Computers & Industrial Engineering 76 (2014) 390–400

Contents lists available at ScienceDirect

Computers & Industrial Engineering

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

Preventive maintenance scheduling of multi-component systems with interval costs

Emil Gustavsson, Michael Patriksson, Ann-Brith Strömberg *, Adam Wojciechowski, Magnus Önnheim

Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Sweden

A R T I C L E   I N F O

Article history:
Received 1 December 2012
Received in revised form 28 January 2014
Accepted 24 February 2014
Available online 4 March 2014

Keywords:
Maintenance scheduling
Integer optimization
Finite horizon
Dynamic grouping
Polyhedral analysis
Case studies

A B S T R A C T

We introduce the preventive maintenance scheduling problem with interval costs (PMSPIC), which is to schedule preventive maintenance (PM) of the components of a system over a finite and discretized time horizon, given a common set-up cost and component costs dependent on the lengths of the maintenance intervals. We present a 0-1 integer linear programming (0-1 ILP) model for the PMSPIC; the model is identical to that presented by Joneja (1990) for the joint replenishment problem within inventory management. We study this model from a polyhedral and exact solutions' point of view, as opposed to previously studied heuristics (e.g. Boctor, Laporte, & Renaud, 2004; Federgruen & Tzur, 1994; Levi, Roundy, & Shmoys, 2006; Joneja, 1990). We show that most of the integrality constraints can be relaxed and that the linear inequality constraints define facets of the convex hull of the feasible set. We further relate the PMSPIC to the opportunistic replacement problem, for which detailed polyhedral studies were performed by Almgren et al. (2012a). The PMSPIC can be used as a building block to model several types of maintenance planning problems possessing deterioration costs. By a careful modeling of these costs, a polyhedrally sound 0-1 ILP model is used to find optimal solutions to realistic-sized multi-component maintenance planning problems. The PMSPIC is thus easily extended by side-constraints or to multiple tiers, which is demonstrated through three applications; these are chosen to span several levels of unmodeled randomness requiring fundamentally different maintenance policies, which are all handled by variations of our basic model.

Our first application considers rail grinding. Rail cracks increase with increasing intervals between grinding occasions, implying that more grinding passes must be performed—thus generating higher costs. We optimize the grinding schedule for a set of track sections presuming a deterministic model for crack growth; hence, no corrective maintenance (CM) will occur between the grinding occasions scheduled. The second application concerns two approaches for scheduling component replacements in aircraft engines. The first approach is bi-objective, simultaneously minimizing the cost for the scheduled PM and the probability of unexpected stops. In the second approach the sum of costs for PM and expected CM—without rescheduling—is minimized. When rescheduling is allowed, the 0-1 ILP model is used as a policy by re-optimizing the schedule at a component failure, which then constitutes an opportunity for PM. The policy manages the trade-off between costs for PM and unplanned CM and is evaluated in a simulation of the engine. The third application considers components' replacement in wind mills in a wind farm, extending the PMSPIC to comprise multiple tiers with joint set-up costs. Due to the large number of components unexpected stops occur frequently, thus calling for a dynamic rescheduling, which is evaluated through a simulation of the system. In each of the three applications, the use of the 0-1 ILP model is compared with age or constant-interval policies; the maintenance costs are reduced by up to 16% as compared with the respective best simple policy. The results are strongest for the first two applications, possessing low levels of unmodeled randomness.

1. Introduction

To ensure that a system stays operational, or to restore a failed system to an operational state, requires maintenance; different system states call for different types of maintenance activities.

* Corresponding author. Tel.: +46 317725378.
E-mail addresses: emilg@chalmers.se (E. Gustavsson), mipat@chalmers.se (M. Patriksson), anstr@chalmers.se (A.-B. Strömberg), wojadam@gmail.com (A. Wojciechowski), onnheimm@chalmers.se (M. Önnheim).

http://dx.doi.org/10.1016/j.cie.2014.02.009
0360-8352/

Maintenance optimization means deciding which maintenance activities to perform, and when, such that one or several objectives are optimized. Maintenance optimization models of systems comprising one or several components, and including repairs and replacements of components, as well as inspections and condition monitoring, are extensively studied in the literature; see the surveys (Dekker, Wildeman, & van der Duyn Schouten, 1997; Nicolai & Dekker, 2008; Wang, 2002). Of particular interest to this article are two types of maintenance activities, often denoted *preventive maintenance* (PM)—performed in order to avoid failure—and *corrective maintenance* (CM)—performed after failure in order to restore the system into an operational state.

This article considers the scheduling of PM activities for a multi-component system using a dynamic finite horizon model. That is, the system to be maintained consists of several components assumed to possess a positive economic dependence such that any maintenance activity generates common *set-up costs* shared by the components. The model is dynamic in order to incorporate unexpected events, i.e., CM activities. In the sequel, we denote by *maintenance occasion* that maintenance occurs for at least one component in the system. Further, *replacement* will denote a generic maintenance activity for a single component, even though in our case studies a PM action is not always an actual replacement.

A common approach to maintenance scheduling—or maintenance decision making—is to use a simple policy, which often contains a number of parameters whose values are optimized either numerically or analytically. The following policies are of interest for the problems studied in this article and will be compared with the optimization model developed. (i) The *constant-interval* policy (CI) (e.g. Tian, Jin, Wu, & Ding, 2011) is to replace all components after a predefined period (the parameter of the policy). (ii) The *age policy*, in which a component is replaced when it reaches a predefined age or at failure, was originally developed for single-component systems. For multi-component systems we consider an age policy with 'soft' and 'hard' component lives (constituting the parameters of the policy), as described by Crocker and Kumar (2000) and summarized as follows: "A maintenance occasion is enforced if the age of any component reaches its hard life or if a component failure occurs. At a maintenance occasion, additionally failed components and components having surpassed their soft life are replaced." That is, the 'hard' life parameter sets a hard limit on the interval between replacements of a component in the system; the 'soft' life is the age parameter after which a component is replaced if the set-up cost has been triggered by some other component. (iii) A policy based on *target built life* (TBL) with hard lives (the TBL and the hard lives constituting the policy parameters) is then considered, as described by Crocker and Sheng (2008): "A maintenance occasion is enforced if the age of any component reaches its hard life or if any component fails. Given a maintenance occasion, components are replaced until the expected number of component failures before the TBL is reached is below one." Note that the constant-interval policy corresponds to a fixed schedule, while the age and TBL policies do not.

The scheduling problem considered in this article is an extension of the *opportunistic replacement problem* (ORP) studied by Almgren et al. (2012a) and described as follows: "The system consists of a set of components. The time between two consecutive replacements of a component may not exceed its assigned maximum replacement interval. To each time point in the planning period corresponds a fixed maintenance set-up cost and replacement costs for each component. The problem is to schedule the component replacements over a finite set of time points in order to minimize the total maintenance cost." Systems consisting of safety critical components should be maintained according to this principle. For each component in such a system the maximum replacement interval corresponds to a technical life which is assigned based on safety criteria. For other types of systems, however, a failure might be a mere inconvenience. Further, a failure may correspond to a signal from a condition monitoring system indicating that a threshold value is surpassed, and that a repair or replacement action is necessary for the system to stay in operation. In Almgren et al. (2012a), a 0-1 integer linear programming (0-1 ILP) model yields significant reductions of the maintenance costs as compared with simpler policies of the types (i)–(iii). Patriksson, Strömberg, and Wojciechowski (2014) consider the stochastic ORP, which extends the ORP to allow for uncertain maximum replacement intervals and—given a failure of one component—to decide whether additional components should be replaced, by using a two-stage stochastic programming model. That setting, however, presumes identical costs for unexpected and scheduled maintenance stops. In this article PM is scheduled, but instead of enforcing a maximum replacement interval, a deterioration cost is assigned to the length of the time interval between scheduled PM actions. We will demonstrate by means of case studies that this provides a rich and promising framework for PM scheduling.

The idea of assigning a deterioration cost to a maintenance interval is not new. The *standard indirect grouping* model for PM, reviewed by Dekker et al. (1997), is also based on this idea and contains a fixed *maintenance occasion cost*, a *preventive maintenance cost*, and a *deterioration cost function* for each component. A maintenance stop occurs every T time units and component i is replaced every k_iT time units. A closed form expression of the average maintenance cost is obtained and values for the parameters $T \in \mathbb{R}_+$ and $k_i \in \mathbb{N}$ are chosen by numerical optimization. Since the average cost is minimized, a static infinite horizon model is obtained.

As discussed in Dekker (1995), varying the form of the deterioration cost function yields a large variety of maintenance problems including *optimal block replacement*, *minimal repair*, and *standard inspection*, as well as inventory problems, such as the *joint replenishment problem* (JRP). The JRP has been studied under indirect grouping strategies as well as over a finite horizon (Khouja & Goyal, 2008)—then denoted the *dynamic JRP* (DJRP); it is closely connected to the *preventive maintenance scheduling problem with interval costs* (PMSPIC) considered in this article. Our 0-1 ILP model was introduced by Joneja (1990) for the DJRP (see Section 2 for an in-depth discussion).

