance locations, it cannot be used to systematically optimize the maintenance location choice.

In the area of aviation, Feo and Bard (1989) introduced the problem of assigning aircraft to given flights and simultaneously optimizing the number of maintenance facilities. They solved the problem as a minimum cost multi-commodity flow problem. They come up with a two-phase heuristic approach to solve it: firstly, many possible trip patterns for individual aircraft are computed; secondly, the most possible patterns are combined to solve the complete problem.

The work of Gopalan (2014) is closely related to the work of Feo and Bard (1989). They assumed that a set of daily routes are given for each aircraft, although these routes are not yet assigned to specific aircraft. Each route needs to be connected to a route on the next day in such a way that the routing passes through a maintenance location with some given periodicity. The objective was to minimize the number of maintenance locations (one of the differences from Feo and Bard, 1989, that considers cost minimization). They proposed four heuristics to solve their problem.

The aforementioned research focused on maintenance location choice. Another area of research in the field of rolling stock maintenance is maintenance scheduling, which is relevant due to the fact that the optimal maintenance location choice typically depends on the schedules for the rolling stock units that need to be maintained. In the literature, maintenance scheduling is often combined with the assignment of rolling stock units (or aircraft) to train trips (or flight legs) such that maintenance constraints are satisfied. Some examples in railways are Herr et al. (2017), Andrés et al. (2015), Wagenaar and Kroon (2015) and Van Hövell (2019) and some examples in aviation are Clarke et al. (1997), Gopalan and Talluri (1998), Sarac et al. (2006) and Gopalan (2014). Those problems may be solved using Mixed Integer Linear Programming (MILP) models (e.g. Herr et al., 2017; Andrés et al., 2015; Wagenaar and Kroon, 2015), sometimes extended with solution approaches such as column generation (e.g. Clarke et al., 1997). Other problems are solved using heuristics (e.g. Gopalan and Talluri, 1998). The model developed in this paper also contains important components of maintenance scheduling that incorporate the timely performance of maintenance activities. In particular, time slots are defined during which maintenance actions shall be performed, i.e. maintenance opportunities (introduced in Section 3), while the exact start time and end time of actions are not considered. The way in which the current contribution includes these maintenance scheduling components differs from what is typically found in the literature in at least two ways. Firstly, the current contribution incorporates these maintenance scheduling components in conjunction with location choice. Secondly, since the proposed problem is solved on a tactical level, the maintenance scheduling component is addressed assuming a fixed rolling stock circulation.

Table 1 summarizes the discussed papers. It shows for each paper whether it is written in the aviation (A) or in the railway (R) context, whether it considers maintenance constraints, whether it considers locations choice, whether it considers timing of maintenance for every individual moving unit (MU, can be either rolling stock or aircraft), the adopted solution approach and the planning stage to which the problem relates.

Some more explanation may be necessary on the column indicating maintenance timing of individual rolling stock units. Even if a paper considers maintenance (for example by requiring that each trip path shall not exceed some specified length), it need not provide information on where and when each maintenance activity shall be performed. When a paper indeed includes the latter aspect, it is considered to address maintenance timing of individual MUs.

This literature review indicates that several aspects have not been addressed in the currently existing literature. First, although variants of problems relating to rolling stock maintenance location have been investigated before, this problem has not been tackled at a tactical level, i.e. deciding on which maintenance locations to open for given rolling stock circulation. Second, no research is available that explicitly distinguishes between daytime maintenance and nighttime maintenance. This distinction is relevant for the situation in the Netherlands, where nighttime maintenance is standard and recent developments have led companies to investigate daytime maintenance as well. Third, although maintenance location choice for the Netherlands has been addressed from a simulation perspective in Zomer (2019), no systematic method has been developed yet to find a set of locations for daytime maintenance to minimize the number of nighttime maintenance activities.

Table 1
Overview of the aviation (A) and railway (R) maintenance literature discussed in Section 2.

	A/R	Maint. considered	Location choice	Indiv. maint. timing	Solution approach	Planning stage
Tönissen et al. (2019)	R	x	x		Optimization	Strategic
Tönissen and Arts (2018)	R	x	x		Optimization	Strategic
Tönissen and Arts (2020)	R	x	x		Optimization	Strategic
Canca and Barrena (2018)	R		x		Heuristic	Strategic
Feo and Bard (1989)	A	x	x		Heuristic	Strategic
Gopalan (2014)	A	x	x		Heuristic	Strategic
Van Hövell (2019)	R	x		x	Optimization	Tactical
Zomer (2019)	R	x	x		Simulation	Tactical
Current	R	x	x	x	Optimization	Tactical



Fig. 1. Example of a rolling stock circulation.

Table 2
Example of the MOs corresponding to the rolling stock circulation example from Fig. 1.

MO	Day	Location	Start time	End time
1	1	Hrl	10:41	16:19
2	1	Ekz	19:52	20:09
3	1	Mt	23:31	00:01
4	2	Ehv	01:06	05:34

3. Problem description

The problem that is addressed in the current research is defined as the *Maintenance Location Choice Problem (MLCP)*. The input of the **MLCP** consists of a rolling stock circulation, a set of potential maintenance locations and a set of maintenance activities that need to be scheduled. The circulation describes the planned movements of the rolling stock. The maintenance locations are locations where maintenance activities can be carried out. These maintenance activities each have a given duration and constraints on the time between consecutive ones. The goal of the **MLCP** is to find an optimal choice of locations used for maintenance, from the set of potential maintenance locations, such that all maintenance activities can be completed in time. Given that rolling stock movements and out of service time in the network can have a direct influence on where each rolling stock unit can be maintained, it is important to consider the rolling stock circulation and some scheduling aspects. The relevant terms and concepts for the **MLCP** are given in the following.

Rolling stock circulation. A rolling stock circulation contains for each rolling stock unit a list of trips that a rolling stock unit is scheduled to perform. These trips include an origin and destination position (which are often train stations), and the corresponding planned departure and arrival times. An example of a rolling stock circulation is found in Fig. 1. Squares represent arrivals or departures of rolling stock. Inside the squares, the station abbreviation and the departure or arrival time at this station is given. Solid lines represent time intervals where a rolling stock unit is used for a train service. For example, from 07.09 to 10.41, the depicted rolling stock unit is planned to be used for a train service between Ekz (Enkhuizen) and Hrl (Heerlen). During the day the rolling stock unit is also planned to visit Mt (Maastricht) and Eindhoven (Ehv). Dashed lines represent time intervals where a rolling stock unit is not in service. For example, between 10.41 and 16.19, the depicted rolling stock unit is not in service and will be standing still (i.e. parked) at Hrl. In practice, when a rolling stock unit is not in service during daytime, it is parked at a maintenance location. Therefore, performing more daytime maintenance activities has no or limited impact on operations and the rolling stock circulation itself will not be modified.

Maintenance opportunities. Maintenance can be carried out during *maintenance opportunities (MOs)*. A MO represents the time between two consecutive train services. During MOs, rolling stock units are typically parked at maintenance locations close to stations. Note that some maintenance tasks can be performed at those parking tracks (e.g. interior cleaning, technical checks), while others may need extra shunting movements (e.g. external washing). Table 2 indicates the MOs corresponding to the rolling stock circulation from Fig. 1. All MOs (even short ones) can potentially be used for maintenance activities. When a maintenance activity is assigned to a MO, this implies that the maintenance activity needs to be performed between the start and end time of this MO. The exact scheduled time of the maintenance activity is not determined in the scope of this research.

