Aggregate overhaul and supply chain planning for rotatables

Citation for published version (APA):

Arts, J. J., & Flapper, S. D. P. (2015). Aggregate overhaul and supply chain planning for rotatables. *Annals of Operations Research*, 224(1), 77-100. <https://doi.org/10.1007/s10479-013-1426-0>

DOI:

[10.1007/s10479-013-1426-0](https://doi.org/10.1007/s10479-013-1426-0)

Document status and date:

Published: 01/01/2015

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Aggregate overhaul and supply chain planning for rotables

Joachim Arts · Simme Douwe Flapper

© Springer Science+Business Media New York 2013

Abstract We consider the problem of planning preventive maintenance and overhaul for modules that are used in a fleet of assets such as trains or airplanes. Each type of module, or rotable, has its own maintenance program in which a maximum amount of time/usage between overhauls of a module is stipulated. Overhauls are performed in an overhaul workshop with limited capacity. The problem we study is to determine aggregate workforce levels, turn-around stock levels of modules, and overhaul and replacement quantities per period so as to minimize the sum of labor costs, material costs of overhaul, and turn-around stock investments over the entire life-cycle of the maintained asset. We prove that this planning problem is strongly $\mathcal{NP}$ -hard, but we also provide computational evidence that the mixed integer programming formulation can be solved within reasonable time for real-life instances. Furthermore, we show that the linear programming relaxation can be used to aid decision making. We apply the model in a case study and provide computational results for randomly generated instances.

Keywords Maintenance · Aggregate planning · Life cycle costs · $\mathcal{NP}$ -hard · Repairable parts · Reverse logistics

1 Introduction

The primary processes of manufacturing and service companies rely on the availability of equipment. When this equipment represents a significant financial investment, it is usually referred to as a capital asset or capital good. Examples of such capital assets include trains, airplanes, MRI-scanners, and military equipment. While the acquisition cost of capital assets is substantial, the costs associated with maintenance and downtime over the lifetime of

J. Arts (✉) · S.D. Flapper
School of Industrial Engineering, Eindhoven University of Technology, P.O. Box 513, 5600 MB,
Eindhoven, The Netherlands
e-mail: j.j.arts@tue.nl

S.D. Flapper
e-mail: s.d.p.flapper@tue.nl

the asset are typically 3 to 4 times the acquisition price, even when the future costs of maintenance and downtime are discounted (Öner et al. 2007). Accordingly, there has been much focus and research on what is called life cycle costing (LCC); see Gupta and Chow (1985) and Asiedu and Gu (1998). The LCC approach to decision making in asset acquisition, maintenance, and disposal stipulates that the consequences of decisions should be accounted for over the entire lifetime of the asset in question.

Another factor influencing maintenance is the modular design of many technical systems. Usually, a capital asset is not maintained in its entirety at any one time. Instead, different modules of the system are dismantled from the asset and replaced by *ready-for-use* modules. After replacement, the module can be overhauled while the capital asset is up and running again. Exchanging modules, rather than maintaining them on the spot, increases the availability of capital assets, as assets are only down for the time it takes to replace a module. After overhaul, the module is ready for use again and can be used in a similar replacement procedure for another asset. To make this system work, some spare modules are needed, and they form a so called *turn-around stock*.

In this paper, we consider the replacement of modules that have their own maintenance program. The maintenance program stipulates a maximum amount of time/usage a module is allowed to be operational before it needs to be overhauled. We refer to this time allowance as the *maximum inter-overhaul time* (MIOT), and we assume that there is a direct relation between the time a module has been in the field and its usage. Due to safety regulations, or contracts with the original equipment manufacturer (OEM), the MIOT is usually quite conservative and so most modules are almost exclusively maintained preventively. We call the practice described in the previous paragraph *maintenance-by-replacement*. Note that this is similar but different from repair-by-replacement, wherein components are replaced in corrective as opposed to preventive maintenance efforts. We refer to the modules involved as *rotables*, because they rotate through a closed-loop supply chain. At this point, we emphasize that rotables differ from repairables as they are studied in much of the spare parts inventory control literature (e.g., Sherbrooke 2004). Repairables do not have a maintenance program of their own, and, consequently, the need for replacement of repairables is usually characterized by stochastic models such as the (compound) Poisson process. By contrast, rotables do have their own maintenance program, and so replacements and overhauls of rotables are planned explicitly by a decision maker.

This paper is motivated by a maintenance-by-replacement system in place at NedTrain, a Dutch company that performs maintenance of rolling stock for several operators on the Dutch railway network. Below we describe several characteristics and constraints of maintenance-by-replacement systems and their implications for planning.

In a maintenance-by-replacement system, replacements and overhauls are subject to the following two constraints respectively. A replacement may not occur unless a ready-for-use rotatable is available to replace the rotatable that requires overhaul, so that the asset can immediately return to operational condition. An overhaul cannot occur unless there is available capacity in the overhaul workshop. Since the result of an overhaul is a ready-for-use rotatable, these constraints are connected.

The maintenance programs of rotables also impose constraints on maintenance-by-replacement systems. For each rotatable type, the maintenance program stipulates a MIOT, the maximum amount of time a rotatable is allowed to be operational before it needs to be overhauled. Note that the decision to replace a rotatable in some period t directly implies that the replacing rotatable needs to be replaced before time $t + \text{MIOT}$.

With respect to the timing of rotatable overhauls and replacements, the LCC perspective offers opportunities. In traditional maintenance models, the focus is on postponing maintenance as long as possible, thereby taking advantage of the technical life of the unit to be

maintained. This approach may not lead to optimal decisions over finite lifetimes of assets. To see why, consider the following example based on practice at NedTrain. The typical lifetime of a rolling stock unit is 30 years. Bogies are important rotables in a train, with MIOTs that range from 4 to 10 years. Suppose the MIOT of two types of bogies is 7 years, and both types of bogies belong to the same type of train. Then, if replacements are planned to occur just in time, bogie replacements occur 4 times during the life cycle of this train type, namely in years 7, 14, 21, and 28. Another plan that is feasible with respect to overhaul deadlines is to replace in years 6, 12, 19, and 25. Note that it is possible to replace rotables earlier than technically necessary, i.e., throwing away some of the useful life of the equipment, *without* increasing the number of replacements (and overhauls) that are needed during the lifetime of an asset. To smooth the workload of the overhaul workshop, it may be possible to overhaul the first type of rotatable according to the first schedule, and the second type of rotatable according to the second. In general, the flexibility in the exact timing of replacements and overhauls can be used to smooth the workload of the overhaul workshop and utilize other resources more efficiently *without* losing efficiency by throwing away remaining useful life of rotables. In effect, we are not and should not be concerned with minimizing the amount of useful lifetime on rotables that is wasted. Rather, we should minimize the cost of maintenance and overhaul that rotables incur over the lifetime of the asset they serve, which is finite. The renewal reward theorem (e.g., Ross 1996) that has proven beneficial in many reliability and maintenance engineering applications (e.g., Ebeling 2010) cannot be applied in this setting. The reason for this is that the horizon we consider is not infinite (not even by approximation). To see this, consider again the example above. Only a few renewals (4 in the example) occur during the time a rotatable is in the field, and the last renewal has very different characteristics from the other renewals in that the last renewal ends with replacing the asset for which the rotatable is used, rather than overhauling the rotatable itself.

In this paper, we study a model for the aggregate planning of rotatable replacements and overhaul for multiple rotatable types that use the same resources in an overhaul workshop. We adopt the LCC perspective and take the finite life cycle of assets into consideration. Our aggregate planning model supports decisions regarding overhaul workshop capacity levels, sizing of turn-around stocks of rotables, and overhaul and replacement quantities per period. The model we present should be implemented in a rolling horizon, i.e., the model generates decision for the next 30 or so years, but only the decisions for the coming few months should be implemented. As time progresses, estimates of input parameters for our model become more accurate, and the model should be solved again to generate decisions that are based on these more accurate estimates.

This paper is structured as follows. In Sect. 2, we review the literature on maintenance and aggregate supply chain planning. We provide and analyze our model in Sect. 3. Computational results based on a real life case are presented in Sect. 4. In Sect. 5, we present computational results for a large test bed of randomly generated instances. Finally, conclusions are offered in Sect. 6.

2 Literature review and contribution

Aggregate planning is performed in many contexts and businesses. We review the literature on maintenance planning in Sect. 2.1. Since our model also deals with the rotatable supply chain, we review aggregate planning models in the context of production and supply chain in Sect. 2.2. In Sect. 2.3, we explain our contribution relative to the literature discussed.

2.1 Preventive maintenance and capacity planning

Wagner et al. (1964) are among the first to consider the joint problem of preventive maintenance and capacity planning. They consider a setting where a set of preventive maintenance tasks is to be planned, while fluctuations in work-force utilization are to be kept at a minimum. The objective is approximately met by formulating the problem as a binary integer program and using rounding procedures to find feasible solutions. Paz and Leigh (1994) give an overview of many different issues involved with maintenance planning and review much of the literature from before 1993. They identify manpower as the critical resource that has to be reckoned with in maintenance planning.