Grigoriev, Van De Klundert, and Spieksma (2006) consider the *periodic maintenance problem* (PMP), which includes deterioration costs. The PMP is periodic in that, at the end of the time horizon the maintenance schedule starts over. Since the deterioration cost is deterministic, no rescheduling is needed and the solution obtained is static. The system consists of a set of machines among which at most one at a time may be maintained. Hence, the maintenance occasions typically are spread out over time in contrast to the PMSPIC, for which the component replacements typically are coordinated at fewer time points. Grigoriev et al. also presents a 0-1 ILP model for the PMP, based on a network flow formulation, which resembles our basic model for the PMSPIC (see Section 2.2). Since periodicity may simplify the integration of maintenance and staff planning, periodic maintenance is often desired as output from maintenance policies; we show in this article how periodicity can be incorporated in the PMSPIC through side-constraints.

The remainder of this article is organized as follows. In Section 2 we define the PMSPIC, present a 0-1 ILP model based on a multi-commodity flow formulation, and establish some important properties of the model. Sections 3–5 present three industrial applications of the model. Section 3 considers the grinding of railway tracks, presuming a deterministic model of crack growth. Section 4 considers preventive component replacements in an aircraft engine module using two approaches: (a) the bi-objective minimization of the cost for the scheduled PM and the probability of an unexpected stop and (b) the minimization of the sum of the costs for the scheduled PM and the expected CM costs, which are assigned as the deterioration costs. Section 5 presents an extension of the model for component replacement decisions in a wind farm—again assigning the expected CM cost as the deterioration cost. In Section 6 we draw conclusions and present suggestions for future research.

All computational tests are performed on an Intel 2.80 GHz dual core Linux PC with 4 GB RAM. The mathematical optimization models are implemented in AMPL (version 12.1) and solved using CPLEX (version 12.1).

2. The problem definition and the 0-1 ILP model

We first provide a formal description of the PMSPIC and show that it generalizes the DJRP—defined by Boctor et al. (2004)—and the ORP—defined by Almgren et al. (2012a). We then present a 0-1 ILP model of the PMSPIC (which was introduced for the DJRP by Joneja, 1990) and discuss its relation to models of the DJRP and the ORP. We show that the integrality restrictions on a majority of the variables can be relaxed, and that the inequality constraints define facets of the convex hull of the set of feasible solutions. The most important consequence of these properties is that the models are—relatively speaking—easy to solve. We then introduce a representation of time-independent costs—to be employed in the case studies—and model periodicity constraints, such that solutions conform to the class of standard indirect grouping strategies.

2.1. The PMSPIC

The preventive maintenance scheduling problem with interval costs (PMSPIC) is defined as follows.

Definition 1 (*PMSPIC*). Consider a system with a set $\mathcal{N} := \{1,\dots,n\}$ of components and a set $\mathcal{T} := \{1,\dots,T\}$ of time steps at which maintenance can be performed, T defining the planning horizon. A PM schedule consists of a set of scheduled replacement times in $\mathcal{T}$ for each component. All n components are assumed to obtain PM at the times 0 and $T+1$. A maintenance occasion (defined by the PM of at least one component) at time $t \in \mathcal{T}$ generates the set-up/maintenance occasion cost d_t. If PM of component $i \in \mathcal{N}$ is scheduled at the times $s \in \mathcal{T} \cup \{0\}$ and $t \in \{s+1,\dots,T+1\}$, but *not* in the (possibly empty) time interval $\{s+1,\dots,t-1\}$, then the *maintenance interval*, denoted (s,t), generates the *interval cost* c^i_{st}. Find a PM schedule which minimizes the sum of maintenance occasion and interval costs. □

Definition 1 assumes that PM is performed at times 0 and $T+1$. This is a valid assumption if the system is new at time 0 and scrapped at the end of the planning horizon. A component i possessing a history without PM at time 0 can be included in the PMSPIC by modifying the interval costs $c^i_{0t},\ t \in \mathcal{T} \cup \{T+1\}$. Similarly, a discount on the interval costs $c^i_{s,T+1},\ s \in \mathcal{T} \cup \{0\}$, can account for components obtaining PM near the end of the planning horizon and which can thus continue operating without maintenance beyond it. Note that we consider maintenance to be instantaneous, i.e., if maintenance is planned at time t then—besides unexpected failures and planned maintenance—the system is assumed to be operational at times $t-1$ and $t+1$.

We next show that a special case of the PMSPIC coincides with the inventory problem DJRP (Boctor et al., 2004; Khouja & Goyal, 2008). Consider the replenishment of n products. Let S_t be a fixed cost for ordering products at time $t \in \mathcal{T}$, s_{it} the individual ordering cost (independent of the number of units ordered) of item $i \in \mathcal{N}$ at time $t \in \mathcal{T}$, h_{it} the unit inventory holding cost for item of type i at time t, and $\widehat{d}_{it}$ the demand of item i at time t. The DJRP is to schedule the replenishment of items such that the sum of ordering and inventory costs are minimized and no shortage arises. As observed by Boctor et al. (2004), an optimal schedule exists such that a replenishment of item i at the times u and v, and not in-between these time points, implies that $\sum_{t=u}^{v-1} \widehat{d}_{it}$ units of item i are ordered at the time u. Consider an instance of the PMSPIC such that the set-up cost corresponds to the joint ordering cost of the DJRP instance, and the interval cost of the PMSPIC corresponds to single-item ordering and inventory holding costs of the DJRP. That is, $c^i_{uv} := s_{iu} + \sum_{t=u+1}^{v-1}\left(\sum_{r=u}^{t-1} h_{ir}\right)\widehat{d}_{it},\ \ u \in \mathcal{T} \cup \{0\},\ \ v \in \{u+1,\dots,T+1\}$, and $d_t := S_t,\ t \in \mathcal{T}$. Solving the corresponding instance of the PMSPIC yields a solution to the DJRP. Since the DJRP is NP-hard (Arkin, Joneja, & Roundy, 1989) this implies that also the PMSPIC is NP-hard. That the PMSPIC is a strict generalization of the DJRP follows by dimensionality: not every set $\{c^i_{uv}\}$ of costs can be realized as the image of a set $\{s_{iu}, h_{ir}\}$ of ordering and inventory costs under the above correspondence.

Another problem arising as a special case of the PMSPIC is the ORP (see Section 1), which models PM of components for a multi-component system with a set-up cost d_t at each time $t \in \mathcal{T}$. To each component $i \in \mathcal{N}$ is assigned a maximum replacement interval T_i and a replacement cost c_{it} at time t. The ORP is to find a minimum cost replacement schedule over the time interval defined by the set $\mathcal{T}$. Consider an instance of the PMSPIC defined by the set-up costs d_t, and the interval costs $c^i_{st} := c_{it}$ when $t-s \leqslant T_i$, and $c^i_{st} := T(\max_{u\in\mathcal{T}}\{d_u\} + n \cdot \max_{j\in\mathcal{N},\ u\in\mathcal{T}}\{c_{ju}\}) + 1$ when $t-s \geqslant T_i+1$. (The schedule defined by the replacement of all the components at each time step then possesses a lower cost than any schedule that includes a maintenance interval longer than T_i for any component i.) Any optimal solution to this PMSPIC instance is thus optimal to the ORP, which is an NP-hard problem (Almgren et al., 2012a). Note that the maximum replacement interval T_i corresponds to a maximum inventory capacity for the DJRP; hence, capacitated versions of the DJRP are special cases of the PMSPIC.

2.2. The 0-1 ILP model

We present a 0-1 ILP model for the PMSPIC and discuss its relation to other models in the literature. We define the set $\mathcal{I} := \{(s,t) | 0 \leqslant s < t \leqslant T+1;\ s,\ t \in \mathbb{Z}\}$ of replacement intervals and introduce the variables

$$x^i_{st} = \begin{cases} 1, & \text{if component } i \text{ receives PM at the times } s \text{ and } t, \\ & \text{and not in-between}, \\ 0, & \text{otherwise}, \end{cases} \qquad i \in \mathcal{N},\ (s,t) \in \mathcal{I},$$

and

$$z_t = \begin{cases} 1, & \text{if maintenance occurs at time } t, \\ 0, & \text{otherwise}, \end{cases} \qquad t \in \mathcal{T}.$$

The PMSPIC is now modeled as the problem to

$$\text{minimize} \quad \sum_{t\in\mathcal{T}} d_t z_t + \sum_{i\in\mathcal{N}} \sum_{(s,t)\in\mathcal{I}} c^i_{st} x^i_{st}, \tag{1a}$$

$$\text{subject to} \quad \sum_{s=0}^{t-1} x^i_{st} \leqslant z_t, \qquad i \in \mathcal{N}, \quad t \in \mathcal{T}, \tag{1b}$$

$$\sum_{s=0}^{t-1} x^i_{st} = \sum_{r=t+1}^{T+1} x^i_{tr}, \qquad i \in \mathcal{N}, \quad t \in \mathcal{T}, \tag{1c}$$

$$\sum_{t=1}^{T+1} x^i_{0t} = 1, \qquad i \in \mathcal{N}, \tag{1d}$$

$$x^i_{st} \in \{0,1\}, \qquad i \in \mathcal{N}, \quad (s,t) \in \mathcal{I}, \tag{1e}$$

$$z_t \in \{0,1\}, \qquad t \in \mathcal{T}. \tag{1f}$$

The model (1) is based on a model for the uncapacitated fixed-charge multi-commodity network design problem (e.g. Balakrishnan, Magnanti, & Wong, 1989), each component i corresponding to a commodity and the set $\mathcal{T} \cup \{0, T+1\}$ corresponding to the set of nodes in a directed graph with the arcs $(s,t) \in \mathcal{I}$ directed "forward in time". The fixed charges are assigned to the nodes. For each component i, one unit of flow is to be sent from node 0 to node $T+1$. A maintenance occasion at time t corresponds to *opening* all the arcs in the set $\{(s,t) | s \in \{0,\dots,t-1\}\}$. The objective (1a) is to minimize the sum of all set-up and interval costs. The constraints (1b) ensure that if a maintenance interval for component i ends at time t, then maintenance occurs at time t. For each component i, the constraints (1c) ensure that the same number of maintenance intervals end and start at time t. The constraints (1d) ensure that a maintenance interval of component i starts at time $t = 0$.

