

Normal-form preemption sequences for an open problem in scheduling theory

Bo Chen¹ · Ed Coffman² · Dariusz Dereniowski³ · Wiesław Kubiak⁴

Published online: 28 August 2015

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Structural properties of optimal preemptive schedules have been studied in a number of recent papers with a primary focus on two structural parameters: the minimum number of preemptions necessary, and a tight lower bound on *shifts*, i.e., the sizes of intervals bounded by the times created by preemptions, job starts, or completions. These two parameters have been investigated for a large class of preemptive scheduling problems, but so far only rough bounds for these parameters have been derived for specific problems. This paper sharpens the bounds on these structural parameters for a well-known open problem in the theory of preemptive scheduling: Instances consist of in-trees of n unit-execution-time jobs with release dates, and the objective is to minimize the total completion time on two processors. This is among the current, tantalizing “threshold” problems of scheduling theory: Our literature survey reveals that any significant generalization leads to an NP-hard prob-

lem, but that any significant, but slight simplification leads to tractable problem with a polynomial-time solution. For the above problem, we show that the number of preemptions necessary for optimality need not exceed $2n - 1$; that the number must be of order $\Omega(\log n)$ for some instances; and that the minimum shift need not be less than 2^{-2n+1} . These bounds are obtained by combinatorial analysis of optimal preemptive schedules rather than by the analysis of polytope corners for linear-program formulations of the problem, an approach to be found in earlier papers. The bounds immediately follow from a fundamental structural property called *normality*, by which minimal shifts of a job are exponentially decreasing functions. In particular, the first interval between a preempted job’s start and its preemption must be a multiple of $1/2$, the second such interval must be a multiple of $1/4$, and in general, the i -th preemption must occur at a multiple of 2^{-i} . We expect the new structural properties to play a prominent role in finally settling a vexing, still-open question of complexity.

✉ Dariusz Dereniowski
deren@eti.pg.gda.pl

Bo Chen
b.chen@warwick.ac.uk

Ed Coffman
coffman@cs.columbia.edu

Wiesław Kubiak
wkubiak@mun.ca

¹ Centre for Discrete Mathematics and Its Applications (DIMAP) and Warwick Business School, University of Warwick, Coventry, UK

² Departments of Electrical Engineering and of Computer Science, Columbia University, New York, USA

³ Faculty of Electronics, Telecommunications and Informatics, Gdańsk University of Technology, Gdańsk, Poland

⁴ Faculty of Business Administration, Memorial University, St. John’s, Canada

Keywords Preemption · Parallel machines · In-tree · Release date · Scheduling algorithm · Total completion time

1 Introduction

We study structural properties of optimal preemptive schedules of a classic problem of scheduling unit execution time (UET) jobs with precedence constraints and release dates on two processors. Optimal nonpreemptive schedules for this and related problems have been well researched in the literature for various objective functions and restrictions. Fujii et al. (1971) present a matching-based algorithm, and Coffman and Graham (1972) devise a job-labeling algorithm for minimum-makespan nonpreemptive schedules. Garey and

Johnson introduce $O(n^2)$ and $O(n^{2.81})$ time algorithms for minimizing maximum lateness for jobs, respectively, without release dates (Garey and Johnson 1976), and with release dates (Garey and Johnson 1977). Gabow (1982) designed an almost linear-time algorithm for the minimum-makespan problem. Leung et al. (2001) and Carlier et al. (2014) extend these results to precedence delays. Baptiste and Timkovsky (2004) focus on minimization of total completion time and present an $O(n^9)$ time shortest-path optimization algorithm for scheduling jobs with release dates. They also conjecture that there always exist so-called *ideal* schedules that minimize both maximum completion time and total completion time for jobs with release dates. This has been known to hold true for equal release dates without preemptions (Coffman and Graham 1972) and with preemptions (Coffman et al. 2003). Coffman et al. (2012) prove the Baptiste–Timkovsky conjecture and give an $O(n^3)$ algorithm for the minimization of total completion time for jobs with release dates—a major improvement over the $O(n^9)$ time algorithm in Baptiste and Timkovsky (2004).

Optimal preemptive schedules have proven more challenging to compute efficiently, especially for jobs with release dates and the total-completion-time criterion. Coffman et al. (2012) prove that these schedules are not ideal, that is, for some instances any schedule minimizing total completion time will be longer than the schedule minimizing maximum completion time. That holds even for in-tree precedence constraints. This last result serves as a point of departure for this paper, with its focus on in-tree precedence constraints, release dates, and the criterion of total-completion-time, the problem $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$ in the well-known three-field notation. Despite numerous efforts, the computational complexity of the problem remains open: reducing the number of processors to $m = 1$ renders the problem polynomially solvable (Baptiste et al. 2004); and so does dropping the precedence constraints (Herrbach and Leung 1990); dropping the release dates (Coffman et al. 2003); and assuming out-trees instead of in-trees (Baptiste and Timkovsky 2001; Brucker et al. 2003; Huo and Leung 2005). With this background in mind, we focus on key structural properties of optimal preemptive schedules for the problem $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$. The complexity of the problem studied in this paper remains also open for $m = 3$ machines. This holds for both nonpreemptive and preemptive cases and for both makespan and total-completion-time objectives.

Sauer and Stone (1987) study the problem with no release dates and maximum-completion-time (makespan) minimization. They show that, for every optimal preemptive schedule, there is an optimal preemptive schedule with at most n preemptions, where preemptions occur at multiples of $1/2$, and go on to define a *shift* that is the duration of an interval between two consecutive time points, each of which is a job

start, a job end, a job resumption, or a job preemption. The shortest necessary shift in an optimal schedule is then called its *resolution*.¹ The minimum resolution over all instances of a given preemptive scheduling problem is called the *problem resolution*. Following Sauer and Stone (1987), the minimum number of preemptions and the minimum resolution necessary for optimal schedules have become two main structural parameters in preemptive scheduling. They have been investigated for a large class of preemptive scheduling problems by Baptiste et al. (2011) who give general bounds for these parameters.

Coffman et al. (2015) provide bounds on the resolutions of various scheduling problems—we refer the reader to their work for a comprehensive overview. In particular, they show upper bounds of $m^{-n/(m+1)}$ and $m^{-(n-1)/(m+1)}$ on resolutions for problems $P|pmtn, in-tree, r_j, p_j = 1|C_{\max}$ and $P|pmtn, in-tree, r_j, p_j = 1|\sum C_j$, respectively, where n is the number of jobs and m is the number of processors. Thus, for the problem $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$ studied in this paper one immediately obtains an upper bound of $2^{-(n-1)/3}$ on its resolution. As for lower bounds, Coffman et al. (2015) shows that the resolution of $P|pmtn, prec, r_j|\sum w_i C_j$ is at least $(m+n)^{-(2n+1)/2}$.

The papers of Sauer and Stone (1987), Baptiste et al. (2011), and Coffman et al. (2015) obtain their resolution bounds by analyzing the corners of feasibility regions of linear programs designed for specific problems. Our approach is combinatorial and does not make use of the theory of linear programming. It yields a lower bound of 2^{-2n-1} on the problem resolution of $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$, which is a significant improvement for in-trees over the lower bound of $(n+2)^{-(2n+1)/2}$ that can be derived directly from Coffman et al. (2015).

We introduce in this paper the concept of *normal* schedules where shifts decrease as a function of time: The first shift is a multiple of $1/2$, the second one is a multiple of $1/4$, and in general, the i -th shift is a multiple of 2^{-i} . We prove that there exist optimal schedules that are normal for in-trees. However, we conjecture that this is no longer the case for arbitrary precedence constraints, i.e., there are instances for which no optimal schedule is normal. The normality of a schedule implies that each shift is a multiple of 2^{-2n+1} , which is a much stronger claim than the usual requirement that all shifts are no shorter than the problem resolution. Normality also implies that there exists an optimal schedule with a finite number (in particular, a number not exceeding $2n-1$) of events which are times when jobs start, end, or are preempted. Thus, $2n-1$ is an upper bound on the number of preemptions necessary for optimality. We also observe that a job may be required to preempt at a point which is neither a start nor the end of another job in order to ensure optimality.

¹ The term resolution was introduced in Coffman et al. (2015).

These preemption events unrelated to job starts or completions seem to be confined to rather contrived instances; they are more the exception than the rule in preemptive scheduling. We also prove that there exists a sequence of problem instances indexed by n for which the number of preemptions in the corresponding optimal schedules is $\Omega(\log n)$. Thus, a tight upper bound on the number of preemptions required for optimality must be at least logarithmic in n .

2 Our approach and results: a general overview

In Sect. 3, we exhibit a sequence of problem instances indexed by n where the number of preemptions of a single job is of order $\Omega(\log n)$. In the remaining sections we focus on proving our lower bound on the minimum shift. We start by showing that an optimal schedule is a concatenation of *blocks*, each with at most three jobs. No job starts or completes inside a block but there is at least one job start at the beginning of a block, and/or at least one job completion at the end of a block. This is done in Sects. 4.2 and 4.3. A block is called l -normal if each job duration in the block is a multiple of $1/2^{l+1}$, and the block length is a multiple of $1/2^l$. In a normal schedule the first block must be 1-normal, the second 2-normal and so on. These concepts are introduced in Sect. 4.4, where it is verified that, in a normal schedule with q blocks, each preemption occurs at a multiple of $1/2^{q+1}$, where $q \leq 2n - 1$. Our goal is to show that there exists an optimal schedule that is normal. Our proof is by contradiction. We begin by assuming an optimal schedule that is also maximal in the sense that it has a latest possible abnormality point i , i.e., a latest block i which is not i -normal. We show that such a block must have exactly three jobs. One completes at the end of the block and has an $(i + 1)$ -normal duration, but the durations of the other two are not $(i + 1)$ -normal, as shown by Lemma 12. These two jobs then trigger an alternating chain of jobs to which they also belong, as shown in Sect. 6. The completion times of the jobs in the chain are not $(i + 1)$ -normal, which makes it possible under normal-block circumstances to either extend the chain by one job or prove that the abnormality point must exceed i ; this is our main result in Theorem 3. Thus, we get a contradiction in either case since the number of jobs is finite and the schedule is maximal. The normal-block circumstances here mean that the alternating chain does not end with a certain structure that we call an *A-configuration*, a configuration that prevents us from extending the alternating chain. However, we show that there always exists a maximal schedule that does not include an A-configuration. This is done in Sect. 5, where the key result is Theorem 2. The main result of the paper follows and states that there is a normal schedule that is optimal for $P2|pmtn, in-tree, r_j, p_j = 1| \sum C_j$.




