Computers and Operations Research 93 (2018) 123–134

Contents lists available at ScienceDirect

Computers and Operations Research

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

Improved algorithms to minimize workload balancing criteria on identical parallel machines

Stefan Schwerdfeger, Rico Walter*

Friedrich Schiller University Jena, Chair for Management Science, Carl-Zeiß-Straße 3, Jena D-07743, Germany

ARTICLE INFO

Article history:
Received 6 March 2017
Revised 10 January 2018
Accepted 27 January 2018
Available online 31 January 2018

Keywords:
Scheduling
Identical parallel machines
Workload balancing
NSSWD-criterion
Exact algorithms
Heuristics

ABSTRACT

In Schwerdfeger and Walter (2016), we proposed a subset sum based improvement procedure (denoted by LS) for solving the problem of minimizing the normalized sum of squared workload deviations on m identical machines. In its core, the algorithm builds on an exact procedure (denoted by WB3) for the case of $m = 3$ machines. Although this approach proved to be superior to existing procedures in terms of solution quality, we identified room for further improvements.

Building upon our prior work, in this follow-up paper we suggest an enhanced version of WB3 that leads to a significant speed-up of several orders of magnitude and we considerably improve on the performance of LS on difficult instances where the ratio of the number of jobs to the number of machines is small. Moreover, we investigate a simple surrogate balancing measure that can also be optimized by our algorithms with only a slight modification. Results of a comprehensive computational study on a large set of benchmark as well as random test instances demonstrate the effectiveness of the improved algorithms.

1. Introduction

Balancing the workload among a set of resources (machines or human workers) is an important objective in many practical applications. For instance, Khouja and Conrad (1995) report on a balancing problem encountered by a mail order firm. To handle order taking and billing and to ease order tracking, customer groups should be assigned almost equally to the employees. From an employee's point of view, such an assignment is viewed as fair. It also helps to avoid dissatisfaction due to unequal treatment (cf. Cossari et al., 2013). Considering production systems, balanced workloads usually imply an efficient utilization of resources (Ho et al., 2009) and reduce idle times as well as work in process (Ouazene et al., 2014). Generally, the task of the production planner is to assign a set $\mathcal{J}$ of n jobs to a set $\mathcal{I}$ of m resources as "equally" as possible, i.e. to find a schedule that optimizes the workload balance. In our basic scheduling scenario, we assume the jobs to be independent and indivisible, the processing times p_j of the jobs $j = 1, \ldots, n$ to be deterministic and positive integers ($j = 1, \ldots, n$), and the machines to be identical and parallel. To assess and measure the balance or "equality" of the workload, different criteria can be applied (cf. Cossari et al., 2013). In case of identical parallel machines, a well-established criterion is the normalized sum of squared workload deviations (abbreviated as NSSWD). It has been introduced by Ho et al. (2009) and is defined as $\text{NSSWD} = \frac{1}{\mu}\left[\sum_{i=1}^{m}(C_i - \mu)^2\right]^{1/2}$ where C_i denotes the completion time of machine i ($i = 1, \ldots, m$) and $\mu = \sum_{j=1}^{n} p_j/m$ represents the average machine completion time. A schedule with a smaller NSSWD-value is assumed to be more balanced than a schedule with a larger NSSWD-value so that we seek to minimize NSSWD and denote the corresponding $\mathcal{NP}$-hard workload balancing problem by $P||\text{NSSWD}$.

Due to its complexity status it is unlikely to find an algorithm that optimally solves any instance of $P||NSSWD$ within a reasonable amount of time. Therefore, when it comes to benchmark heuristic or approximate solutions, one either has to derive lower bound arguments (cf. Schwerdfeger and Walter, 2016) or one can apply the obvious relation that a schedule whose difference C_Δ between the maximal completion time $C_{max} = \max_{i=1,\ldots,m} C_i$ and the minimal completion time $C_{min} = \min_{i=1,\ldots,m} C_i$ is less than or equal to 1 is NSSWD-optimal. The aforementioned relation raises the natural question whether $C_\Delta = C_{max} - C_{min}$ can even serve as a simple but meaningful surrogate balancing measure. This question is in line with Coffman and Langston (1984) who said that the objective function C_Δ "is perhaps the most direct measure of near-equality". Ouazene et al. (2014) also suggested to use C_Δ as a measure for the workload balance on identical parallel machines. They formulated the problem as a mixed-integer program and solved a few test instances with the help of a state-of-the-art solver but they did not develop a tailored solution procedure. In their recent publication, Ouazene et al. (2016) analyzed the correlation between