More recent research on maintenance planning includes Charest and Ferland (1993), Chen et al. (2010), and Safaei et al. (2011). Safaei et al. (2011) consider short term maintenance scheduling to maximize the availability of military aircraft for the required flying program. The problem is cast as a mixed-integer-program (MIP) in which the required workforce is the most important constraint. Chen et al. (2010) study short-term manpower planning using stochastic programming techniques and apply their model to carriage maintenance in the mass-rapid-transit system of Taipei. They consider a horizon of around one week and their model allows for random maintenance requirements due to break-down maintenance (as opposed to planned, preventive maintenance). Charest and Ferland (1993) study preventive maintenance scheduling where each unit that is to be maintained is fixed to a rigid maintenance schedule with fixed inter-maintenance intervals. They model the problem as a MIP and solve this MIP with various heuristic methods, such as exchange procedures and tabu search.

A closely related problem is the clustering of maintenance activities when a set-up cost is associated with performing maintenance. See Van Dijkhuizen and Van Harten (1997) and the references therein for this stream of literature.

Recently, some attention has also been paid to the availability of ready-for-use rotables as a critical constraint in maintenance planning. Driessen et al. (2010) provide a framework for the planning of spare parts that are used in maintenance. Our work fits partially in their framework, as rotables are a special type of spare part. While Driessen et al. (2010) consider mostly repairable logistics, where maintenance is an exogenous fact, we take a broader view by incorporating the maintenance decisions into our model, as rotables have their own maintenance program. Joo (2009) also explicitly considers the availability of ready-for-use rotables as an essential constraint in their overhaul planning model. Joo (2009) considers a set of rotables of a single type that has to meet an overhaul deadline in the (near) future. The model is set-up such that overhaul is performed as late as possible, but before this deadline and within capacity constraints. The key idea is that the useful life of a rotatable must be used to the fullest extent possible. Joo (2009) uses a recursive scheme to plan rotatable overhaul that is very much akin to dynamic programming.

2.2 Aggregate production and supply chain planning

Aggregate planning in production environments was first proposed by Bitran and Hax (1977) and has been expanded upon by many authors (e.g., Bitran et al. 1981, 1982). Today, aggregate production planning models have found their way into standard textbooks in operations and production management (e.g. Silver et al. 1998, Hopp and Spearman 2001, Nahmias 2009). These aggregate production planning (APP) models are used to plan workforce capacity and production quantities of product families over several periods. Similar models are also used in supply chain planning. These models are described and reviewed in Billington

et al. (1983), Erengüç et al. (1999), De Kok and Fransoo (2003) and Spitter et al. (2005). Although all these models generate a production plan for several periods into the future, it is understood that only the decisions for the upcoming period should be implemented. After this period lapses, more and more accurate information may be available to rerun the model and generate decisions for the next period. The reason to include many periods in the model is to be able to evaluate the impact of the decision in the current period further into the future. This way of working is called rolling horizon planning.

Aggregate maintenance planning differs from aggregate production planning in two fundamental ways. First, while in APP exogenous demand triggers the use of production capacity either implicitly or explicitly, maintenance requirements are necessarily endogenous to the modeling approach. The reason for this is that preventive maintenance needs to be performed within limited time intervals due to safety and/or other reasons. Thus, a decision to maintain a rotatable at some time t also dictates that the replacing rotatable undergo preventive maintenance before time $t + \text{MIOT}$. Here the LCC perspective offers added value. While MIOTs have to be respected, there is considerable freedom in the exact timing of performing maintenance without increasing the number of times preventive maintenance is required during the life cycle of an asset. This flexibility can, however, only be leveraged by considering the entire life cycle in the planning process. When this is done, flexibility can be used to utilize resources such as workforce and turn-around stock efficiently. We already noted that rolling horizon planning considers the impact of decisions in the current period on costs in future periods. In the case of maintenance planning, the relevant planning horizon is the lifetime of the assets that are to be maintained.

Second, maintenance has a fundamentally different capacity restriction in the availability for rotatables for replacement actions. While production capacity levels are not directly influenced by earlier production quantities, the availability of ready-for-use rotatables depends on the number of rotatables that have undergone overhaul in previous periods. Thus the number of rotatables in the closed-loop supply chain form a special type of capacity constraint. For a recent literature review on closed-loop supply chains, see Ilgin and Gupta (2010). A fundamental difference between the closed-loop supply chain studied in this paper and other closed-loop supply chains studied in literature so far, is that in this case a return (replacement) automatically generates another return within some preset fixed maximum period of time, the MIOT.

2.3 Contribution

In the field of preventive maintenance, our model has several contributions to existing literature that we summarize below:

- (a) Our model can be used for tactical decision making in which the effects of decisions over long horizons need to be considered. These long horizons explicitly incorporate LCC considerations into decision making and utilize the flexibility there is with respect to the exact timing of overhaul over the whole life cycle of an asset. However, we do not propose to fix a plan for very long horizons; we do propose accounting for consequences of decisions over long horizons.
- (b) Our model makes the constraints imposed by a finite rotatable turn-around stock explicit by modeling the rotatable supply chain. It also supports the decisions regarding the size of rotatable turn-around stocks.
- (c) Our model considers multiple rotatables types that utilize the same overhaul capacity. For each rotatable type, the model plans multiple overhauls into the future.

- (d) We perform a case study, and show that a linear programming relaxation of our optimization problem yields sufficiently accurate results to aid in decision making. We also provide useful insights about planning for NedTrain, the company involved in the case study. In a numerical experiment where instances are generated randomly, we show that the solution to the LP relaxation is usually sufficiently accurate to aid decision making.

3 Model

We consider an installed base of capital assets and a supply chain of rotables in a maintain-by-replacement system. The rotables in this supply chain go through the same overhaul workshop and their overhaul requires the availability of a fixed amount of resources in the overhaul workshop. Each asset consist of several rotables of possibly different type. For each rotatable type, there is a population of this rotatable type in the field. Each rotatable in the population of a type needs to be overhauled before it has been in the field for the MIOT. For the aggregate planning problem under consideration, we divide time in periods. We let T denote the set of periods in the planning horizon, $T = \{1, \dots, |T|\}$. The length of a period is typically one month while the length of the planning horizon should be at least the length of the life cycle of the assets in which the rotatables function. In this way, the model can capture the entire LCC. For rolling stock and aircraft, this planning horizon is about 25–35 years. We let I denote the set of different types of rotables. The first (last) period in the planning horizon during which rotables of type $i \in I$, are in the field is denoted a_i (p_i), $a_i < p_i$. For most types of rotables $a_i = 1$, meaning that rotables of type i are already in the field when a plan is generated. Rotables always support assets and companies plan the disposal of these assets, as well as their replacement with a newer version. When $a_i > 1$, type i rotables support an asset which the company plans to start using in period a_i . Similarly, when $p_i < |T|$, rotatable type i supports an asset that will be disposed of in period p_i . We let $T_i^I = \{a_i, \dots, p_i\}$ denote the set of periods in the planning horizon during which rotables of type $i \in I$ are active in the field. Furthermore, we let I_t denote the set of rotables that are active in the field during period $t \in T$: $I_t = \{i \in I | a_i \leq t \leq p_i\}$.

We also define a set of aggregated periods, $Y = \{1, \dots, |Y|\}$. Typically an aggregated period is a year. Furthermore, we let T_y^Y denote the set of periods that are contained in the aggregated period $y \in Y$. Table 1 shows an example of how T , Y and T_y^Y relate to each other. The example concerns a horizon of three aggregated periods (e.g., years) and 12 regular periods (e.g., quarters). T_1^Y contains the periods contained in the first aggregated period (e.g., the quarters of the first year).

In the rest of this section we will describe the equations that govern different parts of the system under study.

3.1 Supply chain dynamics

The rotatable supply chain is a two-level closed-loop supply chain as depicted in Fig. 1. There are two stock-points where inventory of rotables that are ready-for-use and rotables requiring

Table 1 Example of regular and aggregated time periods and the set T_y^Y

Time in aggregated periods (Y)	1				2				3			
	T_1^Y				T_2^Y				T_3^Y			
Time in periods (T)	1	2	3	4	5	6	7	8	9	10	11	12

Fig. 1 Rotable supply chain overview

overhaul, respectively, are kept. We let the variables $B_{i,t}$ ($H_{i,t}$) denote the number of ready-for-use (overhaul requiring) rotables of type $i \in I$ in inventory *at the beginning* of period $t \in T_i^I$. We let the decision variable $x_{i,t}$ denote replacements of rotables of type $i \in I$ *during* period $t \in T_i^I$. We assume the time required to replace a rotatable is negligible compared to the length of a period. The overhaul workshop acts as a production unit as defined in supply chain literature (De Kok and Fransoo 2003). This means that when an overhaul order is released at any time t , the rotatable becomes available ready-for-use at time $t + L_i$. Thus, L_i is the overhaul lead time and we assume it is an integer multiple of the period length considered in the problem. We let the decision variable $n_{i,t}$ denote the number of overhaul orders for rotatables of type $i \in I$ released in the course of period t . The supply chain dynamics are described by the inventory balance equations:

$$B_{i,t} = B_{i,t-1} - x_{i,t-1} + n_{i,t-L_i-1}, \quad \forall i \in I, \forall t \in T_i^I \setminus \{a_i\} \quad (1)$$

$$H_{i,t} = H_{i,t-1} + x_{i,t-1} - n_{i,t-1}, \quad \forall i \in I, \forall t \in T_i^I \setminus \{a_i\}. \quad (2)$$