Already Fulkerson (1966) presented a flow model for optimal replacement decisions in a single-component system. For $|\mathcal{N}| = 1$ the model (1) reduces to that presented by Fulkerson (1966).

The model (1) was presented in Joneja (1990) for the DJRP (with ordering costs equivalent to c^i_{st} in Definition 1). Only its continuous relaxation was solved—to obtain a lower bound for the evaluation of the heuristics which constitute the main focus of the article. The model is more closely studied in Joneja (1987, unpublished), but to the best of our knowledge, it has not been further studied in the literature, although other formulations of the DJRP and similar problems are present (see Boctor et al., 2004, for an overview). In particular, a model introduced by Boctor et al. (2004) is strengthened (Narayanan & Robinson, 2006) to a model equivalent to (1), which, however, serves only as a basis for heuristics and lower bounding. To the best of our knowledge, the model (1) has neither been studied from a polyhedral point of view, nor been applied directly as a decision support tool, nor at all within the maintenance community.

2.3. Properties of the 0-1 ILP model

We next investigate the mathematical properties of the model (1). We first show that the integrality restrictions on the variables x^i_{st} can be relaxed, then that all linear inequalities define facets of the convex hull of the set of feasible solutions. Let

$$P^{\text{Conv}} := \text{conv}\Big\{(x,z) \in \{0,1\}^{n(T+1)(T+2)/2+T} \Big| \text{(1b), (1c), and (1d) hold}\Big\}$$

denote the convex hull of the set of feasible solutions to the model (1). Further, for any fixed values of $z_t \in \{0,1\},\ t \in \mathcal{T}$, let

$$P^z := \Big\{x \in [0,1]^{n(T+1)(T+2)/2} \Big| \text{(1b), (1c), and (1d) hold}\Big\}$$

denote the feasible set of the continuous relaxation of the corresponding projection of the feasible set of the model (1) onto the dimension of the variables x.

Proposition 1. *If the variables $z_t,\ t \in \mathcal{T}$, are fixed to binary values then the polyhedron $P^z \subset \mathbb{R}^{n(T+1)(T+2)/2}$ possesses integral extreme points.*

Proof. When the values of the variables z are fixed, the model (1) decomposes over the components $i \in \mathcal{N}$. For each component $i \in \mathcal{N}$ a network flow problem is obtained whose feasible set has integral extreme points (e.g., Nemhauser & Wolsey, 1988, Sect. III.1, Cor. 2.9). □

Proposition 1 implies that the integrality restrictions on the variables x^i_{st} can be relaxed.

Proposition 2. *Each of the constraints (1b) induces a facet of the convex hull of feasible solutions P^{Conv}.*

Proof. We start by investigating the dimension of the convex hull polytope P^{Conv}. The model (1) contains $n(T+1)(T+2)/2 + T$ variables and $n(T+1)$ linearly independent equality constraints. It thus holds that $\dim(P^{\text{Conv}}) \leqslant D := n(T+1)T/2 + T$. Define the *face* $F_{it} \subset P^{\text{Conv}}$ as the set of points that, for fixed values of i and t, satisfy the constraint (1b) with equality, implying that $\dim(F_{it}) \leqslant \dim(P^{\text{Conv}}) - 1$. A face of a polytope is a *facet* if its dimension is one unit less than that of the polytope itself (Nemhauser & Wolsey, 1988, Sect. I.4.3). We will present $D+1$ schedules whose corresponding points, (x,z), are affinely independent and belong to P^{Conv}. Out of these schedules, D will be shown to belong to F_{it}, thus implying that $\dim(F_{it}) \geqslant D-1$ and hence that F_{it} is a facet of P^{Conv}. We assert the affine independence by, for each new schedule, setting a variable value to 1 which was set to 0 in all previous schedules.

First, the schedule of no PM and no maintenance occasions (except for the PM of all the components at the times 0 and $T+1$) is constructed; it belongs to the set F_{it}. For each $u \in \mathcal{T}$, a schedule is constructed with no PM but with a single maintenance occasion at the time u (i.e. $z_u = 1,\ z_v = 0,\ v \in \mathcal{T} \setminus \{u\}$). This yields T additional schedules which all belong to the set P^{Conv} and which all, except for $u = t$, also belong to set F_{it}. The set of all thus constructed schedules is affinely independent.

We next introduce PM of a single component. For $v \in \{0,\dots,T-1\}$ and $w \in \{v+1,\dots,T\}$, a (v,w)-schedule for component i is defined by its PM at the times v and w only; these are also the only two maintenance occasions (except that PM is performed for all the components at the times 0 and $T+1$). Each of these $T(T+1)/2$ schedules extends the set of affinely independent schedules and they all belong to F_{it}. This procedure is applied for all components $j \in \mathcal{N} \setminus \{i\}$, with the addition of a maintenance occasion and a replacement of component i at time t in order to assure that the schedules belong to the set F_{it}.

Summarizing, we obtain $1 + T + nT(T+1)/2 = D+1$ affinely independent schedules, among which D belong to the face F_{it}. Hence, F_{it} is a facet of P^{Conv}, which proves the proposition. □

2.4. Time-independence and periodicity

The applications of the model (1) presented in Sections 3–5 consider PMSPIC instances with time-independent costs, for which we here introduce a representation. By additional appropriate constraints the model (1) can enforce solutions conforming to the class of standard indirect grouping solutions, which may, however, be non-optimal for the PMSPIC.

The time-independent cost structure is defined as follows. To each maintenance occasion is assigned the set-up cost d; hence, $d_t = d$ for all $t \in \mathcal{T}$. Further, each scheduled PM of component $i \in \mathcal{N}$ is assigned the *component cost* c^{PM}_i. For component i, the *deterioration cost* function $M_i : \mathcal{T} \to \mathbb{R}$ is such that each PM interval of length $u \in \mathcal{T}$ generates the deterioration cost $M_i(u)$ in addition to the component cost. This means that

$$c^i_{st} := c^{PM}_i + M_i(t-s), \qquad s \in \mathcal{T}, \quad t \in \{s+1,\dots,T+1\}, \tag{2a}$$

and

$$c^i_{0t} := M_i(t), \quad t \in \{1,\dots,T+1\}. \tag{2b}$$

Given a schedule represented by the solution (x,z) to the model (1), the PM cost is defined as

$$C^{\text{PM}}_{(x,z)} := \sum_{t\in\mathcal{T}} dz_t + \sum_{i\in\mathcal{N}}\sum_{s\in\mathcal{T}}\sum_{t=s+1}^{T+1} c^{\text{PM}}_i x^i_{st}. \tag{3a}$$

Correspondingly, since $x^i_{st} = 1$ if and only if (s,t) is an active maintenance interval, the total deterioration cost corresponding to the solution (x,z) is defined as

$$C^D_{(x,z)} := \sum_{i\in\mathcal{N}} \sum_{(s,t)\in\mathcal{I}} M_i(t-s)x^i_{st}. \tag{3b}$$

The cost structure of the time-independent PMSPIC equals that of the independent standard grouping strategy reviewed by Dekker et al. (1997). However, we do not enforce periodic replacement, and we consider a discretized and finite planning horizon. As discussed in Dekker (1995), different choices of the deterioration cost functions M_i lead to a variety of models.

We now extend the model (1) by a set of constraints such that only solutions fulfilling the periodicity requirement of standard indirect grouping (see Section 1) are feasible. Periodicity of components' replacement is imposed by the constraints

$$x^i_{st} = x^i_{2s-t,s}, \qquad i \in \mathcal{N}, \quad (s,t) \in \{(u,v) \in \mathcal{I} | v \leqslant 2u\}, \tag{4a}$$

which imply that if a replacement interval is scheduled between the times s and t, then a replacement interval must be scheduled also between the times $s-(t-s)$ and s.