* Corresponding author.
E-mail addresses: stefan.schwerdfeger@uni-jena.de (S. Schwerdfeger), rico.walter@uni-jena.de (R. Walter).

https://doi.org/10.1016/j.cor.2018.01.018
0305-0548/

```
for j = 1 : n
    if p[j] ≤ μ
        S[p[j]] = 1; Z[p[j], nor[p[j]]] = j; nor[p[j]]++
        if p[j] > T
            T = p[j]
        end
        for s = 0 : r − 1
            h = R[s] + p[j]
            if h ≤ μ
                S[h] = 1; Z[h, nor[h]] = j; nor[h]++
                if h > T
                    T = h
                end
            end
        end
        r = 0
        for s = p[j] : T
            if S[s] > 0
                R[r] = s; r++
            end
        end
    end
end
```

Fig. 1. Pseudocode of the dynamic programming procedure to determine and store all representations.

NSSWD and C_Δ theoretically and established the following relation: NSSWD/$C_\Delta \leq m^2 / \sum_{j=1}^{n} p_j$.

Apart from these publications, minimizing C_Δ has gained rather little attention in the literature on parallel machine scheduling problems. With regards to solution procedures, we are not aware of any other specific algorithm next to the differencing method of Karmarkar and Karp (1982). Exact algorithms for the $\mathcal{NP}$-hard C_Δ-minimization problem on identical parallel machines (denoted by $P||C_\Delta$) also have not been proposed until now. For further contributions to workload balancing problems on identical parallel machines, we refer to the literature review in Schwerdfeger and Walter (2016).

The contribution of this follow-up paper is manifold. Regarding $P||$NSSWD, we propose an enhanced version of Schwerdfeger and Walter's exact WB3-algorithm for the case of $m = 3$ machines that leads to a significant speed-up of several orders of magnitude. We also extend the aforementioned algorithm to an effective exact algorithm for the case of $m = 4$ machines. Using the exact three machine procedure as a subroutine within a local search method, we are also able to considerably improve on the performance of Schwerdfeger and Walter's local search algorithm LS on difficult instances with small ratios of n to m.

With regards to $P||C_\Delta$, we slightly adapt our exact procedure for $P3||$NSSWD in order to obtain an exact algorithm for $P3||C_\Delta$ which is then used as a subroutine within a local search algorithm to minimize C_Δ on $m \geq 4$ identical parallel machines. Moreover, we also address the question whether C_Δ can serve as a meaningful surrogate measure for workload balance. Therefore, we investigate the relationship between minimizing NSSWD and minimizing C_Δ and present results of a statistical analysis on the correlation between the two balancing objectives that attest to their similarity.

Laying the focus on the development of enhanced solution procedures and their computational performance, the present paper is structured as follows. In Section 2, we describe improved algorithmic approaches to the solution of the balancing problem $P||$NSSWD. Afterwards, in Section 3, we are concerned with modifications of our algorithms to solve another balancing problem ($P||C_\Delta$). The computational performance of our algorithms is tested in a comprehensive computational study on a large set of benchmark as well as randomly generated instances (Section 4). Finally, Section 5 concludes the paper.

2. Algorithms for solving $P||$NSSWD

While Ho et al. (2009) and Cossari et al. (2012, 2013) proposed heuristic algorithms that make use of the traditional pairwise exchange of jobs, in Schwerdfeger and Walter (2016) we developed an innovative three machine neighborhood approach that proves beneficial in minimizing different workload balancing criteria. In its core, our local search improvement procedure (denoted by LS) builds on an exact procedure (denoted by WB3) for the case of three machines. The procedure WB3 itself requires solving a sequence of subset sum problems (SSP) and, thus, entails certain computational costs. Since WB3 is repeatedly called within the local search framework, the required computation time of LS heavily depends on the runtime of WB3. Therefore, even slight improvements in WB3's computation time lead to a speed-up of the whole LS-procedure.

In what follows, we develop an enhanced version of WB3 (see Section 2.1). Afterwards, we explain how the underlying idea can easily be transferred to exactly solve the case of four machines and we also describe our modified local search procedure (see Section 2.2).