



Preemptive scheduling with variable profile, precedence constraints and due dates

Zhen Liu, Eric Sanlaville

► To cite this version:

Zhen Liu, Eric Sanlaville. Preemptive scheduling with variable profile, precedence constraints and due dates. [Research Report] RR-1622, INRIA. 1992. inria-00074939

HAL Id: inria-00074939

<https://inria.hal.science/inria-00074939>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

IRIA

UNITÉ DE RECHERCHE
IRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Domaine de Voluceau
Rocquencourt
B.P.105
78153 Le Chesnay Cedex
France
Tél.: (1) 39 63 55 11

Rapports de Recherche

1 9 9 2



ème

anniversaire

N° 1622

Programme 1

*Architectures parallèles, Bases de données,
Réseaux et Systèmes distribués*

PREEMPTIVE SCHEDULING WITH VARIABLE PROFILE, PRECEDENCE CONSTRAINTS AND DUE DATES

**Zhen LIU
Eric SANLAVILLE**

Février 1992



★ R R . 1 6 2 2 ★

Ordonnancement préemptif avec profil variable, contraintes de précédence et dates d'échéance

Zhen LIU

Eric SANLAVILLE

INRIA
Centre Sophia Antipolis
2004 Route des Lucioles
BP 109, 06561 Valbonne
FRANCE

Laboratoire MASI
Université Pierre et Marie Curie
4, place Jussieu
75252 Paris Cedex 05
FRANCE

28 Janvier, 1992

Résumé

Dans cet article nous étudions le problème de l'ordonnancement préemptif de tâches liées par des contraintes de précédence. Nous cherchons à minimiser soit le retard maximal, soit la durée totale d'exécution des tâches. Le nombre de processeurs parallèles disponibles varie avec le temps. Pour la minimisation du retard maximal, nous considérons le cas où des algorithmes de liste du type "Première Echéance d'Abord" sont optimaux parmi les ordonnancements non préemptifs, pour des tâches de durée unité. Nous montrons alors que l'algorithme préemptif "Plus Petite Laxité d'Abord" est optimal parmi les ordonnancements préemptifs de tâches de durées quelconques. Concernant la minimisation de la durée totale, ce résultat implique que si certains algorithmes de liste du type "Plus Haut Niveau d'Abord" sont optimaux, alors l'algorithme "Plus Long Chemin Résiduel d'Abord" est optimal. Ces résultats permettent de résoudre quatre problèmes spécifiques d'ordonnancement avec profil variable.

Keywords: ordonnancement préemptif, ordonnancement de liste, ordonnancement de priorité, profil variable, contraintes de précédence, échéance, retard, durée totale.

Preemptive Scheduling with Variable Profile, Precedence Constraints and Due Dates

Zhen LIU

Eric SANLAVILLE

INRIA
Centre Sophia Antipolis
2004 Route des Lucioles
BP 109, 06561 Valbonne
FRANCE

Laboratoire MASI
Université Pierre et Marie Curie
4, place Jussieu
75252 Paris Cedex 05
FRANCE

January 28, 1992

Abstract

This paper is concerned with the problem of scheduling preemptive tasks subject to precedence constraints in order to minimize the maximum lateness and the makespan. The number of available parallel processors is allowed to vary in time. It is shown that when an Earliest Due Date first algorithm provides an optimal nonpreemptive schedule for Unit Execution Time tasks, then the preemptive priority scheduling algorithm, referred to as Smallest Laxity First, provides an optimal preemptive schedule for real-execution-time tasks. When the objective is to minimize the makespan, we get the same kind of result between Highest Level First schedules solving nonpreemptive tasks with Unit Execution Time and the Longest Remaining Path first schedule for the corresponding preemptive scheduling problem with real-execution-time tasks. These results are applied to four specific profile scheduling problems and new optimality results are obtained.

Keywords: Preemptive Scheduling, List Schedule, Priority Schedule, Variable Profile, Precedence Constraints, Due Dates, Lateness, Makespan.

1 Introduction

We consider the preemptive profile scheduling of partially ordered tasks. The tasks are modeled by a directed acyclic graph where vertices represent tasks and arcs represent precedence relations between the tasks. These tasks are executed, subject to precedence constraints, on some identical parallel processors. The number of processors available to these tasks, referred to as the profile, may vary with time. The executions may be preempted and resumed without any penalty. At any time, a task can be assigned to at most one processor, and a processor can execute at most one task. Each task is associated with a due date. The objective is to minimize the maximum lateness and the makespan (when due dates are not taken into consideration).

The preemptive and nonpreemptive schedulings of partially ordered tasks on multiprocessor systems have been studied by various authors in the literature (see [9] for an extensive survey). The minimizations of the maximum lateness and the makespan are NP-hard problems in general. Polynomial algorithms exist only in some specific cases of the task graphs or the number of processors. List scheduling algorithms form an important class of nonpreemptive scheduling algorithms. They provide optimal solutions for some particular cases (see [2] for studies of their properties). In preemptive scheduling, there is a corresponding class of algorithms, referred to as the priority scheduling algorithms. For three particular scheduling problems, Lawler [8] constructed preemptive priority algorithms, based on the optimal nonpreemptive list algorithms in the case of unit-execution-time (UET) tasks, and showed that these priority algorithms provide optimal preemptive schedules for real-execution-time (RET) tasks.

The notion of profile scheduling was first introduced by Ullman [15] and later by Garey et al. [7] in the complexity analysis of deterministic scheduling algorithms. Dolev and Warmuth [4,5,6] carried out various studies on the nonpreemptive profile scheduling with UET tasks. They obtained polynomial algorithms that minimizes the makespan for specific profiles (e.g., zigzag profile, bounded profile, etc.) and specific task graphs (e.g., in-forest, out-forest, opposing forest, flat graph, etc.). Liu and Sanlaville [10] analyzed the stochastic preemptive scheduling problems. Simple optimal schedules were provided for the stochastic minimization of the makespan of interval-order task graphs, in-forests, and uniform out-forests.

In this paper, we investigate the preemptive version of the deterministic profile scheduling problems. We show that if some list scheduling algorithms of the type Earliest Due Date (EDD) are optimal within the class of nonpreemptive schedules for UET tasks, then the preemptive priority scheduling algorithm, referred to as Smallest Laxity First (SLF), is optimal within the class of preemptive schedules for RET tasks. When the minimization of makespan is under consideration, such a result implies that if list scheduling algorithms of the type Highest Level First (HLF) are optimal within the class of nonpreemptive schedules for UET tasks, then the preemptive priority scheduling algorithm, referred to as Longest Remaining Path (LRP), is optimal within the class of preemptive schedules for RET tasks.

These results are applied to four specific profile scheduling problems and following new preemptive profile scheduling results are obtained:

- SLF (defined on some modified due dates) minimizes the maximum lateness (with respect to the original due dates) when the task graph is an in-forest and the profile is increasing zigzag;
- LRP minimizes the makespan when the task graph is a union of chains;
- LRP minimizes the makespan when the task graph is an in-forest and the profile is increasing zigzag, or when the task graph is an out-forest and the profile is decreasing zigzag;
- LRP minimizes the makespan when the task graph is arbitrary and the profile is bounded by two.

Our results strengthen the relationship between the optimality of the nonpreemptive list algorithms for UET tasks and the optimality of the preemptive priority algorithms for RET tasks. However, our approach is different from that of Lawler [8]. We establish the optimality of the priority algorithms by using directly the optimality of the corresponding list algorithms, whereas in [8], different proofs were necessary in order to obtain the optimality of the priority algorithms for solving different scheduling problems.

The paper is organized as follows. In Section 2, we define the notation and we present some preliminary results on the list and priority schedules and on the graph expansions introduced in the paper. In Section 3, we prove the main results of the paper which relate the optimality of the preemptive SLF (resp. LRP) schedule to those of the nonpreemptive EDD (resp. HLF) schedules. In Section 4, we apply the main theorems to the four profile scheduling problems and we obtain optimal preemptive schedules. When necessary, the nonpreemptive counterparts are first studied. In particular, we extend the optimality of EDD schedules (defined on modified due dates), obtained by Brucker, Garey and Johnson [1] for in-forests with UET tasks and constant profiles, to the increasing zigzag profiles. We also prove that HLF schedules are optimal for the minimization of makespan of chains with UET tasks and arbitrary profiles within the class of nonpreemptive schedules.

2 Preliminaries

2.1 Problem Description

There are n tasks to be processed by a multiprocessor system. The executions of these tasks are constrained by some precedence relations between the tasks. A task graph $G = (V, E)$ is used to describe these relations, where $V = \{1, 2, \dots, n\}$ is the set of vertices representing the tasks, and E is the set of arcs representing the precedence relations between tasks. It is assumed that

G is a directed acyclic graph (dag) and that it contains no transitivity arcs. Denote by $\pi(i)$ and $\sigma(i)$ the sets of immediate predecessors and successors of task i , respectively. Let $\pi^*(i)$ and $\sigma^*(i)$ be the sets of (not necessarily immediate) predecessors and successors of i , respectively. A task without successor (resp. predecessor) is called a final (resp. an initial) task. Task $i \in V$ has a processing requirement p_i and a due date d_i .