Fig. 1 The precedence constraints and release dates for $\mathcal{J}_p, p \geq 2$

3 How many preemptions of a job is required?

In this section we show that, for any given number n of jobs, it is sometimes necessary to preempt a job $p = \Omega(\log n)$ times. Let $A_i, i \geq 0$, be a set of four jobs $a_1^i, a_2^i, a_3^i, a_4^i$ such that $r(a_j^i) = 2i$ for each $j \in \{1, 2, 3\}$, $r(a_4^i) = 2i + 1$ and $a_j^i < a_4^i$ for each $j \in \{1, 2, 3\}$. Then, define

$$\mathcal{J}_p = \bigcup_{i=0}^p A_i,$$

where $a_j^i < a_4^{i+1}$ for each $i \in \{0, \dots, p - 1\}$. (See Fig. 1.)

We prove that the job a_4^p should complete exactly at $2p + 3 - 1/2^{p+1}$ in any optimal schedule. This is done by first proving that no feasible schedule (optimal or not) can complete a_4^p earlier (cf. Proposition 1), and then by proving that starting a_4^p later leads to a schedule that cannot be optimal (cf. Proposition 3).

Proposition 1 *Let $p \geq 0$ be any integer. If \mathcal{P} is a preemptive schedule for \mathcal{J}_p , then the total length of the job a_4^p executing in $[2p + 2, +\infty)$ is at least $1 - 1/2^{p+1}$.*


Proof We prove the lemma by induction on p . Let $p = 0$. (Note that $\mathcal{J}_p = A_0$.) Executing less than $1/2$ units of a_4^0 in $[2, +\infty)$ implies that a_4^0 completes at $5/2 - \epsilon$, for some $\epsilon > 0$. This, however, requires completing each job in $A_0 \setminus \{a_4^0\}$ at $3/2 - \epsilon$ or earlier, which is not possible.

Suppose that the lemma holds for integers smaller than p and we prove it for p . Let \mathcal{P} be a preemptive schedule for \mathcal{J}_p . We consider all jobs that must execute in time interval $I = [2p, s(\mathcal{P}, a_4^p)]$. Each job in $A_p \setminus \{a_4^p\}$ executes in this interval, because $r(a) = 2p$ and $a < a_4^p$ for each $a \in A_p \setminus \{a_4^p\}$. By induction hypothesis and by the facts that $a < a_4^p$ for each $a \in \mathcal{J}_{p-1}$, we obtain that a part of a_4^{p-1} that executes in I is of length at least $1 - 1/2^p$. Thus, the total length of all jobs that execute in I in \mathcal{P} is at least $4 - 1/2^p$. Therefore,

$$s(\mathcal{P}, a_4^p) \geq 2p + |I|/2 = 2p + 2 - 1/2^{p+1}.$$

Thus, at least $1 - 1/2^{p+1}$ units of a_4^p execute in $[2p + 2, +\infty)$ as required. \square

Fig. 2 **a** The execution of the jobs in A_i in \mathcal{P}_p ; **b** the schedule \mathcal{P}_p for \mathcal{J}_2



We iteratively construct a schedule \mathcal{P}_p . Let \mathcal{P}_0 be such that the jobs a_1^0, a_2^0, a_3^0 form a 3/2 schedule in interval $[0, 3/2]$ and a_4^0 executes in $[3/2, 5/2]$. For $p > 0$, first take \mathcal{P}_{p-1} and then execute the jobs in A_p as follows:

$$\begin{aligned}
 & a_1^p \text{ in } [2p, 2p + 1], \\
 & a_2^p \text{ in } [2p + 1 - 1/2^p, 2p + 1 - 1/2^{p+1}] \\
 & \cup [2p + 1, 2p + 2 - (1/2^p - 1/2^{p+1})] \\
 & = [2p + 1 - 1/2^p, 2p + 1 - 1/2^{p+1}] \\
 & \cup [2p + 1, 2p + 2 - 1/2^{p+1}], \\
 & a_3^p \text{ in } [2p + 1 - 1/2^{p+1}, 2p + 2 - 1/2^{p+1}], \\
 & a_4^p \text{ in } [2p + 2 - 1/2^{p+1}, 2p + 3 - 1/2^{p+1}].
 \end{aligned}$$

(See Fig. 2.)

Note that, when $p > 0$, $C(\mathcal{P}_p, a_4^{p-1}) = C(\mathcal{P}_{p-1}, a_4^{p-1}) = 2p + 1 - 1/2^p$ and hence \mathcal{P}_p is feasible. This gives the following.

Proposition 2 *Let $p \geq 0$ be any integer. There exists a schedule for \mathcal{J}_p that completes a_4^p at $2p + 3 - 1/2^{p+1}$. □*

Proposition 3 *Let $p \geq 0$ be any integer. Each optimal schedule for \mathcal{J}_p completes a_4^p at $2p + 3 - 1/2^{p+1}$ and satisfies the following: the total length of idle time in $[0, 2(j + 1)]$ is $1 - 2^{j+1}$ for each $j \in \{0, \dots, p\}$, and there is no time interval contained in $[0, 2(p + 1)]$ in which both processors are idle.*

Proof We prove the lemma by contradiction, i.e., suppose that, for some $p \geq 0$, there exists an optimal schedule \mathcal{P}'_p such that $C(\mathcal{P}'_p, a_4^p) \neq 2p + 3 - 1/2^{p+1}$. Let, without loss of generality, p be the minimum integer for which this holds. One can verify the lemma for $p = 0$ and hence $p > 0$. By Proposition 1,

$$C(\mathcal{P}'_p, a_4^p) > 2p + 3 - 1/2^{p+1}. \tag{1}$$

By Proposition 2 and by the minimality of p , there exists an optimal schedule \mathcal{P}_{p-1} for \mathcal{J}_{p-1} that executes at least $1 - 1/2^p$ units of a_4^{p-1} in $[2p, +\infty)$.

We argue that the total length of a_4^{p-1} , denoted by x , that executes in \mathcal{P}'_p in interval $[2p, +\infty)$ equals exactly $1 - 1/2^p$.

By Proposition 1,

$$x \geq 1 - 1/2^p. \tag{2}$$

Define a schedule \mathcal{P}' that equals \mathcal{P}_{p-1} in the interval $[0, 2p]$ and equals \mathcal{P}'_p in the interval $[2p, +\infty)$. Note that \mathcal{P}' is not feasible only if the total length of a_4^p executing in \mathcal{P}' (that equals $1/2^p + x$) is greater than 1. However,

$$\begin{aligned}
 C(\mathcal{P}', a) &= C(\mathcal{P}'_p, a) \text{ for each } a \in A_p \text{ and} \\
 \sum_{a \in \mathcal{J}_{p-1} \setminus \{a_4^{p-1}\}} C(\mathcal{P}', a) &\leq \sum_{a \in \mathcal{J}_{p-1} \setminus \{a_4^{p-1}\}} C(\mathcal{P}'_p, a). \tag{3}
 \end{aligned}$$

We then obtain a schedule \mathcal{P} by removing the total length of $x - 1 + 1/2^p$ of a_4^{p-1} from \mathcal{P}' in a way that minimizes the completion time of a_4^{p-1} in \mathcal{P} , which gives $C(\mathcal{P}, a_4^{p-1}) \leq C(\mathcal{P}', a_4^{p-1}) - (x - 1 + 1/2^p)$. The schedule \mathcal{P} is feasible and $C(\mathcal{P}, a) = C(\mathcal{P}', a)$ for each $a \in \mathcal{J}_p \setminus \{a_4^{p-1}\}$. Hence, by (3),

$$\sum_{a \in \mathcal{J}_p} C(\mathcal{P}, a) \leq -(x - 1 + 1/2^p) + \sum_{a \in \mathcal{J}_p} C(\mathcal{P}, a).$$

Thus, by the optimality of \mathcal{P}'_p , $x - 1 + 1/2^p = 0$, i.e., $x = 1 - 1/2^p$ as required.

In the schedule \mathcal{P}'_p , the jobs that are executed in time interval $[2p, C(\mathcal{P}'_p, a_4^p)]$ are the ones in A_p and x units of a_4^{p-1} . By a case analysis one can prove that

$$\sum_{a \in (A_p \cup \{a_4^{p-1}\}) \setminus \{a_4^p\}} C(\mathcal{P}'_p, a) \geq 4 \cdot 2p + 4 + 2(1 - 1/2^p).$$

Note that, by construction,

$$\sum_{a \in (A_p \cup \{a_4^{p-1}\}) \setminus \{a_4^p\}} C(\mathcal{P}_p, a) = 4 \cdot 2p + 4 + 2(1 - 1/2^p).$$

Moreover, $x = 1 - 1/2^p$ and the minimality of p imply that

$$\sum_{a \in \mathcal{J}_{p-1} \setminus \{a_4^{p-1}\}} C(\mathcal{P}_p, a) \leq \sum_{a \in \mathcal{J}_{p-1} \setminus \{a_4^{p-1}\}} C(\mathcal{P}'_p, a).$$

This, (1) and $C(\mathcal{P}_p, a_4^p) = 2p + 3 - 1/2^{p+1}$ imply that \mathcal{P}'_u is not optimal. This gives the desired contradiction.

Note that it follows that in each optimal schedule \mathcal{P} for \mathcal{J}_p , $C(\mathcal{P}, a_4^j) = 2p + 3 - 1/2^{j+1}$. Since a_4^j is a successor of all jobs in \mathcal{J}_j , we obtain that there is, in \mathcal{P} , idle time on exactly one processor in interval $(2(j + 1) - 1/2^{j+1}, 2(j + 1))$. Also note that, for each $j \in \{0, \dots, p\}$, no idle time is possible in the interval $I = (2j, 2(j + 1) - 1/2^{j+1})$ because, for $j > 0$, a_1^j, a_2^j, a_3^j and a part of a_4^{j-1} of length $1 - 2^{-j}$ must execute in I , and for $j = 0$, a_1^0, a_2^0 and a_3^0 must execute in I . This completes the proof of the claim. \square

Note that for any given $p \geq 0$, the number of jobs in \mathcal{J}_p equals $4(p + 1)$. Proposition 3 gives us that each optimal schedule for \mathcal{J}_p has a job that completes at a time point that is a multiple of $1/2^{p+1}$ but is not a multiple of $1/2^p$. This gives that there exists a set of n jobs \mathcal{J} such that there exists no optimal solution to $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$ for \mathcal{J} in which each job start, completion and preemption occurs at a time point that is a multiple of $1/2^{n/4-1}$. We note that this upper bound on the resolution of this problem is slightly weaker than the bound $2^{-(n-1)/3}$ proved in Coffman et al. (2015). Our main result of this section is the following lower bound on the number of preemption of one job in an optimal schedule.

