



Open-Shop Scheduling for Unit Jobs Under Precedence Constraints

An Zhang^{1,2}, Yong Chen^{1,2}, Randy Goebel², and Guohui Lin²(✉)

¹ Department of Mathematics, Hangzhou Dianzi University, Hangzhou, China
{anzhang, chenyon}@hdu.edu.cn

² Department of Computing Science, University of Alberta, Edmonton, AB, Canada
{rgoebel, guohui}@ualberta.ca

Abstract. We study open-shop scheduling for unit jobs under precedence constraints, where if one job precedes another job then it has to be finished before the other job can start to be processed. For the three-machine open-shop to minimize the makespan, we first present a simple $5/3$ -approximation based on a partition of the job set into agreeable layers using the natural layered representation of the precedence graph. We then show a greedy algorithm to reduce the number of singleton-job layers, resulting in an improved partition, which leads to a $4/3$ -approximation. Both approximation algorithms apply to the general m -machine open-shops too.

Keywords: Open-shop scheduling · Precedence constraint
Directed acyclic graph · Approximation algorithm

1 Introduction

Machine scheduling with precedence constraints on the jobs has received much attention in the past few decades, and several algorithmic techniques such as the *critical path method* and the *project evaluation and review technique* [9] have been developed from the line of research. Job precedence constraints are common in construction and manufacturing industries, for example, the bicycle assembly problem is an earliest precedence constrained scheduling application introduced by Graham [7].

Precedence constraints describe the job processing order in a way that one or more jobs have to be finished before another job is allowed to start its processing. Such relationships together are usually represented as a *directed acyclic graph* (DAG) $G = (V, E)$, called the *precedence graph*, where V is the set of jobs and an edge $(v_i, v_j) \in E$ states that the job v_i precedes the job v_j , that is, v_i needs to be finished before v_j can start to be processed.

In this paper, we discuss the open-shop scheduling environment and use Om to denote the m -machine open-shop for some constant m , and O to denote the open-shop in which the number of machines is part of the input. In either Om or O , every job needs to be processed non-preemptively by each machine, in any

machine order, and it is *finished* (or said *completed*) when it has been processed by all the machines. Note that the usual scheduling rules apply to a feasible schedule, that is, at any time point, a job can be processed by at most one machine and each machine can be processing at most one job. The makespan of the schedule is the maximum job completion time. The open-shop scheduling to minimize the makespan is denoted as $Om \parallel C_{\max}$ or $O \parallel C_{\max}$, which has received much study [6, 9, 11, 12, 15]. In particular, $O2 \parallel C_{\max}$ is solvable in $O(n)$ -time, where n denotes the number of jobs [6, 9]; $Om \parallel C_{\max}$ becomes weakly NP-hard when $m \geq 3$ [6] but admits a polynomial-time approximation scheme (PTAS) [11, 12]; $O \parallel C_{\max}$ is strongly NP-hard and cannot be approximated within 1.25 [15].

Open-shop scheduling with precedence constraints, denoted as $Om \mid prec \mid C_{\max}$ or $O \mid prec \mid C_{\max}$, is more difficult than its classical counterparts, which can be considered as scheduling without precedence constraints. Several special classes of precedence graphs have been investigated in the literature. If every job has at most one predecessor and at most one successor, the precedence graph is referred to as *chains*. If every job has at most one successor (one predecessor, respectively), the precedence graph is referred to as an *intree* (an *outtree*, respectively). The fact that the precedence graph belongs to a particular class may change the computational complexity of the scheduling problem. In general, one can expect that the precedence constraints increase the problem complexity. For example, $O2 \mid chains \mid C_{\max}$ becomes NP-hard [13]. For more complexity results on precedence constrained scheduling, the interested readers can refer to Lenstra and Rinnooy Kan [8], or Prot and Bellenguez-Morinea [10].