Two time windows are considered: daytime and nighttime. A MO can be during daytime or nighttime. This division is especially relevant for NS, since it is currently considering a transition from performing maintenance only during nighttime to performing

maintenance during both nighttime and daytime (Zomer, 2019; Van Hövell, 2019). The time window for daytime maintenance is set at 07.00–19.00, the time window for nighttime maintenance corresponds to the interval 19.00–07.00. Some MOs may be partly during daytime and partly during nighttime (this occurs, for example, when a MO lasts from 18.00 to 20.00). For these MOs it is not straightforward whether maintenance would be carried out during the day or during the night (i.e. whether the MO should be considered to be during daytime or during nighttime). For these reasons, the following simplification is applied in the current problem: a MO is marked to be *during daytime* if and only if both its start time and its end time are between 07.00 and 19.00 of the same day; a MO is marked to be *during nighttime* in all other cases. Note that, although this assumption is reasonable in most cases, there are some occasions where it leads to misclassification. For example, a MO starting at 11.00 (during daytime) and ending at 19.01 (just after the start of the nighttime time window) would be classified to be a nighttime MO whereas maintenance scheduled in it would probably be performed during daytime. In this case, the MO is classified as a nighttime MO but the part of this MO during the nighttime time window is so small that no maintenance activity can be performed during nighttime. However, it must be noted that MOs like these are very rare. In practice, most rolling stock units are in service during peak hours, and hence there is only a limited number of MOs that are partially during daytime and partially during nighttime. The simple rule to classify MOs as suggested above yields appropriate results in almost all cases.

Maintenance activities. It is assumed that a set of regular maintenance types of arbitrary size is known, and that for each maintenance type a fixed duration and a fixed maximum time interval between subsequent maintenance activities is given. The time necessary for shunting between the maintenance location and the station is included in the maintenance duration. Those shunting movements, necessary to park rolling stock units, are often already included within the rolling stock planning and hence add no or limited additional pressure on the capacity of the railway network. The maximum intervals are determined in terms of time (as opposed to distance), which is common for regular maintenance types at NS. For each rolling stock unit, all maintenance types need to be applied within the respective specified intervals. Specifically, maintenance activities are assigned to maintenance opportunities, which is referred to as the *maintenance scheduling*. It is assumed that maintenance activities are carried out in a subsequent manner and cannot overlap. Furthermore, it is assumed that the number of hours at the start of the planning horizon since the last maintenance activity for each rolling stock unit is given. Also, the set of maintenance activities is assumed to be the same for all rolling stock units, which may in practice be true for at least some maintenance types (such as technical checks).

It is not necessarily the case that for each rolling stock circulation a feasible maintenance schedule can be found in which all maintenance activities are planned within the required intervals. The railway operator is responsible for providing a rolling stock circulation in such a way that a feasible maintenance schedule exists. If a solution exists with night-time-only maintenance, then a solution with the additional option for daytime maintenance is also guaranteed to exist.

Maintenance locations. Maintenance can be carried out at a potential maintenance location. The set of potential maintenance locations is given. For each of these locations, it can be decided whether a location should be opened or not. This decision is referred to as the *location choice*. Each location can be opened during nighttime or during daytime, that is, for each maintenance location, there are four possible outcomes: the location is not used at all, the location is used for daytime maintenance, the location is used for nighttime maintenance, or the location is used for both. It is assumed that each maintenance type can be applied at every location.

It is assumed that all maintenance activities can be performed at all maintenance locations. This is mostly true in The Netherlands for regular maintenance activities like interior cleaning and technical checks.

The reachability of maintenance locations is incorporated implicitly since the current research takes the given rolling stock circulation as an input. Maintenance activities can take place only at those locations where rolling stock units are located according to the rolling stock circulation.

One may argue that, for a correct assignment of maintenance activities to MOs, the available capacity at maintenance locations should be taken into account, entailing a precise planning of the time each maintenance activity is performed and by which maintenance team. However, in the current research the available number of maintenance teams is not restricted. This is justifiable in the light of the current way of working at NS, where a similar approach prevails: the assignment of maintenance activities to MOs is performed first and separately from a planning of maintenance teams at the maintenance location level. Moreover, given the significant capacity issues during nighttime, NS will appreciate any work that can be moved to daytime if that helps in reducing the capacity problems during nighttime, with the capacity during daytime being subordinate to that goal.

Planning horizon. The planning horizon in the current research is equal to the planning horizon in the rolling stock circulation. In other words, a maintenance schedule is determined for the entire time horizon of the rolling stock circulation, and as a result, the optimal location choice is valid for the length of this time horizon as well. In The Netherlands, rolling stock circulations are available for periods of eight weeks. This implies that the planning horizon in the current research is also fixed at eight weeks.

Objective. The objective of the **MLCP** is to assign maintenance activities to MOs and determine for each location whether it is open during daytime and/or nighttime, satisfying the intervals between maintenance activities and the other constraints (such as the maximum number of locations that can be opened during daytime), in such a way that the number of nighttime maintenance activities is minimized. This goal is relevant in practical contexts where the capacity of maintenance locations at nighttime is under pressure. As a technical aside, a small penalty applies for any maintenance activity, to avoid the situation that more daytime maintenance activities than necessary are planned.

4. Model development

The current section provides a Mixed Integer Linear Programming (MILP) model to solve the **MLCP**. For models formulated as a MILP, sophisticated solvers exist that deliver optimal solutions, and such models have been used for similar problems in the past (e.g. Tönissen et al., 2019). A downside of the use of MILP in general is the fact that the computation time can be relatively high. The aim therefore is to provide a MILP model that contains not too many decision variables. Furthermore, given the fact that the **MLCP** is a problem on the tactical planning level, moderate computation times can be acceptable when looking for optimal solutions, as opposed to applications in for instance traffic management (e.g. Wagenaar and Kroon, 2015), where short computational times are the main goal. An overview of all mathematical symbols used in the current section is given in Table 3.

4.1. Mathematical notation

Let I be the set of rolling stock units considered in the current problem and let $i \in I$ be the index used to indicate a specific rolling stock unit. Let $T > 0$ be the length of the planning horizon in hours. Let L denote the set of potential maintenance locations.

Maintenance opportunities. Let J_i denote the set of MOs for rolling stock unit $i \in I$. The location of a rolling stock unit i at MO $j \in J_i$ is denoted by $l_{ij} \in L$. The start time of MO $j \in J_i$ is denoted by $s_{ij} \in \mathbb{R}$ and the end time is denoted by $e_{ij} \in \mathbb{R}$, where time is given as the hours that passed since midnight of the first day in the planning horizon. For example, a MO j of a rolling stock i on the fourth day starting at 9:30 and ending at 13:15 would be represented by $s_{ij} = 81.5$ and $e_{ij} = 85.25$.

Each day is divided into two non-overlapping time windows: the daytime time window and the nighttime time window. Each MO is classified as either a daytime or a nighttime MO. Let d_{ij} indicate whether a MO is classified to be during daytime or during nighttime: let $d_{ij} = 1$ if MO $j \in J_i$ for rolling stock unit i is during daytime and let $d_{ij} = 0$ if a MO $j \in J_i$ for rolling stock unit $i \in I$ is nighttime. Let δ^D be the hour of the day when the daytime maintenance window starts and let δ^N be the hour of the day when the nighttime maintenance window starts. Unless stated otherwise, $\delta^D = 7.00$ and $\delta^N = 19.00$. A MO is classified to be during daytime if and only if both its start and end time are during daytime of the same day, i.e.

$$d_{ij} = \begin{cases} 1 & \text{if } \delta^D \leq e_{ij} \bmod 24 < \delta^N \\ 0 & \text{else} \end{cases}.$$

Maintenance activities. Maintenance activities can be of different maintenance types. Let K be the set of maintenance types. For each maintenance type $k \in K$, let $v_k \in \mathbb{R}^+$ be its duration in hours and let $o_k \in \mathbb{R}^+$ be the maximum interval between two consecutive maintenance activities of maintenance type k in hours.