There are $m \geq 1$ identical parallel processors in the system with speed 1 (so that the processing time of a task equals its processing requirement). The number of processors available to the execution of task graph G varies with time. Define $M = \{a_r, m_r\}_{r=1}^\infty$ as the profile of the system, where $0 = a_1 < a_2 < \dots < a_r < \dots$ are the epochs when the number of available processors changes, and m_r is the number of available processors during $[a_r, a_{r+1})$. Without loss of generality, we assume that $m_r \geq 1$ for all $r = 1, 2, \dots$. Since the processors are identical, we can assume that processor 1 is always available. We will also assume that the profile is not changed infinitely often during any finite time interval. Under these assumptions, there is a finite \bar{r} such that $a_{\bar{r}} > \sum_{i \in V} p_i$. Without loss of generality, we can assume that there is at least one task running at any time unless all the tasks have finished. Thus, we will only consider the truncated sequence $M = \{a_r, m_r\}_{r=1}^{\bar{r}}$. A special case of the variable profile is the constant profile where $m_1 = m$ and $a_2 = \infty$.

A scheduling algorithm decides when an enabled task, i.e., an unassigned unfinished task all of whose predecessors have finished, should be assigned to one of the available processors. At any time, a task can be assigned to at most one processor, and a processor can execute at most one task. A schedule is feasible if these constraints (i.e., the precedence relation, the variable profile, the nonredundancy of the task assignment) are verified. Scheduling can be either preemptive, i.e. the execution of a task can be stopped and later resumed on any processor without penalty, or nonpreemptive: once begun, the execution of a task continues on the same processor until its completion.

Let S be an arbitrary feasible schedule of task graph G under profile M . Let $C_i(S)$ be the completion time of task i under S . The lateness of task i is defined as $L_i(S) = C_i(S) - d_i$, and the maximum lateness of schedule S as $L_S(G, M) = \max_{i \in V} L_i(S)$. Denote by $L_p^*(G, M)$ and $L_{np}^*(G, M)$ the smallest maximum latenesses of task graph G obtained by the nonpreemptive schedules and preemptive schedules, respectively, under profile M . Since nonpreemptive schedules are special cases of preemptive ones, it is trivial that $L_p^*(G, M) \leq L_{np}^*(G, M)$.

When the due dates are set to zero, the maximum lateness becomes the makespan. Denote by $C_S(G, M) = \max_{i \in V} C_i(S)$ the makespan of (G, M) obtained by schedule S . Let $C_p^*(G, M)$ and $C_{np}^*(G, M)$ be the smallest makespans of task graph G obtained by the nonpreemptive schedules and preemptive schedules, respectively, under profile M .

The goal of this paper is to find preemptive feasible schedules that minimize the maximum lateness and the makespan. More precisely, we will establish a relation between some optimal

nonpreemptive schedules and some optimal preemptive schedules.

2.2 List and Priority Schedules

List algorithms are often used in nonpreemptive scheduling. With such algorithms, there is a (static or dynamic) list of tasks. As soon as a processor is available, the enabled task that is closest to the head of the list is assigned to that processor. The *Earliest Due Date* first (EDD) algorithms form a well known subclass of list algorithms, where the tasks are ordered increasingly by their due dates. The schedules generated by EDD algorithms differ in the way that ties are broken. Let $\mathcal{E}(G, M)$ denote the family of schedules obtained by the EDD algorithms for task graph G and profile M .

In accordance with the list algorithms, (dynamic) priority algorithms are used in the preemptive scheduling. At any time, enabled tasks are assigned to available processors according to a priority list which may change in time and may depend on the partial schedule already constructed. A general description is given below (cf. [8,13]):

- At any time t , enabled tasks are ordered according to their priorities, thus forming subsets V_1, \dots, V_k , where all tasks of V_j have the same priority and greater priority than tasks in V_{j+1} .
- Suppose that tasks in V_1, \dots, V_{r-1} , $r \leq k$, are assigned. Let $\tilde{m}_r(t)$ be the number of remaining free processors. If $\tilde{m}_r(t) \geq |V_r|$, then one processor is assigned to each of the tasks in V_r , and the algorithm deals with the next subset. Otherwise, the $\tilde{m}_r(t)$ processors are shared by the tasks of V_r so that each task in V_r is executed at speed $v_r = \tilde{m}_r(t)/|V_r|$.
- This assignment remains unchanged until one of the following events occurs:
 1. A task completes;
 2. The priority order of the tasks is changed;
 3. The profile changes.

At such moments the processor assignment is re-computed.

In the above scheme, the processor sharing can be achieved by McNaughton's wrap-around algorithm (cf. [11]) which is linear in the number of tasks scheduled in each time interval. An example is illustrated in Figure 1 where three tasks are executed at speed $2/3$ on two processors during a unit length interval.

Note that other processor sharing schemes can be used. However, they will generate the same latenesses of the tasks provided the corresponding task processing speeds are the same in these processor sharing schemes. Therefore, we will not make any difference between them.




Figure 1: An example of processor sharing with McNaughton's algorithm

Denote by $p_i^S(t)$ the remaining processing requirement of task i at time t in a schedule S . Define the laxity of task i at time t in this schedule to be $l_i^S(t) = d_i - p_i^S(t)$. The dynamic priority algorithm based on the *Smallest Laxity First* rule is referred to as *SLF*. The schedule it produces for (G, M) is denoted by $SLF(G, M)$.

List and priority schedules are also used for the minimization of makespan. Some simple algorithms are optimal under certain conditions. We will consider the *Highest Level First* (HLF) and the *Longest Remaining Path* (LRP) schedules.

Let h_i be the height of task i , defined as the length of the longest path between i and a final task in G . This length is computed as the summation of the processing requirements of the tasks (excluding i) in the path. Note that in the case of Unit-Execution-Time (UET) tasks (i.e., $p_1 = \dots = p_n = 1$), h_i is also referred to as the level of task i , where, by convention, the level of a final task is 0.

In a Highest Level First (HLF) list algorithm, the tasks are ordered decreasingly by their heights or levels. Let $\mathcal{H}(G, M)$ denote the family of schedules obtained by the HLF algorithms for task graph G and profile M . As in the case of EDD schedules, the schedules in $\mathcal{H}(G, M)$ differ in the way that ties are broken.

In the preemptive case, define the length of the remaining longest path at time t in a preemptive schedule S to be $r_i^S(t) = h_i + p_i^S(t)$. In a Longest Remaining Path first (LRP) schedule, the tasks are decreasingly ordered by the lengths of their remaining longest path.

It follows from the above definitions that

Lemma 2.1 *For any couple (G, M) , if $d_i = -h_i$ for all $i \in V$, then the EDD (resp. SLF) rule coincides with the HLF (resp. LRP) rule.*

2.3 Commensurability and Graph Expansion

In the paper we will deal with commensurable real numbers. The real numbers $x_1, \dots, x_r \in \mathbb{R}$ are said to be mutually commensurable if there exist $w \in \mathbb{R}$ and $\alpha_1, \dots, \alpha_r \in \mathbb{N}^+ = \{0, 1, 2, \dots\}$

such that $x_i = \alpha_i w$ for all $i = 1, \dots, r$. The real number w is called a commensurability factor of x_1, \dots, x_r . Note that if w is a commensurability factor of x_1, \dots, x_r , then for all $k \in \mathbb{N}^* = \{1, 2, \dots\}$, w/k is also a commensurability factor of x_1, \dots, x_r .

A task graph G is said to have commensurable timing with commensurability factor w if the processing times and the due dates of the tasks of G are mutually commensurable with commensurability factor w . The couple (G, M) is said to have commensurable timing with commensurability factor w if the task processing times, the due dates and the profile changing epochs are mutually commensurable with commensurability factor w . In what follows, w always denotes the commensurability factor. A special case of commensurable timing is the UET tasks and integer due dates.

Let $G = (V, E)$ have commensurable timing with commensurability factor w :

$$\forall i \in V, \exists \alpha_i, \beta_i \in \mathbb{N}^+ : \quad p_i = \alpha_i w, \quad d_i = \beta_i w.$$

For all $k \in \mathbb{N}^*$, we define an operation on G , called the k -th expansion of G , as follows: The expanded task graph $G_{w/k} = (V_{w/k}, E_{w/k})$ is obtained by replacing each vertex i in G with a chain of $\alpha_i k$ vertices such that

$$\begin{aligned} V_{w/k} &= \{i_j \mid i \in V, j = 1, \dots, \alpha_i k\}; \\ E_{w/k} &= \{(i_j, i_{j+1}) \mid i \in V, j = 1, \dots, \alpha_i k - 1\} \cup \{(i_{\alpha_i k}, i'_1) \mid (i, i') \in E\}; \\ p_{i_j} &= w/k, \quad j = 1, \dots, \alpha_i k; \\ d_{i_j} &= d_i - (\alpha_i k - j)w/k, \quad j = 1, \dots, \alpha_i k. \end{aligned}$$

For sake of simplicity, $G_{w/1}$ will be denoted by G_w .