Theorem 1 *Given any positive integer p , there exists an instance of the problem $P2|pmtn, in-tree, r_j, p_j = 1|\sum C_j$, such that in any optimal schedule for the instance, this is a job that is preempted at least $p = \Omega(\log |\mathcal{J}|)$ times, where \mathcal{J} is the job set of the instance.*

Proof Let \mathcal{P}_p be an optimal schedule for \mathcal{J} . Thus \mathcal{P}_p satisfies the conditions in Proposition 3. By Theorem 4, we may assume that \mathcal{P}_p is normal. Define $l = C(\mathcal{P}_p, a_4^p) \cdot 2^c$, where $c = 2|\mathcal{J}_p| + 3$.

Take $\mathcal{J} = \{a\} \cup \mathcal{J}_p \cup \{b_1, \dots, b_l\}$. The precedence constraints between the jobs in \mathcal{J}_p are as in Fig. 1. We extend the precedence relation to \mathcal{J} by additionally enforcing:

$$a_4^p \prec b_1 \prec b_2 \prec \dots \prec b_l \text{ and } a \prec b_1.$$

We first construct a schedule \mathcal{P}' as follows. Take \mathcal{P}_p and extend it by executing a so that $C(\mathcal{P}', a) = 2(p+1) + 1/2^{p+1}$ and executing b_1, \dots, b_l so that $C(\mathcal{P}', b_i) = C(\mathcal{P}', b_{i-1}) + 1$ for each $i \in \{1, \dots, l\}$, where $b_0 = a_4^p$. By Proposition 3, such a schedule \mathcal{P}' exists and a is preempted p times in \mathcal{P}' .

Let \mathcal{P} be an optimal normal schedule for \mathcal{J} . By Theorem 4, such a schedule exists. Suppose for a contradiction that a is preempted at most $p - 1$ times in \mathcal{P} . By Proposition 3, we have $C(\mathcal{P}, a) < C(\mathcal{P}', a)$ and $C(\mathcal{P}, a_4^p) > C(\mathcal{P}', a_4^p)$. The number of events in \mathcal{P} (respectively in \mathcal{P}') in interval $[0, C(\mathcal{P}, a_4^p)]$ (respectively, $[0, C(\mathcal{P}', a_4^p)]$) is at most

$2|\mathcal{J}_p| + 3$ because each event either equals 0 or is the start or completion time of a job. Since both schedules are normal,

$$C(\mathcal{P}, a_4^p) - C(\mathcal{P}', a_4^p) \geq 1/2^{2|\mathcal{J}_p|+3}.$$

Thus, by definition of l ,

$$\begin{aligned} \sum_{i=1}^l C(\mathcal{P}, b_i) &\geq \sum_{i=1}^l C(\mathcal{P}', b_i) + l/2^{2|\mathcal{J}_p|+3} \\ &= \sum_{i=1}^l C(\mathcal{P}', b_i) + C(\mathcal{P}_p, a_4^p). \end{aligned} \tag{4}$$

By Proposition 3, \mathcal{P} restricted to the jobs in \mathcal{J}_p is not optimal for \mathcal{J}_p . This in particular implies

$$\sum_{x \in \mathcal{J}_p} C(\mathcal{P}, x) > \sum_{x \in \mathcal{J}_p} C(\mathcal{P}', x).$$

This, together with (4), implies

$$\begin{aligned} \sum_{x \in \mathcal{J}} C(\mathcal{P}, x) &= C(\mathcal{P}, a) + \sum_{x \in \mathcal{J}_p} C(\mathcal{P}, x) + \sum_{i=1}^l C(\mathcal{P}, b_i) \\ &> \sum_{x \in \mathcal{J}_p} C(\mathcal{P}', x) + \sum_{i=1}^l C(\mathcal{P}', b_i) + C(\mathcal{P}_p, a_4^p) \end{aligned}$$

Since, by construction of \mathcal{P}' ,

$$C(\mathcal{P}_p, a_4^p) = C(\mathcal{P}', a_4^p) \geq C(\mathcal{P}', a),$$

we obtain that the total completion time of \mathcal{P}' is strictly smaller than that of \mathcal{P} , which gives a required contradiction.

Finally, note that $C(\mathcal{P}_p, a_4^p) \leq |\mathcal{J}_p| \leq c$ and hence $l \leq 2^{2c}$. Thus, $|\mathcal{J}| = |\mathcal{J}_p| + 1 + l = 2^{O(p)}$ because $c = O(p)$ and $|\mathcal{J}_p| = O(p)$. This implies that $p = \Omega(\log |\mathcal{J}|)$ as required. \square

4 Optimal, normal and maximal schedules

4.1 Preliminaries

Let \mathcal{J} be a set of n unit UET jobs. The *release date* for job a , denoted by $r(a)$, is the earliest start time for a in any feasible schedule of \mathcal{J} . We assume that $r(a)$ is an integer for all $a \in \mathcal{J}$.

For two jobs a and b , we say that a is a *predecessor* of b , and that b is a *successor* of a , if all feasible schedules require that b not start until a has finished. We write $a \prec b$ to denote this relation. In contrast, $a \not\prec b$ means that b can start prior to the completion time of a . Two jobs a and b are

said to be *independent* if $a \not\prec b$ and $b \not\prec a$. For $B \subseteq \mathcal{J}$, we say that the jobs in B are *independent* if each pair of jobs in B is independent. This work deals with *in-tree* precedence constraints, i.e., for each job a there exists at most one job b such that $a \prec b$.

The symbol \mathbb{R}_+ denotes the set of nonnegative real numbers. Given a schedule \mathcal{P} and a job $a \in \mathcal{J}$, define $s(\mathcal{P}, a)$ and $C(\mathcal{P}, a)$ to be the start and completion times of a in \mathcal{P} , respectively. A job is called *release date pinned* in \mathcal{P} if it starts at its release date in \mathcal{P} . The *total completion time* of a schedule \mathcal{P} of \mathcal{J} is given by $\sum_{a \in \mathcal{J}} C(\mathcal{P}, a)$. We say that a preemptive schedule \mathcal{P} is *optimal* if the sum of its job completion times is minimum among all preemptive schedules for \mathcal{J} .

4.2 Events, partitions and basic schedule transformations

For a given schedule \mathcal{P} , define a vector $\mathbf{e} = (e_1, \dots, e_q)$, where $0 = e_1 < e_2 < \dots < e_q$, such that

$$\{e_1, \dots, e_q\} = \{0\} \cup \{s(\mathcal{P}, a) \mid a \in \mathcal{J}\} \cup \{C(\mathcal{P}, a) \mid a \in \mathcal{J}\}.$$

The elements of \mathbf{e} are called the *events* of \mathcal{P} . The part of \mathcal{P} in time interval $[e_i, e_{i+1}]$ is called the *i -th block* of \mathcal{P} , or simply a *block* of \mathcal{P} , $i \in \{1, \dots, q-1\}$. Given $i \in \{1, \dots, q-1\}$, let $\xi_i: \mathcal{J} \rightarrow \mathbb{R}_+$ be a function such that for each $a \in \mathcal{J}$, $\xi_i(a)$ is the total length of a executed in the i -th block of \mathcal{P} . Then, $(\xi_1, \dots, \xi_{q-1})$ is called the *partition* of \mathcal{P} . Denote by $(\mathcal{P}, \mathbf{e}, \boldsymbol{\xi})$ the schedule \mathcal{P} with events \mathbf{e} and partition $\boldsymbol{\xi}$. Unless specified otherwise, it is understood that \mathbf{e} has q components. For each $a \in \mathcal{J}$, $\tau_{\mathcal{P}}(a)$ is the integer $i \in \{1, \dots, q-1\}$ such that $C(\mathcal{P}, a) = e_{i+1}$. In other words, the $\tau_{\mathcal{P}}(a)$ -th block is the last block in which job a appears. Whenever \mathcal{P} is clear from context we will simply write $\tau(a)$. For any function $f: \mathcal{J} \rightarrow \mathbb{R}_+$, let

$$\mathcal{J}(f) = \{a \in \mathcal{J} \mid f(a) \neq 0\}.$$

In the following we will analyze schedules by investigating their events and partitions. Informally speaking, the events and the partition of a schedule \mathcal{P} are insufficient to uniquely reconstruct the schedule \mathcal{P} but they suffice to build a schedule with the same total completion time as \mathcal{P} . The schedules built from a list of events and a partition may differ in how pieces of jobs are executed within the blocks. The main advantage of our approach is that in order to construct a block in $[e_i, e_{i+1}]$ one only needs to solve the problem $P2|p_j, pmtn|C_{\max}$ where the execution time of a job a is $\xi_i(a)$; the proof of Lemma 1 gives more details. We formalize this observation in the next two lemmas.

Proposition 4 *Let $(\mathcal{P}, \mathbf{e}, \boldsymbol{\xi})$ be a schedule. The following conditions hold for each $i \in \{1, \dots, q-1\}$:*

- (i) *For each $a \in \mathcal{J}(\xi_i)$: $r(a) \leq e_i$;*
- (ii) *For each $a \in \mathcal{J}(\xi_i)$: $\xi_i(a) \leq e_{i+1} - e_i$ and $\sum_{a \in \mathcal{J}(\xi_i)} \xi_i(a) \leq 2(e_{i+1} - e_i)$;*
- (iii) *For each $a \in \mathcal{J}(\xi_i)$ and $b \in \mathcal{J}(\xi_j)$, where $i \leq j < q$: $b \not\prec a$.*

Proof Condition (i) follows from the fact that no job in $\mathcal{J}(\xi_i)$ starts or completes in (e_i, e_{i+1}) , $i \in \{1, \dots, q-1\}$. (Note that $r(a) > e_i$ is not possible for $a \in \mathcal{J}(\xi_i)$ because then we would have $s(\mathcal{P}, a) \in (e_i, e_{i+1})$ which would contradict e_i and e_{i+1} being two consecutive events of \mathcal{P} .) Conditions (ii) and (iii) follow directly from the fact that \mathcal{P} is a feasible schedule for \mathcal{J} . (Note that (iii) in particular implies that the jobs in $\mathcal{J}(\xi_i)$ are independent.) \square

We often rely on rearrangements of the events \mathbf{e} of a schedule \mathcal{P} which result in new schedules \mathcal{P}' with events that differ from those in \mathbf{e} . The resulting schedule \mathcal{P}' , however, may still be analyzed in the time intervals $[e_i, e_{i+1}]$, $i \in \{1, \dots, q-1\}$ defined by the original \mathbf{e} . For this analysis, we need the following lemma, in which vectors of increasing real numbers beginning with 0 are regarded as sequences of *time points*.