Unlike most past results which are on computational complexity, in this paper we aim to develop algorithmic positive results for open-shop scheduling with precedence constraints, from the approximation algorithm perspective. We focus on the problems restricted to unit jobs, that is, the jobs have the same processing times on all the machines (*i.e.*, $p_{ij} = 1$); most of these problems remain NP-hard, or their complexity are still open. To name a few, for an arbitrary precedence graph, the problem $O \mid p_{ij} = 1, prec \mid C_{\max}$ was shown to be strongly NP-hard by Timkovsky [14]; when the precedence graph is an out-tree, then the problem $O \mid p_{ij} = 1, outtree \mid C_{\max}$ becomes polynomially solvable [1]; for a more general objective of minimizing the maximum lateness, Timkovsky proved that $O \mid p_{ij} = 1, outtree \mid L_{\max}$ is weakly NP-hard [14], while the problem $O \mid p_{ij} = 1, intree \mid L_{\max}$ is polynomial solvable [2, 3]. We note that, however, there are polynomial time algorithms for $O2 \mid p_{ij} = 1, prec \mid L_{\max}$, even if the jobs have different release times [2, 3].

The problem we study in this paper is the m -machine open-shop for unit jobs under arbitrary precedence constraints, $Om \mid p_{ij} = 1, prec \mid C_{\max}$, where $m \geq 3$. For this fundamental problem in scheduling theory, there is no known computational complexity result in the literature. In fact, even when $m = 3$, whether or not $O3 \mid p_{ij} = 1, prec \mid C_{\max}$ is NP-hard is an open question explicitly listed in the websites maintained by Brucker and Knust [4] and Dürr [5], and in the survey paper by Prot and Bellenguez-Morinea [10].

We first introduce a natural layered representation for the precedence graph in Sect. 2, based on which we can construct a partition of the job set into agreeable subsets. We then construct a schedule using the partition and show that it is a $5/3$ -approximation for the problem $O3 \mid p_{ij} = 1, prec \mid C_{\max}$. In Sect. 3, we propose a greedy algorithm to reduce the number of singleton-job subsets in the earlier partition, resulting in an improved partition, which leads to a $4/3$ -approximation. We also show that both approximation algorithms apply to the general m -machine open-shops.

2 Preliminaries

We study the problem $O3 \mid p_{ij} = 1, prec \mid C_{\max}$, in which the unit jobs should be processed under the given precedence constraints. These precedence constraints are described as a directed acyclic graph (DAG), the *precedence graph*, in which a vertex corresponds to a job and a directed edge represents a precedence relationship between a pair of jobs. In the rest of the paper, we use a job and a vertex interchangeably. Due to all jobs having unit processing times, we assume without loss of generality that in any feasible schedule the starting processing time of every job is an integer.

Let $V = \{v_1, v_2, \dots, v_n\}$ be the given set of unit jobs. If v_i precedes v_j , that is, we can start processing the job v_j only if the job v_i is finished by the three-machine openshop $O3$, then there is a directed path beginning from v_i and ending at v_j . Such a directed path is a directed edge (v_i, v_j) in the simplest case, in the DAG $G = (V, E)$.

A subset $X \subseteq V$ of jobs is *agreeable* if none of the jobs of X precedes another. In particular, two jobs are *agreeable* if none of them precedes the other, and thus they can be processed concurrently on different machines in a feasible schedule.

Lemma 1. *An agreeable subset $X \subseteq V$ of jobs can be processed by the three-machine openshop $O3$ in $|X|$ units of time if $|X| \geq 3$, or in 3 units of time if $|X| = 1, 2$.*

Proof. Let the jobs of X be v_1, v_2, \dots, v_k . When $k = 1$, at any time point T , v_1 can be processed on the first machine M_1 (the second machine M_2 , the third machine M_3 , respectively) starting at T ($T + 1$, $T + 2$, respectively), and thus finished within 3 units of time.

When $k = 2$, at any time point T , v_1 can be processed on the first machine M_1 (the second machine M_2 , the third machine M_3 , respectively) starting at T ($T + 1$, $T + 2$, respectively); v_2 can be processed on the third machine M_3 (the first machine M_1 , the second machine M_2 , respectively) starting at T ($T + 1$, $T + 2$, respectively). Thus both of them are finished within 3 units of time.

When $k \geq 3$, at any time point T , for $j = 1, 2, \dots, k - 2$, v_j can be processed on the first machine M_1 (the second machine M_2 , the third machine M_3 , respectively) starting at $T + j - 1$ ($T + j$, $T + j + 1$, respectively); v_{k-1} can be processed on the third machine M_3 (the first machine M_1 , the second machine M_2 , respectively) starting at T ($T + k - 2$, $T + k - 1$, respectively); v_k can be

processed on the second machine M_2 (the third machine M_3 , the first machine M_1 , respectively) starting at T ($T + 1$, $T + k - 1$, respectively). See Fig. 1 for an illustration. Thus all of them are finished within k units of time. \square




Fig. 1. An sub-schedule to process an agreeable subset $X \subseteq V$ of jobs in $|X|$ units of time when $k = |X| \geq 3$.