Let S be a feasible preemptive schedule on (G, M) with commensurable timing. Construct another preemptive schedule S_w on (G_w, M) as follows: $\forall i \in V$, S_w executes one subtask i_j at speed v at time t if and only if S executes the j -th portion (with duration w) of task i at speed v at time t , $j = 1, \dots, \alpha_i$. It is easily seen that S_w is a feasible schedule of (G_w, M) . The schedule S_w is referred to as the expansion of schedule S .

Lemma 2.2 Assume that (G, M) has commensurable timing with commensurability factor w . Then

$$L_S(G, M) = L_{S_w}(G_w, M)$$

holds for all feasible preemptive schedule S .

Proof. By definition of S_w , $C_{i_{\alpha_i}}(S_w) = C_i(S)$ for all $i \in V$. Hence,

$$L_{S_w}(G_w, M) = \max_{i \in V} \max_{1 \leq j \leq \alpha_i} (C_{i_j}(S_w) - d_{i_j}) \geq \max_{i \in V} (C_{i_{\alpha_i}}(S_w) - d_{i_{\alpha_i}}) = \max_{i \in V} (C_i(S) - d_i) = L_S(G, M).$$

On the other hand, for all $1 \leq j \leq \alpha_i - 1$,

$$C_{i,j}(S_w) - d_{i,j} \leq C_i(S) - w(\alpha_i - j) - d_{i,j} = C_i(S) - d_{i,j}.$$

Hence the result. ■

Lemma 2.3 *Assume that (G, M) has commensurable timing with commensurability factor w . Then*

$$L_p^*(G, M) = L_p^*(G_w, M).$$

Proof. Let U and U_w be the optimal preemptive schedules of (G, M) and (G_w, M) , respectively. Applying Lemma 2.2 to U implies that there is a feasible preemptive schedule U'_w of (G_w, M) such that

$$L_p^*(G, M) = L_U(G, M) = L_{U'_w}(G_w, M) \geq L_p^*(G_w, M).$$

Let now U' be a preemptive schedule of (G, M) which is constructed from U_w in such a way that U' executes the j -th portion (with duration w) of task i at speed v at time t if and only if U_w executes subtask i_j at speed v at time t , $j = 1, \dots, \alpha_i$, $i \in V$. It is easy to verify that U' is a feasible preemptive schedule of (G, M) and that $C_i(U') = C_{i,\alpha_i}(U_w)$, $\forall i \in V$. As $d_{i,\alpha_i} = d_i$, $\forall i \in V$, we get

$$L_p^*(G_w, M) = L_{U_w}(G_w, M) \geq L_{U'}(G, M) \geq L_p^*(G, M).$$

Therefore,

$$L_p^*(G, M) = L_p^*(G_w, M). \quad \blacksquare$$

Now we specifically consider *SLF* schedules.

Lemma 2.4 *Assume that (G, M) has commensurable timing with commensurability factor w . Let SLF_w be the expanded schedule issued from $SLF(G, M)$. Then SLF_w is identical to the schedule obtained by directly applying the *SLF* rule to (G_w, M) :*

$$SLF_w \equiv SLF(G_w, M).$$

Proof. The proof is by induction on the events in SLF_w . Let $0 = t_0 < t_1 < t_2 < \dots$ be the time epochs when the assignment decisions of SLF_w are made. By definition of the expanded schedule,

$$\forall i \in V, \quad \forall t \in \mathbb{R}^+ : \quad p_i^{SLF}(t) = \sum_{j=1}^{\alpha_i} p_{i,j}^{SLF_w}(t) \quad (1)$$

Let $SLF' = SLF(G_w, M)$. We will prove that:

$$p_i^{SLF}(t_s) = \sum_{j=1}^{\alpha_i} p_{i_j}^{SLF'}(t_s), \quad i \in V, \quad s = 0, 1, 2, \dots, \quad (2)$$

which, together with equation (1), will imply the assertion of the lemma.

It is trivial that equation (2) holds when $s = 0$ as $p_i(0) = p_i = \sum_{j=1}^{\alpha_i} w$. Suppose it is true for some $s \geq 0$. Let i be an enabled task at time t_s under SLF . By induction hypothesis, there is a task i_j of G_w which is enabled at t_s under SLF' . Using (1) one gets $p_i^{SLF}(t_s) = p_{i_j}^{SLF'}(t_s) + w(\alpha_i - j)$ so that

$$l_{i_j}^{SLF'}(t_s) = d_i - w(\alpha_i - j) - p_{i_j}^{SLF'}(t_s) = d_i - p_i^{SLF}(t_s) = l_i^{SLF}(t_s).$$

In words, the laxity of i in $SLF(G, M)$ at time t_s is the same as that of i_j in $SLF'(G_w, M)$ at time t_s . If i' is enabled at t_s in $SLF(G, M)$, and i'_j enabled at t_s in SLF' , then i has greater priority than i' in $SLF(G, M)$ if and only if i_j has greater priority than i'_j in SLF' . Hence i is scheduled at the same speed in $SLF(G, M)$ as i_j in SLF' during $[t_s, t_{s+1})$, so the equality remains true at t_{s+1} . By induction, (2) holds for all $s = 0, 1, 2, \dots$. ■

3 Main Results

In this section we will establish the tight relation between the optimality of the EDD algorithms in the nonpreemptive case with UET tasks and that of the SLF algorithm in the preemptive case with Commensurable-Execution-Time (CET) tasks.

Theorem 3.1 *Assume that (G, M) has commensurable timing with commensurability factor w . If*

$$\forall k \in \mathbb{N}^*, \quad \exists S \in \mathcal{E}(G_{w/k}, M) : \quad L_S(G_{w/k}, M) = L_{np}^*(G_{w/k}, M),$$

then $SLF(G, M)$ is an optimal preemptive schedule:

$$L_{SLF}(G, M) = L_p^*(G, M).$$

In other words, if for any expansion of G , there is an EDD list algorithm which minimizes the maximum lateness within the class of nonpreemptive schedules, then the priority schedule SLF is optimal within the class of preemptive schedules.

The proof of this theorem is somewhat tedious and is forwarded to Appendix A.

Let us consider the problem of makespan minimization. Consider first the following lemma.

Lemma 3.1 *For every feasible (preemptive or nonpreemptive) schedule S for (G, M) , if $d_i = -h_i$ for all $i \in V$, then*

$$L_S(G, M) = C_S(G, M).$$

Proof. Let $\tilde{V} \subseteq V$ be the set of tasks whose latenesses are equal to the maximum lateness:

$$\tilde{V} = \{i \in V \mid L_S(G, M) = C_i(S) - d_i = C_i(S) + h_i\}.$$

Let k be a task in \tilde{V} with the smallest height: $h_k \leq h_i, \forall i \in \tilde{V}$.

If k is not a final task, then there is $j \in \sigma(k)$ such that $h_k = h_j + p_j$. As $C_j(S) \geq C_k(S) + p_j$, it follows that $C_j(S) + h_j \geq C_k(S) + h_k = L_S(G, M)$, which implies that $j \in \tilde{V}$. This last fact contradicts the assumption that task k has the smallest height within \tilde{V} . Therefore, k is necessarily a final task.

Thus, $h_k = 0$ so that $L_S(G, M) = C_k(S) \leq C_S(G, M)$. As trivially

$$L_S(G, M) = \max_{i \in V} (C_i(S) + h_i) \geq \max_{i \in V} C_i(S) = C_S(G, M),$$

the lemma is proved. ■

Theorem 3.2 *Assume that (G, M) has commensurable timing with commensurability factor w . If*

$$\forall k \in IN^*, \quad \exists S \in \mathcal{H}(G_{w/k}, M) : \quad C_S(G_{w/k}, M) = C_{np}^*(G_{w/k}, M),$$

then $LRP(G, M)$ is an optimal preemptive schedule:

$$C_{LRP}(G, M) = C_p^*(G, M).$$

Proof. Setting the due dates of G in such a way that $d_i = -h_i$ for all $i \in V$. It then follows from Lemma 2.1 that an HLF (resp. LRP) schedule coincides with an EDD (resp. SLF) schedule. Since the maximum lateness and the makespan are identical in such a case (cf. Lemma 3.1), an application of Theorem 3.1 yields the desired result. ■

Theorem 3.2 states that if for any expansion of G , there is an HLF list algorithm which minimizes the makespan within the class of nonpreemptive schedules, then the priority schedule LRP is optimal within the class of preemptive schedules.

Remark: Theorems 3.1 and 3.2 actually hold in a more general case where the tasks are associated with release dates: a task is executable only after its release date. However, there are few applications with non-zero release dates so that they will not be considered in the paper.

In the remainder of this paper, we will apply Theorems 3.1 and 3.2 to four profile scheduling problems. New optimality results of SLF and LRP schedules are obtained.