Lemma 1 *If there exist q time points $e_1 < \dots < e_q$ and $q-1$ functions $\xi_i: \mathcal{J} \rightarrow \mathbb{R}_+$ ($i = 1, \dots, q-1$) such that for each $a \in \mathcal{J}$, $\sum_{i=1}^{q-1} \xi_i(a) = 1$ and conditions (i)–(iii) in Proposition 4 are satisfied, then there exists a schedule \mathcal{P} such that for each $i \in \{1, \dots, q-1\}$ and for each $a \in \mathcal{J}$ the total length of all pieces of a executed in $[e_i, e_{i+1}]$ equals $\xi_i(a)$.*


Proof For any given $i \in \{1, \dots, q-1\}$, it is enough to construct the part of schedule \mathcal{P} , denoted by \mathcal{P}_i , in the time interval $[e_i, e_{i+1}]$. By (i) and (ii), this is equivalent to solving the problem $P2|p_j, pmtn|C_{\max}$ where the execution time of each job a is $\xi_i(a)$. It is easy to see that such a schedule \mathcal{P}_i exists if and only if the duration of $[e_i, e_{i+1}]$ is at least the larger of the maximum of the execution times $\xi_i(a)$ and the sum of these times averaged over the two processors, i.e.,

$$e_{i+1} - e_i \geq \max \left\{ \frac{1}{2} \sum_{a \in \mathcal{J}} \xi_i(a), \max \{ \xi_i(a) \mid a \in \mathcal{J} \} \right\}.$$

Thus, (ii) guarantees that \mathcal{P}_i exists. Finally note that (iii) guarantees that the precedence constraints between jobs in different blocks are met. \square

We close this section by introducing two basic transformations of a given schedule $(\mathcal{P}, \mathbf{e}, \boldsymbol{\xi})$: the *cyclic shift* and the *swapping* of two jobs. Let $\varepsilon > 0$ and $j > 0$. Let $B = \{a_1, \dots, a_j\} \subseteq \mathcal{J}$ be j different jobs and let $\{i_1, \dots, i_j\} \subseteq \{1, \dots, q-1\}$ be j blocks of \mathcal{P} such that $\xi_{i_k}(a_k) \geq \varepsilon$ and $\xi_{i_{k+1}}(a_k) \leq e_{i_{k+1}+1} - e_{i_{k+1}} - \varepsilon$ for $k \in$

Fig. 3 Swapping a and a' , where $I = \{k - 4, k - 3\}$, $k = \tau(a)$, leads in this case to a schedule \mathcal{P}' with a smaller total completion time



$\{1, \dots, j\}$, where $i_{j+1} = i_1$. We define a *cyclic shift* of B by ε on $\{i_1, \dots, i_j\}$ in \mathcal{P} , or just a cyclic shift if it is clear from context, as follows. Let

$$(e', \xi') = \langle e, \xi, \varepsilon, (i_1 \overset{a_1}{\rightsquigarrow} i_2 \overset{a_2}{\rightsquigarrow} \dots \overset{a_{j-1}}{\rightsquigarrow} i_j \overset{a_j}{\rightsquigarrow} i_1) \rangle$$

be the events and the partition, respectively, obtained by replacing a piece of a_{k+1} of length ε in block j_{k+1} of \mathcal{P} with a piece of a_k of length ε for each $k \in \{1, \dots, j\}$, where $i_{j+1} = i_1$. This transformation may not result in a feasible schedule because the precedence constraints or release dates may be violated. However, if neither is violated, then the assumptions of Lemma 1 are met for \mathcal{P}' with the events e' , and the partition ξ' exists. If \mathcal{P}' exists, then in addition we assume that the blocks of \mathcal{P}' enforce the following restrictions:

- For each $a_k \in B$, if $C(\mathcal{P}, a_k) = e_{i_{k+1}}$ and $i_{k+1} < i_k$ (taking $i_{j+1} = i_1$), then $C(\mathcal{P}', a_k) = e_{i_{k+1}} - \varepsilon$, which reduces the completion time of job a_k by as much as possible with respect to the cyclic shift.
- If $C(\mathcal{P}, a_k) \leq e_{i_{k+1}}$ (taking $i_{j+1} = i_1$), then $C(\mathcal{P}', a_k) = e_{i_{k+1}} + \varepsilon$, which increases the completion time of job a_k by as little as possible with respect to the cyclic shift.

Note that, in general, e does not consist of the events of \mathcal{P}' , and the number of events of \mathcal{P}' may be different than the number of events of \mathcal{P} .

Finally, we introduce the notion of swapping of two jobs which is used in Sects. 4.4 and 6 to reduce total completion time of a schedule by applying the shortest processing time (SPT) rule to two jobs that complete in consecutive blocks. Let \mathcal{P} be a schedule with events e and partition ξ . Let a and a' be two jobs such that $C(\mathcal{P}, a') = e_{\tau(a)}$, $s(\mathcal{P}, a) \leq s(\mathcal{P}, a')$ and a' is independent of any job in $\mathcal{J}(\xi_{\tau(a)})$. We define a transformation of *swapping a and a'* that results in a new schedule \mathcal{P}' as follows (see Fig. 3). Find a set of indices $I \subseteq \{1, \dots, \tau(a) - 1\}$ such that for each $j \in I$,

$$0 < \varepsilon(j) \leq \min \{e_{j+1} - e_j - \xi_j(a), \xi_j(a')\},$$

$\varepsilon(\max I)$ is minimum and $\sum_{j \in I} \varepsilon(j) = \xi_{\tau(a)}(a)$. Such a set I exists because of the constraints imposed on a and a' . The

schedule \mathcal{P}' is obtained by performing the following three steps:

- For each $j \in I$, remove a piece of a' of length $\varepsilon(j)$ from the j -th block of \mathcal{P} .
- Remove the piece of a executing in the $\tau(a)$ -th block and add a piece of a' of length $\xi_{\tau(a)}(a)$ to the $\tau(a)$ -th block of \mathcal{P} .
- Add a piece of a of length $\varepsilon(j)$ to the j -th block of \mathcal{P} for each $j \in I$.

Lemma 2 Given schedule (\mathcal{P}, e, ξ) , let a, a' be two jobs such that $C(\mathcal{P}, a') = e_k$, $s(\mathcal{P}, a) \leq s(\mathcal{P}, a')$ and a' is independent of any job in $\mathcal{J}(\xi_k)$, where $k = \tau_{\mathcal{P}}(a)$. Then, the schedule \mathcal{P}' obtained by swapping a and a' in \mathcal{P} is feasible and $\sum_{a'' \in \mathcal{J}} C(\mathcal{P}', a'') \leq \sum_{a'' \in \mathcal{J}} C(\mathcal{P}, a'')$ with the inequality being strict when $s(\mathcal{P}, a) < s(\mathcal{P}, a')$ and $\xi_{k-1}(a) < e_k - e_{k-1}$.

Proof The fact that \mathcal{P}' is feasible follows directly from its construction. Suppose that $s(\mathcal{P}, a) < s(\mathcal{P}, a')$ and $\xi_{k-1}(a) < e_k - e_{k-1}$. If $k - 1 \notin I$, then $C(\mathcal{P}', a) \leq e_{k-1} + \xi_{k-1}(a) < e_k$. Otherwise, the restriction on taking $\varepsilon(\max I) = \varepsilon(k - 1)$ to be minimum implies, due to $s(\mathcal{P}, a) < s(\mathcal{P}, a')$, that $\xi_{k-1}(a) + \varepsilon(k - 1) < e_k - e_{k-1}$ and hence $C(\mathcal{P}', a) = e_{k-1} + \xi_{k-1}(a) + \varepsilon(k - 1) < e_k$. Thus, the total completion time of \mathcal{P}' is strictly smaller than that of \mathcal{P} as required. \square

4.3 Properties of optimal schedules

We now give some key properties of optimal schedules and describe three configurations that are forbidden in optimal schedules. These results will be used in subsequent sections. The following lemma states that if a job a completes in the i -th block of an optimal schedule \mathcal{P} , i.e., $\tau(a) = i$, then the part of a that executes in that block spans the block. Such a job a is called a *spanning job in block i* .

Lemma 3 Given schedule (\mathcal{P}, e, ξ) , each job $a \in \mathcal{J}$ is a *spanning job in block $\tau(a)$* , i.e., $\xi_{\tau(a)}(a) = e_{\tau(a)+1} - e_{\tau(a)}$.

