Computers and Operations Research 82 (2017) 1–14

Contents lists available at ScienceDirect

Computers and Operations Research

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

Multi-period technician scheduling with experience-based service times and stochastic customers

Xi Chen[a,c], Barrett W. Thomas[a,*], Mike Hewitt[b]

[a] Department of Management Sciences, Tippie College of Business, University of Iowa, Iowa City, IA, 52242, United States
[b] Information Systems and Operations Management, Quinlan School of Business, Loyola University Chicago, Chicago, IL, 60611, United States
[c] Department of Management Science and Engineering, International Business School, Beijing Foreign Studies University, Beijing, 100089, China

ARTICLE INFO

Article history:
Received 14 November 2016
Accepted 29 December 2016
Available online 10 January 2017

MSC:
00-01
99-00

Keywords:
Dynamic programming
Technician scheduling
Multi-period
Learning

ABSTRACT

This paper introduces the multi-period technician scheduling problem with experience-based service times and stochastic customers. In the problem, a manager must assign tasks of different types that are revealed at the start of each day to technicians who must complete the tasks that same day. As a technician gains experience with a type of task, the time that it takes to serve future tasks of that type is reduced (often referred to as experiential learning). As such, while the problem could be modeled as a single-period problem (i.e. focusing solely on the current day's tasks), we instead choose to model it as a multi-period problem and thus capture that daily decisions should recognize the long-term effects of learning. Specifically, we model the problem as a Markov decision process and introduce an approximate dynamic programming-based solution approach. The model can be adapted to handle cases of worker attrition and new task types. The solution approach relies on an approximation of the cost-to-go that uses forecasts of the next day's assignments for each technician and the resulting estimated time it will take to service those assignments given current period decisions. Using an extensive computational study, we demonstrate the value of our approach versus a myopic solution approach that views the problem as a single-period problem.

1. Introduction

As the global economy recovers, there is growing pressure in the skilled labour markets. According to the Hays Global Skills Index 2014, a statistics-based study designed to assess the dynamics of skilled labour markets across 31 countries, this pressure continues to rise, particularly in economies that are returning to pre-crisis levels such as United States, Germany, and the United Kingdom [1]. Maintaining growth in a pressured labor market requires that companies use their expensive and limited labor resources to their greatest potential. One opportunity is for companies to take advantage of the capacity that is gained as employees learn by experience. Matching the right employee with the right job cannot only help a company meet its current needs, but also build capacity for meeting future demand growth as well as build the flexibility needed to buffer against demand uncertainty. Further, the ability to account for each individual employee's ability to learn allows companies to best deploy workers in the face of strategic growth opportunities giving companies the agility that [2] calls essential to surviving in the volatile modern business environment.

In this paper, we explore the issue of how companies can use immediate employee job assignments to meet current demand and build capacity for the future. We focus on service workers, particularly service technicians. The problem discussed in this paper is a variant of the technician and task scheduling problem (TTSP). In the TTSP, a set of technicians serves a set of customer requests. Customers are associated with certain tasks and different tasks have different skills associated with them. In our version of the problem, technicians have different service times depending on their experience in performing a task as well as each technician's ability to transform that experience into improved productivity. We measure experience in the number of times that the technician has performed the task.

The fact that technician productivity, and really all workers productivity, is linked to experience suggests that what could be modeled as a single-period problem (i.e. focusing solely on making assignments to serve the current day's tasks) should instead be modeled as a multi-period problem. As such, we consider the multi-period technician scheduling problem that accounts for the fact that productivity increases (or service time decreases) as

* Corresponding author.
E-mail address: barrett-thomas@uiowa.edu (B.W. Thomas).

http://dx.doi.org/10.1016/j.cor.2016.12.026
0305-0548/

technicians gain experience. These increases in productivity are often referred to as “learning”. We assume that the time that it takes a technician to complete a task depends on the technician’s experience in the skill associated with the task and how quickly the technician learns. How quickly a technician learns is known as the technician’s learning rate. We assume that we have a set of heterogeneous technicians whose learning rates and initial experience are known. The service time depends on the amount of experience the worker has with the skill required by the task.

We assume that daily demand is not revealed until the day of service. Each day, the technicians serve the day’s demand. In this work, we seek to minimize the expected sum of each day’s total service times over a finite horizon. Reflecting what are often tight labor market conditions for technicians (a job definition that can span multiple industries, including home appliance repair, lab technician, and home health care), we choose an objective that seeks to maximize the capacity of an existing workforce. We call our problem variant the multi-period technician scheduling problem with experience-based service times and stochastic customers (MTSP-ESTSC).