4 Applications

4.1 Maximum Lateness of In-forests

We first apply our results to the class of *in-forests* \mathcal{G}_{if} with *increasing zigzag* profiles \mathcal{M}_{iz} . A task graph $G = (V, E)$ is an in-forest, $G \in \mathcal{G}_{if}$, if $|\sigma(i)| \leq 1$ for all $i \in V$, i.e., a task has at most one successor. When G is an in-forest, the final tasks are also referred to as the roots, and the initial tasks as the leaves. A profile M is increasing zigzag, $M \in \mathcal{M}_{iz}$, if for all $r \in IV^*$, $m_r \geq \max_{1 \leq u \leq r-1} m_u - 1$.

For a given in-forest $G \in \mathcal{G}_{if}$ with processing times p_1, \dots, p_n and due dates d_1, \dots, d_n , we define an in-forest $G' \in \mathcal{G}_{if}$ such that G' has the same set of tasks, the same precedence constraints and the same processing times. The due dates in G' are modified as follows:

$$d'_i = \begin{cases} d_i, & i \text{ is a root;} \\ \min(d_i, d'_{\sigma(i)} - p_{\sigma(i)}), & \text{otherwise,} \end{cases}$$

where, with a harmless abuse of notation, $\sigma(i)$ denotes the successor of $i \in G$ when i is not a root of G .

Such a modification on the due dates has no effect on the maximum lateness:

Lemma 4.1 *Let $G \in \mathcal{G}_{if}$ be an in-forest. Then for any feasible (preemptive or nonpreemptive) schedule S ,*

$$L_S(G, M) = L_S(G', M) \stackrel{\text{def}}{=} \max_{i \in V} (C_i(S) - d'_i).$$

Proof. Consider a given schedule S on (G, M) . It is clear that $d_i \geq d'_i$ for all $i \in V$ so that

$$L_S(G, M) \leq L_S(G', M).$$

Let k be a task such that $C_k(S) - d'_k = L_S(G', M)$ and that for all $j \in \sigma^*(k)$, $C_j(S) - d'_j < L_S(G', M)$. If $d_k > d'_k$, then k has a successor $s \in \sigma(k)$ such that $d'_k = d'_s - p_s < d_k$. Thus,

$$C_s(S) - d'_s \geq C_k(S) + p_s - d'_s = C_k(S) - d'_k = L_S(G', M).$$

This contradicts the assumption that $C_j(S) - d'_j < L_S(G', M)$ for all $j \in \sigma^*(k)$. Therefore, we have necessarily $d_k = d'_k$ so that

$$L_S(G, M) \geq C_k(S) - d_k = C_k(S) - d'_k = L_S(G', M).$$

The proof is thus completed. ■

In [1], it was shown that in the case of constant profile and UET tasks, the EDD schedules defined on the modified due dates $\mathcal{E}(G', M)$ are optimal for the minimization of the maximum lateness of in-forests within the class of nonpreemptive schedules. Such a result is extended to the increasing zigzag profiles below:

Theorem 4.1 (Extension of Theorem 2 of [1]) *Let $G \in \mathcal{G}_{if}$ be an in-forest, and $M \in \mathcal{M}_{iz}$ be an increasing zigzag profile. Assume that the tasks are UET and that the profile changing epochs are integer. Then any EDD schedule $S \in \mathcal{E}(G', M)$ defined on the modified due dates minimizes the maximum lateness within the class of nonpreemptive schedules:*

$$L_S(G, M) = L_{np}^*(G, M).$$

Proof. See Appendix B. ■

In order to get optimal preemptive schedules, we have to show that the due dates modified before and after an expansion are the same. Assume that $G \in \mathcal{G}_{if}$ has commensurable timing with commensurability factor w . Denote by

- $(G')_{w/k}$: the k -th expansion of G' which has the modified due dates of G ;
- $(G_{w/k})'$: the task graph with the modified due dates of the k -th expansion $G_{w/k}$ of G .

Lemma 4.2 *Assume that $G \in \mathcal{G}_{if}$ has commensurable timing with commensurability factor w . Then for all $k \in \mathbb{N}^*$, $(G')_{w/k}$ and $(G_{w/k})'$ are isomorphic in the sense that their topologies are isomorphic and that the corresponding tasks have the same processing times and the same due dates.*

Proof. Clearly, $(G')_{w/k}$ and $(G_{w/k})'$ have the same set of tasks and the same structure as $G_{w/k}$, and the processing times of all the tasks are equal to w/k . Let the tasks in $G_{w/k}$ be indexed in such a way that $i_1, i_2, \dots, i_{\alpha_i k}$ are the subtasks of $i \in G$, where α_i is an integer, and i_j is the predecessor of i_{j+1} , $1 \leq j \leq \alpha_i k - 1$. Let d'_{i_j} be the due date of task i_j in $(G')_{w/k}$ and $(d_{i_j})'$ the due date of task i_j in $(G_{w/k})'$. We will show by induction that for all $i \in G$,

$$d'_{i_j} = (d_{i_j})', \quad 1 \leq j \leq \alpha_i k. \quad (3)$$

First, for all the roots $i \in G$, we have that $d'_{i_{\alpha_i k}} = d_i = (d_{i_{\alpha_i k}})'$, so that

$$(d_{i_j})' = d_i - (\alpha_i k - j)w/k = d'_{i_j}, \quad 1 \leq j \leq \alpha_i k.$$

Therefore, (3) holds for all i such that $\sigma(i) = \emptyset$.

Consider now task i such that (3) holds for $s \in \sigma(i)$. It then follows

$$\begin{aligned}
(d_{i_{\alpha_i k}})' &= \min(d_i, (d_{s_1})' - w/k) \\
&= \min(d_i, d'_{s_1} - w/k) \\
&= \min(d_i, d'_s - (\alpha_s k - 1)w/k - w/k) \\
&= \min(d_i, d'_s - p_s) \\
&= d'_i = d'_{i_{\alpha_i k}}.
\end{aligned}$$

Assume now that for some $2 \leq j \leq \alpha_i k$, $(d_{i_j})' = d'_{i_j}$. Then

$$\begin{aligned}
(d_{i_{j-1}})' &= \min(d_{i_{j-1}}, (d_{i_j})' - w/k) \\
&= \min(d_i - (\alpha_i k - j + 1)w/k, d'_{i_j} - w/k) \\
&= \min(d_i, d'_i) - (\alpha_i k - j + 1)w/k \\
&= d'_i - (\alpha_i k - j + 1)w/k = d'_{i_j}.
\end{aligned}$$

Therefore, by induction, we have shown that for all $1 \leq j \leq \alpha_i k$, $(d_{i_j})' = d'_{i_j}$, so that (3) holds for task i . Hence, (3) holds for all $i \in G$. ■

Theorem 4.2 *Let $G \in \mathcal{G}_{if}$ be an in-forest, and $M \in \mathcal{M}_{iz}$ an increasing zigzag profile. Assume that (G, M) has commensurable timing. Then the SLF' schedule defined on the modified due dates minimizes the maximum lateness within the class of preemptive schedules:*

$$L_{SLF'}(G, M) = L_p^*(G, M).$$

Proof. Consider the task graph G' with the modified due dates of G . Since the class of in-forests \mathcal{G}_{if} is closed under expansion (i.e., $G \in \mathcal{G}_{if}$ implies $G_{w/k} \in \mathcal{G}_{if}$ for all $k \in \mathbb{N}^*$), an application of Theorem 4.1 yields that for all $k \in \mathbb{N}^*$, and all $S \in \mathcal{E}((G_{w/k})', M)$,

$$L_S((G_{w/k})', M) = L_{np}^*((G_{w/k})', M).$$

Owing to Lemma 4.2, $(G_{w/k})'$ and $G'_{w/k}$ are isomorphic so that $\mathcal{E}((G_{w/k})', M) = \mathcal{E}(G'_{w/k}, M)$ and that for all $k \in \mathbb{N}^*$, and all $S \in \mathcal{E}(G'_{w/k}, M)$,

$$L_S(G'_{w/k}, M) = L_{np}^*(G'_{w/k}, M).$$

Applying now Theorem 3.1 implies that

$$L_{SLF}(G', M) = L_p^*(G', M).$$

The above relation together with Lemma 4.1 entail

$$L_{SLF'}(G, M) = L_{SLF}(G', M) = L_p^*(G', M) = L_p^*(G, M).$$

■

Remark: Theorem 4.2 extends Theorem 7.3 of Lawler [8] to the case of increasing zigzag variable profile. It is possible to apply Theorem 3.1 to the case of arbitrary task graph and constant profile with two processors. In such a case, a new proof of Theorem 8.3 of Lawler [8] can be obtained in the case of parallel processors.

4.2 Makespan of Chains

Consider the makespan minimization problem. We first analyze the simplest case of the task graphs: the chains. Let \mathcal{G}_{ch} be the class of unions of chains. A task graph $G \in \mathcal{G}_{ch}$ is a union of chains if for all $i \in V$, $|\sigma(i)| \leq 1$ and $|\pi(i)| \leq 1$.

Theorem 4.3 *Let $G \in \mathcal{G}_{ch}$ be a union of chains with UET tasks, and M be a profile that changes only at integer time epochs. Then every HLF schedule $S \in \mathcal{H}(G, M)$ minimizes the makespan within the class of nonpreemptive schedules:*

$$C_{HLF}(G, M) = C_{np}^*(G, M).$$

Proof. Observe first that all HLF schedules have the same makespan, for the task graph is a union of chains. In order to prove the theorem, we only need to show that there is at least one optimal HLF schedule.