Maintenance locations. A potential maintenance location can be opened during daytime (or not) and it can be opened during nighttime (or not). Let $y_l^D \in \{0, 1\}$ be a binary variable equal to 1 if location $l \in L$ is available for daytime maintenance and 0 otherwise. Let $y_l^N \in \{0, 1\}$ be a binary variable equal to 1 if location $l \in L$ is available for nighttime maintenance and 0 otherwise. The number of potential maintenance locations that can be opened during daytime is restricted. Let L_{max}^D denote the maximum number of potential maintenance locations that can be opened during daytime. The location corresponding to MO j for rolling stock unit i is indicated with $l_{ij} \in L$.

Maintenance schedule. Maintenance activities are assigned to maintenance opportunities. Let $x_{ijk} \in \{0, 1\}$ be a binary variable equal to 1 if maintenance of type k is performed to rolling stock unit $i \in I$ at MO $j \in J_i$, and 0 otherwise. It is required that the total time available at MO j is not exceeded: $\sum_{k \in K} x_{ijk} v_k \leq e_{ij} - s_{ij}$.

Furthermore, a MO j can only be used if the corresponding location is open at the moment of the MO. Therefore, $d_{ij} = 0$ and $y_{l_{ij}}^N = 0$ implies $x_{ijk} = 0, \forall k \in K$. In the same way, $d_{ij} = 1$ and $y_{l_{ij}}^D = 0$ implies $x_{ijk} = 0, \forall k \in K$.

Moreover, the intervals between successive maintenance activities of the same type should satisfy the given criteria. The interval between two MOs $j, j' \in J_i, j \neq j'$ is measured from the end time of the first MO to the start time of another MO: $s_{ij'} - e_{ij}$. If activity $k \in K$ is scheduled for rolling stock unit $i \in I$ in MO j , then the next maintenance activity should be scheduled such that the interval constraints are satisfied. Let $V_{ijk} \subset J_i$ denote the set of maintenance opportunities for rolling stock unit $i \in I$ that start after the end of MO $j \in J_i$ but earlier than o_k hours after the end of MO $j \in J_i$. This set is $V_{ijk} = \{p \in J_i : e_{ij} < s_{ip} \leq e_{ij} + o_k\}$ for $j \in J_i$. It is then required that for all $i \in I, j \in J_i$, the following implication holds:

$$x_{ijk} = 1 \implies \exists j' \in V_{ijk} : x_{ij'k} = 1.$$

Observe that a next maintenance activity only needs to be scheduled if maintenance needs to be carried out within the current planning horizon, that is, if $e_{ij} + o_k \leq T$.

Let b_{ik} be the number of hours since the last maintenance activity of type k for rolling stock unit i at midnight of the first day. Then let $V_{i0k} = \{p \in J_i : s_{ip} \leq o_k + b_{ik}\}$.

The current research assumes that maximum intervals are defined in terms of time. For applications that define the maximum intervals in terms of distance, the sets V_{ijk} and V_{i0k} need to be defined in terms of distance; the rest of the modeling approach remains unchanged.

Table 3
Mathematical notation in the **MLCP** model, for reference.

	Variable	Significance
Sets	I	The set of rolling stock units considered in the current problem.
	J_i	The set of all MOs for rolling stock unit i
	K	The set of maintenance activity types.
	L	The set of potential maintenance locations.
	$V_{jk} \subseteq J_i$	The set of MOs of which at least one should be used for maintenance type k for rolling stock unit i if maintenance type k was performed in MO j for rolling stock unit i .
Indices	$i \in I$	Index used to identify any rolling stock unit
	$j \in J_i$	Index used to identify any MO
	$k \in K$	Index used to identify any maintenance activity type.
	$l \in L$	Index used to identify any potential maintenance location.
	$l_{ij} \in L$	Index used to identify the location corresponding to MO j for rolling stock unit i .
Parameters	$b_{ik} \in \mathbb{R}^+$	The number of hours since the last maintenance activity of type k for rolling stock unit i at midnight of the first day.
	$d_{ij} \in \{0, 1\}$	Binary input parameter used to indicate whether MO i for rolling stock unit j is during the day ($d_{ij} = 1$) or during the night ($d_{ij} = 0$).
	$e_{ij} \in \mathbb{R}$	The end time of MO j for rolling stock unit i .
	$L_{max}^D \in \mathbb{N}$	The maximum number of daytime maintenance locations that may be opened.
	$o_k \in \mathbb{R}^+$	The maximum interval between two subsequent maintenance activities of type k .
	$s_{ij} \in \mathbb{R}$	The start time of MO j for rolling stock unit i
	$\delta^D \in (0, 24)$	The hour when the daytime maintenance window starts.
	$\delta^N \in (0, 24)$	The hour when the nighttime maintenance window starts.
	$v_k \in \mathbb{R}^+$	The duration of maintenance activity of type k in hours.
	y_l^N	Binary input variable equal to 1 if location $l \in L$ is available for nighttime maintenance and 0 otherwise.
	Decision variables	$x_{ijk} \in \{0, 1\}$
y_l^D		Binary decision variable equal to 1 if location $l \in L$ is available for daytime maintenance, and 0 otherwise.

4.2. Model formulation

The **MLCP** model aims to find x_{ijk} and y_l satisfying the above described constraints that minimize the number of maintenance activities during nighttime. In this model, the decision variables are x_{ijk} ($i \in I, j \in J_i, k \in K$) and y_l^D ($l \in L$). The variables y_l^N ($l \in L$) are considered to be given in the input.

The model can then be formulated as follows:

$$\min \sum_{i \in I} \sum_{j \in J_i} \sum_{k \in K} x_{ijk}(1 - d_{ij}) + \epsilon \sum_{i \in I} \sum_{j \in J_i} \sum_{k \in K} x_{ijk} \tag{1}$$

subject to

$$1 \leq \sum_{p \in V_{i0k}} x_{ipk} \quad \forall i \in I, k \in K \tag{2}$$

$$x_{ijk} \leq \sum_{p \in V_{ijk}} x_{ipk} \quad \forall i \in I, j \in J_i, k \in K : e_{ij} + o_k \leq T \tag{3}$$

$$x_{ijk} \leq y_{l_{ij}}^D \cdot d_{ij} + y_{l_{ij}}^N \cdot (1 - d_{ij}) \quad \forall i \in I, j \in J_i, k \in K \tag{4}$$

$$\sum_{k \in K} x_{ijk} v_k \leq e_{ij} - s_{ij} \quad \forall i \in I, j \in J_i \tag{5}$$

$$\sum_{l \in L} y_l^D \leq L_{max}^D \tag{6}$$

$$x_{ijk} \in \{0, 1\}, y_l^D \in \{0, 1\}. \tag{7}$$

The objective function (1) minimizes the number of nighttime maintenance activities. The second term penalizes every maintenance activity with an arbitrarily small (in any case smaller than 1), but strictly positive penalty cost ϵ in order to avoid unnecessary maintenance activities being performed. Constraints (2) and (3) enforce that intervals between successive maintenance activities are satisfied, the first assuring that the first maintenance activity is scheduled in time and the latter assuring that all subsequent maintenance activity are scheduled in time. Constraints (4) ensure that maintenance can only be executed at a location that is opened. Constraints (5) take account of the requirement that the duration of maintenance may not exceed the total time of a MO. The number of locations for daytime maintenance is restricted by constraint (6). Constraints (7) ensure that the integer decision variables are binary.