We next establish that periodicity of the maintenance occasions is imposed by the constraints

$$z_s + z_t \leqslant 1 + z_{2s-t}, \quad (s,t) \in \mathcal{I} : T \geqslant t < 2s > 0, \tag{4b}$$

$$z_s \leqslant z_{ks}, \qquad k \in \{1,\ldots,\lfloor T/s \rfloor\}, \quad s \in \mathcal{T}. \tag{4c}$$

Assume that for some $p \geqslant 2$, $z_p := 1$ and $z_r := 0$, $r \in \{1,\ldots,p-1\}$, hold. From (4c) then follows that $z_{kp} = 1$, $k \in \{1,\ldots,\lfloor T/p \rfloor\}$. In (4b), let $s := kp$ and $t := kp + r$, $k \in \{1,\ldots,\lfloor T/p \rfloor\}$, $r \in \{1,\ldots,p-1\}$ (i.e., s is a multiple of p while t is *not*); the appearance of (4b) then is the inequality $z_{kp+r} \leqslant z_{kp-r}$. By induction with respect to k it then follows that $z_t = 0$, $t \in \mathcal{T} \setminus \{kp \mid k \in \{1,\ldots,\lfloor T/p \rfloor\}\}$, i.e., any set of maintenance occasions satisfying (4b) and (4c) is periodic.

Note that the requirements (4b) and (4c) make the PMSPIC polynomially solvable by evaluating all T possible lengths of the intervals between maintenance occasions; by Proposition 1 each of these evaluations can be performed in polynomial time.

3. Application to rail grinding

Rolling contact fatigue (RCF) defects, such as head checks, squats, corrugations, and wear, may cause severe problems in a railway system. If the natural wear rates of the rail are too low—such that initiated cracks and defects are not worn away quickly enough—artificial procedures for defect removal must be adopted (Grassie, 2005). The most common measure to prevent or remove RCF defects is *rail grinding*: a grinding machine removes a layer of the rail in order to reduce cracks and corrugations and to maintain a correct track geometry. For a thorough description of RCF cracks, see Lewis and Olofsson (2009). We consider planning the rail grinding on a set of track sections.

In the model (1), $\mathcal{N}$ denotes a set of track sections, each characterized by its curve radius. The PM actions considered are *rail grinding* and *replacement* of the track section. Since the rail degrades with time, the cost for rail grinding increases with the length of the interval between two consecutive maintenance occasions. In this application, time represents tonnage, which is measured in mega gross tonnes (MGT; 1MGT means 10^6 tonnes being transported on the rail).

3.1. Rail degradation and maintenance costs

We consider a simple model for RFC crack depth growth considering only head check cracks. The model is based on the two assumptions that the crack depth growth rate (a) depends on the radius of the track section and (b) increases or is constant with time (tonnage); see (INNOTRACK, 2009, Eq. (8)), which explains how the growth rate depends on the radius of the track section for the rail *grade R220*. Since the increase of the crack growth rate with time is uncertain, we consider two degradation models: the crack depth growth rate (i) is constant and (ii) increases by 1% per MGT.

Cracks are removed by performing a number of grindings of the track section. The number of grindings required to remove a crack depends on its depth. We assume that one pass with the grinding machine can remove cracks up to h mm depth, and that the cost of one grinding pass is c_g. A multiple (depending on the time [tonnage] between two consecutive occasions) of c_g is thus paid for each grinding occasion. If the cracks grow too deep—which in our modeling is represented by the time between two consecutive maintenance occasions on track section i exceeding its limit T_i—then the section must be replaced at a cost $c_r \gg c_g$. This cost structure is modeled by the following deterioration cost function (see Section 2.4 and Fig. 1): The crack depth of track section i at t time units after the previous PM action is a function of its curve radius and is denoted $L_i(t)$. The deterioration cost is then defined as $M_i(t) := c_g \left\lceil \frac{L_i(t)}{h} \right\rceil$ if $t \leqslant T_i - 1$ and $M_i(t) := c_r$ if $t = T_i$. The length of the maintenance interval may not exceed T_i, which is incorporated in the model (1) by letting $M_i(t) \gg c_r$ if $t \geqslant T_i + 1$. We let $c_i^{\mathrm{PM}} := 0$, implying that the interval costs in the model (1) are defined as $c^i_{st} = M_i(t-s)$. Finally, each time *any* maintenance is performed the set-up cost d, for utilizing a grinding machine, is paid.

3.2. Numerical results

We consider a set of five track sections (i.e., $n = 5$) with radii 500 m, 700 m, 900 m, 1100 m, and 1300 m, respectively. Whenever several track sections have identical curve radii these will render the same grinding strategy in an optimal solution to the program (1). In a more detailed model, including the geographical locations of the track sections, different strategies can, however, be optimal for different sections having equal radii.

We employ the costs for rail grinding and rail replacement presented in Chattopadhyay, Reddy, and Larsson-Kråik (2005). Since the main objective of our case study is to analyze the grinding strategies for a system of track sections, these costs are normalized with respect to c_g, such that $c_g = 1$ and $c_r = 76$. We assume that each grinding pass removes cracks up to $h = 2.5$ mm, and that the rail needs to be replaced when the crack depth exceeds 10 mm (Banverket, 2004). The time limits are then computed as $T_i = \mathrm{argmin}_t\{t | L_i(t) \geqslant 10\mathrm{mm}\}$, $i \in \mathcal{N}$. Since the set-up cost d is not considered by Chattopadhyay et al., we analyze the rail grinding and replacement model for three magnitudes, $d \in \{0,1,10\}$, of these costs. Time is discretized into intervals corresponding to 2 MGT and the planning horizon comprises $T = 100$ intervals. Fig. 2 illustrates optimal maintenance schedules for the constant and increasing degradation models, respectively. As the maintenance occasion cost d increases, the coordination of the maintenance occasions becomes increasingly important. For the increasing degradation model, an optimal solution consists of more frequent maintenance occasions with fewer grinding passes at each occasion, compared to the constant degradation model. This is due to the fact that, for an increasing crack growth rate, the lengths of the time intervals, in which the deterioration cost function is constant, decrease.

We compared the maintenance schedules obtained by solving the model (1) with three other maintenance scheduling principles: (i) Solve the model (1) extended by the constraints (4a), implying that the maintenance operations on each track section are periodical. (ii) Solve the model (1) extended by the constraints (4b) and

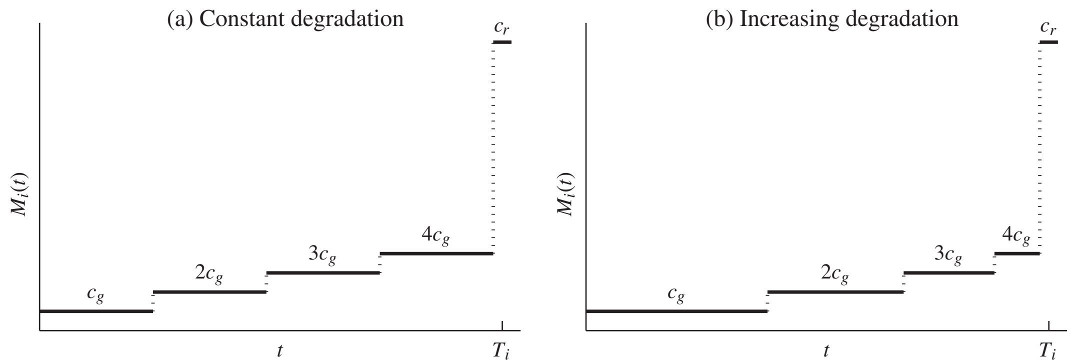

Fig. 1. Deterioration cost functions for PM on a rail section with (a) constant and (b) increasing crack growth rates; an increasing rate corresponds to decreasing lengths of the intervals of constant values of the cost function. At time T_i since the last PM action on a rail section, it must be replaced at a large cost.

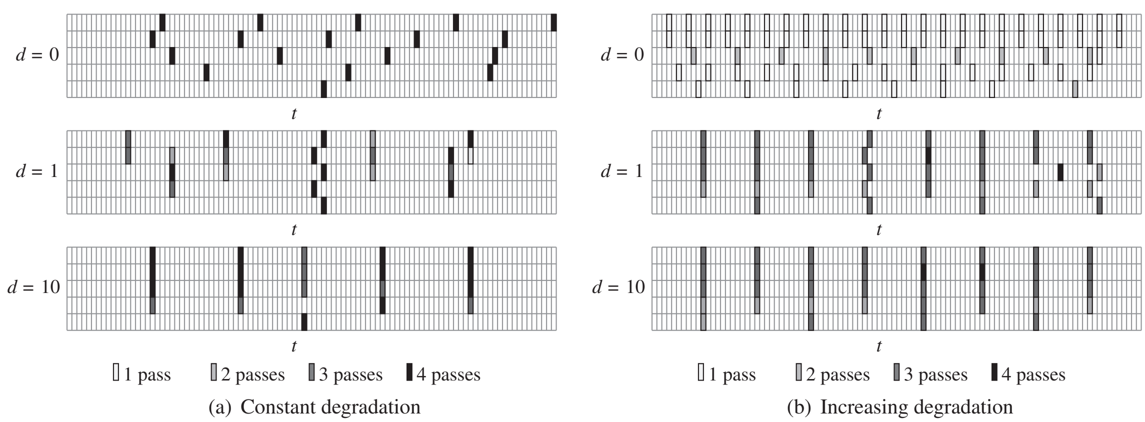

Fig. 2. Rail grinding schedules for the constant and increasing degradation models, and three magnitudes for the values d of the maintenance occasion cost. In each of the schedules, the five rows represent the five track sections.