Proof The proof is by contradiction. There exists $\varepsilon > 0$ such that at most one job executes in $I = [e_{\tau(a)+1} - \varepsilon, e_{\tau(a)+1}]$

on each machine in \mathcal{P} and $\varepsilon \leq e_{\tau(a)+1} - e_{\tau(a)} - \xi_{\tau(a)}(a)$. Let B be the set of the jobs that execute in I . Clearly, $a \in B$ and $1 \leq |B| \leq 2$. There exists a job $b' \in \mathcal{J} \setminus B$ such that $\xi_{\tau(a)}(b') \neq 0$. Indeed, otherwise a could be executed in $[e_{\tau(a)}, e_{\tau(a)} + \xi_{\tau(a)}(a)]$ without making any other changes in the schedule. Since the new schedule completes a earlier (because $\xi_{\tau(a)}(a) < e_{\tau(a)+1} - e_{\tau(a)}$), this would contradict the optimality of \mathcal{P} . Then $C(\mathcal{P}, b') > \tau(a)$ and we can use some of the space of $\xi_{\tau(a)}(b')$ for job a to complete a earlier. More formally, define $\varepsilon' = \min\{\varepsilon, \xi_{\tau(a)}(b')\}$. Let $e' = e_{\tau(a)+1} - \varepsilon'$ and for each job $c \in \mathcal{J}$ let

$$\xi'(c) = \begin{cases} \xi_{\tau(a)}(c), & \text{if } c \notin \{b'\} \cup (B \setminus \{a\}), \\ \xi_{\tau(a)}(c) - \varepsilon', & \text{if } c \in \{b'\} \cup (B \setminus \{a\}), \end{cases}$$

and

$$\xi''(c) = \begin{cases} 0, & \text{if } c \notin \{b'\} \cup (B \setminus \{a\}), \\ \varepsilon', & \text{if } c \in \{b'\} \cup (B \setminus \{a\}). \end{cases}$$

By Lemma 1, there exists a schedule \mathcal{P}' such that for each $t \in \{1, \dots, q-1\} \setminus \{\tau(a)\}$ the total length of all pieces of each job $c \in \mathcal{J}$ executed in $[e_t, e_{t+1}]$ is $\xi_t(c)$, the total length of all pieces of each job c executed in $[e_{\tau(a)}, e']$ equals $\xi'(c)$, and the total length of all pieces of each job c executed in $[e', e_{\tau(a)+1}]$ equals $\xi''(c)$. However, $C(\mathcal{P}', c) = C(\mathcal{P}, c)$ for each $c \in \mathcal{J} \setminus \{a\}$ and $C(\mathcal{P}', a) < C(\mathcal{P}, a)$, which contradicts the optimality of \mathcal{P} . \square

Lemma 4 Given schedule (\mathcal{P}, e, ξ) , if $a \in \mathcal{J}$ is not a spanning job in block i ($i \in \{1, \dots, q-2\}$), $s(\mathcal{P}, a) \leq e_i$ and $C(\mathcal{P}, a) \geq e_{i+1}$, then there is no idle time in the i -th block of \mathcal{P} .

Proof Suppose for a contradiction that there is idle time of length $\varepsilon > 0$ in $[e_i, e_{i+1}]$ on one of the processors in \mathcal{P} . We get a contradiction by obtaining another schedule \mathcal{P}' such that $C(\mathcal{P}, b) = C(\mathcal{P}', b)$ for each $b \in \mathcal{J} \setminus \{a\}$ and $C(\mathcal{P}', a) < C(\mathcal{P}, a)$. Namely, take $\varepsilon' = \min\{\varepsilon, \xi_{\tau(a)}(a), e_{i+1} - e_i - \xi_i(a)\}$. By Lemma 3, $\tau(a) > i$ and hence $\varepsilon' > 0$. By Lemma 1, the desired schedule \mathcal{P}' obtained from \mathcal{P} by moving the piece of a that executes in $[C(\mathcal{P}, a) - \varepsilon', C(\mathcal{P}, a)]$ to the i -th block of \mathcal{P} is feasible. \square

Given schedule (\mathcal{P}, e, ξ) , two jobs a and b with $\tau(a) < \tau(b)$ are said to *interlace* if job b is not spanning in block $\tau(a)$ and there exists $t < \tau(a)$ such that job a is not spanning in block t , $\xi_t(b) > 0$, $r(a) < e_{t+1}$ and a is independent of all jobs in $\mathcal{J}(\xi_t) \cup \dots \cup \mathcal{J}(\xi_{\tau(a)})$. Note that, informally speaking, the above constraints imply that a piece of a executed in $[C(\mathcal{P}, a) - \varepsilon, C(\mathcal{P}, a)]$, for some $\varepsilon > 0$, can be exchanged with a piece of b of length ε executing in the t -th block of \mathcal{P} . We formalize this observation in the next lemma.

Lemma 5 If \mathcal{P} is an optimal schedule, then no two jobs interlace in \mathcal{P} .

Proof Let e and ξ be the events and the partition of \mathcal{P} , respectively. Suppose for a contradiction that two jobs a and b with $\tau(a) < \tau(b)$ interlace and t is the block in the definition. Let

$$\varepsilon = \min\{\xi_t(b), e_{\tau(a)+1} - e_{\tau(a)} - \xi_{\tau(a)}(b), \xi_{\tau(a)}(a), e_{t+1} - e_t - \xi_t(a)\}.$$

Note that $\varepsilon > 0$. By Lemma 1, there exists a schedule \mathcal{P}' with e' and partition ξ' such that

$$(e', \xi') = (e, \xi, \varepsilon, (t \overset{b}{\rightsquigarrow} \tau(a) \overset{a}{\rightsquigarrow} t)).$$

The schedule \mathcal{P}' is feasible for two reasons. First, $r(a) < e_{t+1}$ implies that if $s(\mathcal{P}, a) < e_{t+1}$, then $r(a) \leq e_t$ and if $s(\mathcal{P}, a) \geq e_{t+1}$, then $s(\mathcal{P}', a) \geq e_{t+1} - \varepsilon$ according to the definition of the transformation, which implies that a does not start prior to its release date in \mathcal{P}' . Second, the fact that a is independent of all jobs in $\mathcal{J}(\xi_t) \cup \dots \cup \mathcal{J}(\xi_{\tau(a)})$ implies that a does not violate the precedence constraints in \mathcal{P}' . For each $c \in \mathcal{J} \setminus \{a\}$, $C(\mathcal{P}, c) = C(\mathcal{P}', c)$ and $C(\mathcal{P}, a) > C(\mathcal{P}', a)$. This contradicts the optimality of \mathcal{P} . \square

Lemma 6 Let (\mathcal{P}, e, ξ) be an optimal schedule. Let $I = [x, y]$ be an interval and let $B \subseteq \mathcal{J}$ be such that $a \in B$ if and only if the total length of job a executing in I is strictly between 0 and $y - x$.

If jobs in B are independent, $C(\mathcal{P}, a) \geq y$ and $r(a) \leq x$ for each $a \in B$, then $|B| \leq 2$.

Proof It follows from definition of set B that no job completes in (x, y) . We first argue that

$$C(\mathcal{P}, b) > y \text{ for each } b \in B. \quad (5)$$

Suppose for a contradiction that $C(\mathcal{P}, b) = y$ for some job $b \in B$. Since the total length of b in I is less than $y - x$, there exists a non-empty interval $I' \subseteq I$ such that no part of b executes in I' . We obtain a schedule \mathcal{P}' by exchanging the part of \mathcal{P} that executes in I' with the part of \mathcal{P} that executes in $[y - |I'|, y]$. Since the release date of each job that executes in I is at most x and the jobs whose parts execute in I are independent, we obtain that \mathcal{P}' is indeed a feasible schedule. Then, $C(\mathcal{P}', b) = y - |I'| < y = C(\mathcal{P}, b)$ and $C(\mathcal{P}', a) \leq C(\mathcal{P}, a)$ for each $a \in \mathcal{J} \setminus \{b\}$, which completes the proof of (5).

We now prove the lemma. Suppose for a contradiction that $|B| > 2$. Let b be a job in B with minimum completion time in \mathcal{P} . Since $|B| > 2$, Lemma 3 implies that there exists $b' \in B$ such that $\tau(b) < \tau(b')$ and b' is not a spanning job in chunk $\tau(b)$. Define $\varepsilon = \min\{y - x - p, \xi_{\tau(b)}(b), p', e_{\tau(b)+1} - e_{\tau(b)} - \xi_{\tau(b)}(b')\}$, where p and p' are the total lengths of

b and b' respectively executing in I . Due to the choice of b' , $\varepsilon > 0$. We obtain a schedule \mathcal{P}' by first exchanging the pieces of b' of total length ε executing in I with a piece of b of length ε executing in chunk $\tau(b)$. The resulting \mathcal{P}' may not be feasible in I , however, the McNaughton's rule can readily turn this part into a feasible schedule. This provides a feasible schedule \mathcal{P}' because the release date of each job whose part executes in I is at most x and the jobs that execute in I in \mathcal{P} are independent. By (5), $C(\mathcal{P}', b) = C(\mathcal{P}, b) - \varepsilon$. Note that if a job completes at y in \mathcal{P} , then the total length of this job in I equals $y - x$; otherwise the job would belong to B contradicting (5). Thus, no job completes later in \mathcal{P}' than in \mathcal{P} —a contradiction with the optimality of \mathcal{P} . \square

Lemma 7 *Let schedule (\mathcal{P}, e, ξ) be optimal. If $\mathcal{J}(\xi_i) \neq \emptyset$ ($i \in \{1, \dots, q - 1\}$), then:*

- (i) *There exists a job in \mathcal{J} that is spanning in block i ;*
- (ii) *$|\mathcal{J}(\xi_i)| \leq 3$ and if $|\mathcal{J}(\xi_i)| = 3$, then some job in $\mathcal{J}(\xi_i)$ completes at e_{i+1} in \mathcal{P} .*

Proof Note that $|\mathcal{J}(\xi_i)| > 3$ would lead to a contradiction to Lemma 6 with $I = [e_i, e_{i+1}]$. Moreover, if $i = \tau(a)$ for some $a \in \mathcal{J}$, then by Lemma 3, a is spanning in block i and the lemma holds.

Thus, assume that no job finishes in the i -th block of \mathcal{P} . If $|\mathcal{J}(\xi_i)| \leq 2$, then it remain to prove (i): if no job a is spanning in block i , then by Lemma 4, there is no idle time in the i -th block of \mathcal{P} , which would violate $|\mathcal{J}(\xi_i)| \leq 2$. This completes the proof of case $|\mathcal{J}(\xi_i)| \leq 2$. We prove, by contradiction, that $|\mathcal{J}(\xi_i)| = 3$ is not possible if no job completes at e_{i+1} . Denote $B = \{a \in \mathcal{J} \mid 0 < \xi_i(a) < e_{i+1} - e_i\}$. Clearly, $|B| > 1$. On the other hand, $|B| < 3$ for otherwise the job in B with smallest completion time interlaces with one of the two other jobs in B , which contradicts Lemma 5. Thus, $|B| = 2$. Denote $B = \{b, b'\}$ and assume without loss of generality that $C(\mathcal{P}, b) \leq C(\mathcal{P}, b')$. According to Lemma 3, job b is spanning in block $\tau(b)$. Also job b' is spanning in block $\tau(b)$, since otherwise b and b' interlace, which is not possible according to Lemma 5. The only job, call it c' , in $\mathcal{J}(\xi_i) \setminus \{b, b'\}$ completes in $(e_{i+1}, e_{\tau(b)})$ for otherwise this job and b interlace—again a contradiction with Lemma 5. Thus, in particular, $\tau(b) > i + 1$. This situation is depicted in Fig. 4.