Given two disjoint agreeable subsets X_1 and X_2 , if a job of X_1 precedes a job of X_2 , then we say X_1 precedes X_2 . A collection of mutual disjoint agreeable subsets is *acyclic* if the precedence relations among the subsets do not contain any cycle. A subset of k jobs is called a k -subset, for $k = 1, 2, \dots$. For simplicity, a 1-subset is also called a *singleton*.

Corollary 1. *Let \mathcal{C} be an acyclic partition of V into agreeable subsets, in which there are b 2-subsets and c singletons. Then a schedule π can be constructed to achieve the makespan $C_{\max}^\pi = n + b + 2c$, where $n = |V|$.*

Proof. Using Lemma 1, all the $n - 2b - c$ jobs outside of those 2-subsets and singletons can be finished in $n - 2b - c$ units of time, and each 2-subset and each singleton can be finished in 3 units of time, respectively. Putting them together, we have a schedule π of makespan $C_{\max}^\pi = (n - 2b - c) + 3b + 3c = n + b + 2c$. \square

By Corollary 1, we wish to solve the problem $O3 \mid p_{ij} = 1, prec \mid C_{\max}$ by partitioning the jobs into acyclic agreeable subsets such that the quantity $b + 2c$ is minimized. Our main contribution is an algorithm that produces an acyclic partition achieving a number of singletons no more than the number of isolated jobs (to be defined) in the optimal schedule.

In the rest of the section, we introduce a representation for the DAGs which is used in our algorithm design and analysis.

2.1 A DAG Representation

Let $G = (V, E)$ be the precedence graph describing all the given precedence constraints, where a directed path from v_i to v_j suggests that the job v_i precedes

the job v_j (that is, v_j cannot be processed unless v_i is finished by the three-machine openshop). Through out the paper, we let $n = |V|$ and $m = |E|$.

If $(v_i, v_j) \in E$ and there exists a path from v_i to v_j not involving the edge (v_i, v_j) , then we call (v_i, v_j) a *redundant edge*, in the sense that the precedence constraint between every pair of jobs is still there after we remove the edge (v_i, v_j) from the graph. We may thus simplify the graph G by removing all redundant edges, which can be executed in $O(m)$ time by a *breadth-first-search* (BFS). Afterwards, for each edge $(v_i, v_j) \in E$, we call v_i a *parent* of v_j and v_j a *child* of v_i . Note that a job can have multiple parents, and multiple children as well.

In the following layered representation of the graph $G = (V, E)$, each job will be associated with a level (a positive integer). The first layer consists of all the jobs with in-degree 0, and these are the level-1 jobs. Iteratively, after the level- ℓ jobs are determined, they and the edges (these are out-edges) incident at them are removed from the graph; then the $(\ell + 1)$ -st layer consists of all the jobs with in-degree 0 in the remainder graph, and these are the level- $(\ell + 1)$ jobs. The process terminates when all the jobs of the original graph G have been partitioned into their respective layers. We assume that there are ℓ_{\max} layers in total. The entire layer partitioning process is executed in $O(m)$ time. In the sequel, without loss of generality, a DAG $G = (V, E)$ is always represented in this way, in which every job is associated with a level and L_i denotes the subset of all the level- i jobs, for $i = 1, 2, \dots, \ell_{\max}$. See Fig. 2 for an illustration.




Fig. 2. A layered representation of the precedence graph $G = (V, E)$, in which there are ℓ_{\max} layers (each as a dashed rectangle) in total, $L_1, L_2, \dots, L_{\ell_{\max}}$. U denotes the subset of all the vertices on the longest paths in G , $U_i = L_i \cap U$, for $i = 1, 2, \dots, \ell_{\max}$ (each as a dashed oval), and $S = (U, F)$ denotes the induced subgraph on U .