(4c), implying that the maintenance occasions for the system are periodical. (iii) Apply the age policy with soft and hard lives (see Section 1). The hard lives are set to $a_i^{\text{hard}} := T_i - 1$ and the soft lives to $a_i^{\text{soft}} := T_i - \eta$. The age policy is applied for each of the values of $\eta \in \{1, \ldots, \min_{i \in \mathcal{N}}\{T_i\}\}$; for each problem instance the value yielding the lowest objective value is chosen.

In Table 1, the costs, $C^{\text{PM}}_{(x,y)}$ (3a), resulting from the four maintenance principles are presented for the two cases of *constant* and *increasing degradation*, respectively. When $d = 0$, the inclusion of the constraints (4a) or (4b)–(4c) does not affect the maintenance cost as compared to solving the model (1), since then the problem can be solved for each individual track section, which in turn means that the requirement of a cyclic schedule poses no further restrictions. For $d \in \{1, 10\}$ the cyclic schedules generate higher costs than the model (1) schedule. For both constant and increasing degradation, the cost of the schedule resulting from the model (1), (4a) is never higher than that resulting from the model (1), (4b)–(4c). This is due to the fact that the constraints (4a) follow as a consequence of the constraints (4b) and (4c), (1b). The age policy performs well for the constant degradation case, but results in more expensive schedules for the increasing degradation case. For the most successful instance, the model (1) reduces the costs by 10% as compared to the age policy.

4. Application to maintenance of a low pressure turbine

The low pressure turbine (LPT) is a module of the aircraft engine RM12, which is manufactured and serviced at GKN Aerospace (formerly Volvo Aero Corporation) in Trollhättan, Sweden. The engine RM12 consists of several modules, each comprising several components. When a component is to be replaced, the engine is removed from the aircraft and the corresponding module is disassembled. When an engine is removed for maintenance a replacement engine

Table 1
The total maintenance cost $\left(C^{\text{PM}}_{(x,z)}\right)$ and the number of maintenance occasions (#) resulting from four scheduling principles for rail grinding.

Degradation rate	d	Model (1)		Model (1), (4a)		Model (1), (4b)–(4c)		Age policy	
		$C^{\text{PM}}_{(x,z)}$	#	$C^{\text{PM}}_{(x,z)}$	#	$C^{\text{PM}}_{(x,z)}$	#	$C^{\text{PM}}_{(x,z)}$	#
Constant	0	80	37	80	37	80	26	80	18
Constant	1	89	8	91	9	93	5	92	8
Constant	10	137	5	138	5	138	5	138	5
Increasing	0	98	42	98	42	98	42	98	27
Increasing	1	119	13	119	11	120	9	132	10
Increasing	10	192	8	192	8	194	7	196	7

is hired. We consider the maintenance scheduling of the LPT module only; the data used originate from a case study in Almgren et al. (2012a).

The LPT contains ten components, four classified as *safety critical* and six as *on condition*. A failure of a safety critical component may have a catastrophic outcome; each such component is thus assigned a technical life defining the latest allowed time for its replacement, i.e., a *maximum allowed replacement interval*. An on condition component is replaced after a condition measurement indicates that a threshold value has been reached. We regard the event of a condition measurement enforcing a replacement as a *component failure*, and assume its probability to be Weibull distributed over time. Time units are measured in flight hours (fh). For each component $i \in \mathcal{N}$, the replacement cost c_i^{PM} is composed by its purchase cost and the work cost for its replacement. The set-up cost d consists of the costs for removing the module and hiring a replacement engine. Due to confidentiality, all costs presented are normalized by the set-up cost d. Time is discretized into intervals of 50 fh and the planning horizon comprises $T = 100$ intervals (i.e., 5000 fh).

An optimal schedule for replacement of the safety critical components only could be obtained by solving an instance of Almgren et al. (2012a, model (1)). One may refrain from planning future replacements of the remaining on condition components, but an absent maintenance plan can result in many unplanned production stops, due to maintenance as well as at inconvenient times (as, e.g., when the aircraft is not located at the home base). Any unplanned maintenance stop generates at least the set-up cost d, but the cost may be much higher, however, hard to estimate. Therefore, we first consider the bi-objective optimization problem of minimizing the cost for PM and the probability of at least one unexpected maintenance stop within the planning horizon. Then, we assign a cost to each unexpected maintenance occasion. For both problems, we investigate the cases with and without periodicity restrictions.

Maintenance scheduling of the LPT has previously been studied by Almgren et al. (2012a) and Patriksson et al. (2014). Almgren et al. (2012a) assign maximum replacement intervals (representing the expected lives) to on condition components while unexpected stops are ignored. Almgren et al. (2012b) employ the same methodology for the entire aircraft engine. Patriksson et al. (2014) model maintenance decisions for the LPT assuming equal costs for scheduled and unexpected stops, thus restricting PM to occasions when the replacement of at least one component is triggered by condition measurements. None of these three approaches considers additional costs for unexpected stops.

4.1. The bi-objective approach

Consider the bi-objective minimization of the PM cost, defined by (3a), and the probability of unexpected maintenance stops, defined by (5) below, subject to the constraints (1b)–(1e) and (1f). A schedule is *Pareto optimal* if no other feasible schedule exists such that both the probability of an unexpected stop and the PM cost is lower (see Ehrgott (2005) for the concept of Pareto optimality in multiple objective optimization). We aim at finding Pareto optimal schedules whose objective values are scattered over the set of Pareto optimal solutions, i.e., the *Pareto front*.

4.1.1. The choice of deterioration costs

Assume that each component $i \in \mathcal{N}$ possesses a cumulative failure distribution function $F_i : \mathbb{R} \to \mathbb{R}$. A solution (x, z) fulfilling the constraints (1b)–(1e) and (1f) corresponds to one replacement schedule for each component i. Let q_i denote the number of scheduled replacements of component i and S_i^r the number of time steps between replacements $r-1$ and r, $r \in \{1, \ldots, q_i+1\}$. We assume independence of failures between different components and between different replacement intervals, and presume that the replacement schedule defined by (x, z) is employed. The probability of an unexpected maintenance stop during the time interval $[0, T+1]$ is then calculated as

$$p_{(x,z)}^{\mathrm{stop}} := 1 - \prod_{i \in \mathcal{N}} \prod_{r=1}^{q_i+1} \left(1 - F_i(S_i^r)\right), \tag{5}$$

where $1 - F_i(S_i^r)$ is the probability of no failures occurring for individual r of component i. For $u \in \mathcal{T}$, the deterioration cost function is defined as

$$M_i(u) := -\frac{1-\gamma}{\gamma} \log(1 - F_i(u)), \quad \gamma \in (0, 1], \tag{6}$$

where γ is a weight parameter. That is, $M_i(S_i^r)$ is proportional to the log-probability of failure of component i within its rth replacement interval. The next proposition shows that solving the model (1) yields a Pareto optimal schedule.

Proposition 3. *Let (x, z) be an optimal solution to (1) with the deterioration function defined by (6). The schedule corresponding to (x, z) is Pareto optimal for the minimization of (3a) and (5) subject to (1b)–(1e) and (1f), and the value of (1a) at (x, z) equals*

$$\frac{1}{\gamma}\left(\gamma C_{(x,z)}^{\mathrm{PM}} - (1-\gamma)\log\left(1 - p_{(x,z)}^{\mathrm{stop}}\right)\right). \tag{7}$$

Proof. By the equalities (5) and (6) it holds that

$$-\frac{1-\gamma}{\gamma}\log\left(1 - p_{(x,z)}^{\mathrm{stop}}\right) = \sum_{i \in \mathcal{N}} \sum_{r=1}^{q_i+1} M_i(S_i^r).$$

Using the definitions (2) and (3a) of deterioration and PM costs, for the case when $d_t = d$, $t \in \mathcal{T}$, the objective (1a) can be expressed as (7). Since the function $g : \mathbb{R} \to \mathbb{R}$, defined by $g(u) := -\log(1-u)$, is increasing, the schedule given by (x, z) is Pareto optimal with respect to the objectives (3a) and (5). □

The weight parameter γ controls the balance between the objectives. Letting $\gamma \approx 0$ means minimizing the probability of unexpected failures,[1] while $\gamma = 1$ means minimizing the cost of PM.[2] Since the feasible set of an ILP is in general non-convex, our weighted average approach does not guarantee that all Pareto optimal schedules can be found. In our tests the Pareto optimal solutions found are, however, fairly well spread.

Since the component failures are Weibull distributed (e.g. Murthy, Xie, & Jiang, 2004) with shape parameter β_i and scale parameter α_i, the deterioration function can be simplified as

$$M_i(u) = -\frac{1-\gamma}{\gamma}\log(1 - F_i(u)) = \frac{1-\gamma}{\gamma}\left(\frac{u}{\alpha_i}\right)^{\beta_i}.$$

4.1.2. Results

The weight parameter γ is heuristically varied such that each new point on the Pareto front is generated between the two of the already generated points having the largest gap. We study schedules obtained by solving the two models (1), (4) and (1) (i.e., with, and without, periodicity requirements on both maintenance occasions and replacements). Since the model (1), (4) is a restriction of the model (1), all (Pareto optimal) solutions to (1), (4) are feasible solutions to (1). Hence, any Pareto optimal solution to (1), (4) is either Pareto optimal to (1) or dominated by at least one solution to (1). This property is indicated in Fig. 3, which

[1] For $\gamma \approx 0$ an optimal schedule is to replace all components at all times.
[2] For $\gamma = 1$ an optimal schedule would contain no replacements at all.