Equations (1) and (2) require initial conditions. The stock levels for rotables already in the field in the first planning period ($a_i = 1$) are initialized by the parameters B_i^d and H_i^d respectively; so $B_{i,a_i} = B_i^d$ and $H_{i,a_i} = H_i^d$ if $a_i = 1$. Here, and throughout the remainder of this paper, the superscript d is used for parameters known from data that initialize variables. (Note that B_{i,a_i} is a variable and B_i^d is a parameter known from data.) For rotables that enter the field after the first period, the initial stock level conditions are to start with the entire turn-around stock $S_i \in \mathbb{N}$ consisting of ready-for-use repairables, and no rotables requiring maintenance; so $B_{i,a_i} = S_i$ and $H_{i,t} = 0$ if $a_i > 1$. The initial turn-around stock levels for rotables that are not yet in the field in period 1, S_i , are decision variables. For $t = a_i - L_i + 1, \dots, a_i - 1$, $n_{i,t}$ also has initial conditions: $n_{i,t} = n_{i,t}^d$ for $t \in \{a_i - L_i + 1, \dots, a_i - 1\}$. These initial conditions are known from data if $a_i = 1$ and set to 0 if $a_i > 1$. We assume that when $n_{i,t}$ overhaul orders are released during period t , these releases occur uniformly during that period.

3.2 Workforce capacity and flexibility in the overhaul workshop

The workforce capacity in the workshop is flexible. Workforce is acquired or disposed of at the ending of each aggregated time period $y \in Y$. We let the decision variable W_y denote

the available labor hours during aggregated period $y \in Y$. For example, if the length of an aggregated period is a year, W_y represent the number of labor hours to be worked during that year given the number of contracts with laborers. However, there is flexibility as to when exactly these hours are to be used during the aggregated period (year). If we let the decision variable w_t denote the amount of labor hours used during period $t \in T$, this can be expressed as follows:

$$W_y = \sum_{t \in T_y^Y} w_t, \quad \forall y \in Y. \quad (3)$$

The average number of hours worked during any period $t \in T_y^Y$ is $W_y/|T_y^Y|$. We let the parameters δ_t^l and δ_t^u denote lower and upper bounds on the fraction of $W_y/|T_y^Y|$ that is utilized during period $t \in T_y^Y$:

$$\delta_t^l W_y/|T_y^Y| \leq w_t \leq \delta_t^u W_y/|T_y^Y|, \quad \forall y \in Y, \forall t \in T_y^Y. \quad (4)$$

Thus the flexibility of manpower planning per period is also constrained by (4).

The labor allocated in any period t affects possible overhaul order releases as follows. We let r_i denote the amount of labor hours required to start overhaul of a type $i \in I$ rotatable. Then overhaul order releases must satisfy:

$$\sum_{i \in I_t} r_i n_{i,t} \leq w_t, \quad \forall t \in T. \quad (5)$$

Finally, we note that W_y can be changed from one aggregated period to the next. Such a change from aggregated period y to $y+1$ is bounded from below and above as a fraction of W_y by Δ_y^l and Δ_y^u respectively:

$$\Delta_y^l W_y \leq W_{y+1} \leq \Delta_y^u W_y, \quad \forall y \in \{1, \dots, |Y| - 1\}. \quad (6)$$

Finally we note that W_1 is initialized by the parameter W^d .

3.3 Rotable availability

Since the asset from which the rotatables are to be replaced require high availability, we require that replacements may not occur unless there is a ready-for-use rotatable available to complete the replacement. Similarly, we require that the release of an overhaul order must be accompanied immediately by a rotatable requiring overhaul. Recalling our assumption that the replacements and overhaul order releases during any period are uniformly distributed over that period, rotatable availability can be expressed as

$$n_{i,t} \leq H_{i,t} + x_{i,t}, \quad \forall i \in I, \forall t \in T_i^I, \quad (7)$$

$$x_{i,t} \leq B_{i,t} + n_{i,t-L_i} \quad \forall i \in I, \forall t \in T_i^I. \quad (8)$$

3.4 Overhaul deadlines propagation

Due to safety and possibly other reasons, the maintenance program of rotatables of type $i \in I$ stipulates that any rotatable of type i has to be replaced before it has been operational for q_i periods. Thus for each period in the planning horizon, there are a number of rotatables of type

$i \in I$ that have to be replaced before or in that period, and we denote this quantity $D_{i,t}$ for rotables of type $i \in I$ in period $t \in T_i^I$. For a given rotatable type $i \in I$, these quantities are known for period a_i up to $\min\{a_i + q_i - 1, p_i\}$ and given by $D_{i,t}^d$.

We assume that rotables of any type are replaced in an oldest rotatable first fashion, i.e., whenever a rotatable of any type is to be overhauled, the specific rotatable of that type that has been in the field the longest is always chosen. Thus, from period $a_i + q_i$ onwards

$$D_{i,t} = x_{i,t-q_i}, \quad \forall i \in I, \forall t \in \{a_i + q_i, \dots, p_i\}. \quad (9)$$

It is possible to replace rotables ahead of time, and we let U_i^d denote the number of rotables of type $i \in I$ that have been replaced ahead of time at time $a_i - 1$. To comply with the maintenance program, the replacements have to satisfy:

$$U_i^d + \sum_{t'=a_i}^t x_{i,t'} \geq \sum_{t'=a_i}^t D_{i,t'} \quad \forall i \in I, \forall t \in T_i^I. \quad (10)$$

This constraint can also be described using an auxiliary variable, $U_{i,t}$, that represents the number of replacements of rotables of type i in excess of what is strictly necessary by period t .

Proposition 1 *The set of inequalities (10) is equivalent to the set of constraints:*

$$x_{i,t} \geq D_{i,t} - U_{i,t-1}, \quad \forall i \in I, \forall t \in T_i^I, \quad (11)$$

$$U_{i,t} = x_{i,t} - D_{i,t} + U_{i,t-1}, \quad \forall i \in I, \forall t \in T_i^I \setminus \{p_i\} \quad (12)$$

$$U_{i,a_i-1} = U_i^d, \quad \forall i \in I \quad (13)$$

Proof We show that (12)–(13) imply that

$$U_{i,t} = U_i^d + \sum_{t'=a_i}^t x_{i,t'} - \sum_{t'=a_i}^t D_{i,t'}, \quad \forall i \in I, \forall t \in \{a_i - 1, \dots, p_i - 1\}. \quad (14)$$

Substituting (14) back into (11) yields (10). To verify that (14) and (12)–(13) are equivalent, we use induction. First observe that (13) implies that (14) holds for all $i \in I$ and $t = a_i - 1$. Now suppose that (14) holds for some $i \in I$ and $t - 1 \in \{a_i, \dots, p_i - 1\}$. Then (12) implies that

$$\begin{aligned} U_{i,t} &= x_{i,t} - D_{i,t} + U_{i,t-1} \\ &= x_{i,t} - D_{i,t} + U_i^d + \sum_{t'=a_i}^{t-1} x_{i,t'} - \sum_{t'=a_i}^{t-1} D_{i,t'} \\ &= U_i^d + \sum_{t'=a_i}^t x_{i,t'} - \sum_{t'=a_i}^t D_{i,t'}, \end{aligned} \quad (15)$$

where the second equality holds because of the induction hypothesis. $\square$

The alternative way of writing (10) is useful because it leads to a sparser set of equations that significantly improves the computational feasibility of the model.

3.5 Cost factors

There are four cost factors in our model. Cost per labor hour during aggregated period $y \in Y$ is denoted c_y^W . For rotatables not yet in the field in the first period of the planning horizon, a turn-around stock of rotatables needs to be acquired at the price of c_i^a per rotatable of type $i \in I$. (Note that c_i^a may also include the expected inventory holding cost over the relevant time horizon.) There are also material costs associated with overhaul, and these are denoted $c_{i,t}^m$ for rotatables of type $i \in I$ when the overhaul order was released during period $t \in T_i^I$. Similarly, $c_{i,t}^r$ represent costs of replacing a rotatable of type $i \in I$ during period $t \in T_i^I$. Note that we do not explicitly model the cost of replacing a rotatable earlier than required; these costs can be modeled implicitly through the dependence on time included in all the cost factors. Adding all costs over the relevant horizon we find that the total relevant costs (TRC) satisfy

$$TRC = \sum_{y \in Y} c_y^W W_y + \sum_{i \in I | a_i > 1} c_i^a S_i + \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^m n_{i,t} + \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^r x_{i,t}. \quad (16)$$

3.6 Model remarks

In the above sections, we have given a mathematical description of the planning problem. In this description there are some implicit assumptions that we highlight and justify in this section.