Let $b_i(S)$ denote the time instant when task i is assigned for execution under a nonpreemptive schedule S (so that $C_i(S) = b_i(S) + 1$). Consider an optimal schedule S^* which yields the minimal makespan $C_{np}^*(G, M)$. If S^* is not a HLF type schedule, then there are at least two tasks i and j such that i is at a higher level than j and that i is assigned after j , v.z., $h_i > h_j$ and $b_i > b_j$.

Let $\sigma^*(i) = \{i_1, i_2, \dots, i_{h_i}\}$ and $\sigma^*(j) = \{j_1, j_2, \dots, j_{h_j}\}$ be the sets of (not necessarily immediate) successors of i and j , respectively. Denote by $i_0 = i$ and $j_0 = j$. Assume that $i_u \in \pi(i_{u+1})$, $0 \leq u \leq h_i - 1$, and that $j_v \in \pi(j_{v+1})$, $0 \leq v \leq h_j - 1$. Let k be the largest integer in $\{0, 1, \dots, h_j\}$ such that for all $0 \leq u \leq k$, $b_{i_u} > b_{j_u}$.

Construct a schedule S which differs from S^* only in the assignments of tasks i_0, i_1, \dots, i_k and j_0, j_1, \dots, j_k . In S , the assignments of S^* for tasks i_u and j_u , $0 \leq u \leq k$, are interchanged. It is easy to see that S is a feasible schedule and has the same makespan as S^* . Further, schedule S has at least one less non-HLF decision. If S is still not HLF, then we repeat this interchange procedure on S to reduce the number of non-HLF decisions. After at most $n \times n$ steps of interchange, we will finally obtain an HLF schedule which has the optimal makespan $C_{np}^*(G, M)$. ■

Remark: Theorem 4.3 still holds when the chains have non-zero release dates. The interchange argument of the proof remains valid.

Theorem 4.4 *Let $G \in \mathcal{G}_{ch}$ be a union of chains. If (G, M) has commensurable timing, then the LRP schedule minimizes the makespan within the class of preemptive schedules:*

$$C_{LRP}(G, M) = C_p^*(G, M).$$

Proof. It is clear that the class of unions of chains \mathcal{G}_{ch} is closed under expansion. An application of Theorems 3.2 and 4.3 implies the result. ■

Note that in the preemptive case, scheduling problems for a union of disjoint chains and for a set of independent tasks are equivalent.

4.3 Makespan of Forests

We next study the minimization of the makespan of forests. Apart from the class of in-forests \mathcal{G}_{if} , we will also consider the class of *out-forests* \mathcal{G}_{of} . A task graph $G = (V, E)$ is an out-forest, $G \in \mathcal{G}_{of}$, if $|\pi(i)| \leq 1$ for all $i \in V$, i.e., a task has at most one predecessor. Clearly, the class of out-forests \mathcal{G}_{of} is closed under expansion. When G is an out-forest, the initial tasks are also referred to as the roots, and the final tasks as the leaves.

The out-forests will be scheduled with the class of *decreasing zigzag* profiles \mathcal{M}_{dz} . A profile M is decreasing zigzag, $M \in \mathcal{M}_{dz}$, if for all $r \in \mathbb{N}^*$, $m_r \leq \min_{1 \leq u \leq r-1} m_u + 1$.

Consider first the nonpreemptive schedules. The following lemma is due to Dolev and Warmuth [5].

Lemma 4.3 (Theorems 5.1 and 5.2 in [5]) *Let $(G, M) \in (\mathcal{G}_{if}, \mathcal{M}_{iz}) \cup (\mathcal{G}_{of}, \mathcal{M}_{dz})$ be either an in-forest with increasing zigzag profile or an out-forest with decreasing zigzag profile. Assume that all the tasks have UET and that the profile changes at integer time epochs. Then every HLF schedule $S \in \mathcal{H}(G, M)$ minimizes the makespan within the class of nonpreemptive schedules:*

$$C_{HLF}(G, M) = C_{np}^*(G, M).$$

Theorem 4.5 *Let $(G, M) \in (\mathcal{G}_{if}, \mathcal{M}_{iz}) \cup (\mathcal{G}_{of}, \mathcal{M}_{dz})$ be either an in-forest with increasing zigzag profile or an out-forest with decreasing zigzag profile. If (G, M) has commensurable timing, then the LRP schedule minimizes the makespan within the class of preemptive schedules:*

$$C_{LRP}(G, M) = C_p^*(G, M).$$

Proof. The assertion comes from Theorem 3.2 and Lemma 4.3 together with the fact that both classes of in-forests and out-forests are closed under expansion. ■

Remark: Theorem 4.5 extends the result of Muntz and Coffman [13] to the zigzag variable profiles. The optimality of LRP schedule for the makespan minimization of in-forests with increasing zigzag profiles may also be obtained from Theorem 4.2.

4.4 Makespan of Arbitrary Task Graphs

Consider now the makespan minimization of arbitrary task graphs. In [3], Coffman and Graham proved that if there are constantly two processors and if the tasks have UET, then some special HLF schedules, referred to as the CG schedules in our paper, minimize the makespan within the class of nonpreemptive schedules. In the CG schedules, the list of tasks is determined by a lexicographical order on the sets of successors. Such a result still holds when the profile is bounded by two, i.e., $m_r \leq 2$ for all $r \in \mathcal{IN}^*$.

Lemma 4.4 (Extension of Theorem 1 in [3]) *Let G be an arbitrary task graph with UET tasks, and M be a profile which is bounded by two and changes at integer time epochs. Then every CG schedule minimizes the makespan within the class of nonpreemptive schedules:*

$$C_{CG}(G, M) = C_{np}^*(G, M).$$

Proof. The result can be shown by mimicking the proof provided in [3]. See [14] for details. ■

Theorem 4.6 *Let G be an arbitrary task graph and M be a profile bounded by two. If (G, M) has commensurable timing, then the LRP schedule minimizes the makespan within the class of preemptive schedules:*

$$C_{LRP}(G, M) = C_p^*(G, M).$$

Proof. Since the CG schedules form a subclass of HLF schedules, an application of Theorem 3.2 yields the desired result. ■

Remark: Theorem 4.6 extends the result of Muntz and Coffman [12] to variable profiles bounded by two.

A Proof of Theorem 3.1

The scheme of our proof is similar to that of [13]. We first establish some intermediary results.

Lemma A.1 *Assume that (G, M) has commensurable timing with commensurability factor w . Then there is a constant Δ such that*

$$\forall k \in \mathbb{N}^* : \quad L_p^*(G, M) \leq L_{np}^*(G_{w/k}, M) \leq L_p^*(G, M) + \frac{\Delta}{k}.$$

Proof. The first inequality comes from the fact

$$L_p^*(G, M) = L_p^*(G_{w/k}, M) \leq L_{np}^*(G_{w/k}, M),$$

where we used Lemma 2.3 for the first equality. We now prove the second inequality.

Let S^* be an optimal preemptive schedule, and C_1, \dots, C_n be the completion times of the tasks of G . Without loss of generality, we assume that the tasks of G are labeled in such a way that $C_1 \leq \dots \leq C_n$. Let $C_0 = 0$.

For $j = 1, \dots, n$, let V_j denote the set of tasks that are assigned for execution in the time interval $\theta_j = [C_{j-1}, C_j)$. Note that all the tasks in V_j are enabled at time C_{j-1} so that there is no precedence relation between these tasks. Clearly, $|V_j| \leq n$.

For a given time interval θ_j , we define the *assigned pieces* (i, μ, t^1, t^2) , $i \in V$, $1 \leq \mu \leq m$, such that task i is continuously executed by processor μ during $[t^1, t^2)$, where $C_{j-1} \leq t^1 < t^2 \leq C_j$, and that task i is assigned to processor μ neither at time t^{1-} nor at time t^{2+} . The quantity $t^2 - t^1$ is referred to as the length of the assigned pieces.

By hypothesis, the number of profile changes during each interval θ_j is bounded by \bar{r} . It is possible to transform S^* so that the number of total assigned pieces in each time interval is bounded. Such a transformation can be obtained by for instance considering the tasks in V_j one by one. A task i is executed by an available processor continuously until the profile changes or the total amount of executions of i in S^* during θ_j is reached. Under such a transformation, the number of preemptions in each time interval is bounded by $mn\bar{r}$. Hence, suppose the number of total assigned pieces in S^* is bounded by B in each time interval θ_j , $j = 1, 2, \dots, n$.

We want to construct a nonpreemptive (possibly idling) schedule S_k for $G_{w/k}$ such that the relation

$$L_{S_k}(G_{w/k}, M) \leq L_{S^*}(G, M) + \frac{\Delta}{k}$$

holds for some constant Δ independent of k . For this purpose, we first construct an intermediate schedule S' from S^* .