Lemma 2. *Given a DAG $G = (V, E)$, L_i is agreeable for every i , and a level- i job has at least one level- $(i - 1)$ parent ($i \geq 2$).*

Proof. By how the layers are constructed. □

Lemma 3. *Given a DAG $G = (V, E)$, the partition $\mathcal{C} = \{L_1, L_2, \dots, L_{\ell_{\max}}\}$ is an acyclic collection of agreeable subsets.*

Proof. By how the layers are constructed and Lemma 2, L_i precedes L_j if and only if $i < j$. □

Lemma 4. *Given a DAG $G = (V, E)$, the minimum makespan $C_{\max}^* \geq \max\{n, 3\ell_{\max}\}$.*

Proof. Since we are dealing with unit jobs, $C_{\max}^* \geq n$. Select one job v_i from L_i , for every i , such that v_i is a child of the job v_{i-1} . One clearly sees that in any feasible schedule, the job v_i starts processing after the job v_{i-1} is finished by the three-machine openshop; the makespan of the schedule is thus at least $3\ell_{\max}$. This proves the lemma. □

Theorem 1. *A schedule π can be constructed from the partition $\mathcal{C} = \{L_1, L_2, \dots, L_{\ell_{\max}}\}$ to achieve the makespan $C_{\max}^\pi \leq \frac{5}{3}C_{\max}^*$.*

Proof. Let b and c denote the number of 2-subsets and the number of singletons among $L_1, L_2, \dots, L_{\ell_{\max}}$. By Corollary 1 a schedule π can be constructed from \mathcal{C} to achieve the makespan $C_{\max}^\pi = n + b + 2c$.

Using the trivial bound $\ell_{\max} \geq b + c$ in Lemma 4, we have $C_{\max}^* \geq \max\{n, 3(b + c)\}$. It follows that

$$C_{\max}^\pi = n + b + 2c \leq C_{\max}^* + \frac{2}{3}C_{\max}^* = \frac{5}{3}C_{\max}^*.$$

This proves the theorem. □

Clearly, from the layered representation of the graph $G = (V, E)$, we see that every longest path begins with a level-1 job and ends at a level- ℓ_{\max} job, and it passes through every intermediate layer. That is, every longest path contains exactly ℓ_{\max} jobs (and $\ell_{\max} - 1$ edges). Let U denote the subset of all the jobs on the longest paths and F denote the subset of edges inherited by U (i.e., $F = E[U]$). We call $S = (U, F)$ the *spine* of the graph $G = (V, E)$, and let $H = G[V - U]$ denote the subgraph of G induced on the remaining subset $V - U$ of jobs. See Fig. 2 for an illustration.

We define a connected component in a DAG in the usual way by ignoring the direction of the edges. If the spine $S = (U, F)$ has more than one connected component, then we can safely conclude that every layer of the graph $G = (V, E)$ contains at least two jobs, that is, $|L_i| \geq 2$ for $i = 1, 2, \dots, \ell_{\max}$. Recall that our goal is to partition all the jobs into acyclic agreeable subsets to minimize the number of singletons. We call such partitions the *optimal partitions* or *optimal collections* of acyclic agreeable subsets. We assume in the rest of the paper that

the spine $S = (U, F)$ of the input graph $G = (V, E)$ is connected and there are singleton layers in $S = (U, F)$, as otherwise we trivially achieve an optimal partition without any singletons. Let U_i denote the subset of level- i jobs of U , for $i = 1, 2, \dots, \ell_{\max}$. If $|U_i| = 1$, then the job of U_i , denoted as s_i , is called a singleton job of U .

Lemma 5. *Given a DAG $G = (V, E)$ and its spine $S = (U, F)$, any acyclic partition of agreeable subsets contains at least ℓ_{\max} subsets.*

Proof. Select one job u_i from U_i , for every i , such that u_i is a child of the job u_{i-1} . (For example, these can be the jobs on a single longest path.) One clearly sees that in acyclic partition of agreeable subsets, the jobs u_i and u_j do not belong to a common subset when $i \neq j$. This suggests there are at least ℓ_{\max} subsets in the partition. This proves the lemma. \square

Lemma 6. *Given a DAG $G = (V, E)$ and its spine $S = (U, F)$, a singleton job of U cannot be processed concurrently with any other job of U in any feasible schedule.*

Proof. Because the singleton job is not agreeable with any other job of U . \square