We assume many parameters to be deterministic and known, when in fact they are either random variables whose exact value will only become known later. Consider for example p_i , the period in which type i rotatables become obsolete. This period depends on when the asset for which type i rotatables are used goes out of service. Companies plan the end-of-life of their assets, and so at least estimates of p_i are available in practical situations. We also note that these estimates typically become more accurate as the end-of-life of an asset becomes more imminent. Similar arguments can be made for a_i when $i \in I \setminus I_1$, r_i , q_i , Δ_y^l , Δ_y^u , etc. Since our model should be implemented in a rolling horizon setting, the estimates of these parameters are either very good or deterministic for decisions that need to be implemented in the near future. Obviously, these estimates may have considerable error for periods far into the future. However, these periods are included in the model to account for costs occurring later in the life cycle of the involved assets, but that are affected by current decisions. Furthermore, we note that our model can easily deal with non-stationarity in the input over time, while stochastic models generally cannot. As a final argument, we would like to point out that Dzielinski et al. (1963) and Spitter (2005) (Chap. 6) have tested deterministic rolling horizon models via simulation in dynamic and/or stochastic environments and have shown that these models perform favorably and approach the performance of optimization models that do incorporate stochasticity. Comparisons with stochastic optimization models is only possible, however, in relatively simple environments. In particular it is usually assumed that stochastic quantities have stationary distributions over time, which, in our setting, is unlikely at best.

3.7 Mixed integer programming formulation

The modeling results of the previous sub-sections lead to an optimization problem that we shall call the aggregate rotatable overhaul and supply chain planning (AROSCP) problem. For convenience, all introduced notation is summarized in Table 2, where also a distinction

Table 2 Overview of notation

Sets

I	Set of all types of rotables (not the rotables themselves)
I_t	Set of all types of rotatable types in the field in period $t \in T$, $I_t = \{i \in I a_i \leq t \leq p_i\}$
T	Set of all periods considered in the planning horizon (typically months)
T_i^I	Set of periods during which rotatable $i \in I$ is active in the field ($T_i^I = \{a_i, \dots, p_i\}$)
Y	Set of aggregated periods (typically years)
T_y^Y	Set of periods that are contained in a certain aggregated period $y \in Y$

Input parameters

a_i	First period in the planning horizon that rotables of type $i \in I$ are in the field ($a_i \in T$)
B_i^d	Number of ready-for-use rotables of type $i \in I$ available (on stock) at the beginning of period a_i
c_i^a	Acquisition cost of rotatable $i \in I \setminus I_1$
$c_{i,t}^m$	Material costs associated with releasing an overhaul order
$c_{i,t}^r$	Costs of replacing a rotatable $i \in I$ during period $t \in T_i^I$
c_y^W	Cost per labor hour during aggregated period $y \in Y$
$D_{i,t}^d$	Number of rotables of type $i \in I$ that require overhaul in or before period $t \in \{a_i, \dots, a_i + q_i\}$
H_i^d	Number of <i>non</i> -ready for use rotables of type $i \in I$ on stock at the beginning of period a_i
L_i	The overhaul lead time (in periods) for rotables of type $i \in I$
$n_{i,t}^d$	Number of overhaul order releases of rotables of type $i \in I$ in period $t \in \{a_i - L_i, \dots, a_i - 1\}$
p_i	Last period that rotables of type $i \in I$ are in the field during the planning horizon ($p_i \in T$)
q_i	Inter-overhaul deadline for rotables of type $i \in I$
r_i	Amount of labor hours required to start overhaul of a type $i \in I$ rotatable
U_i^d	Number of replacements of rotables of type i in excess of what is strictly necessary by period $a_i - 1$
W^d	The number of labor hours available in the first aggregate period
$\Delta_y^l(\Delta_y^u)$	Lower (upper) bound on the change in number of labor contracts from aggregated period y to $y + 1$, $y \in \{1, \dots, Y - 1\}$
$\delta_t^l(\delta_t^u)$	Lower (upper) bound on labor hours for rotatable overhaul made available in period $t \in T$ expressed as a fraction of $W_y/ T_y^Y $, for $t \in T$

(Auxiliary) variables

$B_{i,t}$	Number of ready-for-use rotables of type $i \in I$ available at the beginning of period $t \in T_i^I$
$D_{i,t}$	Number of rotables of type $i \in I$ that require overhaul in or before period $t \in T_i^I$
$H_{i,t}$	Number of <i>non</i> -ready for use rotables of type $i \in I$ at the beginning of period $t \in T_i^I$
$U_{i,t}$	Number of replacements of rotables of type i in excess of what is strictly necessary by period t , i.e. $U_{i,t} = \sum_{t'=a_i}^t x_{i,t'} - \sum_{t'=a_i}^t D_{i,t'}$

Decision variables

$n_{i,t}$	Number of overhaul order releases of rotables of type $i \in I$ during period $t \in \{a_i - L_i + 1, \dots, p_i\}$
S_i	Turn-around stock of rotables of type $i \in I$
W_y	Number of labor hours available in aggregated period $y \in Y$
w_t	Number of labor hours for overhaul that are allocated to period $t \in T$
$x_{i,t}$	Number of rotatable replacements of type $i \in I$ during period $t \in T$

is made between sets, parameters, (auxiliary) variables and decision variables. A natural formulation of AROSCP is a mixed integer program, as shown below.

$$\text{minimize } TRC = \sum_{y \in Y} c_y^W W_y + \sum_{i \in I | a_i > 1} c_i^a S_i + \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^m n_{i,t} + \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^r x_{i,t} \quad (17)$$

subject to

$$B_{i,t} = B_{i,t-1} - x_{i,t-1} + n_{i,t-L_i-1} \quad \forall i \in I, \forall t \in T_i^I \setminus \{a_i\} \quad (18)$$

$$H_{i,t} = H_{i,t-1} + x_{i,t-1} - n_{i,t-1} \quad \forall i \in I, \forall t \in T_i^I \setminus \{a_i\} \quad (19)$$

$$B_{i,a_i} = S_i \quad \forall i \in I \setminus I_1 \quad (20)$$

$$B_{i,a_i} = B_i^d \quad \forall i \in I_1 \quad (21)$$

$$H_{i,a_i} = 0 \quad \forall i \in I \setminus I_1 \quad (22)$$

$$H_{i,a_i} = H_i^d \quad \forall i \in I_1 \quad (23)$$

$$n_{i,t} = n_{i,t}^d \quad \forall i \in I, \quad t \in \{a_i - L_i, \dots, a_i - 1\} \quad (24)$$

$$W_y = \sum_{t \in T_y^Y} w_t \quad \forall y \in Y \quad (25)$$

$$\delta_t^l W_y / |T_y^Y| \leq w_t \leq \delta_t^u W_y / |T_y^Y| \quad \forall y \in Y, \forall t \in T_y^Y \quad (26)$$

$$\Delta_y^l W_y \leq W_{y+1} \leq \Delta_y^u W_y \quad \forall y \in \{1, \dots, |Y| - 1\} \quad (27)$$

$$W_1 = W^d \quad (28)$$

$$\sum_{i \in I_t} r_i n_{i,t} \leq w_t \quad \forall t \in T \quad (29)$$

$$n_{i,t} \leq H_{i,t} + x_{i,t} \quad \forall i \in I, \forall t \in T_i^I \quad (30)$$

$$x_{i,t} \leq B_{i,t} + n_{i,t-L_i} \quad \forall i \in I, \forall t \in T_i^I \quad (31)$$

$$x_{i,t} \geq D_{i,t} - U_{i,t-1} \quad \forall i \in I, \forall t \in T_i^I \quad (32)$$

$$U_{i,t} = x_{i,t} - D_{i,t} + U_{i,t-1} \quad \forall i \in I, \forall t \in T_i^I \quad (33)$$

$$U_{i,a_i-1} = U_i^d \quad \forall i \in I \quad (34)$$

$$D_{i,t} = D_{i,t}^d \quad \forall i \in I, \forall t \in \{a_i, \dots, \min\{a_i + q_i - 1, p_i\}\} \quad (35)$$

$$D_{i,t} = x_{i,t-q_i} \quad \forall i \in I, \forall t \in \{a_i + q_i, \dots, p_i\} \quad (36)$$

$$x_{i,t}, n_{i,t} \in \mathbb{N}_0 \quad \forall i \in I, \forall t \in T \quad (37)$$

$$S_i \in \mathbb{N} \quad \forall i \in \{i \in I | a_i > 1\} \quad (38)$$

$$0 \leq n_{i,t}, x_{i,t}, B_{i,t}, H_{i,t}, U_{i,t} \quad \forall i \in I, \forall t \in T \quad (39)$$

$$0 \leq W_y \quad \forall y \in Y \quad (40)$$

$$0 \leq w_t \quad \forall t \in T. \quad (41)$$

Here, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. We remark that it is possible to choose parameter values such that a feasible solution to this MIP does not exist. In particular, infeasibility can be created by

setting the parameters $D_{i,t}^d$ to exceed the available capacity in terms of either work force or rotatable availability.

Because MIPs are hard to solve in general, it is natural to question what the complexity of AROSCP is. In this regard, we offer the following proposition.

Proposition 2 *The aggregate rotatable overhaul and supply chain planning problem (ARO-SCP) is strongly $\mathcal{NP}$ -hard.*

The proof of Proposition 2 uses reduction from BIN-PACKING and is found in Appendix A. In Sect. 4, we provide computational evidence that, despite the computational complexity of the problem, mixed integer programming can still be used to find optimal or close to optimal solutions for instances of real-life size.

3.8 Modeling flexibility

The formulation presented in (17)–(41) still has significant modeling flexibility. We illustrate this flexibility by several examples.