Let us consider the first time interval $\theta_1 = [0, C_1)$. Schedule S' is constructed from S^* by cutting a small portion of execution time from each assigned piece so that its length becomes a multiple of w/k . At time C_1 , an assigned piece of length w/k is added sequentially on processor 1 (which is assumed to be always available) for each of the assigned pieces. To be more precise, let there be $B_1 \leq B$ assigned pieces in θ_1 : $(v_i, \mu_i, t_i^1, t_i^2)$, $1 \leq i \leq B_1$, $v_i \in V$. In S' , each assigned piece $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* is replaced by two pieces $(v_i, \mu_i, t_i^1, t_i^1 + \lfloor (t_i^2 - t_i^1)k/w \rfloor w/k)$ and $(v_i, 1, C_1 + (i-1)w/k, C_1 + iw/k)$. Processor μ_i is idle during the time interval $[t_i^1 + \lfloor (t_i^2 - t_i^1)k/w \rfloor w/k, t_i^2)$. Let

$$C'_1 \stackrel{\text{def}}{=} C_1 + B_1 w/k < C_1 + (B + m)w/k.$$

Clearly, in the time interval $[0, C'_1)$, the lengths of all the assigned pieces of S' are integer multiples of w/k , and S' finishes all the amounts of executions of the tasks in S^* during the time interval $[0, C_1)$.

Assume that for some $j \geq 1$, in the time interval $[C'_{j-1}, C'_j)$, the lengths of all the assigned pieces of S' are integer multiples of w/k , and that S' finishes all the amounts of executions of the tasks in S^* during the time interval $[C_{j-1}, C_j)$.

Consider now the assigned pieces of S^* in time interval θ_{j+1} . There are two cases: $C_j \leq C'_j \leq C_{j+1}$ or $C'_j > C_{j+1}$.

Assume first $C_j \leq C'_j \leq C_{j+1}$. Let there be $B_{j+1} \leq B$ assigned pieces in θ_{j+1} . We split the \bar{m} assigned pieces (v, μ, t^1, t^2) of S^* such that $t^1 < C'_j < t^2$ into two pieces (v, μ, t^1, C'_j) and (v, μ, C'_j, t^2) . Let the $B_{j+1} + \bar{m}$ assigned pieces $(v_i, \mu_i, t_i^1, t_i^2)$, $1 \leq i \leq B_{j+1} + \bar{m}$, be ordered in such a way that $t_i^1 \geq C'_j$ holds for all $1 \leq i \leq B' \leq B_{j+1} + \bar{m}$, and that $t_i^2 \leq C'_j$ holds for all $B' < i \leq B_{j+1} + \bar{m}$. Construct S' as follows:

- For the assigned pieces $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* with $i \leq B'$, we apply the same procedure as in the case $j = 1$, v.z., the assigned piece $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* is replaced by two pieces $(v_i, \mu_i, t_i^1, t_i^1 + \lfloor (t_i^2 - t_i^1)k/w \rfloor w/k)$ and $(v_i, 1, C_{j+1} + (i-1)w/k, C_{j+1} + iw/k)$ in S' .
- For the assigned pieces $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* with $i > B'$, we slightly increase their lengths so that they become integer multiples of w/k and then sequentially assign them at processor 1 at time $C_{j+1} + B'w/k$, i.e., the assigned piece $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* is replaced by $(v_i, 1, t'_i, t''_i)$ in S' , where

$$t'_i = C_{j+1} + B'w/k + \sum_{u=B'+1}^{i-1} \lfloor (t_u^2 - t_u^1)k/w + 1 \rfloor w/k;$$

$$t_i'' = t_i' + \lfloor (t_i^2 - t_i^1)k/w + 1 \rfloor w/k.$$

Let $C'_{j+1} = t''_{B_{j+1}+\bar{m}}$. It is easily verified that in the time interval $[C'_j, C'_{j+1})$ the lengths of all the assigned pieces of S' are integer multiples of w/k , and that S' finishes all the amounts of executions of the tasks in S^* during the time interval $[C_j, C_{j+1})$. Note that

$$\begin{aligned} t''_{B_{j+1}+\bar{m}} &\leq C_{j+1} + B'w/k + m(C'_j - C_j) + (B_{j+1} + \bar{m} - B')w/k \\ &\leq C_{j+1} + m(C'_j - C_j) + (B + m)w/k. \end{aligned}$$

Therefore,

$$C'_{j+1} - C_{j+1} \leq m(C'_j - C_j) + (B + m)w/k.$$

Assume now $C'_j > C_{j+1}$. Let there be $B_{j+1} \leq B$ assigned pieces in θ_{j+1} : $(v_i, \mu_i, t_i^1, t_i^2)$, $1 \leq i \leq B_{j+1}$. We slightly increase their lengths so that they become integer multiples of w/k and then sequentially assign them at processor 1 at time C'_j , i.e., the assigned piece $(v_i, \mu_i, t_i^1, t_i^2)$ of S^* is replaced by $(v_i, 1, t_i', t_i'')$ in S' , where

$$\begin{aligned} t_i' &= C'_j + \sum_{u=1}^{i-1} \lfloor (t_u^2 - t_u^1)k/w + 1 \rfloor w/k; \\ t_i'' &= t_i' + \lfloor (t_i^2 - t_i^1)k/w + 1 \rfloor w/k. \end{aligned}$$

Let $C'_{j+1} = t''_{B_{j+1}}$. As in the previous case, it is easily verified that in the time interval $[C'_j, C'_{j+1})$ the lengths of all the assigned pieces of S' are integer multiples of w/k and that S' finishes all the amounts of executions of the tasks in S^* during the time interval $[C_j, C_{j+1})$. Since

$$\begin{aligned} C'_{j+1} &= t''_{B_{j+1}} \\ &\leq C'_j + B_{j+1}w/k + m(C_{j+1} - C_j) \\ &\leq C'_j + (B + m)w/k + C_{j+1} + (m - 1)C_{j+1} - mC_j \\ &\leq C'_j + (B + m)w/k + C_{j+1} + (m - 1)C'_j - mC_j, \end{aligned}$$

we obtain that

$$C'_{j+1} - C_{j+1} \leq m(C'_j - C_j) + (B + m)w/k.$$

This construction is continued until $j = n$ so that a complete schedule S' is generated. By induction, the lengths of all the assigned pieces of S' are integer multiples of w/k , and for all $i \in V$, the total execution time of task i is at least p_i . Furthermore, an easy computation yields

$$\forall 1 \leq j \leq n : \quad C'_j - C_j \leq \frac{m^j - 1}{m - 1} (B + m)w/k \leq \frac{\Delta}{k},$$

where

$$\Delta = \frac{m^n - 1}{m - 1}(B + m)w.$$

The schedule S_k is now defined for the couple $(G_{w/k}, M)$ as follows: task i_j is running at time t if and only if the j -th portion of length w/k of task i is running at time t under schedule S' , where $1 \leq j \leq \alpha_i k$, and i_j is the j -th task in the chain that replaces task i in the expansion of G to obtain $G_{w/k}$. It is easy to see that for all $i \in V$,

$$C_{i_{\alpha_i k}}(S_k) \leq C'_i \leq C_i + \Delta/k.$$

As S_k can be considered as the expansion of a schedule S'' obtained by restricting S' to the first $\alpha_i k$ portions of length w/k of each task $i \in V$, an application of Lemma 2.2 implies that

$$L_{S_k}(G_{w/k}, M) = L_{S''}(G, M) = \max_{i \in V} (C_{i_{\alpha_i k}}(S_k) - d_i) \leq \max_{i \in V} (C_i - d_i) + \Delta/k = L_p^*(G, M) + \Delta/k.$$

Therefore,

$$L_{np}^*(G_{w/k}, M) \leq L_{S_k}(G_{w/k}, M) \leq L_p^*(G, M) + \Delta/k,$$

which completes the proof. ■

Recall that in a priority scheduling algorithm, there are three types of events: (1) A task completes; (2) The priority of one subset becomes the same as another; (3) The profile changes. Let t_1 be the time epoch when the first event occurs in a priority schedule S applied to (G, M) . Denote by $G_{t_1}(S)$ the remaining graph of G at time t_1 under the schedule S , where the processing times of the tasks of $G_{t_1}(S)$ are the remaining processing times of tasks in G at time t_1 under the schedule S . The due dates remain unchanged. Similarly, denote by $M_{t_1}(S) = \{a_r^1(S), m_r^1(S)\}_{r=1}^{\bar{r}}$ the remaining profile, where

$$a_1^1(S) = 0, \quad a_r^1(S) = a_{r+1} - \mathbf{1}_{\{t_1 = a_2\}} - t_1, \quad r \geq 2,$$

$$m_r^1(S) = m_{r+1} - \mathbf{1}_{\{t_1 = a_2\}}, \quad r \geq 1.$$

Lemma A.2 *Assume that (G, M) has commensurable timing with commensurability factor w . Let $t_1 > 0$ be the first time epoch when an event occurs in the SLF schedule. Then there exists an integer $\gamma_1 \in \mathbb{N}^*$ such that:*

- i) t_1 is an integer multiple of $w_1 = w/\gamma_1$;
- ii) the remaining couple $(G_{t_1}(SLF), M_{t_1}(SLF))$ has commensurable timing with commensurability factor w_1 ;

- iii) for all $k \in \mathbb{N}^*$, and for all EDD schedule $S \in \mathcal{E}(G_{w_1/k}, M)$, the remaining task graph of $G_{w_1/k}$ at t_1 in schedule S , denoted by $G' = (G_{w_1/k})_{t_1}(S)$, is isomorphic to $(G_{t_1}(SLF))_{w_1/k}$ in the sense that both dags are isomorphic, and that the corresponding tasks have the same processing times and the same due dates.

