European Journal of Operational Research 216 (2012) 172–177

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Heuristic solution methods for the stochastic flow shop problem

Kenneth R. Baker [a,*], Dominik Altheimer [b]

[a] Tuck School of Business, Dartmouth College, Hanover, NH, USA
[b] Helmut Schmidt Universitat, Hamburg, Germany

ARTICLE INFO

Article history:
Received 10 December 2010
Accepted 8 July 2011
Available online 23 July 2011

Keywords:
Stochastic scheduling
Flow shop
Makespan
Lognormal distribution

ABSTRACT

We investigate the stochastic flow shop problem with *m* machines and general distributions for processing times. No analytic method exists for solving this problem, so we looked instead at heuristic methods. We devised three constructive procedures with modest computational requirements, each based on approaches that have been successful at solving the deterministic counterpart. We compared the performance of these procedures experimentally on a set of test problems and found that all of them achieve near-optimal performance.

1. Introduction

The flow shop problem plays an important role in the theory of scheduling. The deterministic version was introduced to the literature by Johnson (1954), in what is often identified as the first formal study of a problem in scheduling theory. That article has led to a large number of papers studying variations of the basic model and various algorithmic approaches for finding solutions. For example, Reisman et al. (1997) claimed to have located 170 articles containing contributions to the "subdiscipline" of flow shop scheduling. More recently, Ruiz and Maroto (2005) cited 53 articles in their review paper on heuristic procedures for the permutation flow shop problem with makespan objective. Framinan et al. (2004) cited 76 articles in a review paper on the same topic. Reza Hejazi and Saghafian (2005) cited 176 articles in a review paper on exact and heuristic approaches to the same problem. Clearly, the flow shop scheduling problem has attracted a lot of attention.

On the other hand, progress with the stochastic version of the flow shop problem has been very limited. Few general results have been obtained, and the optimization of basic cases remains a challenge. In this paper, we present a comparative study of heuristic methods for solving the *m*-machine stochastic flow shop problem with the objective of minimizing the expected makespan. We focus on a few relatively simple heuristic approaches that are motivated by the existing literature, and we compare their performance on a set of test problems. Finally, we summarize our results and suggest what questions might guide future research on this subject.

2. Background on the deterministic model

The classical flow shop problem contains *n* jobs and *m* machines, as well as a set of standard assumptions (see, for example Baker and Trietsch, 2009a). The objective is to minimize the length of the schedule or *makespan*. In the case of two machines, we can construct an optimal job sequence by employing Johnson's Rule (Johnson, 1954), which leads to an efficient algorithm. In the case of three or more machines, the flow shop problem is *NP*-hard. Several effective heuristic procedures have been invented for solving problems with three or more machines. Relatively recent reviews of that literature have been compiled by Framinan et al. (2004) and Ruiz and Maroto (2005). We mention two heuristics in particular, as they are adapted to the stochastic model in our work. The first is due to Campbell et al. (1970), known as the CDS heuristic. The second is due to Nawaz et al. (1983), known as the NEH heuristic. Both are *constructive* heuristics. This term means that the algorithms perform a predictable amount of computation and ultimately construct a complete schedule. In contrast, an *improvement* heuristic starts with a given sequence and searches for improvement, but the computational effort is unpredictable. Improvement heuristics are usually based on generic methods such as neighborhood search. Sophisticated forms of improvement heuristics include tabu search, simulated annealing and genetic algorithms.

The CDS algorithm uses Johnson's Rule in a heuristic fashion and creates several schedules from which a "best" schedule is chosen. The algorithm corresponds to a multistage use of Johnson's Rule applied to a two-machine pseudo-problem derived from the original. The NEH algorithm constructs a single sequence, starting with a list of the jobs. The first two jobs on the list are removed, the two possible permutation sequences of those jobs are constructed, and the better of the two is retained (with ties broken arbitrarily).

* Corresponding author. Tel.: +1 603 646 2064.
E-mail address: ken.baker@dartmouth.edu (K.R. Baker).

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.07.021

The relative sequence position of the first two jobs is then fixed. At each succeeding step, a job is removed from the list and placed optimally into the partial sequence retained from the previous step. When the last job is removed from the list, a full sequence is chosen from among the possible insertions at the final step.

The CDS algorithm and the NEH algorithm are computationally efficient. The computational complexity is $O(mn\log n)$ for the CDS algorithm and $O(mn^2)$ for the NEH algorithm. Comparative studies by Park et al. (1984), Widmer and Hertz (1989), Taillard (1990), Ho and Chang (1991), Ponnambalam et al. (2001) and Ruiz and Maroto (2005) have tended to reinforce the proposition that these are the two best and most robust constructive procedures available.

3. The stochastic model

The most common stochastic version of the flow shop problem assumes that the processing times are allowed to be random variables. In particular, we assume that the processing time of job j on machine k follows a probability distribution with mean μ_{kj} and standard deviation σ_{kj} (denoted σ when it applies across all jobs and machines). For convenience, we also assume that the processing times are drawn independently from distributions of a given family, such as the normal or the uniform. As a result, the makespan will also be random, and the objective is to minimize its expected value. This single change from the deterministic version of the problem is sufficient to make the problem quite difficult to solve. In fact, no analytic solution procedure exists for the stochastic version. Little attention has even been paid to finding heuristic procedures for the stochastic flow shop problem, although Portougal and Trietsch (2006) have shown that Johnson's Rule applied to mean values will produce asymptotically optimal expected makespan values in the stochastic case. Our paper essentially presents the first study comparing heuristic procedures for the m-machine stochastic flow shop problem with expected makespan criterion.