In many practical applications the availability of workforce fluctuates with the time of year; particularly during holiday and summer season there is reduced workforce availability. This can be modeled through the bounds on w_t , δ_t^u and δ_t^l .

The cost parameters in (17) depend on t . This dependence can be used to penalize early overhaul of rotatables and to discount future costs, e.g., by taking $c_{i,t}^m = \alpha^t c_i^m$ with $\alpha \in (0, 1]$.

Labor flexibility has taken a very specific form that is congruent with the setting we will describe in Sect. 4. Traditionally, labor flexibility has been modeled by including overtime at extra cost in the model, as has also been done in Bitran and Hax (1977) and the related literature as reviewed in Sect. 2. These modeling constructs are easily incorporated into our MIP formulation.

In our model we have assumed capacity bounds to exist only on labor in the overhaul workshop. The model can easily be extended with capacity constraints on the number of replacements in the maintenance depot and capacity constraints of different types (e.g., on equipment and tools) in both the overhaul workshop and the maintenance depot. Note however that when these constraints are clearly not binding, it is best to avoid the extra modeling and data collection efforts associated with such extensions.

4 Case study

In this section, we report on a case-study at NedTrain, a Dutch company that maintains rolling stock. The fleet maintained by NedTrain consists of some 3000 carriages across 6 main train types. Almost all carriages rest on two bogies. Bogies are rotatables and there are about 30 different types of bogies in the fleet maintained by NedTrain. In the city of Haarlem, NedTrain has a facility dedicated to the overhaul of all types of bogies in the fleet. Bogies are considered important rotatables and this case-study is about the overhaul and supply chain planning of rotatables at NedTrain. An example of a bogie is shown in Fig. 2. The data set used for the case study we present is outlined in considerable detail in the master thesis of Vernooij (2011). Here, we present a high level description of the data. Rolling stock has a planned life cycle of 30 years and our model uses this as the length of the planning horizon. The period length we consider is a month, while the aggregated period length is a year. The instance we study has 56 bogie types, i.e. $|I| = 56$, 30 bogie types of which are

Fig. 2 An example of a bogie

currently in operation and 26 of which belong to new types of trains that will enter the fleet some time in the next 30 years. The population size of any bogie type ranges from 32 to 611 and depends on how many trains there are of a specific type in the fleet, and how often a bogie type appears in any specific trainset. (For instance, bogies with traction engines appear less often than bogies without in most trainsets.) The flexibility of changing capacity from one aggregated period to the next is limited at 10 %, i.e., $\Delta_y^l = 0.9$ and $\Delta_y^u = 1.1$ for all $y \in \{1, \dots, |Y| - 1\}$. The flexibility of allocating labor to specific periods is also limited to 10 %, i.e., $\delta_t^l = 0.9$ and $\delta_t^u = 1.1$ for all $t \in T$. The MIOTs, q_i , range from 72 to 240 months. Overhaul lead times are 1 period for all bogie types ($L_i = 1$ for all $i \in I$). To start overhaul of any bogie (of any type), 200 hours of labor need to be available ($r_i = 200$ for all $i \in I$). For confidentiality reasons, we do not report any real cost figures. Under the MIP formulation in this paper, this instance has 64896 variables (42968 of which are auxiliary variables) and 76378 constraints.

4.1 Computational feasibility

Since the AROSCP is $\mathcal{NP}$ -hard, we first test the computational feasibility of the model. To this end we propose 3 ways to (approximately) solve the problem:

- (i) Solve the MIP formulation while allowing for an optimality gap of 1 %
- (ii) Relax the integrality constraints on $n_{i,t}$ and $x_{i,t}$ and solve the resulting MIP while allowing for an optimality gap of 1 %¹
- (iii) Solve the linear programming relaxation of the MIP formulation.

All these three methods can be readily implemented using several MIP/LP solvers. We did this for four well-known solvers: CPLEX 12.5.0.0², GUROBI 4.6.1.³, CBC 2.7.5⁴, and

¹The reverse option in which the integrality constraints are only relaxed for S_i is not interesting because for integral $x_{i,t}$ and $n_{i,t}$, optimal S_i cannot be fractional.

²CPLEX is a commercial solver that can use multiple CPU cores in parallel. More information on this solver can be found on <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.

³GUROBI is a commercial solver that can use multiple CPU cores in parallel. More information on this solver can be found on <http://www.gurobi.com/>.

⁴CBC stands for Coin Branch and Cut and is an open source solver associated with the COIN-OR initiative. At present, CBC can only use one CPU core. More information on this solver can be found on <http://www.coin-or.org/Cbc/>

Table 3 Computational times for different solvers and solution methods using ‘out of the box’ settings

Solver	Solution Method	Average	Halfwidth of 95 % CI
GUROBI 4.6.1	MIP (MIPGap 1 %)	5128.0	27.83
	Partial MIP relaxation	119.2	0.30
	LP relaxation	85.6	0.60
CPLEX 12.5.0.0	MIP (MIPGap 1 %)	Out of memory after 881 seconds	
	Partial MIP relaxation	186.9	0.12
	LP relaxation	126.8	0.29
CBC 2.7.5	MIP (MIPGap 1 %)	Infeasible after 43200 seconds	
	Partial MIP relaxation	207.4	0.49
	LP relaxation	293.8	0.16
GLPK 4.47	MIP (MIPGap 1 %)	Infeasible after 43200 seconds	
	Partial MIP relaxation	3031.8	2.40
	LP relaxation	138.2	0.09

GLPK 4.47.⁵ We solved the instance of AROSCP described above 5 times for each combination of solver and (approximate) solution method. The average computational times and halfwidths of 95 % confidence intervals based on the t -distribution are shown in Table 3. All experiments were run on a machine with Intel Core Duo 2.93 GHz processor and 4GB RAM. For the solvers, we used the ‘out of the box’ settings.

It is notable that only GUROBI can solve the MIP formulation; the other solvers either run out of memory or time. With a computational time of less than two hours, the performance of GUROBI is quite good. All solvers can solve the Partial MIP relaxation and the LP relaxation. The LP relaxation can be solved in a matter of minutes by any solver. In the next section, we show that the results produced by both the partial MIP relaxation and the LP relaxation are quite good in terms of both the estimated LCC and the decisions that follow from the solution.

4.2 Sensitivity of result to integrality constraints

The most important decisions that follow from the model are the dimensioning of aggregate workforce capacity (W_y) and turn-around stocks (S_i). Figure 3 shows the aggregate capacity plan, W_y , for the planning horizon of 30 years as found by the three (approximate) solution methods proposed in Sect. 4.1. From Fig. 3, it is evident that for tactical decision making, the results of both the Partial MIP relaxation and the LP-relaxation are sufficiently accurate, although the results of the Partial MIP relaxation are somewhat closer to the solution of the original MIP.

Results for the turn-around stock levels are also very close across different solution methods, as shown in Fig. 4. Here the turn-around stocks of the LP-relaxation are determined by rounding up to the nearest integer. We remark that rounding the turn-around stock levels found by the LP-relaxation yields results that are closer to the MIP solution than the Partial MIP solution that does not relax integrality constraints on the turn-around stocks, S_i .

⁵GLPK stands for GNU linear programming kit and is an open source solver. GLPK can only use one CPU core. More information on this solver can be found on <http://www.gnu.org/software/glpk/>.

Fig. 3 Aggregate capacity plan for case-study at NedTrain using different solution methods

Fig. 4 Turn-around stock sizes for case at NedTrain as determined by different solution methods

Table 4 shows the costs found by all three solution methods, normalized so that the solution found by the MIP formulation is exactly 100. It is notable that estimated lower bounds of TRC found by solving either relaxation are very tight. Also Table 4 shows that the division of costs over labor, material, acquisition and replacement costs are almost identical across solution methods, suggesting that the solution of the LP-relaxation does not only provide a tight lower bound, but also a similar solution that allocates costly resources in a similar manner.

In conclusion, we observe that for making good decisions and estimating cost accurately, it is sufficient to solve relaxations of AROSCP. In particular the linear programming relaxation is a good candidate given its computational tractability.

Table 4 Cost break down for different solution methods

	MIP	Partial MIP relaxation	LP-relaxation
Labor cost	55.1	54.9	54.9
Rotable acquisition cost	2.0	1.9	1.8
Overhaul material cost	42.7	42.7	42.7
Replacement cost	0.2	0.2	0.2
Total	100.0	99.7	99.6

4.3 Insights from case-study

From Table 4, we know that labor costs are the most dominant cost factor. Our model allows for labor flexibility through the pairs of parameters Δ_y^u , Δ_y^l and δ_t^u , δ_t^l . The first pair of parameters controls what we call long term labor flexibility, as they model how the size of the workforce can be changed over a longer horizon. The second pair of parameters, δ_t^u , δ_t^l , models the flexibility to allocate labor of the current workforce to different periods within the same aggregated period. For this reason, we say that δ_t^u , δ_t^l model short term labor flexibility. We performed a sensitivity analysis on long term versus short term labor flexibility. In what follows, we say that long (short) term labor flexibility is x % when $\Delta_y^u = 1 + x/100$ and $\Delta_y^l = 1 - x/100$ ($\delta_t^u = 1 + x/100$ and $\delta_t^l = 1 - x/100$) for all $y \in Y$ ($t \in T$). Figure 5 shows how TRC varies with different percentages of long and short term labor flexibility. Here again costs were normalized to 100 for the MIP solution of the original instance with 10 % labor flexibility in both the short and long term. It appears that short term labor flexibility has relatively little effect on costs over the horizon under consideration, while long term labor flexibility can be leveraged quite effectively. An explanation for this is that the greatest gains in planning rotatable overhaul supply chains are often achieved by moving replacements and overhauls more than a year backward in time. Thus, effective planning does not rely on moving labor capacity around in the short term. Rather, gains can be made by planning replacement and overhauls such that exercising short term labor flexibility has only marginal impact. Overhauls and replacements interact with each other on the time scale of the MIOT. Thus, taking the entire life cycle of an asset and not artificially penalizing early overhaul and replacements really pays off.