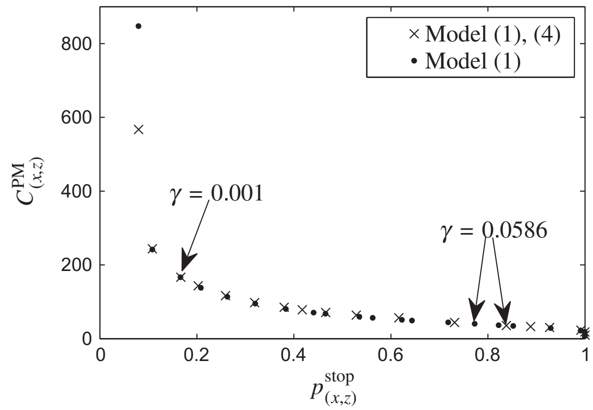

(a) Points on the Pareto front for different values of γ.

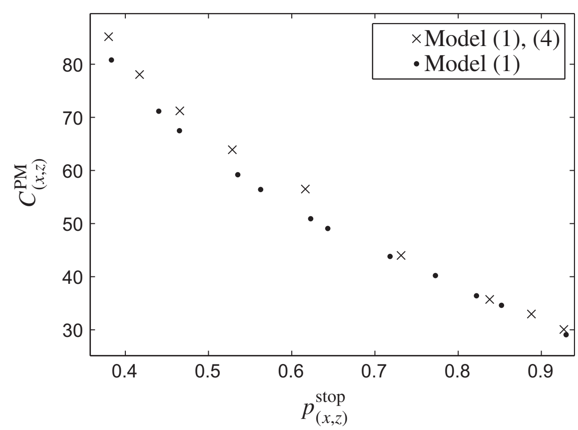

(b) A close-up of the points on the Pareto front.

Fig. 3. Objective values for Pareto optimal schedules with and without the periodicity requirements (4) on the maintenance occasions and the component replacements, for the bi-objective LPT case.

illustrates the results of our computations. The Pareto optimal periodic schedules cost up to 5% more than the corresponding non-periodic schedules with approximately the same probability of an unexpected maintenance stop; see Fig. 3(b). Table 2 provides the objective values and the CPU times required for solving the models for two values of the weight parameter γ (the corresponding points are indicated in Fig. 3(a)). We observe, as expected, that adding the periodicity restriction (4) typically reduces the solution time. The results for $\gamma = 0.0586$ also demonstrate that verifying the optimality of a solution to the model (1) can be quite time consuming; the optimal solution to this instance was found already after 60 s.

4.2. The expected maintenance cost

A cost is assigned to each unexpected maintenance stop and we minimize the sum of the PM costs (3a) and the cost of unexpected maintenance stops over the planning period, i.e., the CM cost, defined in (9) below.

4.2.1. Choosing the deterioration cost

Assume that the components' failure probabilities are independent. If component i fails, then a maintenance stop, at the cost d^{CM}, and a replacement, at the cost c_i^{PM}, must be performed. Consider the renewal function $m_i : \mathbb{R}_+ \to \mathbb{R}$, such that $m_i(u)$ is the expected number of failures of a newly replaced component i during the time period $[0, u]$, provided that it is replaced at failure. The deterioration cost function is given by

$$M_i(u) := \left(c_i^{\mathrm{PM}} + d^{\mathrm{CM}}\right) m_i(u), \quad u \in \mathcal{T}. \tag{8}$$

The calculation of the renewal function for the Weibull distribution is non-trivial (Smith & Leadbetter, 1963). For inclusion into the model (1), the value of the Weibull renewal function is required only at a discrete set of time points. Hence, we simulate a large number of Weibull processes, then used to obtain a numerical approximation of the Weibull renewal function.

Assume that (x, z) is a feasible solution to the model (1), that components are preventively replaced according to the corresponding schedule, and that each unexpectedly failed component is instantly replaced, with no further modification of the schedule. The CM cost is then defined as

$$C_{(x,z)}^{\mathrm{CM}} = \sum_{i \in \mathcal{N}} \sum_{r=1}^{q_i+1} \left(c_i^{\mathrm{PM}} + d^{\mathrm{CM}}\right) m_i(S_i^r), \tag{9}$$

where q_i and S_i^r are defined as in Section 4.1.1. By the relations (8), (9), (3a), and (2) the objective function in (1a) corresponds to $C_{(x,z)}^{\mathrm{PM}} + C_{(x,z)}^{\mathrm{CM}}$. If rescheduling at an unexpected failure of a component is not allowed, then the model (1) yields an optimal replacement schedule. It is, however, natural to consider a rescheduling at a component failure. By re-optimizing the model at a given failure event and rescheduling accordingly, the model (1) can be used as a policy which may yield a lower total maintenance cost than that resulting from the schedule initially computed. Allowing for rescheduling (by the model (1)) at a component failure—which is thus a PM opportunity—means, however, that the deterioration cost function no longer captures the full expected CM cost, and, hence, that the maintenance decisions are not provably optimal to the problem of minimizing the sum of costs for PM and expected CM.

4.2.2. Simulation and results

In order to investigate the effect of rescheduling, we simulate the system according to the following. The term *policy* here means either a simple policy, or a fixed replacement schedule, or a rescheduling using the model (1). A scenario is obtained by—for each component—sampling a sequence of component lives from its life distribution. The policy provides the time of the first scheduled maintenance stop. If a failure of any component occurs before that stop, the unexpected maintenance stop cost d^{CM} arises and the failed component is replaced. If no failure occurs, the maintenance occasion cost d arises and a scheduled maintenance stop is performed. At both scheduled and unscheduled maintenance stops the policy provides decisions for component replacements and for the next scheduled maintenance stop. The replacement of component i generates the cost c_i^{PM}. After a component is replaced, it obtains the life of the next component individual in the scenario

Table 2
Objective values and CPU-times obtained for two values of the weight parameter γ on the bi-objective LPT case.

γ	Model (1)			Model (1), (4)		
	$p_{(x,z)}^{stop}$	$C_{(x,z)}^{\mathrm{PM}}$	CPU-time (s)	$p_{(x,z)}^{stop}$	$C_{(x,z)}^{\mathrm{PM}}$	CPU-time (s)
0.001	0.166	166.7	3.31	0.167	166.3	2.15
0.0586	0.773	40.2	225.22	0.839	35.7	2.19

sequence. All the policies are compared over the same set of scenarios.

We consider three simple policies defined in Section 1. The first is the constant-interval (CI) policy, defined by replacing all components after t_B time units. The value of the parameter t_B is determined by comparing the mean cost from the policy over the set of scenarios for $t_B \in \mathcal{T}$. The second is the age policy based on soft (a_i^{soft}) and hard (a_i^{hard}) lives for all components. For the safety critical components, $a_i^{\text{hard}} = a_i^{\text{soft}} := T_i$, where T_i is the technical life of the component. For the on condition components, $a_i^{\text{hard}} := T_i + \eta^{\text{hard}}$ and $a_i^{\text{soft}} := T_i + \eta^{\text{soft}}$, where the values of η^{soft} and η^{hard} are selected according to the following. First, $\eta^{\text{hard}} := T$ and the value of $\eta^{\text{soft}} \in \{-\min_{i \in \mathcal{N}}\{T_i\}, \dots, T\}$ yielding the lowest mean cost over the simulated scenarios is selected. Then, the value of $\eta^{\text{hard}} \in \{\eta^{\text{soft}}, \dots, T\}$ yielding the lowest mean cost over the simulated scenarios is selected. The third simple policy considered is the target built life (TBL) policy. The TBL value is set by performing an exhaustive search, and the hard lives are set as for the age policy. Note that the CI policy does not consider rescheduling while the others do. Further, policies allowing for rescheduling do not correspond to fixed schedules, since the decisions on PM actions depend on the future expected CM.

We compare the use of simpler policies with the replacement schedules obtained by solving the model (1), and consider problem settings with, and without, the restrictions (4). First, the schedules are implemented without rescheduling: at failure, only the failed component is replaced. Then, rescheduling is implemented by solving the model (1) at every maintenance stop: scheduling for the rest of the planning period and employing the life distributions conditioned on the respective components' ages for their first interval costs. Only the replacement decisions at the current maintenance stop are implemented.

The cost of unexpected maintenance stops is hard to estimate; we compare the performance of the suggested policies for $d^{\text{CM}} \in \{d, 2d, 4d\}$. Table 3 shows the average maintenance costs and the corresponding standard deviations, over the same set of simulation scenarios. The model (1) with rescheduling provides the most cost efficient policy. Out of the simple policies, the age policy performs the best with an average cost that is 4–15% higher than that of the model (1) with rescheduling. The best fixed schedule is computed by the model (1). The periodic schedules are slightly more expensive than the corresponding non-periodic ones. For the case $d^{\text{CM}} = 4d$, the fixed schedule from model (1) is almost as good as its rescheduling analogue; this is due to the fact that unexpected maintenance stops are fairly rare. Fig. 4 shows the periodic and non-periodic fixed schedules resulting from the model (1). The CPU-time required for solving the model (1) was always less than 10 s.

