Computers & Industrial Engineering xxx (2014) xxx–xxx

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

ARTICLE INFO

Article history:
Available online xxxx

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

ABSTRACT

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 multicommodity 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