At NedTrain, it is practice to not plan overhauls and replacements very far into the future. The reason is that the MIOTs are subject to some uncertainty. The engineers that fix the MIOTs try to fix them as late as possible in the hope that they may stretch these MIOTs. The basic idea is that, by stretching the MIOT, a rotatable needs to undergo overhaul less often per time unit and so associated material and labor costs are incurred less often. While this is true for asset with an infinite life cycle and overhaul shops that have capacity available when convenient and not otherwise, it is not necessarily true for assets with a finite life cycle and overhaul shops that provide specialized labor that has to be contracted ahead of time. A result of knowing the MIOT late is that the overhaul workshop does not know how much work to expect, so it plans for the worst case scenario. Especially for the sake of labor costs and workload smoothing, it is much more beneficial to fix MIOTs early and optimize the plan for overhaul and supply chain operations.

Fig. 5 The value of long term versus short term labor flexibility

5 Numerical results for randomly generated instances

The results for the case study indicate that the LP-relaxation of our formulation yields sufficiently accurate results to aid decision making. In this Section, we show that this behavior is typical by generating instances of the planning problem randomly and verifying that similar results are found. In Sect. 5.1, we give an overview of how instances are generated (pseudo) randomly, and in Sect. 5.2, we explain the metrics we use to compare the LP-relaxation to the MIP optimum and discuss the numerical results.

5.1 Random instance generator

We generate instances randomly, but the orders of magnitude from which we generate values for these instances are based on the orders of magnitude observed at NedTrain, the company of the case study in Sect. 4. For a more detailed discussion of what these orders of magnitude are and how they arise we refer to Vernooij (2011). Table 5 shows how instances were generated (pseudo) randomly. A more detailed explanation of how instances are generated randomly is provided in Appendix B. In Table 5 and Appendix B, we use the notation $\mathcal{UD}(a, b)$ to denote a discrete uniform random variable on the integers $a, \dots, b$ and $\mathcal{U}(a, b)$ to denote the (continuous) uniform random variable on the interval (a, b) .

5.2 Results

In Sect. 4, the results of the original MIP and LP-relaxation were quite close, as evidenced by Figs. 3–5 and Table 4. In this experiment, we measure how ‘close’ the solutions of the MIP and LP-relaxation are by eight metrics. In this section, we use the superscripts LP and MIP on variables to denote that their values are obtained by solving the LP-relaxation and MIP formulation respectively. The eight metrics we consider are:

$$\Delta_{TRC} = \frac{TRC^{MIP} - TRC^{LP}}{TRC^{MIP}} \quad (42)$$

Table 5 Overview of how instances are generated randomly

Parameter/set	Generation	Index range
Sets		
I_1	$\{1, 2, 3, \dots, \mathcal{UD}(25, 40)\}$	–
$I \setminus I_1$	$\{ I_1 + 1, \dots, I_1 + \{j \in I_1 : p_j < 360\} \}$	–
I	$I_1 \cup I \setminus I_1$	–
T	$\{1, 2, \dots, 360\}$	–
Y	$\{1, 2, \dots, 30\}$	–
T_y^Y	$\{12 \cdot (y - 1) + 1, \dots, 12 \cdot (y - 1) + 12\}$	$y \in Y$
Rotable characteristics		
a_i	1	$i \in I_1$
p_i	$\min\{360, \mathcal{UD}(11, 460)\}$	$i \in I_1$
q_i	$\mathcal{UD}(72, 240)$	$i \in I_1$
r_i	$\mathcal{UD}(180, 220)$	$i \in I_1$
L_i	1	$i \in I$
a_i	$p \min\{j \in I_1 : \{k \in \{1, \dots, j\} : p_k < 360\} = i - I_1 \} + 1$	$i \in I \setminus I_1$
p_i	360	$i \in I \setminus I_1$
q_i	$q \min\{j \in I_1 : \{k \in \{1, \dots, j\} : p_k < 360\} = i - I_1 \}$	$i \in I \setminus I_1$
r_i	$r \min\{j \in I_1 : \{k \in \{1, \dots, j\} : p_k < 360\} = i - I_1 \}$	$i \in I \setminus I_1$
Initialization and flexibility		
τ_i	$a_i + \mathcal{UD}(10, q_i)$	$i \in I$
D_{i, τ_i}^d	$\mathcal{UD}(30, 600)$	$i \in I$
$D_{i, t}^d$	0	$i \in I, \quad t \in \{a_i, \dots, a_i + q_i\} \setminus \{a_i + \tau_i\}$
U_i^d	0	$i \in I$
$n_{i, t}^d$	0	$i \in I, \quad t \in \{a_i - L_i, \dots, a_i - 1\}$
H_i^d	0	$i \in I$
B_i^d	$\mathcal{U}(0.1, 0.3) \cdot D_{i, \tau_i}^d$	$i \in I$
W^d	150000	–
$\Delta^l(\Delta^u)$	$\mathcal{U}(0.7, 0.95) \quad (\mathcal{U}(1.05, 1.3))$	–
$\Delta_y^l(\Delta_y^u)$	$\Delta^l(\Delta^u)$	$y \in \{1, \dots, 29\}$
$\delta^l(\delta^u)$	$\mathcal{U}(0.7, 0.95) \quad (\mathcal{U}(1.05, 1.3))$	–
$\delta_t^l(\delta_t^u)$	$\delta^l(\delta^u)$	$t \in T$
Cost parameters		
α	0.95	–
c_i^a	$\mathcal{UD}(300000, 400000) \cdot \alpha^{a_i/12-1}$	$i \in I \setminus I_1$
$c_{i, t}^m$	$\mathcal{UD}(4000, 6000) \cdot \alpha^{\lceil t/12 \rceil - 1}$	$i \in I, \quad t \in T_i^I$
$c_{i, t}^r$	$\mathcal{UD}(30, 50) \cdot \alpha^{\lceil t/12 \rceil - 1}$	$i \in I, \quad t \in T_i^I$
c_y^w	$\mathcal{UD}(60, 80) \cdot \alpha^{y-1}$	$y \in Y$

$$\Delta_W = \frac{|C_W^{LP} - C_W^{MIP}|}{C_W^{MIP}} \cdot 100 \%, \quad \text{with } C_W = \sum_{y \in Y} c_y^W W_y \quad (43)$$

$$\Delta_a = \frac{|C_a^{LP} - C_a^{MIP}|}{C_a^{MIP}} \cdot 100 \%, \quad \text{with } C_a = \sum_{i \in I: a_i > 1} c_i^a S_i \quad (44)$$

$$\Delta_m = \frac{|C_m^{LP} - C_m^{MIP}|}{C_m^{MIP}} \cdot 100 \%, \quad \text{with } C_m = \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^m n_{i,t} \quad (45)$$

$$\Delta_r = \frac{|C_r^{LP} - C_r^{MIP}|}{C_r^{MIP}} \cdot 100 \%, \quad \text{with } C_r = \sum_{i \in I} \sum_{t \in T_i^I} c_{i,t}^r x_{i,t} \quad (46)$$

$$\Delta_{\text{capacity}}^{\max} = \max_{y \in Y} \left| \frac{W_y^{LP} - W_y^{MIP}}{W_y^{MIP}} \cdot 100 \% \right| \quad (47)$$

$$\Delta_{\text{capacity}}^{\max(5)} = \max_{y \in \{1, \dots, 5\}} \left| \frac{W_y^{LP} - W_y^{MIP}}{W_y^{MIP}} \cdot 100 \% \right| \quad (48)$$

$$\Delta_S = \frac{\sum_{i \in I \setminus I_1} |\lceil S^{LP} \rceil - S^{MIP}|}{|I \setminus I_1|} \quad (49)$$

Metrics (42)–(46) measure the relative deviation of the objective function and the different terms of the objective function; together they convey the same information as Table 4 does for the case study. Metrics (47)–(48) measure the relative deviation of aggregate capacity decisions, for the long term and the short term. In the case-study, this information is conveyed by Fig. 3. Finally, metric (49) measures the average absolute deviation of sizing the turn-around stock and conveys the information shown by Fig. 4.