If we restrict attention to the two-machine stochastic flow shop problem, it is still the case that no general results are known, but if we restrict ourselves further to the case of exponential distributions, then we have one result, which states: the expected makespan is minimized by sequencing the jobs in nonincreasing order of $(1/\mu_{i1} - 1/\mu_{i2})$. This ordering is known as Talwar's Rule. It was conjectured to be optimal by Talwar (1967) and later proven optimal by Cunningham and Dutta (1973). Thus, sequencing jobs based on the differences in their mean processing rates provides the optimal two-machine solution for one special case.

With the solution to the two-machine case established, we might look next to the m-machine case with exponential distributions, but generalizations of Talwar's Rule have not been developed for three or more machines. One advantage of the exponential assumption is the possibility of analytic calculation of the expected makespan. (Lacking an optimal sequencing rule, we must have such a capability merely to compare one sequence with another.) However, a disadvantage of the exponential distribution is the fact that it has only one parameter: its mean cannot be different from its standard deviation. Pinedo (1982) suggests the following rule of thumb: "Schedule jobs with smaller expected processing times and larger variances in the processing times toward the beginning and the end of the sequence." But for this rule to have meaning, we must deal with distributions that have distinct mean and variance parameters, unlike the exponential. Kalczynski and Kamburowski (2006) heuristically adapted Talwar's result for Weibull distributions, but did not attempt to generalize beyond two machines.

Baker and Trietsch (2010) tested three simple heuristic procedures for the two-machine stochastic model with general probability distributions. They compared Johnson's Heuristic (Johnson's Rule applied to the mean processing times), Talwar's Heuristic (Talwar's Rule applied to the mean processing times), and an Adjacent Pairwise Interchange Heuristic (which swapped adjacent jobs if their sequence, when considered separately, could be improved). Although none of the heuristic procedures dominated the others, Baker and Trietsch found that they all achieved very good performance, providing expected makespan values that, on average, were within 1% of the best value found. In our work, we demonstrate that this same good performance can be achieved in the m-machine case.

For the exponential case, Gourgand et al. (2003) show that the expected makespan calculation can be carried out for m machines analytically using a Markovian approach, but even that method encounters limitations due to problem size. (They proceed no further in making the calculation than medium-size problems of 20 jobs and 5 machines.) They conclude that we must ultimately rely on simulation techniques to evaluate the expected makespan. Thus, to make progress on the model with general probability distributions for processing times, we shall have to rely on (1) heuristic procedures to find good sequences and (2) simulation procedures to calculate expected makespan values.

4. Heuristic procedures

We describe three main heuristic procedures for sequencing jobs in the stochastic flow shop with expected makespan objective. Two of these procedures follow the logic of the CDS algorithm. In other words, they create a series of two-machine pseudo-problems; then those pseudo-problems are solved by a two-machine algorithm (either Johnson's Rule or Talwar's Rule). The procedures are thus referred to as Johnson's Heuristic and Talwar's Heuristic.

Johnson's Heuristic solves a two-machine stochastic flow shop problem by replacing the processing times with their mean values. Then, the resulting deterministic problem is solved by Johnson's Rule to deliver a desired sequence for the jobs. This procedure is heuristic because it solves a deterministic counterpart of the stochastic problem. Talwar's Heuristic solves a two-machine stochastic flow shop problem by applying Talwar's Rule (sorting the jobs by nonincreasing differences of the mean processing rates). This procedure is heuristic because its optimality does not extend to general distributions.

Thus, the first two heuristic procedures might be called the CDS/Johnson Heuristic and the CDS/Talwar Heuristic. Our third procedure is the NEH algorithm, applied to the stochastic problem directly. That is, the procedure finds the best two-job sequence; then, keeping the two jobs in their better order, it finds the best insertion of the third job into the two-job sequence, then the best insertion of the fourth job into the best three-job sequence, and so on. The jobs are considered in the order of nonincreasing total mean processing time.

Each heuristic procedure requires the ability to compare two job sequences and choose the better one. In other words, we must be able to find the expected makespan for each of two sequences in a given stochastic flow shop problem and identify the smaller of the two. For this purpose, we use simulation. Gourgand et al. (2003) assessed the accuracy of simulation by making comparisons in cases for which their Markovian analysis is practicable. They tested different sample sizes on a standard dataset and found, for example, that sample sizes of 200,000 produced 95% confidence intervals on the order of 0.1% and average estimation errors on the order of 0.05%. Those results were based on comparing simulation and analytic calculations for stochastic flow shop problems with exponential processing times. The tests used lexicographic job sequences (i.e., the equivalent of an arbitrary sequence) and indicated that a sample size this large is more than sufficient to obtain useful estimates.