Assume there are in total k singleton jobs in U , which are $s_{i_1}, s_{i_2}, \dots, s_{i_k}$, where $s_{i_j} \in U_{i_j}$ (that is, $|U_{i_j}| = 1$) and $1 \leq i_1 < i_2 < \dots < i_k \leq \ell_{\max}$. Let v_i be a level- i job outside of U , i.e., $v_i \in L_i - U_i$. If $i > i_j$ and v_i is agreeable with s_{i_j} , then none of the jobs of U_{i-1} can be a parent of v_i ; it follows from Lemma 2 that v_i has a parent $v_{i-1} \in L_{i-1} - U_{i-1}$. When $i - 1 > i_j$, v_{i-1} must also be agreeable with s_{i_j} , and we may repeat the above argument to conclude that there is a job v_{i_j} of $L_{i_j} - U_{i_j}$ which is a predecessor of v_i . Since both s_{i_j} and v_{i_j} are in L_{i_j} , they are agreeable (Lemma 2). We thus have proved the following lemma.

Lemma 7. *Given a DAG $G = (V, E)$ and its spine $S = (U, F)$, for a singleton job $s_{i_j} \in U$ if there is a job of $V - U$ agreeable with s_{i_j} , then there is a level- i job of $L_i - U_i$ with $i \geq i_j$ which is agreeable with s_{i_j} .*

3 A 4/3-Approximation for O3 | prec, p_{ij} = 1 | C_{max}

We have shown in Theorem 1 that we can construct a schedule π from the partition $\mathcal{C} = \{L_1, L_2, \dots, L_{\ell_{\max}}\}$ to achieve the makespan $C_{\max}^{\pi} \leq \frac{5}{3}C_{\max}^*$, suggesting that the O3 | prec, p_{ij} = 1 | C_{max} problem admits a linear time 5/3-approximation. In this section, we present an improved 4/3-approximation algorithm.

3.1 Algorithm Description

Our algorithm is mostly based on the above Lemma 7, for each singleton job s_{i_j} of U , to find a job of $V - U$ which is agreeable with s_{i_j} such that they can be

processed concurrently. The algorithm is greedy and iterative, and is denoted as APPROX.

Recall that there are in total k singleton jobs in U , which are $s_{i_1}, s_{i_2}, \dots, s_{i_k}$ (that is, $U_{i_j} = \{s_{i_j}\}$), with $1 \leq i_1 < i_2 < \dots < i_k \leq \ell_{\max}$. There are $k + 1$ iterations in the algorithm APPROX, which together construct an acyclic partition $\mathcal{D} = \{D_{\ell_{\max}}, D_{\ell_{\max}-1}, \dots, D_2, D_1\}$. We initialize $R = V - U$.

In the first iteration, sequentially for $i = \ell_{\max}, \ell_{\max} - 1, \dots, i_k + 1$, we simply let $D_i = L_i$ and remove the jobs of $L_i - U_i$ from R . If $|L_{i_k}| \geq 2$, then we let $D_{i_k} = L_{i_k}$ and remove the jobs of $L_{i_k} - s_{i_k}$ from R . Otherwise, among all the jobs of R , we pick one job that is agreeable with s_{i_k} (i.e., not a predecessor of s_{i_k}) and has the maximum level. Assume this job is $v_i \in L_i - U_i$ such that $i > i_k$. We let $D_{i_k} = \{s_{i_k}, v_i\}$ and remove the job v_i from R . If no job of R is agreeable with s_{i_k} , then we let $D_{i_k} = \{s_{i_k}\}$ and say that s_{i_k} remains as a singleton job in the partition \mathcal{D} . This ends the iteration.

In general, in the j -th iteration ($j = 2, 3, \dots, k$), sequentially for $i = i_{k+2-j} - 1, i_{k+2-j} - 2, \dots, i_{k+1-j} + 1$, we simply let $D_i = L_i$ and remove the jobs of $L_i - U_i$ from R . We remark that here the set L_i might not be the original L_i , since some of its jobs might be picked in earlier iterations and thus have been removed. Nevertheless, since $|U_i| \geq 2$, we conclude that $|D_i| \geq 2$ too. If $|L_{i_{k+1-j}}| \geq 2$, then we let $D_{i_{k+1-j}} = L_{i_{k+1-j}}$ and remove the jobs of $L_{i_{k+1-j}} - s_{i_{k+1-j}}$ from R . Otherwise, among all the jobs of R , we pick one job that is agreeable with $s_{i_{k+1-j}}$ (i.e., not a predecessor of $s_{i_{k+1-j}}$) and has the maximum level. Assume this job is $v_i \in L_i - U_i$ such that $i > i_{k+1-j}$. We let $D_{i_{k+1-j}} = \{s_{i_{k+1-j}}, v_i\}$ and remove the job v_i from R . If no job of R is agreeable with $s_{i_{k+1-j}}$, then we let $D_{i_{k+1-j}} = \{s_{i_{k+1-j}}\}$ and say that $s_{i_{k+1-j}}$ remains as a singleton job in the partition \mathcal{D} . This ends the iteration. A high-level description of such a typical iteration of the algorithm APPROX is depicted in Fig. 3.