In Sect. 3, we already noted that not every instance of AROSCP is feasible. This is true in particular for instances generated randomly, as explained in Sect. 5.1. In this experiment, we generated instances until 200 feasible instances were obtained. To achieve this, a total of 280 instances were generated. For these 200 instances, we solved the MIP formulation (while allowing for an optimality gap of 1 %), and the LP-relaxation and computed metrics (42)–(49). Table 6 reports the results as well as the computation times (in minutes) on a machine with Dual Core 2.9 GHz processor with 4 GB of RAM and GUROBI 5.0 as solver.

First, we note that the relative deviation with respect to the optimal objective value and its separate components is very small. Laying Δ_a aside for a moment, we see that Δ_{TRC} , Δ_W , Δ_m , Δ_r are all well below 1.0 % on average and well below 1.5 % in the worst case. This is remarkable, especially considering that the MIP solution still has a remaining optimality gap somewhere below 1.0 %. The odd one out is Δ_a . As in the case study, the acquisition

Table 6 Accuracy of LP-relaxation in approximating an integer optimal solution ($N = 200$)

	Δ_{TRC}	Δ_W	Δ_a	Δ_m	Δ_r	$\Delta_{\text{capacity}}^{\max}$	$\Delta_{\text{capacity}}^{\max(5)}$	Δ_S	CompTime [min]
Avg	0.73	0.23	5.78	0.15	0.14	7.54	2.11	0.19	39.4
Min	0.33	0.00	0.48	0.00	0.00	1.93	0.02	0.00	1.3
Max	1.37	1.27	16.57	0.49	0.44	20.17	9.44	0.50	334.0
Stdev	0.23	0.21	2.46	0.11	0.10	3.39	1.63	0.11	61.7

costs are a relatively small part of the total costs (as evidenced here by the fact that Δ_{TRC} is considerably smaller than Δ_a). Furthermore, the direct comparison of turn-around stocks provided by the metric Δ_S is again quite favorable, probably also because S^{LP} is rounded up for the purpose of comparison. All in all, these results indicate that the LP-relaxation can be used to perform sensitivity analyses with respect to costs, thus saving considerable computation time. Furthermore, it is possible to use the dual variables provided by solving the LP-relaxation using the simplex method to streamline the sensitivity analysis. Consider, for example, the sensitivity of the model with respect to the bounds on capacity flexibility Δ_y^l and Δ_y^u . In the case study, the sensitivity to these bounds was investigated by repeatedly solving the problem, but via dual variables it is possible to investigate the sensitivity of the optimal solution to these bounds around some operating point.

The measures $\Delta_{\text{capacity}}^{\max}$ and $\Delta_{\text{capacity}}^{\max(5)}$ appear quite high. We note however that $\Delta_{\text{capacity}}^{\max} = 9.01\%$ and $\Delta_{\text{capacity}}^{\max(5)} = 0.35\%$ for the case study instance. Thus, performance is comparable to the instance in the case-study.

In closing, we note that the computation times for solving the MIP formulation are considerable, even going up to over 5.5 hours. For sensitivity and scenario analyses, it seems the LP-relaxation provides a good substitute with more acceptable computation times.

6 Conclusion

In this paper, we have presented a model for the aggregate planning of rotatable overhaul and supply chain operations. The model has many realistic features and incorporates LCC considerations in planning decisions when it is used in a rolling horizon setting. Despite the fact that solving the presented model to optimality is $\mathcal{NP}$ -hard in general, we have provided evidence to suggest that a linear programming relaxation of the problem supplies the user with useful information that aids in decision making and even yields decisions that are close to optimal for large instances of AROSCP, as found in practice. In the context of a real life case study, we have argued that it is beneficial to fix MIOTs relatively early so that an effective plan for overhaul and supply chain operations can be made that utilizes the flexibility of overhaul planning that exists only when considering the entire life cycle of an asset.

Acknowledgements The authors would like to thank Bob Huisman of NedTrain for introducing this research topic to the authors and many stimulating discussions. The authors also thank Karin Vernooij who conducted an initial study in this subject, and who collected some of the data that we use. The authors also thank anonymous referees, Geert-Jan van Houtum and Matthew Reindorp for their constructive feedback.

Appendix A: Proof of Proposition 2

We show that being able to solve the AROSCP will enable one to decide the BIN-PACKING decision problem, i.e. we reduce BIN-PACKING to AROSCP. The following decision problem, known as BIN-PACKING, is strongly $\mathcal{NP}$ -complete (e.g. Garey and Johnson 1979): Given positive integers $\alpha_1, \dots, \alpha_m$, β , and κ , is there a partition of $\{1, \dots, m\}$ into disjoint sets $\Upsilon_1, \dots, \Upsilon_\kappa$ such that $\sum_{j \in \Upsilon_i} \alpha_j \leq \beta$ for $i = 1, \dots, \kappa$?

Now suppose we are given an instance of BIN-PACKING. Without loss of generality, we may assume that $\sum_{i=1}^m \alpha_i \leq \kappa\beta$ and $\alpha_i \leq \beta$ for all $i \in \{1, \dots, m\}$. From this instance of BIN-PACKING, we will show how to create an instance of AROSCP in polynomial time such that the answer to this instance of BIN-PACKING is yes, if and only if the optimal

objective value of the corresponding instance of AROSCP equals 0. The basic idea behind this reduction is the following. By setting the initial number of non ready-for-use rotables sufficiently high, the release of overhaul orders is constrained only by available workforce capacity, i.e. by (29). This workforce capacity can be kept constant at β across periods by constraints (27) and (28). Now the problem can be viewed as packing overhaul order releases into several one period bins of fixed size β . By penalizing these order releases in all but κ periods, the objective becomes to pack as many order releases as possible in the κ periods in which the order releases are not penalized. If the optimal objective of AROSPC is 0, then it was possible to pack all overhaul order releases in κ periods and so the instance of BIN-PACKING is a yes-instance.

More formally, the reduction is as follows. Set $Y = \{1, \dots, m+1\}$ and $T_y^Y = \{y\}$ for all $y \in Y$; thus, aggregated and regular periods coincide. Set $W^d = \beta$, and $\Delta_y^l = \Delta_y^u = \delta_y^l = \delta_y^u = 1$. This ensures capacity is identical across periods. Set $I = \{1, \dots, m\}$ and set $a_i = 1$, $p_i = m+1$, $q_i = m+1$, $L_i = 1$, $H_i^d = 1$, $B_i^d = 0$, $n_{i,0}^d = 0$, $D_{i,m+1}^d = 1$ and $r_i = \alpha_i$ for all $i \in I$. Furthermore, set $D_{i,t}^d = 0$ for all $i \in I$ and $t \in \{1, \dots, m\}$. Thus, each type of rotatable needs to be replaced exactly once before or in the last period of the planning horizon. This instance of AROSCP is feasible because the following is a feasible solution: $n_{i,i} = 1$ for $i \in I$ and $n_{i,t} = 0$ otherwise, $x_{i,m+1} = 1$ for all $i \in I$ and $x_{i,t} = 0$ otherwise. (Note that all other variables are set by constraints). There are no acquisitions ($a_i = 1$ for all $i \in I$) so c_i^a does not need to be set. Most other cost parameters are set to 0; $c_y^w = 0$ for all $y \in Y$ and $c_{i,t}^r = 0$ for all $i \in I$ and $t \in T_i^l$. However, we set $c_{i,t}^m = 1$ for all $i \in I$ and $t \in \{1, \dots, m-\kappa\}$ and set $c_{i,t}^m = 0$ otherwise. Note that $m-\kappa \geq 1$ because, by assumption, $\sum_{i=1}^m \alpha_i \leq \kappa\beta$ and $\alpha_i \leq \beta$ for all $i \in \{1, \dots, m\}$. The objective function now reduces to $\sum_{i \in I} \sum_{t=1}^{m-\kappa} c_{i,t}^m n_{i,t}$. Let OPT denote the optimal solution to this instance of AROSCP. If $OPT = 0$ then, necessarily $n_{i,t} = 0$ for all $i \in I$ and $t \in \{1, \dots, m-\kappa\}$. Furthermore, by constraint (25), $\sum_{i \in I} r_i n_{i,t} \leq w_t$ for all $t \in T$, which, by our choice of parameter values, implies $\sum_{i \in I} \alpha_i n_{i,t} \leq \beta$ for all $t \in \{m-\kappa+1, \dots, m\}$. All rotatables in this instance of AROSCP have to be overhauled exactly once in or before period m because of constraints (31), (32) and (35). Therefore, for each $i \in I$ there is some $t \in \{m-\kappa+1, \dots, m\}$ such that $n_{i,t} = 1$. when $OPT = 0$. Now it follows that a partition that satisfies the requirement of the original BIN-PACKING problem is given by:

$$\Upsilon_j = \{i \in I | n_{i,m-\kappa+j} = 1\}, \quad j \in \{1, \dots, \kappa\}.$$

In an analogous manner, it is possible to construct an optimal solution with objective 0 to an instance of AROSCP if the corresponding instance and truth certificate of BIN-PACKING is given, by setting all $x_{i,t}$ and $n_{i,t}$ to 0, except $x_{i,m+1} = 1$ for all $i \in I$ and $n_{i,m-\kappa+j} = 1$ if $i \in \Upsilon_j$. Thus, we have shown that an instance of BIN-PACKING is a yes-instance if and only if the corresponding AROSCP problem has an optimal objective of 0. Observing further that the reduction above is a pseudo-polynomial reduction as defined by Garey and Johnson (1979, p. 101), and that BIN-PACKING is strongly $\mathcal{NP}$ -complete, completes the proof. $\square$

Appendix B: Details on the random instance generator

B.1 Rotable characteristics

First, we generate the number of different rotatable types that are already in the field at the beginning of the planning horizon, $|I_1|$, from $\mathcal{UD}(25, 40)$. Then for each $i \in I_1$, we draw p_i ,

q_i , r_i and L_i as shown in Table 5. Note that for $i \in I_1$, $a_i = 1$ by definition and needs not be generated randomly.