5. Results

The main goal of the experiments performed in the current section is to investigate the validity of the model by studying the outcomes under various inputs and to make practical recommendations regarding the number and the choice of maintenance locations to be opened. Also, some numbers regarding computational efficiency are reported.

5.1. Experimental design

The proposed **MLCP** model is tested on real data from the Dutch railway network (NS, 2018). The data set contains planned rolling stock movements of trips that are operated by NS. A rolling stock circulation covers an 8-week period, which represents a basic planning period, and is issued several weeks before the start of such a period. A data set for such a period is referred to as Basic Day update (in Dutch *BasisDag update*, BDU). This data is an appropriate choice for the current research since it offers a complete view of all rolling stock movements as expected several weeks before operations. In the current research, 10 different BDUs have been used. Each of these BDUs are individual data sets on which the **MLCP** model can be run. The period covered by these 10 BDUs starts on 10-12-2017 and ends on 1-9-2019. This data set includes all rolling stock units operated by NS, meaning it is of a realistic size and regards a large part of the Dutch rail transport network. The number of rolling stock units in each BDU varies between 820 and 993. Due to the acquisition of new rolling stock units, an increasing trend in the number of rolling stock units operating on the network is visible, leading to the fact that each BDU does not need to contain the same number of rolling stock units.

NS operates rolling stock units of different types such as, for example, VIRM (used mainly for intercity services) and FLIRT (used mainly for local services). In reality, all rolling stock units have different characteristics, meaning that maintenance types may differ for various rolling stock types, but as indicated in Section 3 of the current research, we ignore the differences between the rolling stock maintenance times. Two maintenance types are considered ($K = \{1, 2\}$): maintenance type 1 has a duration of 0.5 h ($v_1 = 0.5$) and a maximum interval of 24 h between successive maintenance activities ($o_1 = 24$), and maintenance type 2 having a duration of 1 h ($v_2 = 1$) and a maximum interval of 48 h between successive maintenance activities ($o_2 = 48$). Rolling stock units are assumed to be as-good-as-new at the start of the planning horizon ($b_{ik} = 0$ for all $i \in I, k \in K$). The start of the daytime time window δ^D is fixed at $\delta^D = 07.00$ and the start of the nighttime maintenance window δ^N is fixed at $\delta^N = 19.00$. The technical parameter ε is fixed at $\varepsilon = 0.001$.

The following parameters are varied to construct the experiments investigated in Section 5.2.

- We consider 10 BDUs $\beta \in \{1, \dots, 10\}$, each representing a period of approximately 8 weeks. Unless stated otherwise, BDU set 10 is used, which is valid in the period 9-6-2019 to 1-9-2019.
- The set of rolling stock units v , indicating which rolling stock units are considered in the analysis, is varied. Smaller-sized sets are used to avoid large computation times and thus to allow validating the model by solving multiple instances; larger-sized sets are more realistic which are used to demonstrate the practical performance of the model. In addition, the use of sets of various sizes also enables to investigate the effect of the number of rolling stock units on model outcomes. In most sets, rolling stock units of type VIRM are used since these rolling stock units typically spread out over the entire country visiting many different maintenance locations and therefore tend to provide most interesting insights on the network level. Specifically, the following sets of rolling stock units are considered:
 - Set 1, with 10 rolling stock units of type VIRM4. The 10 rolling stock units that occur first in the rolling stock schedule are selected.
 - Set 2, with 20 rolling stock units of type VIRM4. The 20 rolling stock units that occur first in the rolling stock schedule are selected.
 - Set 3, with all rolling stock units of type VIRM4 (in BDU 10: 92 rolling stock units)
 - Set 4, with all rolling stock units of type VIRM4 (in BDU 10: 92 rolling stock units) and VIRM6 (in BDU 10: 68 rolling stock units)
 - Set 5, with all rolling stock units used for intercity services (in BDU 10: 360 rolling stock units)
- The planning horizon $\tau \in \{7, 21, 42\}$ measured in days. For meaningful interpretation, τ should not exceed the length of the period for which the BDU considered is valid (i.e. in general eight weeks). Observe that $T = 24\tau$, where T is the planning horizon in hours (see Section 4).
- The maximum number of daytime maintenance locations $L_{max}^D \in \{0, 1, 2, 3, 5, 10, 20\}$.

Section 5.2 presents four experiments to investigate the performance of the **MLCP** model provided in Section 4. Experiment 1 focuses on the maximum number of locations for daytime maintenance: when more locations can be opened during daytime, supposedly more activities can be performed during daytime and hence the goal of the **MLCP** model (minimizing the number of nighttime maintenance activities) may be better achieved. Experiment 2 focuses on the influence of the planning horizon. If a shorter planning horizon yields a similar location choice, a shorter time horizon may suffice to determine good maintenance locations, which in turn leads to a shorter running time. Experiment 3 compares 10 different BDUs, leading to 10 different scenarios that apply for different time periods, allowing to investigate whether results on the location choice are consistent over multiple time periods. Experiment 4 includes all rolling stock units that serve intercity (long-distance) lines, while the results in the first three experiments considered only the rolling stock types VIRM4 and VIRM6. This allows to investigate the performance and results of the **MLCP** model

Table 4

Number of constraints and variables in the **MLCP** model for different combinations of rolling stock set ν and planning horizon τ . An increasing number of rolling stock units and length of the planning horizon leads to higher numbers of constraints and variables.

ν	$\tau = 7$		$\tau = 21$		$\tau = 42$	
	# constraints	# variables	# constraints	# variables	# constraints	# variables
1	650	293	2,078	883	4289	1,771
2	1,355	596	4,185	1,742	8445	3,445
3	6,348	2,715	19,108	7,819	38,248	15,475
4	10,562	4,479	31,797	12,973	63,642	25,711
5	28,947	12,305	88,217	36,013	177,107	71,569

in a scenario of realistic size. More experiments, including experiments elaborating upon the sensitivity of various parameters and the computation time of the **MLCP** model, can be found in Zomer (2020).

It must be noted that, from the ten available BDUs, five resulted in infeasible solutions when applied in combination with rolling stock sets 4 and 5. This becomes relevant in Experiment 3, where all 10 BDUs are considered. These five scenarios are infeasible due to the fact that the corresponding data set contained some rolling stock units for which it is not possible to create a feasible maintenance schedule, for example, no two suitable successive MOs could be found within 24 h to schedule maintenance activities of type 1. Even if this is the case for only one rolling stock unit, no feasible solution can be found. In reality, this problem could be overcome by manually removing the rolling stock units that caused the infeasibility and running the model again. However, in the current research, results are reported for the five BDUs for which a feasible solution was found. This is expected to provide sufficient insights as to whether results of one period carry over to multiple periods.