Fig. 4 The proof of Lemma 7: the positioning of jobs b, b' and c'




Fig. 5 **a** The illustration of Lemma 8; **b** the illustration of Lemma 9

Let $Y = \{a \in \mathcal{J} \setminus \mathcal{J}(\xi_i) \mid e_{i+1} \leq s(\mathcal{P}, a) \leq e_{\tau(b)}\}$. Since e_{i+1} is an event of \mathcal{P} , $Y \neq \emptyset$. By Lemma 5, if $e_{i+1} \leq C(\mathcal{P}, a) \leq e_{\tau(b)}$, then $a \in Y$ or $a = c'$. If there exists $c \in Y$ such that $C(\mathcal{P}, c) = e_{\tau(b)}$, then we obtain a schedule \mathcal{P}' by swapping b and c . By Lemma 2, \mathcal{P}' is feasible. Moreover, if job b is non-spanning in block $\tau(b) - 1$, then the total completion time of \mathcal{P}' is smaller than that of \mathcal{P} , which completes the proof. If, on the other hand, job b is spanning in block $\tau(b) - 1$, then $C(\mathcal{P}', b) = e_{\tau(b)}$ and $\xi'_{\tau(b)-1}(b) = 0$, where ξ' is the partition of \mathcal{P}' , in which case b and b' interlace in \mathcal{P}' —a contradiction with Lemma 5. Thus, it remains to consider the situation when no such c exists. This, since $e_{\tau(b)}$ is an event of \mathcal{P} , implies that c' ends at $e_{\tau(b)}$ in \mathcal{P} . Moreover, $\mathcal{J}(\xi_{\tau(c')}) \subseteq \{c', b, b'\}$ for otherwise \mathcal{P} would not be optimal. Thus, some job $c \in Y$ ends at $e_{\tau(c')}$ because $e_{\tau(c')}$ is an event of \mathcal{P} and no job in Y can start at $e_{\tau(c')}$. Therefore, one of jobs $\{c', b\}$ must be non-spanning in block $\tau(c)$. Swapping this job with c gives, by Lemma 2, a schedule with smaller total completion time than that of \mathcal{P} , which provides the required contradiction and completes the proof of the lemma. \square

The following two lemmas describe additional configurations that cannot be present in an optimal schedule. The first situation is depicted in Fig. 5a, while the statement of Lemma 9 is shown in Fig. 5b.

Lemma 8 *Given schedule (\mathcal{P}, e, ξ) , let e_j be an event in \mathcal{P} and jobs c, c' , and d be such that*

- (i) $C(\mathcal{P}, c) = C(\mathcal{P}, c') = e_j$;
- (ii) $C(\mathcal{P}, d) = e_{j+1}$ and $s(\mathcal{P}, d) < e_j$;
- (iii) *Jobs in $\{c, c'\} \cup \mathcal{J}(\xi_j)$ are independent.*

Then, \mathcal{P} is not optimal.

Proof By Lemma 7 (ii), one of jobs c or c' , say c , satisfies $\xi_t(c) = 0$, where $e_t = s(\mathcal{P}, d)$. We then have $s(\mathcal{P}, d) < s(\mathcal{P}, c)$ for otherwise c and d would interlace, thus we get a contradiction by Lemma 5. Therefore, we can swap jobs c and d . By Lemma 2, the resulting schedule is feasible and has smaller total completion time than \mathcal{P} , as required. \square

Lemma 9 *Let schedule (\mathcal{P}, e, ξ) be optimal and jobs a and a' be such that*

- $\xi_j(a) > 0$, $\xi_{j'}(a) > 0$, $j < j' - 1$ and job a is spanning in block t for each $t \in \{j + 1, \dots, j' - 1\}$;
- $\xi_j(a') = 0$ and $C(\mathcal{P}, a') = e_{j'}$;
- No successor of a' starts at $e_{j'}$.

Then, $s(\mathcal{P}, a') \geq e_{j+1}$ and $\tau(a) > j'$.

Proof If $s(\mathcal{P}, a') < e_{j+1}$, then due to $\xi_j(a') = 0$, $s(\mathcal{P}, a') < e_j$. But then, a and a' would interlace, which is not possible in an optimal schedule according to Lemma 5.

By assumption, a' is independent of any job in $\mathcal{J}(\xi_{j'})$. Also, $s(\mathcal{P}, a) \leq e_{j+1} \leq s(\mathcal{P}, a')$. Then, $\tau(a) > j'$ follows from an observation that otherwise swapping a and a' in \mathcal{P} would produce, by Lemma 2, a schedule with smaller total completion time than that of \mathcal{P} . \square

4.4 Abnormality points and maximal schedules

We now define normal schedules, abnormality points and maximal schedules. In particular Lemma 12 gives key necessary conditions for an abnormality point.

For any $x \in \mathbb{R}_+$ and nonnegative integer l , we say that x is l -normal if $x = l'/2^l$ for some integer l' . We say that a block of a schedule \mathcal{P} is l -normal if the length of the block is l -normal and the total execution time of each job in the block is $(l + 1)$ -normal. A preemptive schedule \mathcal{P} with q events is normal if the i -th block of \mathcal{P} is i -normal for each $i \in \{1, \dots, q - 1\}$. If a schedule \mathcal{P} with q events e and partition ξ is not normal, then the minimum index $i \in \{1, \dots, q - 1\}$ such that the i -th block of \mathcal{P} is not i -normal is called the *abnormality point* of \mathcal{P} . If a schedule is normal, then its abnormality point is denoted by ∞ for convenience. We have the following simple observations.

Observation 1 If x is l -normal, then x is l' -normal for each $l' \geq l$. \square

Observation 2 If $i \neq \infty$ is the abnormality point of a schedule \mathcal{P} with events e , then e_i is $(i - 1)$ -normal. \square

According to our definition, if an i -th block of a schedule \mathcal{P} is i -normal, then $\xi_i(a)$ is $(i + 1)$ -normal for each $a \in \mathcal{J}$, however, this does not necessarily imply that job preemptions occur at $(i + 1)$ -normal time points in the i -th block of \mathcal{P} . Such job preemptions can possibly take place only strictly between e_i and e_{i+1} since both e_i and e_{i+1} are i -normal by assumption. By the next observation, we may assume without loss of generality that i -normal blocks have job preemptions only at $(i + 1)$ -normal time points.

Observation 3 If the i -th block of a schedule \mathcal{P} is i -normal, then there exists a schedule \mathcal{P}' with the same events, partition and total completion time as that of \mathcal{P} , in which each preemption, resumption, job start and job completion in the i -th block occurs at $(i + 1)$ -normal time point.

Proof It follows from the McNaughton's algorithm. \square

Let us introduce a partial order, denoted by \trianglelefteq , to the set of all schedules. For schedules \mathcal{P} and \mathcal{P}' , we write $\mathcal{P} \trianglelefteq \mathcal{P}'$ if and only if one of the following holds:

- $\mathcal{P} = \mathcal{P}'$;
- \mathcal{P}' is optimal, while \mathcal{P} is not;
- Both \mathcal{P} and \mathcal{P}' are optimal and, additionally, \mathcal{P}' is normal while \mathcal{P} is not;
- Both \mathcal{P} and \mathcal{P}' are optimal, but neither is normal. Additionally $i \leq i'$, where i and i' are the abnormality points of \mathcal{P} and \mathcal{P}' , respectively.

Any element in \mathcal{J} that is maximal under the partial order is called a *maximal* schedule.

Lemma 10 Let schedule (\mathcal{P}, e, ξ) have abnormality point $i \neq \infty$. For each $a \in \mathcal{J}$ and for each $i' \leq i$, $\sum_{j=i'}^{q-1} \xi_j(a)$ is i' -normal.

Proof Let $a \in \mathcal{J}$ be selected arbitrarily. Note that

$$\sum_{j=i'}^{q-1} \xi_j(a) = 1 - \sum_{j=1}^{i'-1} \xi_j(a).$$

Since $i \geq i'$ is the abnormality point of \mathcal{P} , Observation 1 implies that $\xi_j(a)$ is i -normal for each $j \in \{1, \dots, i' - 1\}$. \square

The next lemma, informally speaking, allows us to further consider only those maximal schedules with abnormality point $i \neq \infty$ in which the abnormality of the i -th block is due to the length of the jobs in this block, and not due to the length of this block.

Lemma 11 Let \mathcal{P} be a maximal schedule with the events e_1, \dots, e_q . If $i \neq \infty$ is the abnormality point of \mathcal{P} , then there exists a maximal schedule \mathcal{P}' with abnormality point i such that $e_1, \dots, e_i, e'_{i+1}, \dots, e'_q$ are its events and $e'_{i+1} - e_i$ is i -normal.

Proof If $i = \tau(a)$ for some $a \in \mathcal{J}$, then by Lemma 3, $e_{i+1} = e_i + \xi_i(a)$. By Lemma 10, $\xi_i(a)$ is i -normal. Thus, $\mathcal{P}' = \mathcal{P}$ is the required schedule, which proves the lemma. Hence, $i \neq \tau(a)$ for each $a \in \mathcal{J}$, i.e., no job completes at e_{i+1} . Note that $e_{i+1} = s(\mathcal{P}, a)$ for some $a \in \mathcal{J}$. Suppose for a contradiction that $e_{i+1} - e_i$ is not i -normal. Thus, e_{i+1} is not i -normal. Lemma 4 implies that there is no idle time in the i -th block of \mathcal{P} . Thus, $|\mathcal{J}(\xi_i)| \geq 2$ and therefore there exists $d \in \mathcal{J}(\xi_i)$ that is non-spanning in block $i + 1$ because a starts at e_{i+1} .