To solve the problem, we propose an approximate dynamic programming (ADP) approach that, at each stage, solves a mixed integer program (MIP) to assign technicians to tasks. In addition to recognizing the resulting service times in the current period, the MIP approximates the impact of those assignments on future technician service times (the “cost-to-go”) with a forecast of each technician’s task assignments in the next period. Assignment decisions in the next period are partially driven by each technician’s service times on each type of task, which are in turn driven by their experience level on each type of task.ÊÊTo capture this fact, the forecasting model embedded in the MIP is a function of the assignment decisions in the current period.

One of the challenges associated with solving optimization models that recognize that humans learn is that the quantitative models of human learning proposed by the psychology community are non-linear. As a result, to solve a MIP at each stage of the ADP, we adapt an exact reformulation method from the literature that relies on the fact that the function we use to map experience to service time has a finite domain.

In this study, we make the following research contributions. First, we present the first model that explicitly models the impact of individualized, experience-based learning on the technician scheduling problem. Such a model will facilitate organizational productivity improvement by allowing for more effective workforce management. Further, we discuss how the presented Markov decision process (MDP) model can be adapted to handle cases of worker attrition and new task types. Second, we introduce a method for approximating the future value of today’s workforce assignments. With the addition of technician learning to the model, decisions today affect tomorrow’s productivity. Thus, the approximation method allows us to do so. Third, using the approximation, we demonstrate how the approximate Bellman equation can be transformed into a linear, mixed integer program. This transformation is significant because the nonlinearity of the learning functions naturally lead to a nonlinear integer programming formulation. Our formulation allows the approximate Bellman equation to be solved by a standard implementation of a commercial integer programming solver.

We demonstrate the value of the proposed solution approach with three experiments. In one experiment, the set of technicians and the set of task types remain the same over the problem horizon. In the second, we introduce a workforce disruption in the middle of the horizon in which one technician leaves the workforce and a new technician is added. The third variant adds an additional task type in the middle of the horizon. For each of the three problem variants, we compare the proposed solution approach to a myopic solution approach that views the problem as a single-period problem, ignoring the impact of current period decisions on future service times. Our comparisons demonstrate that the proposed solution approach leads to higher-quality solutions by better positioning technicians to meet future demands.

The remainder of this paper is organized as follows. Section 2 reviews the literature related to the MTSP-ESTSC. Section 3 presents a model for the problem. Section 4 describes the solution approach. Section 5 discuss the design of the experiments, and Section 6 presents our computational results. Finally, Section 7 concludes this work and suggests areas of future research.

2. Literature review

We review the literature of technician scheduling and routing problems as well as for experience-based learning.

2.1. Individual learning and its applications

That humans learn as they gain experience, “learning-by-doing” is a well known phenomenon. The learning effect was first examined on a scientific basis by [3], who quantified learning curves with the observation that in the aircraft industry the working costs per unit declined with an increasing production output. Subsequent empirical studies confirmed the existence and importance of learning effects (see for example [4–7]). In 2016, the concept has become mainstream enough that it is now included in textbooks on operations management ([8,9]).

The mathematical descriptions of learning are often called learning curves. Reviews of the literature on learning curves can be found in [10–13]. Because of the availability of distributions from which to generate workforces, in this research, we use the hyperbolic learning model described in [14]. We note that, while we employ the hyperbolic learning model, most learning curves have similar shapes and would support conclusions similar to those discussed in Section 6. We note that there also exists an extensive literature focusing on organizational learning [15]. provides an excellent reference.

Work that explicitly models individual learning and the associated heterogeneity of the workforce demonstrates the value of capturing learning [16]. shows that simply modeling worker heterogeneity without considering learning improves system performance versus assuming uniform workforce productivity in flow-line production. [17] extend the analysis of [16] to demonstrate the impact of heterogeneous learning and forgetting curves on system productivity in a assembly-line setting [18]. confirms the results of [17] for technician routing. In addition to flow lines, assembly lines, and technician routing, the value of modeling learning has also been found in call centers ([19]), departmental assignment ([20]), machine scheduling (see [21] for a review), project selection ([22,23]), and vehicle routing ([24]).

One of the challenges of much of the workforce planning literature that models individual learning is that the nonlinearity of the learning curves creates challenges. For this reason, work such as [19] and [25], simplify the model of individual learning to avoid the nonlinearities. Work such as [26] exploits structural properties of the optimal solution to increase the size of the problem that can be solved. However, such approaches do not generalize. Work such as [27–29] are limited to solving small problems [30]. introduce a linear and integer reformulation of the learning curve that takes advantage of the fact that most work is assigned in time intervals. The reformulation allows much larger problems to be solved than had been previously. We take advantage of the reformulation in this work as well. Other examples of the reformulation