To report results concerning locations, the locations are abbreviated by their standard code: Ah (Arnhem), Alm (Almere), Amr (Alkmaar), Asd (Amsterdam Centraal), Bkh (Binckhorst), Bkd (Bokkeduinen), Ddr (Dordrecht), Dv (Deventer), Ehv (Eindhoven), Es (Enschede), Gn (Groningen), Gvc (Den Haag Centraal), Hdr (Den Helder), Hfdo (Hoofddorp Opstel), Hlm (Haarlem), Hrl (Heerlen), Llso (Lelystad Opstel), Lw (Leeuwarden), Mt (Maastricht), Rsd (Roosendaal), Rtd (Rotterdam Centraal), Vl (Venlo), Vs (Vlissingen), Zd (Zaandam), Zl (Zwolle).

Table 4 shows the number of constraints and variables for instances with different combinations of rolling stock set ν and of the planning horizon τ . All these numbers relate to BDU 10, but it is noted that no significant differences across different BDUs apply. The increasing number of constraints and variables usually leads to higher computation times. To reduce the computation time, only those MOs that have a length of at least the minimum maintenance duration of 30 min are considered. All other MOs are irrelevant since these are not suitable for any type of maintenance and are therefore discarded.

5.2. Results

The current section provides results on the **MLCP**. The **MLCP** model is implemented using Python version 3.7 and solved using Gurobi version 9.0.1. The results were generated using a computer with Windows 10 Enterprise 64, 8.00 GB of RAM and with an Intel(R) Core(TM) i7-6600U CPU 2.60 GHz processor. Zomer et al. (2020) provides the implementation code of the **MLCP** model used to generate the results in this section. Due to confidentiality, the actual data cannot be provided, but a synthetic data set is made available instead.

5.2.1. Experiment 1: influence of the maximum number of daytime maintenance locations

For varying numbers of the maximum number of daytime maintenance locations, this section presents two measures. Firstly, it presents the *mean daytime share*, which is the percentage of hours of activity performed during daytime, averaged per day, relative to the mean total hours of activity performed per 24 h, averaged per day. A high mean daytime share indicates that relatively more work is performed during daytime, and this is desirable, since it is an indication of a decrease in the pressure on maintenance location capacity during nighttime. However, to check that not only the relative amount of work but also the absolute amount of work during nighttime decreases, it is necessary to verify that the total hours of activity remains more or less constant. Therefore, secondly, this section reports the *mean total hours of activity* per 24 h, averaged per day, for multiple scenarios.

Mean daytime share. Scenarios have been run for all combinations of L_{max}^D , ν and τ , with $L_{max}^D \in \{0, 1, 2, 3, 5, 20\}$, $\nu \in \{1, 2, 3, 4\}$ and $\tau \in \{7, 21, 42\}$ and mean daytime shares for each of these scenarios have been computed. Since the mean daytime shares showed to be relatively invariant for different values of ν and τ , Table 5 shows *average* mean daytime shares, averaged over all 12 combinations of ν and τ , for various values of L_{max}^D . It shows that the average mean daytime share is increasing with L_{max}^D .

The fact that the mean daytime shares are increasing in L_{max}^D is expected. When more locations are opened for daytime maintenance, more maintenance can be performed during daytime and hence the mean daytime share increases. This is particularly interesting since the goal of NS is to move hours of activity from nighttime to daytime. Apparently, when opening 20 locations, up to on average 42.0% of the work can be performed during daytime. It must be noted that when more maintenance locations are opened, the utilization of additional maintenance locations decreases. This can be observed from the fact that with 5 maintenance locations, a day share of 22.3% can be attained, whereas in order to achieve a day share of 42.0% (i.e. twice as high) no less than 20 locations (i.e. 4 times as much locations) are necessary. See also Section 5.2.2.

Table 5
The mean daytime share increases for increasing values of L_{max}^D .

L_{max}^D	Average mean daytime share
0	0.0%
1	4.2%
2	8.9%
3	13.3%
5	22.3%
20	42.0%

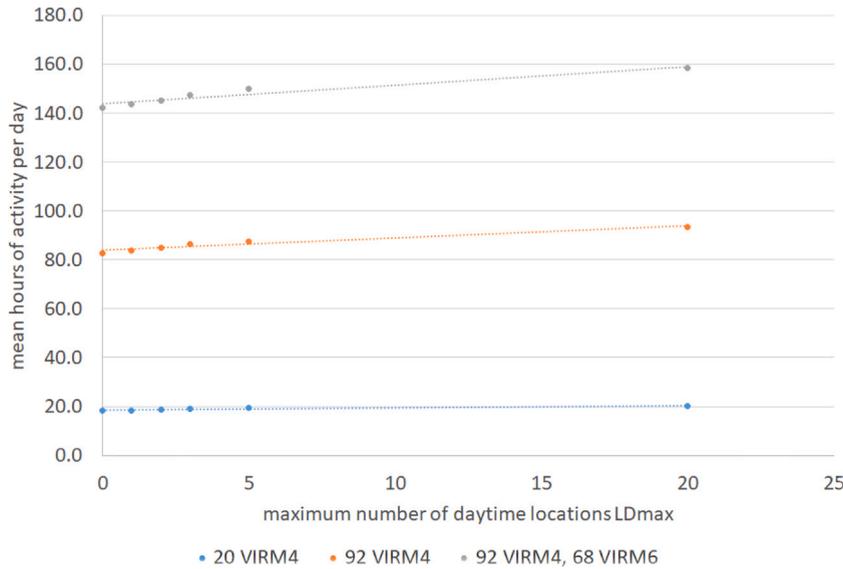


Fig. 2. The mean total hours of activity increases when the number of daytime maintenance locations L_{max}^D increases.

Mean total hours of activity. Scenarios have been run for all combinations of L_{max}^D , ν and τ , with $L_{max}^D \in \{0, 1, 2, 3, 5, 20\}$, $\nu \in \{2, 3, 4\}$ and $\tau \in \{7, 21, 42\}$. The set $\nu = 1$ yields similar results but has been excluded from the graph for visualization purposes. The mean total hours of activity for each of these scenarios has been computed. Since the mean total hours of activity showed to be relatively invariant over different values of the planning horizon τ , Fig. 2 shows mean total hours of activity per day, averaged over all 3 choices of τ , for various values of L_{max}^D and ν .

Fig. 2 shows the mean total hours of activity for various sets of rolling stock units $\nu \in \{2, 3, 4\}$. Numbers are averaged over scenarios with different values of the planning horizon $\tau \in \{7, 21, 42\}$. It shows that indeed the mean total hours of activity per day increases when the maximum number of daytime maintenance locations increases: up to about 15% when 20 daytime maintenance locations are opened. This implies that, when more possibilities for daytime maintenance arise, this leads to more maintenance in total. From this figure it also becomes clear that this holds for various numbers of rolling stock units, and (evidently) that the mean total hours of activity also increases when more rolling stock units are added to the analysis. The increase of the mean total hours of activity is however convincingly outweighed by the increase of the mean daytime share (cf. Table 5), and hence the application of daytime maintenance decreases the pressure on the capacity of maintenance locations during nighttime.

5.2.2. Experiment 2: influence of the planning horizon

The goal of the second experiment is to investigate the sensitivity of the solutions provided by the MLCP model to various choices of the planning horizon τ . A shorter planning horizon requires less computation time. If a shorter planning horizon yields similar location choice, then the required computation time to determine the optimal location choice could be significantly reduced.

Scenarios have been run for all combinations of L_{max}^D , ν and τ , with $L_{max}^D \in \{5, 20\}$ and $\tau \in \{7, 21, 42\}$ with $\nu = 4$ (i.e. all VIRM4 and VIRM6), which means there are three scenarios for each choice of L_{max}^D .