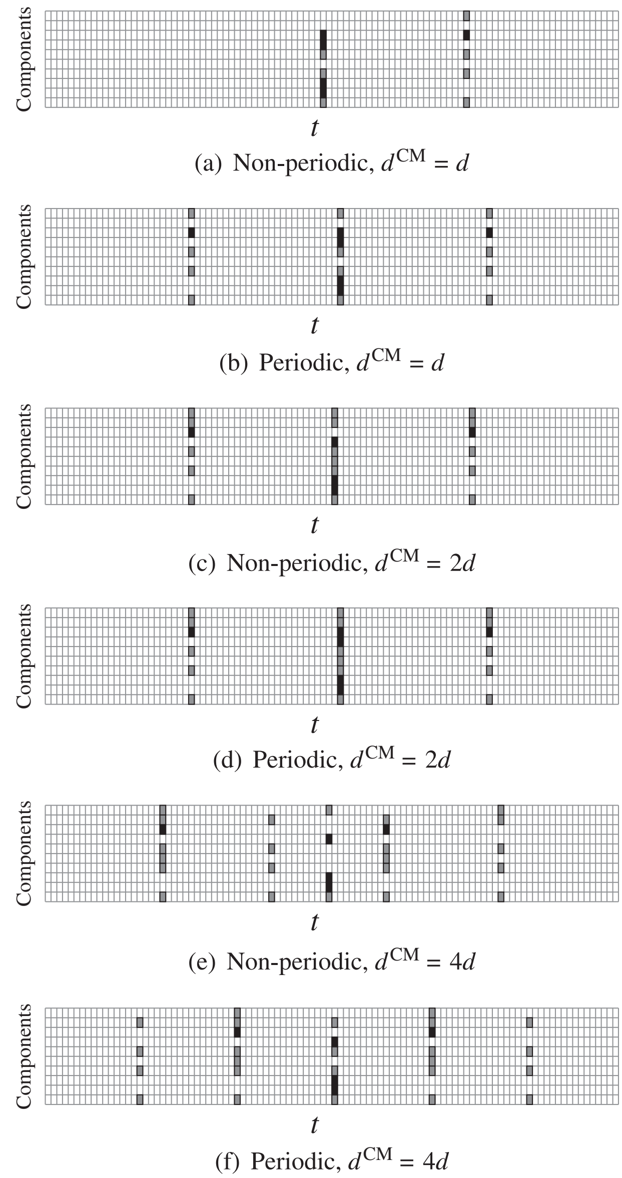

Fig. 4. Fixed schedules obtained as solutions to model (1) with or without periodicity restrictions (4) for the LPT case. Black and gray boxes correspond to replacements of safety critical and on condition components, respectively.

5. Application to maintenance of a farm of wind turbines

We consider the scheduling of PM in a farm comprising $M = 5$ identical wind turbines. This problem was studied by Tian et al. (2011) in the context of condition based maintenance (CBM). In this setting, maintenance costs are time-independent and composed by component replacement costs, and set-up costs for individual wind turbine maintenance occasions and for visits to the wind turbine farm. The probability of a component failure is assumed to increase with its age according to a Weibull distribution (Tian et al., 2011). Let $\mathcal{M}$ denote the set of identical wind turbines, each with a set $\mathcal{N}$ of components to be maintained during a planning period, defined by the set $\mathcal{T}$, and let the set $\mathcal{I}$ be defined as in Section 2.2. The cost for visiting the wind farm is denoted f and the

Table 3
Comparison of maintenance policies for the LPT. The mean total maintenance cost (standard deviation) for each policy applied to 100 simulation scenarios is listed. The cost and standard deviation is presented in units of the maintenance occasion cost d.

d^{CM}	Allowing for rescheduling			Fixed schedule		
	Model (1)	Age	TBL	Model (1)	Model (1), (4)	CI
d	27.4 (3.2)	28.5 (2.8)	29.2 (3.2)	29.1 (3.2)	30.9 (3.3)	31.2 (3.3)
$2d$	31.8 (4.1)	34.9 (4.1)	36.7 (4.4)	34.8 (5.8)	35.7 (5.8)	39.5 (5.5)
$4d$	39.0 (7.3)	44.7 (8.8)	50.4 (8.4)	40.5 (7.8)	41.4 (6.8)	48.3 (8.7)

set-up cost for individual wind turbine maintenance is denoted d. For component $i \in \mathcal{N}$ the cost of a preventive (corrective) replacement is denoted c_i^{PM} ($c_i^{\mathrm{CM}} > c_i^{\mathrm{PM}}$). For component i and replacement interval $(s,t) \in \mathcal{I}$, the interval cost c_{st}^i is given by the deterioration cost function M_i, defined in (2). We introduce the variables

$$x_{st}^{mi} = \begin{cases} 1, & \text{if component } i \text{ in turbine } m \text{ is replaced} \quad m \in \mathcal{M}, \\ & \text{at times } s \text{ and } t, \text{and not in} - \text{between}, \quad i \in \mathcal{N}, \\ 0, & \text{otherwise}, \quad (s,t) \in \mathcal{I}, \end{cases}$$

$$z_t^m = \begin{cases} 1, & \text{if turbine } m \text{ is maintained at time } t, \quad m \in \mathcal{M}, \\ 0, & \text{otherwise}, \quad t \in \mathcal{T}, \end{cases}$$

and

$$y_t = \begin{cases} 1, & \text{if the farm is visited at time } t, \\ 0, & \text{otherwise}, \end{cases} \quad t \in \mathcal{T}.$$

The model (1) is then extended to the problem to

$$\text{minimize} \quad \sum_{t\in\mathcal{T}} f y_t + \sum_{t\in\mathcal{T}}\sum_{m\in\mathcal{M}} d z_t^m + \sum_{m\in\mathcal{M}}\sum_{i\in\mathcal{N}}\sum_{(s,t)\in\mathcal{I}} c_{st}^i x_{st}^{mi}, \tag{10a}$$

$$\text{subject to} \quad \sum_{s=0}^{t-1} x_{st}^{mi} \leqslant z_t^m, \quad i \in \mathcal{N}, \quad m \in \mathcal{M}, \quad t \in \mathcal{T}, \tag{10b}$$

$$\sum_{s=0}^{t-1} x_{st}^{mi} = \sum_{s=t+1}^{T+1} x_{ts}^{mi}, \quad i \in \mathcal{N}, \quad m \in \mathcal{M}, \quad t \in \mathcal{T}, \tag{10c}$$

$$\sum_{s=1}^{T+1} x_{0s}^{mi} = 1, \quad i \in \mathcal{N}, \quad m \in \mathcal{M}, \tag{10d}$$

$$z_t^m \leqslant y_t, \quad m \in \mathcal{M}, \ t \in \mathcal{T}, \tag{10e}$$

$$x_{st}^{mi} \in \{0,1\}, \quad i \in \mathcal{N}, \ m \in \mathcal{M}, \ (s,t) \in \mathcal{I}, \tag{10f}$$

$$z_t^m \in \{0,1\}, \quad m \in \mathcal{M}, \quad t \in \mathcal{T}, \tag{10g}$$

$$y_t \in \{0,1\}, \quad t \in \mathcal{T}. \tag{10h}$$

We will utilize the model (10) to minimize the sum of PM and expected CM costs, analogous to the ideas presented in Section 4.2. The performance of the model will be evaluated for the case when rescheduling at failure is allowed.

5.1. The choice of deterioration cost functions

In addition to the deterioration cost functions employed in Section 4.2, we include a term approximating the effect of rescheduling as follows. Let $S_1, S_2, \ldots, S_k$ be i.i.d. random variables drawn from the failure time distributions of component i. The k:th failure time of component i is defined as $J_k := \sum_{l=1}^{k} S_l$. A natural deterioration cost function is the expectation

$$M_i(t) := \mathbb{E}\left[\sum_{k=1}^{\infty} \mathbb{1}_{\{J_k \leqslant t\}} G(J_k, t)\right],$$

where $\mathbb{1}$ denotes the indicator function, and the value of $G(s,t)$ measures the effective cost of a failure occurring at time s provided that a replacement is scheduled at time $t > s$, and a replacement has been performed at time 0. A simple model for the function G to account for rescheduling is to assume that whenever a failure occurs long before a planned preventive replacement, i.e., when $s \ll t$, a dynamic rescheduling is likely to result in a pure CM for the failed component. The marginal failure cost in this scenario is thus fairly well represented by the sum $(c_i^{\mathrm{CM}} + d + f)$. On the other hand, if a component fails just before a planned PM, i.e., when $s \approx t$, a dynamic rescheduling is likely to move any soon-to-be-performed planned PM to the present time. The marginal failure cost in this scenario may thus be represented by the sum $(c_i^{\mathrm{CM}} - c_i^{\mathrm{PM}})$. We link these two scenarios by modeling the effective cost of a failure as $G(s,t) := (c_i^{\mathrm{CM}} + d + f) - \left(\frac{s}{t}\right)^{\lambda}(c_i^{\mathrm{PM}} + d + f)$, where $\lambda > 0$.