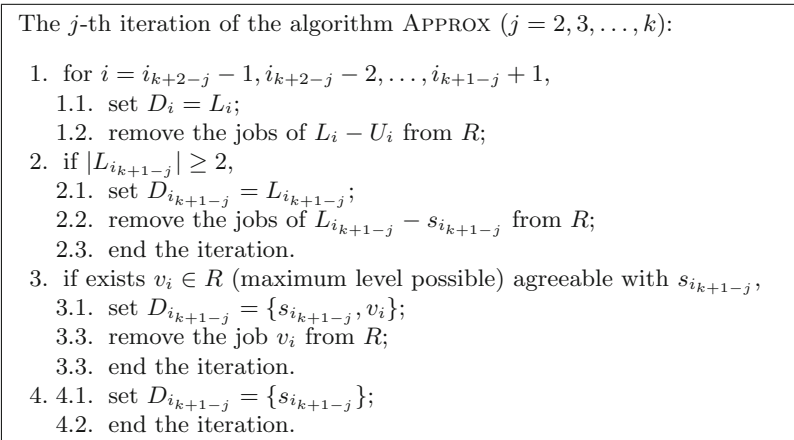


Fig. 3. A high-level description of a typical iteration of the algorithm APPROX.

In the last (the $(k + 1)$ -st) iteration, sequentially for $i = i_1 - 1, i_1 - 2, \dots, 2, 1$, we simply let $D_i = L_i$ and remove the jobs of $L_i - U_i$ from R . Again, we know that here the set L_i might not be the original L_i , since some of its jobs might be picked in earlier iterations. Nevertheless, since $|U_i| \geq 2$, we conclude that $|D_i| \geq 2$ too. This ends the last iteration and the construction of \mathcal{D} is complete. See Fig. 4 for an illustration on \mathcal{D} achieved on the graph $G = (V, E)$ shown in Fig. 2.




Fig. 4. An illustration on the acyclic partition $\mathcal{D} = \{D_{\ell_{\max}}, D_{\ell_{\max}-1}, \dots, D_2, D_1\}$ achieved on the precedence graph $G = (V, E)$ shown in Fig. 2. The ℓ_{\max} layers $L_1, L_2, \dots, L_{\ell_{\max}}$ are shown as dashed rectangles and each subset D_i is shown as a dashed oval.

3.2 Performance Analysis

The main result in this section is the following theorem.

Theorem 2. *The schedule π constructed from the partition $\mathcal{D} = \{D_1, D_2, \dots, D_{\ell_{\max}}\}$ has a makespan $C_{\max}^\pi \leq \frac{4}{3}C_{\max}^*$.*

Proof. We prove first that the partition \mathcal{D} is acyclic, in a way that D_i precedes D_{i+1} for $i = 1, 2, \dots, \ell_{\max} - 1$. Suppose to the contrary D_i precedes D_j but $i > j$; then D_i precedes D_{i-1} . Note that D_i (D_{i-1} , respectively) consists of a subset of jobs of L_i (L_{i-1} , respectively) and possibly a job v_r with a smaller level $r \leq i - 1$. It follows that $i = i_j$ for some j (that is, s_{i_j} is a singleton job of U), and v_r precedes a job of D_{i-1} , denoted as v_t of level t . Thus we have $r < t \leq i - 1$. If v_t is agreeable with s_{i_j} , then by the algorithm description v_t should be picked into D_{i_j} , a contradiction. Hence v_t precedes s_{i_j} , which implies

that v_r precedes s_{i_j} too, again a contradiction. These contradictions together prove that for any $i > j$, D_i doesn't precede D_j .