Some of the rotatables $i \in I_1$ may belong to assets that will be disposed of before the end of the planning horizon, i.e., possibly $p_i < 360$ for some $i \in I_1$. If this is the case, we assume that this asset will be replaced by a new type of asset, which consists of rotatables with identical characteristics that will remain current for the remainder of the planning horizon. For example, if $1 \in I_1$, $p_1 = 270$, $a_1 = 1$, $q_1 = 120$, $r_1 = 200$, $L_1 = 1$, and $|I_1| = 32$, then we add rotatable type 33 to I and set $a_{33} = p_1 + 1 = 271$, $p_{33} = p_1 = 360$, $q_{33} = q_1 = 120$, $r_{33} = r_1 = 200$, and $L_{33} = L_1 = 1$. This procedure is shown formally in Table 5 using set expressions and the fact that I is generated to contain a sequence of integers. We note that it is also possible that a rotatable type that is replaced some time during the planning horizon is replaced with a rotatable type that has different characteristics. In particular, new rotatable types are likely to be more reliable due to technological advancements. The models can also accommodate these scenarios. However, we make the conservative assumption that rotatable types are replaced rotatable types with identical characteristics.

B.2 Initial conditions and flexibility

For each type of rotatable $i \in I$, there are revisions already due. We assume the worst case scenario that the first upcoming revision of any one rotatable type are due in a single period. For rotatable type $i \in I$, this single period is $a_i + \tau_i$ and τ_i is generated as $\tau_i = \mathcal{UD}(10, q_i)$. The number of revisions due in period $a_i + \tau_i$ is drawn from $\mathcal{UD}(30, 600)$. This means that for each $i \in I$, $D_{i,t}^d = 0$ if $t \neq \tau_i$ and $D_{i,\tau_i}^d = \mathcal{UD}(30, 600)$.

Again as a worst case scenario, $U_i^d = 0$ and $n_{i,t}^d = 0$ for all relevant i and t , meaning that there are no recent order releases and that there have been no replacements ahead of time. Initially, for each $i \in I_1$ there is no stock of non-ready for use rotatables, i.e., $H_i^d = 0$ for all $i \in I_1$. The initial ready-for-use stock of $i \in I_1$ is generated as a fraction of the first peak number of revisions due in period $a_i + \tau_i$: $B_i^d = \lceil \mathcal{U}(0.1, 0.3) \cdot D_{i,\tau_i}^d \rceil$, where $\lceil x \rceil$ is x rounded up to the nearest integer.

The bounds Δ_y^l and Δ_y^u are obtained by generating Δ^l (Δ^u) as $\mathcal{U}(0.7, 0.95)$ ($\mathcal{U}(1.05, 1.3)$) and setting $\Delta_y^l = \Delta^l$ and $\Delta_y^u = \Delta^u$ for all $y \in Y \setminus \{|Y|\}$. Similarly δ_t^l (δ_t^u) are obtained by generating δ^l (δ^u) as $\mathcal{U}(0.7, 0.95)$ ($\mathcal{U}(1.05, 1.3)$) and setting $\delta_t^l = \delta^l$ and $\delta_t^u = \delta^u$ for all $t \in T$. Finally, in each case, W^d is set as 150000. We do not generate this parameter randomly, because the r_i are already generated randomly.

B.3 Costs parameters

The costs parameters are discounted on a yearly basis by the discount parameter α . Within a year however, there is no discounting, so that for a period $t \in T$, the corresponding year is $\lceil t/12 \rceil$.

References

- Asiedu, Y., & Gu, P. (1998). Product life cycle cost analysis: state of the art review. *International Journal of Production Research*, 36(4), 883–908.
- Billington, P., McClain, J., & Thomas, L. (1983). Mathematical programming approaches to capacity-constrained MRP systems: review, formulation and problem reduction. *Management Science*, 29(10), 1126–1141.

- Bitran, G., & Hax, A. (1977). On the design of hierarchical production planning systems. *Decision Sciences*, 8, 28–55.
- Bitran, G., Haas, E., & Hax, A. (1981). Hierarchical production planning: a single stage system. *Operations Research*, 29(4), 717–743.
- Bitran, G., Haas, E., & Hax, A. (1982). Hierarchical production planning: a two stage system. *Operations Research*, 30(2), 232–251.
- Charest, M., & Ferland, J. (1993). Preventive maintenance scheduling of power generating units. *Annals of Operations Research*, 41, 185–206.
- Chen, C., Yan, S., & Chen, M. (2010). Short-term manpower planning for MRT carriage maintenance under mixed deterministic and stochastic demands. *Annals of Operations Research*, 181, 67–88.
- De Kok, T., & Fransoo, J. (2003). Planning supply chain operations: definition and comparison of planning concepts. In A. de Kok & S. Graves (Eds.), *Supply chain management: design, coordination and operation* (pp. 597–675). Amsterdam: Elsevier.
- Driessen, M., Arts, J., Van Houtum, G., Rustenburg, W., & Huisman, B. (2010). *Maintenance spare part planning and control: a framework for control and agenda for future research*. Beta working paper 325, Eindhoven University of Technology, Eindhoven, the Netherlands.
- Dzielinski, B., Baker, C., & Manne, A. (1963). Simulation tests of lot size programming. *Management Science*, 9(2), 229–258.
- Ebeling, C. (2010). *Introduction to reliability and maintainability engineering* (2nd ed.). Long Grove: Waveland Press.
- Erengüç, E., Simpson, N., & Vakharia, A. (1999). Integrated production/distribution planning in supply chains: an invited review. *European Journal of Operational Research*, 115(2), 219–236.
- Garey, W., & Johnson, D. (1979). *Computers and intractability: a guide to the theory of NP-completeness*. New York: W.H. Freeman and company.
- Gupta, Y., & Chow, W. (1985). Twenty-five years of life cycle costing—theory and applications: a survey. *International Journal of Quality and Reliability Management*, 2(3), 51–76.
- Hopp, W., & Spearman, M. (2001). *Factory physics*. New York: McGraw-Hill.
- Ilgin, M., & Gupta, S. (2010). Environmentally conscious manufacturing and product recovery (ECMPRO): a review of the state of the art. *Journal of Environmental Management*, 91, 563–591.
- Joo, S. (2009). Scheduling preventive maintenance for modular designed components: a dynamic approach. *European Journal of Operational Research*, 192, 512–520.
- Nahmias, S. (2009). *Production and operations analysis* (6th ed.). New York: McGraw-Hill.
- Öner, K., Franssen, R., & Kiesmüller, G. (2007). Life cycle costs measurement of complex systems manufactured by an engineer-to-order company. In R. Qui, D. Russel, W. Sullivan, & M. Ahmad (Eds.), *The 17th international conference on flexible automation and intelligent manufacturing* (pp. 569–589). Amsterdam: Elsevier.
- Paz, N., & Leigh, W. (1994). Maintenance scheduling: issues, results and research needs. *International Journal of Operations & Production Management*, 14(8), 47–69.
- Ross, S. (1996). *Stochastic processes* (2nd ed.). Hoboken: Wiley.
- Safaei, N., Banjevic, D., & Jardine, A. (2011). Workforce constrained maintenance scheduling for military aircraft fleet: a case study. *Annals of Operations Research*, 186, 295–316.
- Sherbrooke, C. (2004). *Optimal inventory modeling of systems: multi-echelon techniques* (2nd ed.). Boston: Kluwer.
- Silver, E., Pyke, D., & Peterson, R. (1998). *Inventory management and production planning and scheduling*. Hoboken: Wiley.
- Spitter, J. (2005). *Rolling schedule approaches for supply chain operations planning*. PhD thesis, Eindhoven University of Technology, Department of Technology Management, Eindhoven, the Netherlands. <http://alexandria.tue.nl/extra2/200511140.pdf>.
- Spitter, J., Hurkens, C., De Kok, A., Lenstra, J., & Negenman, E. (2005). Linear programming models with planned lead times for supply chain operations planning. *European Journal of Operational Research*, 163, 706–720.
- Van Dijkhuizen, G., & Van Harten, A. (1997). Optimal clustering of frequency-constrained maintenance jobs with shared set-ups. *European Journal of Operational Research*, 99, 552–564.
- Vernooij, K. (2011). *An aggregate planning for preventive maintenance of bogies by NedTrain*. Master's thesis, Eindhoven University of Technology, School of Industrial Engineering, the Netherlands. http://alexandria.tue.nl/extra2/afstversl/tm/Vernooij_2011.pdf.
- Wagner, H., Giglio, R., & Glaser, R. (1964). Preventive maintenance scheduling by mathematical programming. *Management Science*, 10(2), 316–334.