European Journal of Operational Research 256 (2017) 102–115

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

A multi-item approach to repairable stocking and expediting in a fluctuating demand environment

Joachim Arts*

School of Industrial Engineering, Eindhoven University of Technology, PO box 513, 5600MB, Eindhoven, the Netherlands

ARTICLE INFO

Article history:
Received 27 May 2015
Accepted 2 June 2016
Available online 11 June 2016

Keywords:
Inventory
Spare parts
Column generation
Maintenance
Markov modulated Poisson process

ABSTRACT

We consider a single inventory location where multiple types of repairable spare parts are kept for service and maintenance of several different fleets of assets. Demand for each part is a Markov modulated Poisson process (MMPP). Each fleet has a target for the maximum expected number of assets down for lack of a spare part. The inventory manager can meet this target by stocking repairables and by expediting the repair of parts. Expedited repairs have a shorter lead time. There are multiple repair shops (or departments) that handle the repair of parts and the load imposed on repair shops by expedited repairs is constrained. A dual-index policy makes stocking and expediting decisions that depend on demand fluctuations for each spare part type. We formulate the above problem as a non-linear non-convex integer programing problem and provide an algorithm based on column generation to compute feasible near optimal solutions and tight lower bounds. We show how to use the MMPP to model demand fluctuations in maintenance and other settings, including a moment fitting algorithm. We quantify the value of lead time flexibility and show that effective use of this flexibility can yield cost reductions of around 25 percent.

1. Introduction

Service and manufacturing operations rely heavily on the availability of equipment such as aircraft, MRI-scanners, trains, and manufacturing equipment. The owners of such assets need to keep their equipment up and running as efficiently as possible. This is usually done by replacing defective components with ready for use components. The defective component is often expensive. Therefore a defective component is usually repaired and put back on stock. Such components are called *repairables*, and the working method described above is called *repair-by-replacement*. A sufficiently large number of spare repairables are needed to make such a system work, in particular to ensure a sufficiently high level of availability of capital assets. Buying sufficient repairables of all types needed to maintain a fleet of equipment is a major investment decision for firms with capital assets, because repairables are expensive.

The amount of spare repairables to buy is not the only major decision that affects the availability of capital assets. The repair of parts is usually done in house by several repair shops organized according to technical disciplines. For example, large airlines have repair shops for, amongst others, mechanical parts, avionics, and pneumatics. Our research has been instigated by a project we conducted at NedTrain, the maintenance division for rolling stock for the Dutch railways. The model we present in this paper has been used in a case study at NedTrain and was conducted by Van Aspert (2014) (see also Van Aspert, 2013). NedTrain also invests in repairable spare parts and has repair shops for mechanical parts, compressors, pneumatics, low voltage electronics, and high voltage electrical systems amongst others. The repair operations in these shops affect the availability of several fleets of trains. At the operational level, the inventory and repair shop planners coordinate to make sure repair priority is given to parts for which the inventory is most likely to run out in the near future.

The objective of this paper is to present a tractable optimization model that assists decision makers in answering the following questions

1. How many spare parts should we buy of each repairable type?
2. When should we expedite the repair of a given repairable type?

We assume the decision maker has to make these decisions for several fleets of equipment (e.g. local trains and long distance trains), and across parts that use different repair resources (e.g. pneumatics and electronics). The objective of the decision maker is to minimize the costs involved with purchasing or holding repairable spare parts while:

- meeting a service level in the form of a maximum average number of backorders for each fleet, and

* Tel.: +31 402474983.
E-mail address: j.j.arts@tue.nl

http://dx.doi.org/10.1016/j.ejor.2016.06.003
0377-2217/

- keeping the load imposed on each repair resource due to expedited orders below a set target level.

Note that this stocking problem cannot be resolved for each fleet separately because repairables that belong to different fleets (may) use the same resources for repair. We consider a setting where repair resources are flexible and model this through the possibility to request regular repair or expedited repair when sending a defective part to the repair shop. Expedited repairs have a shorter lead time than regular repairs. Since the flexibility of a repair resource is limited, there is a constraint on the amount of repair work that can be expedited per time unit for each repair resource. We refer to the amount of work that a repair resource handles per time unit as the load. Repairables from different fleets compete for the opportunity to load a repair resource with expedited orders.

Demand for a single type of repairable spare part usually fluctuates over time. These demand fluctuations arise for several reasons such as periodic inspections, usage patterns of equipment over time and the season of year. Slay and Sherbrooke (1988) observe empirically that demand for aircraft parts is non-stationary, and our experience with NedTrain also shows demand for many parts is non-stationary. When the reasons for demand fluctuations are understood, expediting decisions can be made to anticipate these fluctuations and to make effective use of repair resources.