Next consider an optimal schedule π^* that achieves the minimum makespan C_{\max}^* , and assume without loss of generality that the makespan is achieved at the first machine M_1 . For a singleton job s_{i_j} of U , Lemma 6 states that it cannot be processed concurrently with any other job of U in π^* . Therefore, there are at most two distinct jobs of $V - U$, such that for each of them, when the machine M_1 is processing it, one of the other machines M_2 and M_3 is processing s_{i_j} . We say that these two jobs of $V - U$ are *associated* with the singleton job s_{i_j} . It is important to note that a job of $V - U$ associated with a singleton job cannot be associated with another singleton job, for otherwise the two singleton jobs were processed concurrently in π^* (contradicting Lemma 6).

Either there is one or two jobs of $V - U$ associated with the singleton job s_{i_j} , we pick one randomly. If the picked job has a level less than or equal to i_j , then we use t_{i_j} to denote it. If the picked job has a level greater than i_j , then we apply Lemma 7 to locate one of its predecessor jobs with level i_j and use t_{i_j} to denote this predecessor. One sees that all these t_{i_j} 's, if exist, are distinct.

If there is no job of $V - U$ associated with the singleton job s_{i_j} , we say s_{i_j} is *isolated* in π^* .

Recall that in the partition \mathcal{D} , when $i \notin \{i_1, i_2, \dots, i_k\}$, $|D_i| \geq 2$. If $|D_{i_j}| = 1$, that is, $D_{i_j} = \{s_{i_j}\}$, then we say s_{i_j} is *isolated* in the schedule π constructed from \mathcal{D} . We prove in the following the most important property that the number of isolated jobs in π is not greater than the number of isolated jobs in π^* (though the two meanings of "isolated" are different).

Assume s_{i_j} is isolated in π . We find a path from s_{i_j} to an isolated job in π^* as follows: If s_{i_j} is isolated in π^* , then the path has length 0. If s_{i_j} is not isolated in π^* , that is, we have a job t_{i_j} associated with s_{i_j} , then t_{i_j} should have been picked by the algorithm APPROX in an earlier iteration, since otherwise in this $(k + 1 - j)$ -th iteration the singleton job s_{i_j} wouldn't be left alone in the set D_{i_j} . Therefore, we identify another singleton job $s_{i_{j'}}$, where $j' > j$, which is not isolated in π because in the $(k + 1 - j')$ -th iteration the algorithm APPROX picked up t_{i_j} to accompany the singleton job $s_{i_{j'}}$. Our path extends from s_{i_j} to $s_{i_{j'}}$. If $s_{i_{j'}}$ happens to be isolated in π^* , then our path ends; otherwise, we continue to use its associated job $t_{i_{j'}}$ to locate a third singleton job $s_{i_{j''}}$, where $j'' > j$ too, which is not isolated in π , and our path extends to $s_{i_{j''}}$. Due to the finitely many singleton jobs, our path ends at a singleton job $s_{i_{j^*}}$, which is isolated in π^* .

One sees that we have used the associated jobs t_{i_j} 's, which are distinct from each other, to locate an isolated job in π^* for each isolated job in π . Therefore, an isolated job in π^* wouldn't be discovered by multiple isolated jobs in π . In other words, the number of isolated jobs in π is not greater than the number of isolated jobs in π^* , denoted as c^* . Suppose there are b 2-subsets and c singletons in the partition \mathcal{D} ; then there are c isolated jobs in π . We have

$$c \leq c^*. \tag{1}$$

In the optimal schedule π^* , the machine M_1 processes nothing while each of the other two machines is processing an isolated job. That is, the machine M_1 idles for at least $2c^*$ units of time before the makespan. Since the load of M_1 is n , we have

$$C_{\max}^* \geq n + 2c^*. \tag{2}$$

On the other hand, we still have $\ell_{\max} \geq b + c$ and $C_{\max}^* \geq 3\ell_{\max}$; therefore,

$$C_{\max}^* \geq \max\{n + 2c^*, 3(b + c)\}, \tag{3}$$

which is a better lower bound than the one in Lemma 4. It follows that

$$C_{\max}^\pi = n + b + 2c = (n + 2c) + b \leq C_{\max}^* + \frac{1}{3}C_{\max}^* = \frac{4}{3}C_{\max}^*.$$

This proves that the performance ratio for the algorithm APPROX is $4/3$.