Proof. Let V_1, V_2, \dots, V_u be the subsets of the tasks that are assigned for execution under the SLF schedule. Assume that laxities of the tasks in the same subset are identical and that the tasks in V_i have (strictly) smaller laxity than those in V_{i+1} , $i = 1, \dots, u-1$.

If $|V_1| + |V_2| + \dots + |V_u| \leq m_1$ (recall that m_1 is the number of available processors at time 0), then all these tasks are executed at speed 1. In this case, we define $\gamma_1 = 1$. Otherwise, $|V_1| + |V_2| + \dots + |V_u| > m_1$. Let $a = m_1 - (|V_1| + |V_2| + \dots + |V_{u-1}|)$ and $b = |V_u|$. According to the definition of priority schedules, we have $1 \leq a < b$. In this case, the tasks in V_1, V_2, \dots, V_{u-1} are executed at speed 1 and those in V_u at speed a/b . Define $\gamma_1 = ab(b-a)$.

i) Let $q_i \in \{0, 1, \frac{a}{b}\}$ be the speed of task $i \in V$, where $q_i = 0$ if task i is not assigned for execution. There are three possible cases:

- If at time t_1 , an event of type 1 occurs, then there is a task i (with $q_i > 0$) which completes at time t_1 . Thus $t_1 = p_i/q_i$.
- If an event of type 2 occurs at time t_1 , then at least two tasks i and j which had different laxities at time 0 become of the same priority at time t_1 :

$$d_i - (p_i - q_i t_1) = d_j - (p_j - q_j t_1).$$

Note that this is possible only if $(q_i - q_j)[(d_i - p_i) - (d_j - p_j)] < 0$, i.e., the speed of the task with smaller laxity at time 0 is (strictly) greater than that of the task with larger laxity at time 0. Therefore,

$$t_1 = \frac{1}{q_i - q_j} [(d_j - p_j) - (d_i - p_i)].$$

- If now the event is of type 3, then $t_1 = a_2$.

Due to the fact that (G, M) has commensurable timing with commensurability factor w , we obtain that in all these three cases, t_1 is an integer multiple of $w_1 \stackrel{\text{def}}{=} w/\gamma_1$.

ii) By definition, γ_1 is an integer multiple of the speed q_i , $i \in V$. Using the fact that (G, M) has commensurability factor w (which implies the commensurability of factor w/γ_1), one readily gets that all the remaining processing times in the task graph $G_{t_1}(SLF)$ are integer multiples of w_1 . Using further the fact that t_1 is an integer multiple of w_1 implies that all the profile

changing epochs of $M_{t_1}(SLF)$ are integer multiples of w_1 . Therefore, the remaining couple $(G_{t_1}(SLF), M_{t_1}(SLF))$ has commensurable timing with commensurability factor w_1 .

iii) Since G has commensurability factor w_1 , we assume that $p_i = \alpha_i w_1$, where α_i is an integer, $i \in V$. Consider the expansion $G_{w_1/k} = (V_{w_1/k}, E_{w_1/k})$, where the subtasks of $i \in V$ are indexed by $i_1, i_2, \dots, i_{\alpha_i k}$. For all $i \in V$, let $T_i = \{i_1, i_2, \dots, i_{\alpha_i k}\}$.

Owing to Lemma 2.4, we have that the expanded graph of $G_{t_1}(SLF)$, denoted by $(G_{t_1}(SLF))_{w_1/k}$, is identical to the remaining graph $(G_{w_1/k})_{t_1}(SLF)$ of the expanded graph $G_{w_1/k}$ at time t_1 under the SLF schedule. Therefore, we only have to show the isomorphism between G' and the remaining graph $(G_{w_1/k})_{t_1}(SLF)$ of $G_{w_1/k}$ at time t_1 under the schedule SLF.

For all $i \in V$, let $T_i(S) \subseteq T_i$ (resp. $T_i(SLF) \subseteq T_i$) be the set of tasks of T_i that remain at time t_1 in the EDD schedule S (resp. SLF schedule). As the remaining processing times in $G_{t_1}(SLF)$ are integer multiples of w_1 , it suffices to prove that for all $i \in V$, $T_i(S) = T_i(SLF)$, or simply $|T_i(S)| = |T_i(SLF)|$.

Consider the subsets of tasks of G that are assigned for execution under SLF: V_1, V_2, \dots, V_u . Assume that $|V_1| + |V_2| + \dots + |V_u| > m_1$ so that $\gamma_1 = ab(b-a)$. (The case $|V_1| + |V_2| + \dots + |V_u| = m_1$ can be shown analogously and is thus omitted). Hence in the SLF schedule, the tasks in V_1, V_2, \dots, V_{u-1} are executed at speed 1 and those in V_u at speed a/b . It is shown in i) that t_1 is an integer multiple of w_1 . The interested reader can verify that t_1 is in fact an integer multiple of bw_1 . Thus, $t_1 = \eta bw_1$, where η is an integer. It then follows that

$$|T_i(SLF)| = \begin{cases} (p_i - t_1)k/w_1 = (\alpha_i - \eta b)k, & i \in V_1 \cup \dots \cup V_{u-1}; \\ (p_i - t_1 a/b)k/w_1 = (\alpha_i - \eta a)k, & i \in V_u; \\ p_i/w_1 = \alpha_i k, & i \in V - (V_1 \cup \dots \cup V_u). \end{cases} \quad (4)$$

Consider now the nonpreemptive EDD schedule S for $G_{w_1/k}$. Since all the tasks in $G_{w_1/k}$ have the same processing time (w_1/k) , a task i_j has smaller laxity than task i'_j , if and only if i_j has smaller due date than task i'_j . Furthermore, for all $1 \leq j \leq \eta bk = t_1 k/w_1$,

- task i_j has the same due date as i'_j for all $i, i' \in V_f$, $1 \leq f \leq u$;
- task i_j has (strictly) smaller due date than i'_j for all $i \in V_f$, $i' \in V_{f'}$, $1 \leq f < f' \leq u$;
- task i_j has (strictly) smaller due date than i'_j , for all $i \in V_1 \cup \dots \cup V_u$, $i' \in V - (V_1 \cup \dots \cup V_u)$, $j' \geq 1$.

Therefore, in the nonpreemptive EDD schedule S for $G_{w_1/k}$, during each of the first $\eta bk = t_1 k/w_1$ time intervals of length w_1/k , one task from each of the sets T_i , $i \in V_1 \cup \dots \cup V_{u-1}$, and

a enabled tasks having the smallest due dates among those in $\bigcup_{i \in V_u} T_i$ are chosen for execution. Since there are in total ηbk intervals of length w_1/k in $[0, t_1)$, the EDD schedule S finishes ηbk tasks from each of the sets T_i , $i \in V_1 \cup \dots \cup V_{u-1}$, and finishes ηak tasks from each of the sets T_i , $i \in V_u$. Thus,

$$|T_i(S)| = \begin{cases} \alpha_i k - \eta bk, & i \in V_1 \cup \dots \cup V_{u-1}; \\ \alpha_i k - \eta ak, & i \in V_u; \\ \alpha_i k, & i \in V - (V_1 \cup \dots \cup V_u). \end{cases} \quad (5)$$

The equations (4) and (5) readily imply assertion iii). ■

Informally, Lemma A.2 means that if G is sufficiently “sliced”, all EDD schedules behave analogously to SLF(G, M) during $[0, t_1)$. Using this property, we establish the following lemma.

Lemma A.3 *Assume that (G, M) has commensurable timing with commensurability factor w . Then there exists an integer γ such that for all $k \in \mathbb{N}^*$, and all EDD schedule $S_k \in \mathcal{E}(G_{w_0/k}, M)$, where $w_0 = w/\gamma$,*

$$L_{S_k}(G_{w_0/k}, M) = L_{SLF}(G, M).$$

Proof. Observe first that in the schedule $SLF(G, M)$, the number of events is bounded by $2n + \bar{r}$. Indeed, the number of tasks, the number of subsets of tasks having the same priorities, and the number of profile changes are bounded by $n + n + \bar{r}$. Whenever an event occurs in $SLF(G, M)$, either the number of tasks is decreased by one (event type 1), or the number of subsets is decreased by one (event type 2, in which case at least two subsets merge), or the number of profile changes is decreased by one (event type 3).