Case 1 There is an i -normal number in $(e_i, e_{i+1}]$. Let x be the maximal i -normal number in $(e_i, e_{i+1}]$. Then, $r(a) \leq x$

because there is no i -normal number in $(x, e_{i+1}]$ and $r(a) \leq e_{i+1}$ is i -normal by Observation 1. Let

$$0 < \varepsilon \leq \min \{ \xi_{i+1}(a), e_{i+1} - x, \xi_i(d), e_{i+2} - e_{i+1} - \xi_{i+1}(d) \}.$$

No job completes at e_{i+1} and therefore the jobs in $\mathcal{J}(\xi_i) \cup \mathcal{J}(\xi_{i+1})$ are independent. Thus, by Lemma 1, there exists a schedule \mathcal{P}' with events e' and partition ξ' , where

$$(e', \xi') = (e, \xi, \varepsilon, (i \overset{d}{\rightsquigarrow} i + 1 \overset{a}{\rightsquigarrow} i)).$$

Moreover, due to the McNaughton’s rule, one can assume that $s(\mathcal{P}', a) = x$. By Observation 2, e_i is $(i - 1)$ -normal and hence, by Observation 1, $x - e_i$ is i -normal. Since the first $i - 1$ blocks are identical in \mathcal{P} and \mathcal{P}' , and $e'_{i+1} = x$, \mathcal{P}' is the desired schedule, which completes the proof in this case.

Case 2 There is no i -normal number in $(e_i, e_{i+1}]$. By Observation 1, there is no $(i - 1)$ -normal number in $(e_i, e_{i+1}]$. Let $x > e_i$ be the minimum $(i - 1)$ -normal number. Since $i + 1 < q$, more than one block intersects (e_i, x) .

Suppose first that exactly two blocks intersect (e_i, x) , and there exists a job b such that $\xi_i(b) + \xi_{i+1}(b) = e_{i+2} - e_i$. One of the two blocks is of length at least $(x - e_i)/2$. By Observation 2, $e_i + (x - e_i)/2$ is i -normal. Hence, due to the condition in Case 2, this must be the $(i + 1)$ -st block. However, the schedule with events $(e_1, \dots, e_i, e_i + e_{i+2} - e_{i+1}, e_{i+2}, \dots, e_q)$ and partition $(\xi_1, \dots, \xi_i, \xi_{i+2}, \xi_{i+1}, \xi_{i+3}, \dots, \xi_{q-1})$ would satisfy the assumption in Case 1. This allows us to construct the desired schedule \mathcal{P}' as in Case 1.

Suppose now that exactly two blocks intersect (e_i, x) and there exists no job b such that $\xi_i(b) + \xi_{i+1}(b) = e_{i+2} - e_i$. Lemma 6 applied to $I = [e_i, e_{i+2}]$ gives a contradiction. Observe that the corresponding set $\{a, b, d\} \subseteq B$ in Lemma 6 is of size at least 3. Moreover, since no job completes in (e_i, e_{i+2}) , B contains only independent jobs.

Finally, we consider the case when more than two blocks intersect (e_i, x) . Thus, the job a does not complete before x . Moreover, no job completes at e_{i+2} because otherwise either \mathcal{P} is not optimal or e_{i+2} is i -normal by Lemma 10. Since $e_{i+2} < x$ we get a contradiction in either case. Therefore, there is a job a' that starts at e_{i+2} . Clearly, a' does not complete before x . Thus, Lemma 6 for $I = [e_i, \min\{x, e_{i+3}\}]$ again gives a contradiction. Observe that the corresponding set $\{a, a', d\} \subseteq B$ in Lemma 6 is of size at least 3. Moreover, since no job completes in (e_i, e_{i+3}) , B contains only independent jobs. \square

Given schedule (\mathcal{P}, e, ξ) , for $i \in \{1, \dots, q - 1\}$ define

$$A_i(\mathcal{P}) = \{a \in \mathcal{J} \mid \xi_i(a) \text{ is not } (i + 1)\text{-normal}\}.$$

Lemma 12 *Let \mathcal{P} be a maximal schedule. If $i \neq \infty$ is the abnormality point of \mathcal{P} , then $|A_i(\mathcal{P})| = 2$ and $|\mathcal{J}(\xi_i)| = 3$.*

Proof Let e and ξ be the events and the partition of \mathcal{P} , respectively. By Lemma 11, $e_{i+1} - e_i$ is i -normal. We have $|\mathcal{J}(\xi_i)| > 2$ because otherwise by Lemmas 4 and 3, $\xi_i(a) = e_{i+1} - e_i$ for each $a \in \mathcal{J}(\xi_i)$, which would contradict the fact that $e_{i+1} - e_i$ is i -normal. Lemma 7 implies that $|\mathcal{J}(\xi_i)| = 3$ and there exists $a \in \mathcal{J}(\xi_i)$ such that $\xi_i(a) = e_{i+1} - e_i$. Thus, $a \notin A_i(\mathcal{P})$, and we have that $|A_i(\mathcal{P})| \leq 2$ because $A_i(\mathcal{P}) \subseteq \mathcal{J}(\xi_i)$. Also, $|A_i(\mathcal{P})| > 1$ by Lemma 4. \square

Finally, the following observation shows that the abnormality point of a schedule does not decrease after a certain type of schedule modifications. The observation directly follows from the definition of normality.

Lemma 13 *Let \mathcal{P} be a schedule and $\varepsilon = l'/2^l$ for some integers l and l' . If \mathcal{P}' is a schedule that is obtained from \mathcal{P} by a sequence of modifications, each modification being a removal of a piece of length that is a multiple of ε from a j -th block and an insertion of this piece into a j' -th block, where $\min\{j, j'\} \geq l - 1$, then the abnormality point of \mathcal{P}' is not smaller than that of \mathcal{P} . \square*

We close this section with a brief summary of key properties of optimal and maximal schedules obtained in this section.

An optimal schedule is a concatenation of blocks. The structure of any block is as follows. Each non-empty block has a spanning job, and pieces of at most three jobs can be present in a block; if two jobs are present in a block, then both are spanning (cf. Lemmas 4 and 7) while if three jobs are present in a block, then one of them completes in the block (cf. Lemma 7) and it is the only spanning job in this block (cf. Lemma 3). The latter property from Lemma 3 is general—each job that ends in a block is spanning. We also emphasize that interlacing of jobs is not allowed in optimal schedules (see Lemma 5). This observation is heavily used throughout the paper.

Later on in Sect. 4.4 we switch our focus to optimal schedules that are also maximal. Lemma 11 shows that we may assume without loss of generality that, in a maximal schedule with abnormality point $i \neq \infty$, the length of block i is i -normal and thus there exists a job a in block i such that the length of a in this block is not i -normal. Then, Lemma 12 says that there must be another job b , different from a , that has its total length in block i not i -normal either. Note that, due to Lemmas 3, 4 and 7, this implies that neither a nor b is

spanning in block i , none of them finishes in this block and there must be a third spanning job that finishes in block i .

5 A-configurations

In this section we first define an undesirable structure that may appear in a schedule; we refer to this structure as an *A-configuration*. Our proof that there exists a normal optimal schedule for each \mathcal{J} , given in Sect. 6, relies on the key assumption that there exists a maximal schedule without A-configurations, or *A-free* maximal schedule, for each set of jobs \mathcal{J} . The proof proceeds by contradiction. The abnormality point implies the existence of an *alternating chain* in a maximal schedule (Sect. 6.1 for definition of alternating chains). In Sect. 6.3, we show that if A-configuration is not present in the part of the schedule that follows the alternating chain, then we are able to extend this chain by adding one more job. Thus we arrive at the required contradiction since this method allows us to extend the alternating chain *ad infinitum* in the presence of a finite number of jobs.


Therefore, the A-configurations are undesirable and the main goal of this section is to prove that an A-free maximal schedule exists for each \mathcal{J} . Our proof is by contradiction: informally speaking, we take a maximal schedule having an A-configuration as early as possible, and, after some schedule transformations, we either obtain a new schedule with smaller total completion time or with an earlier A-configuration. In the former case we clearly obtain a contradiction. In the latter case, a contradiction occurs only if the new schedule is maximal, i.e., its abnormality point is not smaller than that of the initial schedule. For this reason, while performing the initial schedule transformations we must ensure that they do not change the abnormality point in the latter case. The proof works for in-trees, however, it does not for general precedence constraints. The question whether there is an A-free maximal schedule for each \mathcal{J} and general precedence constraints remains open.

Let $(\mathcal{P}, \mathbf{e}, \xi)$ be a schedule. We say that \mathcal{P} has an *A-configuration of length ℓ ($\ell > 0$) starting at e_j* if there exist two jobs a and b such that

- $C(\mathcal{P}, a) = e_j$ and $C(\mathcal{P}, b) = e_{j'}$ for some $j' > j$;
- $[e_j - \ell, e_j]$ is a maximal interval where a executes non-preemptively;
- b executes non-preemptively in $[e_j, e_{j'}]$, and b does not execute in $[e_j - \ell, e_j]$;
- $s(\mathcal{P}, b) < e_j - \ell$;
- a and each job in $\mathcal{J}(\xi_j) \cup \dots \cup \mathcal{J}(\xi_{j'})$ are independent.

We also say that the jobs a and b form the A-configuration. See Fig. 6 for an exemplary A-configuration.

Fig. 6 An example of A-configuration



If no pair of jobs form an A-configuration in \mathcal{P} , then \mathcal{P} is called *A-free*. For any time interval I , if for any $x \in I \cap \{e_i : i = 1, \dots, q\}$ there is no A-configuration at x in \mathcal{P} , then \mathcal{P} is *A-free in I* . The main result of this section is the following theorem.

Theorem 2 *If any maximal schedule has an abnormality point $i \neq \infty$, then there exists an A-free maximal schedule.*

We first provide several technical lemmas before presenting our proof of Theorem 2. A schedule \mathcal{P} of abnormality point i is said to be *A-maximal* if it is maximal and, unless $i = \infty$, one of the following two statements is true:

- \mathcal{P} is A-free;
- any maximal schedule is not A-free, and \mathcal{P} has the earliest starting A-configuration among maximal schedules.

Proposition 5 *Let \mathcal{P} be A-maximal. If a and b form an A-configuration in \mathcal{P} with $C(\mathcal{P}, a) < C(\mathcal{P}, b)$, then $s(\mathcal{P}, a) \leq s(\mathcal{P}, b)$.*

Proof Suppose for a contradiction that $s(\mathcal{P}, a) > s(\mathcal{P}, b)$. Then, swapping jobs a and b in \mathcal{P} produces, by Lemma 2, a schedule with smaller total completion time than that of \mathcal{P} —a contradiction. \square

The first of the following two lemmas describes a situation that guarantees an A-configuration, while the second a situation that cannot happen in an A-maximal schedule with an A-configuration.