Fig. 3 indicates for each location that was opened during daytime in at least one of the scenarios, in how many of the investigated scenarios the location was chosen. Moreover, for each location the mean hours of daytime activity assigned to each location is indicated. This provides insight in the location consistency and the workload assigned to each location. Blue bars represent the number of scenarios (out of 3) in which the respective location was open during daytime. Orange bars represent the average daily number of hours of daytime activity assigned to the respective location. The locations are presented in descending order of average hours of daytime activity.

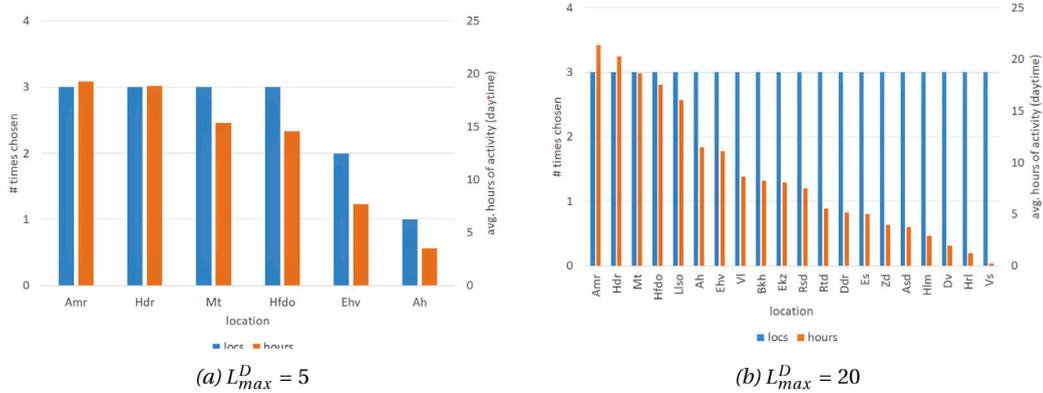


Fig. 3. Location choice is consistent over various choices for the planning horizon. Moreover, significant workloads are assigned to various locations.

For $L_{max}^D = 5$ (blue lines in Fig. 3(a)), the locations Amr, Hdr, Hfdo and Mt are chosen in all three scenarios. This justifies that these four locations are good choices for daytime maintenance in the given subset, but more importantly it shows that the location choice is robust under various time horizons. For $L_{max}^D = 20$ (blue lines in Fig. 3(b)), the location choice is for all three values of the planning horizon τ exactly equal, i.e. each time the same 20 locations are chosen, again demonstrating robustness under various time horizons. This offers the important practical insight that, in order to determine a good maintenance location choice, shorter time horizons suffice, requiring shorter computation times.

Moreover, Fig. 3 shows in orange color the mean hours of *daytime* activity, averaged over all 3 combinations of τ . When $L_{max}^D = 5$ and all rolling stock units are included, the hours of daytime activity is higher than 2 h per day for these locations that are chosen in all 3 scenarios. In other words, for the locations for which the choice is most consistent, also the highest workload is found.

When comparing Figs. 3(a) and 3(b), the locations Amr, Hdr, Mt and Hfdo appear to be four good candidate locations to be opened. They are chosen in all three scenarios, and the workload assigned to these locations is the highest. Hence, also when more locations can be opened (20 in this case), the locations found for a smaller maximum number of daytime maintenance locations (5 in this case, see Fig. 3) still remain good candidates. This is important in practical situations where the number of locations is gradually expanded: the locations found in a stage where only a few locations can be opened, still remain good candidates even when more locations can be opened in a later stage. This is not an obvious finding, it may have been true that the location choice would change drastically for an increasing value of L_{max}^D . Moreover, from Fig. 3 the importance of the four locations mentioned above can be deduced from the high utilization; the other locations have lower and similar utilization.

Observe that, especially in the case with $L_{max}^D = 20$ (Fig. 3(b)), the mean hours of activity for some locations, such as Dv, Hrl and Vs, is low. This means that those locations are not utilized very well, and hence it is questionable whether it is worthwhile to open these locations in terms of efficient use of human resources. In practice, this could lead to an advice to rerun the model for a slightly smaller value of L_{max}^D .

5.2.3. Experiment 3: consistency over different BDUs

Scenarios have been run for $L_{max}^D = 20$, $v = 4$ (i.e. all VIRM4 and VIRM6) and $\tau = 42$. As indicated, out of the 10 available BDUs, only 5 resulted in a feasible solution and hence only these have been considered here.

Fig. 4, shows for each location that was opened during daytime in at least one of scenarios (with a feasible solution), in how many of these five scenarios this location is opened (blue color) and the mean hours of daytime activity at each opened location, averaged over the scenarios in which the location was opened (orange color).

It shows that most of the locations (14 in total) are chosen in all five scenarios (blue lines). This implies that the optimal locations are consistent across different BDUs. From a managerial perspective, this is important, since it means that once a location is chosen, it will most likely still be optimal to choose this location in a next period, i.e. for a new rolling stock circulation plan. This means that the applicability of the results, obtained by applying the MLCP model, are in this case not only valid for one BDU, but also carry over to multiple BDUs and hence to multiple periods of time.

In addition, in some cases, a location was chosen in only a limited number of scenarios: for example, Vs and Zl were chosen in only one out of five scenarios. However, in general, these locations are not the locations to which high workloads are assigned. The influence of choosing not to open such a location may therefore be limited.

Looking at the mean hours of daytime activity of the opened locations (orange lines), it can be observed that there are four locations with on average more than three hours of activity per day (Mt, Hdr, Amr and Ah) and that these locations are consistently chosen over the multiple scenarios. This is positive since it indicates that the choice on whether a location is opened consistently or not correlates with the workload assigned to that location. In other words, if a location is opened consistently throughout multiple scenarios, this usually also means that a relatively high workload is assigned to this location. Like in Fig. 3(b), from Fig. 4 the importance of the four locations mentioned above can be deduced from the high utilization; the other locations have lower and similar utilization.

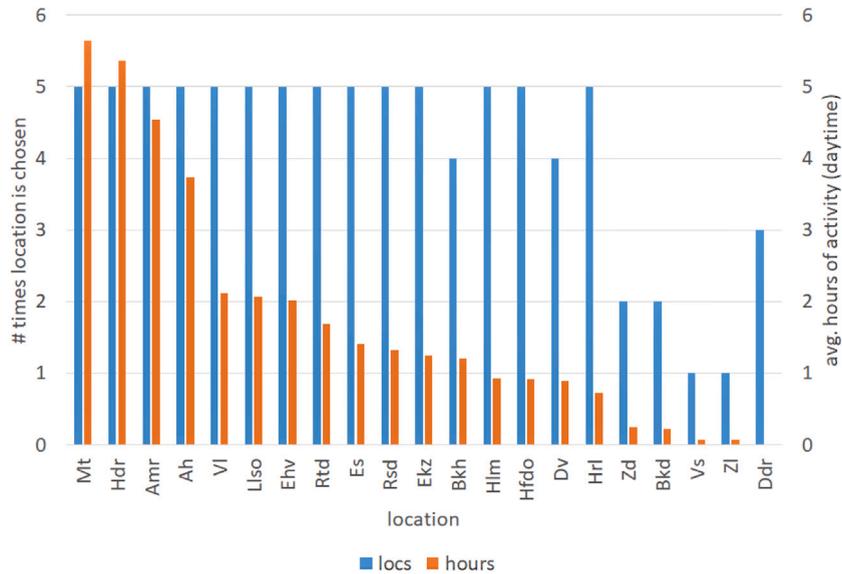


Fig. 4. Location choice is consistent over various BUs.