Let N and t_1, t_2, \dots, t_N be the number and the time epochs, respectively, of events in the SLF schedule for (G, M) . Let $G^{(j)} = G_{t_j}(SLF)$ and $M^{(j)} = M_{t_j}(SLF)$ be the remaining graph and the remaining profile at time t_j , $1 \leq j \leq N$, where $G^{(N)} = \emptyset$.

It is readily shown by induction using Lemma A.2 that there are integers $\gamma_1, \gamma_2, \dots, \gamma_N$ such that t_j is an integer multiple of $w/(\gamma_1 \dots \gamma_j)$, $j = 1, 2, \dots, N$. Let

$$\gamma = \gamma_1 \gamma_2 \dots \gamma_N, \quad w_0 = w/\gamma.$$

Thus, for all $k \in \mathbb{N}^*$, and all EDD schedule $S_k \in \mathcal{E}(G_{w_0/k}, M)$, an induction on $j = 1, 2, \dots, N$ using Lemma A.2 implies that $(G_{w_0/k})_{t_j}(S_k)$ is isomorphic to $(G^{(j)})_{w_0/k}$ for all $j = 1, 2, \dots, N$. Thus,

$$L_{S_k}(G_{w_0/k}, M) = L_{SLF}(G_{w_0/k}, M) = L_{SLF}(G, M),$$

where we used Lemma 2.4 for the last equality. ■

We are now in a position to prove the main theorem.

Proof of Theorem 3.1. Let (G, M) verify the conditions of the theorem. Applying Lemma A.3 to the couple (G_w, M) implies that there is an integer γ such that for all $k \in \mathbb{N}^*$, and all EDD schedule $S_k \in \mathcal{E}(G_{w_0/k}, M)$, where $w_0 = w/\gamma$,

$$L_{S_k}(G_{w_0/k}, M) = L_{SLF}(G, M).$$

Let S_k^* be the EDD schedule which is optimal within the class of nonpreemptive schedules for the minimization of maximum lateness of $(G_{w_0/k}, M)$. Then, according to Lemma A.1

$$L_p^*(G, M) \leq L_{np}^*(G_{w/(\gamma k)}, M) = L_{S_k^*}(G_{w_0/k}, M) = L_{SLF}(G, M) \leq L_p^*(G, M) + \frac{\Delta}{\gamma k}.$$

Letting k tend to infinity immediately entails the desired result. ■

B Proof of Theorem 4.1

Instead of proving Theorem 4.1, we show the following Theorem which, by a standard argument, immediately implies Theorem 4.1.

Theorem B.1 (Extension of Theorem 1 in [1]) *Let $G \in \mathcal{G}_{if}$ be an in-forest, and $M \in \mathcal{M}_{iz}$ be an increasing zigzag profile. Assume that the tasks have UET and that the profile changing epochs are integer. Then any EDD schedule $S \in \mathcal{E}(G', M)$ defined on the modified due dates meets all the original due dates if and only if such a feasible nonpreemptive schedule exists.*

Proof. The proof is similar to that of Theorem 1 in [1] except that we have to consider the variable profile. It suffices to show that, if an EDD schedule $S \in \mathcal{E}(G', M)$ defined on the modified due dates does not meet all the original due dates, then there is no feasible nonpreemptive schedule which achieves this.

Owing to Lemma 4.1, a given schedule meets all the original due dates if and only if it meets all the modified due dates. Assume there is an EDD schedule $S \in \mathcal{E}(G', M)$ defined on the modified due dates which does not meet all the original due dates. Let u be the task having the smallest modified due date among the tasks for which S fails to meet their modified due dates, i.e.,

$$C_u(S) > d'_u, \quad \text{and} \quad d'_u \leq d'_i, \quad \text{if } C_i(S) > d'_i.$$

Let $b = C_u(S) - 1$ be the time at which task u is assigned to a processor. We will show by contradiction that there is no idle processor in the time interval $[0, b)$ under S and that all the tasks assigned for execution under S during the time interval $[0, b)$ have strictly smaller modified due dates. These facts trivially imply that there is no feasible schedule that meets all the modified due dates.

Assume that these facts are not true. Then, there is an integer $0 \leq t \leq b - 1$ such that during the time interval $[t, t + 1)$ of schedule S , there is an available processor which is either idle or executing a task j with strictly larger modified due date than task u ($d'_j > d'_u$). Let t be the largest such integer. Then during the time interval $[t + 1, b)$ of schedule S , all the available processors are busy and are executing tasks whose modified due dates are smaller or equal to d'_u .

According to the definition of due date modifications, all the predecessors of a task $i \in G$ have strictly smaller modified due dates than i . Therefore, in the EDD schedule S defined on the modified due dates, if task i is assigned for execution before task j which has strictly smaller modified due date ($d'_j < d'_i$), then there is no precedence constraint between i and j .

Thus, if $t = b - 1$, then there is at least a predecessor task i of u such that i is executed during $[t, t + 1)$ (otherwise task u would be assigned by time $t = b - 1$). This implies $C_i(S) = C_u(S) - 1 > d'_u - 1 \geq d'_i$, which contradicts the assumption on task u . Hence, $t = b - 1$ is impossible.

Therefore, $t < b - 1$. Assume without loss of generality that the profile is specified every unit of time, i.e., $a_r = r - 1$, $r = 1, 2, \dots$. Since all the m_{t+1} tasks which start execution at time $t + 1$ have (non-strictly) smaller modified due dates than u , all these tasks have predecessors executing during the time interval $[t, t + 1)$. As the precedence graph is an in-forest, two tasks can not have the same predecessor. Hence at least m_{t+1} tasks which have (non-strictly) smaller modified due dates than u start execution at t . Thus, $m_t \geq m_{t+1} + 1$.

Since the profile is zigzag increasing, we have necessarily $m_t - 1 = m_{t+1} \stackrel{\text{def}}{=} \hat{m}$. Moreover, for all $t' \geq t + 1$, $m_{t'} \geq \hat{m}$. By the definitions of u and t , all the tasks running during the time interval $[t + 1, b)$ under schedule S have at least one (not necessarily immediate) predecessor which is assigned for execution at time t . Therefore, these tasks form \hat{m} chains, one of which contains predecessors of u . Let i be the immediate predecessor of u such that $C_i(S) = b$. It then follows that $C_i(S) = C_u(S) - 1 > d'_u - 1 \geq d'_i$, which, again, contradicts the assumption on task u . The proof is thus completed. ■

References

- [1] P. Brucker, M. R. Garey and D. S. Johnson, "Scheduling Equal-length Tasks under Treelike Precedence Constraints to Minimize Maximum Lateness", *Math. of Oper. Res.*, **2** (1977), pp. 275-284.
- [2] E. G. Coffman, Jr. (ed.) *Computer and Job-Shop Scheduling Theory*, Wiley, 1976.
- [3] E. G. Coffman, Jr. and R. L. Graham, "Optimal Scheduling for Two-Processor Systems", *Acta Informatica*, **1** (1972), pp. 200-213.
- [4] D. Dolev, M. K. Warmuth, "Scheduling Precedence Graphs of Bounded Height", *J. of Algorithms*, **5** (1984), pp. 48-59.
- [5] D. Dolev and M. K. Warmuth, "Scheduling Flat Graphs", *SIAM J. on Comput.*, **14** (1985), pp. 638-657.
- [6] D. Dolev and M. K. Warmuth, "Profile Scheduling of Opposing Forests and Level Orders", *SIAM J. Alg. Disc. Meth.*, **6** (1985), pp. 665-687.
- [7] M. R. Garey, D. S. Johnson, R. E. Tarjan, M. Yannakakis, "Scheduling Opposite Forests", *SIAM J. Alg. Disc. Meth.*, **4** (1983), pp. 72-93.
- [8] E. L. Lawler, "Preemptive Scheduling of Precedence Constrained Scheduling", in *Deterministic and Stochastic Scheduling*, Dempster et al. (editors), Reidel, 1982, pp. 101-123.
- [9] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan and D. B. Shmoys, "Sequencing and Scheduling: Algorithms and Complexity", Report BS-R8909, CWI, 1989.
- [10] Z. Liu, E. Sanlaville, "Stochastic Scheduling with Variable Profile and Precedence Constraints". Rapport de Recherche INRIA, No. 1525, 1991. Submitted.
- [11] R. McNaughton, "Scheduling with Deadlines and Loss Functions", *Manag. Sci.*, **6** (1959), pp. 1-12.
- [12] R. R. Muntz and E. G. Coffman, Jr., "Optimal Preemptive Scheduling on Two-Processor Systems", *IEEE Trans. on Comp.*, **C-18** (1969), pp. 1014-1020.
- [13] R. R. Muntz and E. G. Coffman, Jr., "Preemptive Scheduling of Real-Time Tasks on Multiprocessor Systems", *J. of the ACM*, **17** (1970), pp. 325-338.
- [14] E. Sanlaville, "Problèmes d'ordonnancement à profils variables : état de l'art et extensions", Tech. Rep. IBP/MASI 90-39, Univ. Paris VI, France, September 1990.
- [15] J. D. Ullman, "NP-Complete Scheduling Problems" *J. Comput. System Sci.*, **10** (1975), pp. 384-393.

ISSN 0249-6399