In this paper, we provide a mathematical model for the decision problem described above. This model has been conceived with an application at NedTrain in mind. We emphasize however that the applicability of the model and results in this paper extend to other companies that maintain their own equipment. We will illustrate the need, as well as the application of the model, using an example that runs throughout this entire paper. This example is about a fictitious railway company and is big enough to capture all model facets, yet small enough to inspect results in detail to gain insights and intuition. We finish this introduction by starting this example. The rest of the paper is organized as follows. Section 2 reviews related literature and positions the contribution of this paper with respect to existing literature. The mathematical model is provided in Section 3. The analysis of the model is in Section 4. Computational results of the model for industrial size instances are provided in Section 5 and concluding remarks are offered in Section 6.

Example 1. The railway company Thomas&Co needs new trains to replace locomotives with pulled carriages. They decide to buy 100 trains from Liam Engineering Inc., and plan to use those for the next 30–40 years on long distance train services. Along with this order of 100 trains, Liam Engineering Inc. offers the possibility to buy (repairable) spare parts at a considerably discounted price. Thomas&Co would like to buy repairable spare parts at this discounted price and is taking this opportunity to decide on the stocking levels of repairables for the new fleet, as well as to reconsider the stocking levels for repairables of other fleets. ◇

2. Literature review and contribution

Multi-item repairable inventory models are abundant in literature. We refer the reader to the books of Sherbrooke (2004), Muckstadt (2005) and Van Houtum and Kranenburg (2015), and review papers by Guide and Srivastava (1997), Kennedy, Patterson, and Fredendall (2002), and Basten and Van Houtum (2014) for a broad overview. In this section, we briefly discuss literature with similar modeling assumptions and literature that expounds on or uses similar solution methods as those used in this paper. On the modeling side, the main contributions of this paper are the fluctuating demand model and the use of a dynamic expediting policy that depends on demand fluctuations. On the analysis side, we decompose the problem per item via a column generation algorithm. Therefore, this section is organized around three main topics: fluctuating demand (Section 2.1), repair expediting and scheduling policies (Section 2.2), and decomposition and column generation algorithms (Section 2.3).

2.1. Fluctuating demand

Demand for repairables that fluctuates over time has been considered before in a series of models developed by the RAND corporation under the name Dyna-METRIC (Carillo, 1989; Hillestad, 1982; Isaacson & Boren, 1993). Initially, these models were based on an extension of Palm's theorem for non-stationary Poisson processes, but these efforts eventually developed into simulation models that do not allow efficient optimization. In the Dyna-METRIC approach, demand is a non-stationary Poisson process, but the Poisson demand rate is a deterministic function of time. Rather than performing steady-state analysis, the Dyna-METRIC approach is to perform a transient analysis at some particular point in time that is chosen by the modeler. The Dyna-METRIC model does not include the possibility to expedite repair. Demand fluctuations are therefore only buffered by holding inventory.

A similar approach is followed by Lau and Song (2008) with two exceptions: They also model the finite repair capacity using queueing approximations and they evaluate the transient behavior of the system at several points of interest rather than only one. For their extensions to Dyna-METRIC, they take heuristic and approximative approaches.

Our work differs from these contributions because demand fluctuations are modeled by a Markov modulated Poisson process. This resembles practice more closely as the intensity of demand over time behaves as a stochastic process rather than a deterministic function. Additionally, our model deals with these demand fluctuations not only by holding repairable inventory, but also by using the possibility to expedite repair. Our modeling also allows us to evaluate our system exactly and compute tight lower bounds on optimal system performance. The use of the Markov modulated Poisson process to model demand for inventory systems has already been advocated by Song and Zipkin (1993). However, no practical fitting algorithms have been provided for modeling demand. (There are, however, practical fitting algorithms for the MMPP process in the context of communication networks; see e.g. Heffes and Lucantoni (1986), Meier-Hellstern (1987), Yoshihara, Kasahara, and Takahashi (2001), Nelson and Gerhardt (2010)). We provide two practical fitting procedures. The first is based on the maintenance and repair setting and uses information from maintenance planning. The second procedure is a moment fitting procedure that fits on demand over the lead-time.

2.2. Expediting and repair scheduling policies

The possibility to either expedite repair or prioritize the scheduling of repairs in the repair shop has been considered many times, mostly under the assumption of fixed given turn-around stock levels (Hausman & Scudder, 1982; Liang, Balcıoğlu, & Svaluto, 2013; Pyke, 1990; Scudder, 1986; Scudder & Chua, 1987; Tiemessen & Van Houtum, 2012). In these contributions, the repair shop is modeled by a finite server queue. Given a limited capacity, the question becomes: How should limited repair capacity be allocated to repair jobs of various types, i.e., which repair jobs deserve priority?

As observed by Tiemessen and Van Houtum (2012), even for fixed given turn-around stock levels, computing optimal priority rules, or evaluating a given rule, requires computation times that grow exponentially in the number of different repairable types. Also the derivation of structural properties of optimal policies or