Table 6

In a realistic case, a substantial amount of work can be performed during daytime.

τ	L_{max}^D	Mean total hrs. activity	Mean daytime share	Running time (s)
7	10	305.9	22.6%	204
	20	310.6	30.9%	201
42	10	328.8	22.2%	11,522
	20	334.9	30.1%	7,885

5.2.4. Experiment 4: Performance of scenario of realistic size

Experiments 1–3 have investigated the performance of the **MLCP** model under various parameter settings. The largest data set considered so far, included rolling stock units of one train type, VIRM. To assess the working of the **MLCP** model in practice, it is necessary to consider a larger and more representative data set. To this end, this fourth experiment considers all rolling stock units used for intercity services combining multiple train types (rolling stock set $\nu = 5$).

Table 6 gives the most important results for the scenarios with 10 and 20 locations available: it provides the mean hours of daytime activity, the mean daytime share, the number of variables and constraints and the running time of the four different scenarios investigated in the current experiment. It shows that, for $L_{max}^D = 20$, a mean daytime share of over 30% can be attained. The mean daytime share appears to be higher for the scenarios with a maximum number of 20 locations opened for daytime maintenance, which is consistent with the findings in Experiment 1 (in Section 5.2.1). Also, it can be observed that the mean total hours of activity may increase when more opportunities for daytime maintenance emerge. In the **MLCP** model, any increase in the number of daytime maintenance activities is considered desirable as long as it leads to a decrease in nighttime maintenance activities. As a result, the increase in daytime maintenance activities may be stronger than the decrease in nighttime maintenance activities, leading to a higher number of hours of activity. Another reason for this increase lies in the fact that maintenance types with durations of 24 and 48 h are considered; therefore, any switch from nighttime maintenance and daytime maintenance (and back) includes some inefficiency.

Furthermore, it can be observed that, when the model grows, the running time increases significantly. The largest running time observed is 11,522 s or approximately 3.2 h for the largest model run in the current research. In the light of the fact that the **MLCP** concerns a planning problem at the tactical level, this running time is deemed acceptable for practical purposes.

Fig. 5 provides the mean hours of daytime activity per location, for a planning horizon width $\tau = 42$, the set of rolling stock units $\nu = 5$, for $L_{max}^D = 10$ (blue) and $L_{max}^D = 20$ (orange). It shows that, in the situation with a maximum number of locations for daytime maintenance of 10, at least 5 locations have a workload of more than 8 h per day (Gvc, Bkd, Bkh, Dv and Gn), which can be considered substantial since it is enough to provide work to one maintenance team. Moreover, it shows that the addition of maintenance locations does not seem to reduce the average workload on any of the initial 10 locations. Hence, the initial 10 locations are still good choices, even when daytime maintenance is possible at more locations. Instead, it can be assumed that some additional maintenance work was allocated to these newly opened locations. The added locations, however, are assigned a much lower workload than the initial 10 locations. It may be therefore questionable whether the addition of these locations is worthwhile.

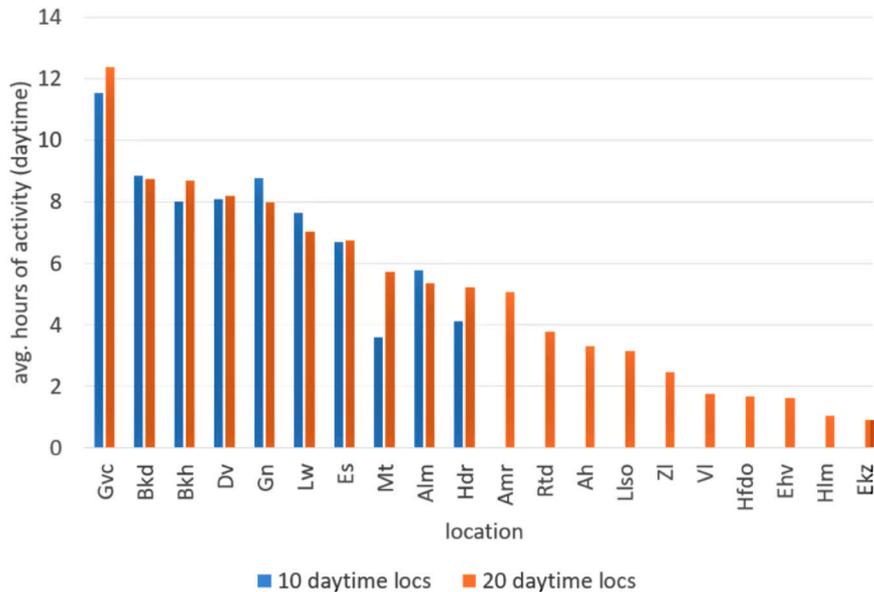


Fig. 5. Good candidate locations for daytime maintenance found in the scenario of realistic size.

6. Conclusions

This paper introduces the Maintenance Location Choice Problem (MLCP) and models it as a mixed integer program, taking as its main input a rolling stock circulation and provides for this rolling stock circulation an optimal maintenance location choice that minimizes the total number of maintenance activities during nighttime, thereby reducing the capacity pressure during the night. For the NS case with all rolling stock units of type VIRM4 and VIRM6, up to 22.3% of the work can be performed during daytime if five locations are opened, to even 42.0% if 20 locations are opened during daytime. The location choice is consistent for different lengths of planning horizons and for different input data sets. The four locations with the highest assigned workload (in hours of daytime activity) are Alkmaar, Den Helder, Maastricht and Hoofddorp Opstel. For the largest scenario (including all rolling stock units used for intercity services), a mean daytime share of at least 30.1% of the activities can be achieved if 20 maintenance locations are opened. In this case, Den Haag Centraal, Bokkeduinen, Binckhorst, Deventer show to be the locations with the highest workloads. Moreover, the model is efficient for planning purposes: for the largest scenario run (with all rolling stock units used for intercity services) for a planning horizon of 42 days, the computation time was 3.2 h. For operational purposes, however, this running time is not sufficiently efficient.

Some limitations to the current approach can be identified. First, the current research uses planned rolling stock data. Although this is a common practice in the railway industry and therefore a logical choice, it does not take into account the potential occurrence of disruptions. Due to disruptions, the rolling stock circulation is often altered during operations and as a result, the provided optimal maintenance schedule and optimal maintenance location choice may not be the best anymore at the moment of execution. The current model can provide a basis for more robust decisions on daytime maintenance locations and/or the online adjustment of a maintenance schedule. Second, the current research uses a simplified definition of maintenance types. Their durations and intervals may not exactly correspond to practical situations. Moreover, in the current research these are assumed to be equal for each rolling stock unit, whereas in reality these values often differ per rolling stock type. We expect that enriching the formulation with these types would not be very complex (it would entail replacing the general set of maintenance types K for a set of specific maintenance types for each rolling stock unit, K_i , and replace the general variables on duration o_k and interval v_k for specific ones o_{ki} and v_{ki} , respectively). Third, the capacity of maintenance locations is not restricted. According to expert judgment at NS, this is acceptable: currently, the primary goal of NS is to reduce the pressure on maintenance locations during nighttime and hence it is not relevant to restrict the capacity of maintenance locations during nighttime. Also, the work assigned to daytime maintenance locations seems to be not excessive and hence it is deemed unnecessary to restrict the capacity of daytime maintenance locations. However, there may be applications in which it is desirable to restrict the capacity of maintenance locations as well, which is an interesting direction for further work.