Table 4
The mean total maintenance costs (standard deviation) per day and average scenario statistics for the three policies for the wind farm case.

	CI	CM	Model (10)
Cost (std) ($/day)	820 (81)	1200 (80)	770 (71)
Wind farm visits	24.9	54.6	19.6
Visits/turbine	9.1	11.5	8.1
Replacements/component	24.1	11.6	24.6
Failures/scenario	19.9	54.6	18.8

5.2. Results

Data for component lives and PM and CM costs are taken from Tian et al. (2011). The parameter value $\lambda = 3$ is chosen by trial-and-error and the planning period of 8900 days is discretized into $T = 90$ time steps. The simulation set-up described in the first paragraph of Section 4.2.2 (augmented by the farm set-up cost f) is used. We compare the use of the model (10)—re-optimized at failure—with a CI policy (see Sections 1 and 4.2.2) as implemented by Tian et al., as well as with a pure CM policy. As shown in Table 4, the model (10) results in the lowest average total maintenance cost; the average cost of the CI policy is 6% higher, while the cost of the CM policy is 67% higher. The benefit of the model (10) over the CI policy is only slight and not at the level of the improvement obtained by considering CBM as discussed by Tian et al. The CBM strategy, however, requires access to condition monitoring equipment, which the strategies considered here do not. It should be noted that the initial schedule resulting from the model (10) equals that obtained by the optimal CI policy of Tian et al. Hence, the advantage of the model (10) over the CI policy is purely due to the rescheduling. Comparing the model (10) (with rescheduling) with the two policies reveals a clear reduction of the number of visits to the wind farm, and a slight decrease in the average number of failures. Coupled with the 13.5% reduction of the standard deviation of the total maintenance cost, this signifies an increase in reliability of the wind farm by using the model (10) with rescheduling.

6. Conclusions and future research

We introduce the preventive maintenance scheduling problem with interval costs (PMSPIC), which is an NP-hard problem. We demonstrate that the 0-1 ILP model (1), originally developed for the joint replenishment problem, provides a mathematical model for the PMSPIC. We show that the integrality restrictions on a majority of the variables in the model can be relaxed, and that the inequality constraints define facets of the convex hull of feasible solutions.

The model (1) is evaluated through case studies from the aircraft, railway, and wind power industries. It is designed for PM scheduling, but can also be dynamically used in a setting allowing for rescheduling. Using the model (1) reduces the maintenance costs by up to 16% as compared with the best simple policy investigated. Modeling the expected cost for CM between consecutive maintenance occasions by an interval cost reveals a limitation of our approach: Interval costs for separate components are expected to be accurate models only in cases where an unexpected CM event has only a weak influence on the maintenance schedules for the system. In the aircraft engine study, comprising relatively few CM events, this seems to be the case. However, for the wind farm study, due to the large number of components unexpected stops will be frequent, hence the separate component interval cost does not accurately capture the effects on the system of unexpected CM events.

In future research, we will study problems comprising several types of maintenance decisions; inspections, and repairs at different levels. We also intend to develop the model (1) to include, e.g., production and staff planning.

Acknowledgements

We thank the editor and two anonymous reviewers for comments and suggestions that made it possible to substantially improve the presentation of the article. We further acknowledge the financial support from the Swedish Energy Agency (through Grant No. 2010-004592 and Chalmers Energy Area of Advance) and the Swedish Governmental Agency for Innovation Systems (VINNOVA, through Chalmers Transport Area of Advance).

References

Almgren, T., Andréasson, N., Patriksson, M., Strömberg, A.-B., Wojciechowski, A., & Önnheim, M. (2012a). The opportunistic replacement problem: Theoretical analyses and numerical tests. *Mathematical Methods of Operations Research, 76*(3), 289–319.

Almgren, T., Andréasson, N., Palmgren, M., Patriksson, M., Strömberg, A.-B., Wojciechowski, A., et al. (2012b). Optimization models for improving periodic maintenance schedules by utilizing opportunities. In *Proceedings of 4th production and operations management world conference 2012*. <http://www.pom2012.org/cms/view/51>.

AMPL (2012). V. 12.1, AMPL Optimization LLC. <http://www.ampl.com>.

Arkin, E., Joneja, D., & Roundy, R. (1989). Computational complexity of uncapacitated multi-echelon production planning problems. *Operations Research Letters, 8*(2), 61–66.

Balakrishnan, A., Magnanti, T., & Wong, R. (1989). A dual-ascent procedure for large-scale uncapacitated network design. *Operations Research, 37*, 716–740.

Banverket (2004). Slipning av räler och spårväxlar, Handbok BVH 825.10 (in Swedish). ida2004.banverket.se.

Boctor, F., Laporte, G., & Renaud, J. (2004). Models and algorithms for the dynamic-demand joint replenishment problem. *International Journal of Production Research, 42*(13), 2667–2678.

Chattopadhyay, G., Reddy, V., & Larsson-Kråik, P.-O. (2005). Decision on economical rail grinding interval for controlling rolling contact fatigue. *International Transactions in Operational Research, 12*(6), 545–558.

CPLEX (2012). CPLEX Optimizer v. 12.1, IBM ILOG. <http://www.ilog.com/products/cplex>.

Crocker, J., & Kumar, U. (2000). Age-related maintenance versus reliability centred maintenance: A case study on aero-engines. *Reliability Engineering & System Safety, 67*(2), 113–118.

Crocker, J., & Sheng, Q. (2008). Is prevention better than cure? *Journal of Simulation, 2*(1), 53–60.

Dekker, R. (1995). Integrating optimisation, priority setting, planning and combining of maintenance activities. *European Journal of Operational Research, 82*(2), 225–240.

Dekker, R., Wildeman, R., & van der Duyn Schouten, F. (1997). A review of multi-component maintenance models with economic dependence. *Mathematical Methods of Operations Research, 45*(3), 411–435.

Ehrgott, M. (2005). *Multicriteria optimization*. *Lecture notes in economics and mathematical systems* (Vol. 491). New York, NY, USA: Springer.

Federgruen, A., & Tzur, M. (1994). The joint replenishment problem with time-varying costs and demands: Efficient, asymptotic and ε-optimal solutions. *Operations Research, 42*(6), 1067–1086.

Fulkerson, D. (1966). Flow networks and combinatorial operations research. *The American Mathematical Monthly, 73*(2), 115–138.

Grassie, S. (2005). Rolling contact fatigue on the British railway system: Treatment. *Wear, 258*(7–8), 1310–1318.

Grigoriev, A., Van De Klundert, J., & Spieksma, F. (2006). Modeling and solving the periodic maintenance problem. *European Journal of Operational Research, 172*(3), 783–797.

INNOTRACK (2009). D4.1.4 Rail degradation algorithms. <http://www.innotrack.net>.

Joneja, D. (1987). *Multi-echelon and joint replenishment production and distribution systems with non-stationary demand*. Tech. Rep. 731, School of OR & IE, Cornell University, Ithaca, New York, NY, USA.

Joneja, D. (1990). The joint replenishment problem: New heuristics and worst case performance bounds. *Operations Research, 38*(4), 711–723.

Khouja, M., & Goyal, S. (2008). A review of the joint replenishment problem literature: 1989–2005. *European Journal of Operational Research, 186*(1), 1–16.

Levi, R., Roundy, R. O., & Shmoys, D. B. (2006). Primal-dual algorithms for deterministic inventory problems. *Mathematics of Operations Research, 31*(2), 267–284.

Lewis, R., & Olofsson, U. (2009). *Wheel–rail interface handbook*. Cambridge, UK: Woodhead publishing. ISBN 978-1-84569-412-8.

Murthy, D., Xie, M., & Jiang, R. (2004). *Weibull models* (Vol. 358). Hoboken, NJ, USA: Wiley-Interscience.

Narayanan, A., & Robinson, E. (2006). More on "Models and algorithms for the dynamic-demand joint replenishment problem". *International Journal of Production Research, 44*(2), 383–397.

Nemhauser, G., & Wolsey, L. (1988). *Integer and combinatorial optimization*. New York, NY, USA: John Wiley & Sons.

Nicolai, R. P., & Dekker, R. (2008). Optimal maintenance of multi-component systems: A review. In K. A. H. Kobbacy & D. N. P. Murthy (Eds.), *Complex system maintenance handbook*. *Springer series in reliability engineering* (pp. 263–286). Berlin, Germany: Springer-Verlag.

Patriksson, M., Strömberg, A.-B., & Wojciechowski, A. (2014). The stochastic opportunistic replacement problem, Part II: A two stage solution approach. *Annals of Operations Research*. http://dx.doi.org/10.1007/s10479-012-1134-1.

Smith, W. L., & Leadbetter, M. R. (1963). On the renewal function for the Weibull distribution. *Technometrics, 5*(3), 393–396.

Tian, Z., Jin, T., Wu, B., & Ding, F. (2011). Condition based maintenance optimization for wind power generation systems under continuous monitoring. *Renewable Energy, 36*(5), 1502–1509.

Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research, 139*(3), 469–489.