For the running time, the algorithm APPROX maintains the precedence relationships and updates the subsets L_i 's for constructing the partition \mathcal{D} . The most time is spent for locating an agreeable job for accompanying a singleton job of U , which might take $O(n)$ time. Therefore, it is safe to conclude that the total running time of the algorithm APPROX is $O(n^2)$. This finishes the proof of the theorem. \square

Corollary 2. *The problem $Om \mid p_{ij} = 1, prec \mid C_{\max}$ admits an $O(n^2)$ -time $(2 - \frac{2}{m})$ -approximation algorithm.*

Proof. Basically we can construct from the acyclic partition \mathcal{D} a schedule with makespan $C_{\max} \leq n + (m - 2)b + (m - 1)c$. While the lower bounds in Eq. (3) are updated as $C_{\max}^* \geq \max\{n + (m - 1)c^*, m(b + c)\}$. Since we still have $c \leq c^*$, these two inequalities imply that $C_{\max} \leq (1 + (m - 2)/m)C_{\max}^* = (2 - \frac{2}{m})C_{\max}^*$. \square

4 Concluding Remarks

We studied the open-shop scheduling problem for unit jobs under precedence constraints. The problem has been shown to be strongly NP-hard when the number of machines is part of the input [14], but left as an open problem when the number m of machines is a fixed constant greater than 2, since 1978 [8]. We approached this problem by proposing a $(2 - \frac{2}{m})$ -approximation algorithm, for $m \geq 3$. Addressing the complexity and designing better approximations are both challenging and exciting.

Acknowledgements. This research is partially supported by the NSFC Grants 11571252, 11771114 and 61672323, the China Scholarship Council Grant 201508330054, and the NSERC Canada.

References

1. Bräsel, H., Kluge, D., Werner, F.: A polynomial algorithm for the $[n/m/0, t_{ij} = 1, tree/C_{\max}]$ open shop problem. *Eur. J. Oper. Res.* **72**, 125–134 (1994)
2. Brucker, P.: *Scheduling Algorithms*. Springer, New York (2007). <https://doi.org/10.1007/978-3-540-69516-5>
3. Brucker, P., Jurisch, B., Jurisch, M.Z.: Open shop problems with unit time operations. *Oper. Res.* **37**, 59–73 (1993)
4. Brucker, P., Knust, S.: Complexity results for scheduling problems (2009). <http://www2.informatik.uni-osnabrueck.de/knust/class/>
5. Dürr, C.: *The scheduling zoo* (2016). <http://schedulingzoo.lip6.fr>
6. Gonzalez, T., Sahni, S.: Open shop scheduling to minimize finish time. *J. ACM* **23**, 665–679 (1976)
7. Graham, R.L.: Combinatorial scheduling theory. In: Steen, L.A. (ed.) *Mathematics Today Twelve Informal Essays*. Springer, New York (1978). https://doi.org/10.1007/978-1-4613-9435-8_8
8. Lenstra, J.K., Rinnooy Kan, A.H.G.: Complexity of scheduling under precedence constraints. *Oper. Res.* **26**, 22–35 (1978)
9. Pinedo, M.L.: *Scheduling: Theory, Algorithm and Systems*. Springer, New York (2016). <https://doi.org/10.1007/978-3-319-26580-3>
10. Prot, D., Bellenguez-Morinea, O.: A survey on how the structure of precedence constraints may change the complexity class of scheduling problems. *J. Sched.* **21**, 3–16 (2018)
11. Sevastianov, S.V., Woeginger, G.J.: Makespan minimization in open shops: a polynomial time approximation scheme. *Math. Program.* **82**, 191–198 (1998)
12. Sevastianov, S.V., Woeginger, G.J.: Linear time approximation scheme for the multiprocessor open shop problem. *Discrete Appl. Math.* **114**, 273–288 (2001)
13. Tanaev, V.S., Sotskov, Y.N., Strusevich, V.A.: *Scheduling Theory: Multi-Stage Systems*. Springer, Heidelberg (1994). <https://doi.org/10.1007/978-94-011-1192-8>
14. Timkovsky, V.G.: Identical parallel machines vs. unit-time shops and preemptions vs. chains in scheduling complexity. *Eur. J. Oper. Res.* **149**, 355–376 (2003)
15. Williamson, D.P., et al.: Short shop schedules. *Oper. Res.* **45**, 288–294 (1997)