The MLCP model is intended for use on a tactical level (i.e. months before operations), since it assumes that a rolling stock circulation is already given. Hence, it answers the question at which already existing locations personnel needs to be stationed during daytime and/or night time. However, when applied on multiple rolling stock circulations, the MLCP model can also serve more strategic goals (i.e. years before operations), regarding for example the question at what positions new locations could best be built.

Several directions for future research can be recommended. First, in the current research the rolling stock circulation is assumed to be given and fixed. The disintegration between rolling stock circulation and maintenance scheduling has many advantages in terms of computational efficiency, but it must be recognized that more opportunities may arise when the rolling stock circulation would have been allowed to be adjusted. An example of such opportunities is the concept of rolling stock exchanges (Van Hövell, 2019), which is based on the idea of dividing long daytime MOs (with more than enough time to do maintenance) into several shorter MOs (each with just enough time to do maintenance) to enable more daytime maintenance opportunities. Also, it may be interesting to consider small time adjustments to the timetable and further, also to incorporate maintenance tasks and location selection directly in the rolling stock circulation problem. Second, an interesting next research topic is how to improve the computational performance of the **MLCP**. Using the structure of the problem, the problem may potentially be decoupled into multiple smaller sub-problems that are much easier to solve, e.g. solving rolling stock types independently, or considering a rolling horizon framework, thereby first optimizing a few days in ahead and iteratively adding more days. Third, as indicated, some of the input data sets appeared to result in infeasible solutions. Since the model is not capable of allowing (minor) violations of constraints, for some input sets no solution can be found at all. However, in some contexts it may be better to obtain a slightly violated solution than obtaining no solution at all. This could for instance be achieved by accepting and penalizing constraint violations such as reducing the time for some maintenance activities by a small amount. Fourth, in the current research, no restrictions regarding the capacity at maintenance locations are taken into account. Although this is acceptable regarding the current way of working at NS, other applications may require that maintenance location capacity be explicitly addressed. It requires more research to investigate how this can be incorporated within the current modeling framework.

Fifth, several extensions to the current model may be interesting. For instance, it may be interesting to investigate the effects of applying a different, potentially more complex weighting scheme in the objective of the **MLCP** model to better reflect preferences for specific assignments of maintenance activities to MOs. Also, in some practical settings it may not be possible to move as many tasks from nighttime to daytime. In such cases it is worthwhile to investigate the benefits of having a smart way of selecting a limited number of the tasks to be performed during daytime. Moreover, more research is required in order to carefully determine the required buffer times that need to be included into account to maintenance tasks in order to guarantee robust and resilient solutions.

Sixth, the **MLCP** model can be applied to other domains as well. Examples of such extensions include the applicability of the research to serve other railway undertakings with potentially different objectives and policies, the investigation of other cost structures and the inclusion of exchange opportunities of rolling stock units so that more rolling stock units can be maintained during daytime.

Acknowledgments

We thank SMA und Partner AG for supporting one of the authors in this research and providing an opportunity to implement the **MLCP** model to the planning tool Viriato. Also, thanks to NS for providing the data required. Finally, we thank the three anonymous reviewers for their suggestions that have improved the quality of this paper.

References

- Andrés, J., Cadarso, L., Marín, A., 2015. Maintenance scheduling in rolling stock circulations in rapid transit networks. *Transp. Res. Procedia* (ISSN: 23521465) 10 (July), 524–533. <http://dx.doi.org/10.1016/j.trpro.2015.09.006>.
- Canca, D., Barrena, E., 2018. The integrated rolling stock circulation and depot location problem in railway rapid transit systems. *Transp. Res. E* (ISSN: 13665545) 109 (May 2017), 115–138. <http://dx.doi.org/10.1016/j.tre.2017.10.018>.
- Clarke, L., Johnson, E., Nemhauser, G., Zhu, Z., 1997. The aircraft rotation problem. *Ann. Oper. Res.* 69, 33–46.
- Dinmohammadi, F., Alkali, B., Shafiee, M., 2016. A risk-based model for inspection and maintenance of railway rolling stock. In: *European Safety and Reliability Conference, ESREL*, pp. 2016–2029.
- Feo, T.A., Bard, J.F., 1989. Flight scheduling and maintenance base planning. *Manage. Sci.* (ISSN: 0025-1909) 35 (12), 1415–1432. <http://dx.doi.org/10.1287/mnsc.35.12.1415>.
- Gopalan, R., 2014. The aircraft maintenance base location problem. *European J. Oper. Res.* (ISSN: 03772217) 236 (2), 634–642. <http://dx.doi.org/10.1016/j.ejor.2014.01.007>.
- Gopalan, R., Talluri, K.T., 1998. The aircraft maintenance routing problem. *Oper. Res.* 46 (2), 260–271.
- Herr, N., Nicod, J.-M., Varnier, C., Zerhouni, N., Malek Cherif, F., 2017. F., Joint optimization of train assignment and predictive maintenance scheduling. In: *7th International Conference on Railway Operations Modelling and Analysis*. Raillille 2017, April, pp. 699–708.
- NS. VVGB NedTrain, 2018. Various microsoft access data sheets containing all planned rolling stock movements for NS, for 10 different periods of approximately 8 weeks each, with start dates between 10-12-2017 and 9-6-2019.
- Sarac, A., Batta, R., Rump, C.M., 2006. A branch-and-price approach for operational aircraft maintenance routing. *European J. Oper. Res.* (ISSN: 03772217) 175 (3), 1850–1869. <http://dx.doi.org/10.1016/j.ejor.2004.10.033>.
- Tönissen, D.D., Arts, J.J., 2018. Economies of scale in recoverable robust maintenance location routing for rolling stock. *Transp. Res. B* (ISSN: 01912615) 117, 360–377. <http://dx.doi.org/10.1016/j.trb.2018.09.006>.
- Tönissen, D., Arts, J., 2020. The stochastic maintenance location routing allocation problem for rolling stock. *Int. J. Prod. Econ.* 230, 107826.
- Tönissen, D.D., Arts, J.J., Shen, Z.-J.M., 2019. Maintenance location routing for rolling stock under line and fleet planning uncertainty. *Transp. Sci.* (ISSN: 0041-1655) 53 (5), 1252–1270. <http://dx.doi.org/10.1287/trsc.2018.0866>.
- Van Hövell, M., 2019. Increasing the Effectiveness of the Capacity Usage at Rolling Stock Service Locations. Delft University of Technology.
- Wagenaar, J., Kroon, L.G., 2015. Maintenance in railway rolling stock rescheduling for passenger railways. *SSRN Electron. J.* 1–38. <http://dx.doi.org/10.2139/ssrn.2566835>.
- Zomer, J., 2019. Verkenning Mogelijkheden Overdag Behandelen. Technical Report September, NS.
- Zomer, J., 2020. Simultaneous optimization of rolling stock maintenance scheduling and rolling stock maintenance location choice (M.Sc. thesis). Delft University of Technology, Delft, The Netherlands.
- Zomer, J., Bešinović, N., de Weerd, M.M., Goverde, R.M.P., 2020. Demo code for the maintenance location choice problem model. GitHub repository. <https://github.com/jzomergit/MLCP/>.