Proposition 6 *Given maximal schedule $(\mathcal{P}, \mathbf{e}, \xi)$, let e_j be an event in \mathcal{P} and jobs a, c, c_1 , and d be such that*

- (i) $C(\mathcal{P}, c_1) = e_j$, $C(\mathcal{P}, c) = e_{j+1}$, and $C(\mathcal{P}, d) = e_{j+2}$;
- (ii) $\mathcal{J}(\xi_{j-1}) = \{c, c_1\}$, $\mathcal{J}(\xi_j) = \{a, c\}$ and $\mathcal{J}(\xi_{j+1}) = \{a, d\}$;
- (iii) $s(\mathcal{P}, d) < e_{j-1}$.

Then, jobs c and d form an A-configuration. (See Fig. 7a for an illustration.)

Proof Let $k < j - 1$ be the maximum index such that in block $k - 1$ job c is non-spanning but spanning in block t for each $t \in \{k, \dots, j - 1\}$. Note that by (i), (ii) and Lemma 4, k is well defined. We prove, by induction on $t \in \{1, \dots, j - k\}$, that

$$\xi_{j-t}(c) = \xi_{j-t}(c_t) = e_{j-t+1} - e_{j-t} \wedge C(\mathcal{P}, c_t) = e_{j-t+1} \text{ for some } c_t \in \mathcal{J} \setminus \{a, d\} \wedge s(\mathcal{P}, c) < e_{j-t}, \quad (6)$$




Fig. 7 a Illustration of Proposition 6. b Illustration of Proposition 7

which immediately follows from (i), (ii) and Lemma 4 for $t = 1$. So assume inductively that the claim holds for some $t - 1 \geq 1$ and we prove it for t . It suffices to argue that some job c_t completes at e_{j-t+1} since then Lemma 3 implies (c_t) is spanning in block $j - t$. By the induction hypothesis and the fact that all jobs have the same execution time, neither c_{t-1} nor c start at e_{j-t+1} . Since e_{j-t+1} is an event of \mathcal{P} , some job c_t completes at e_{j-t+1} as required. We have $s(\mathcal{P}, c) < e_{j-t}$ for otherwise we can swap jobs c and d in $[s(\mathcal{P}, c), C(\mathcal{P}, d)]$. The resulting schedule is feasible and has smaller total completion time than \mathcal{P} . Thus \mathcal{P} is not optimal—contradiction. This proves (6).

If $0 < \xi_{k-1}(c) < e_k - e_{k-1}$, then by modifying the schedule in block $k - 1$ we may without loss of generality assume that c resumes at e_k . Thus, (6) implies that c and d form an A-configuration of length $e_{j+1} - e_k$ at e_{j+1} . \square

Proposition 7 Let (\mathcal{P}, e, ξ) be an A-maximal schedule. Suppose that jobs a and b form an A-configuration at e_j in \mathcal{P} with $C(\mathcal{P}, a) < C(\mathcal{P}, b)$. Then there exists no $e_t \leq s(\mathcal{P}, b)$ such that (see Fig. 7b for an illustration, where it is possible that job c ends at the start of b):

- (i) $r(a) < e_t$, $\xi_t(a) = e_{t+1} - e_t$, job a is non-spanning in block $t - 1$;
- (ii) Jobs in $\mathcal{J}(\xi_{t-1}) \cup \{a\}$ are independent;
- (iii) Some job c in $\mathcal{J}(\xi_{t-1})$ satisfies $C(\mathcal{P}, c) \geq e_{j'}$, $\xi_{j'}(c) = 0$, and if $C(\mathcal{P}, c) = e_{j'}$, then the jobs in $(\mathcal{J}(\xi_{j'}) \setminus \{b\}) \cup \{c\}$ are independent, where $e_{j'} = s(\mathcal{P}, b)$;
- (iv) The abnormality point of \mathcal{P} is not in $\{t, \dots, j'\}$.

Proof Suppose for a contradiction that such an e_t exists. Let $\ell > 0$ be the length of the A-configuration formed by a and b . Define

$$\varepsilon = \min \{ \xi_{t-1}(c), e_t - r(a), e_t - e_{t-1} - \xi_{t-1}(a), \xi_{j'}(b), \ell \}.$$

By (i) and (iii), we have $\varepsilon > 0$. Let \mathcal{P}' be a schedule obtained by moving a piece of c of length ε from the $(t - 1)$ -st block




Fig. 8 Schedule transformation in the proof of Proposition 7




Fig. 9 An illustration of Proposition 8

to the j' -th block, a piece of b of length ε from the j' -th block to $[C(\mathcal{P}, a) - \varepsilon, C(\mathcal{P}, a)]$, and a piece of a from $[C(\mathcal{P}, a) - \varepsilon, C(\mathcal{P}, a)]$ to the $(t - 1)$ -st block. By (i), (ii), (iii) and Lemma 1, the schedule \mathcal{P}' is feasible. This transformation is shown in Fig. 8 when $\varepsilon = e_t - e_{t-1} - \xi_{t-1}(a)$ and $j' = t$. Clearly, $C(\mathcal{P}, d) = C(\mathcal{P}', d)$ for each $d \in \mathcal{J} \setminus \{a, c\}$ and, by (iii), $C(\mathcal{P}', c) \leq C(\mathcal{P}, c) + \varepsilon$.

If $\varepsilon = \ell$, then the total completion time of \mathcal{P}' is strictly smaller than that of \mathcal{P} , because a resumes at $C(\mathcal{P}, a) - \ell$ in \mathcal{P} , i.e., $C(\mathcal{P}', a) < C(\mathcal{P}, a) - \ell$. We get a contradiction since \mathcal{P} is optimal.

Otherwise, if $\varepsilon < \ell$, then a and b form an A-configuration in \mathcal{P}' at $C(\mathcal{P}, a) - \varepsilon$. Also, $C(\mathcal{P}', a) = C(\mathcal{P}, a) - \varepsilon$. Let i be the abnormality point of \mathcal{P} . If $i \leq t - 1$, then the fact that \mathcal{P} and \mathcal{P}' are the same in $[0, e_{t-1}]$, we obtain that the abnormality point of \mathcal{P}' equals i and \mathcal{P}' is A-maximal. If $i > t - 1$, then by (iv), $i > j'$ and hence ε is t -normal and, by Lemma 13, \mathcal{P}' is A-maximal. Therefore, we obtain a contradiction in both cases, which proves the lemma. \square

The following lemma describes how two jobs that form an A-configuration start in an A-maximal schedule. See Fig. 9 for an illustration of the two possible cases: the two jobs can start at different time points, or at the same time.

Proposition 8 Suppose that each A-maximal schedule has an A-configuration. Then, there exists an A-maximal schedule (\mathcal{P}, e, ξ) such that the earliest A-configuration formed by a and b with $C(\mathcal{P}, a) < C(\mathcal{P}, b)$ satisfies the following properties:

- (i) $\mathcal{J}(\xi_{j'}) = \{a, b\}$, $|\mathcal{J}(\xi_j)| = 2$ and some job completes at $e_{j'}$, where $e_{j'} = s(\mathcal{P}, b)$ and $e_j = s(\mathcal{P}, a)$,

- (ii) $0 \leq j' - j \leq 1$, and
- (iii) $e_{j'+1}$ is an integer.

Proof Let \mathcal{P} be A-maximal. Let a and b form the earliest A-configuration in \mathcal{P} . By Proposition 5, $s(\mathcal{P}, a) \leq s(\mathcal{P}, b) = e_{j'}$. Without loss of generality assume that $s(\mathcal{P}, b)$ is as late as possible.

Then job a is spanning in block j' since otherwise a and b interlace and we get a contradiction by Lemma 5. Moreover,

$$\mathcal{J}(\xi_{j'}) = \{a, b\}, \tag{7}$$

since otherwise, by Lemma 7, it holds $\mathcal{J}(\xi_{j'}) = \{a, b, c\}$ and $C(\mathcal{P}, c) = e_{j'+1}$. The latter implies, by Lemma 3, that job c is spanning in block j' , which contradicts that job a is spanning in block j' and proves (7).

We prove (iii) first. Suppose for a contradiction that $e_{j'+1}$ is not an integer and let h be the greatest integer smaller than $e_{j'+1}$. Since $e_{j'+1}$ is an event and, by definition of A-configuration, none of the jobs a and b ends at $e_{j'+1}$, (7) implies that some job c starts at $e_{j'+1}$.

We show that $\xi_{j'+1}(b) = 0$, which will make our first transformation in (9) feasible. This holds for $j' + 1 = \tau(a)$, since by definition of A-configuration $b \notin \mathcal{J}(\xi_{\tau(a)})$. For $j' + 1 < \tau(a)$, we have $\xi_{j'+1}(b) = 0$ or job a is spanning in block $j' + 1$ for otherwise a and b interlace and we get a contradiction by Lemma 5. However, $\xi_{j'+1}(b) > 0$ and job a is spanning in block $j' + 1$, which imply $\mathcal{J}(\xi_{j'+1}) = \{a, b, c\}$. Thus, by Lemma 7, some job must finish at $e_{j'+1}$ and since $j' + 1 < \tau(b) < \tau(b)$, this job must be c . By Lemma 3, job c is spanning in block $j' + 1$ —a contradiction. Therefore,

$$\xi_{j'+1}(b) = 0. \tag{8}$$

Since $\tau(a) > j' + 1$, we obtain by (7) and definition of A-configuration that no job ends at $e_{j'+2}$ and hence Lemma 7 implies that job c is spanning in block $j' + 1$. Now take

$$\varepsilon = \min \{ \xi_{j'+1}(c), e_{j'+1} - \max \{ h, e_{j'} \} \}$$

and let \mathcal{P}' be a schedule with events e' and partition ξ' , where

$$(e', \xi') = (e, \xi, \varepsilon, (j' \overset{b}{\rightsquigarrow} j' + 1 \overset{c}{\rightsquigarrow} j')). \tag{9}$$

Figure 10a illustrates the transformation from \mathcal{P} to \mathcal{P}' for $\varepsilon = e_{j'+1} - h$, when $h > e_{j'}$. Observe that (8) and $\xi_{j'}(b) \geq \varepsilon$ ensure the feasibility of \mathcal{P}' . Also, by (7) and Lemma 12, we have $i \neq j'$, where i is the abnormality point of \mathcal{P} possibly equal ∞ . Clearly, if $i < j'$, then \mathcal{P} and \mathcal{P}' have the same abnormality point i since the two schedules are identical in $[0, e_{j'}]$. If $i > j'$, then by Lemma 11, $\xi_{j'+1}(c)$ is $(j' + 1)$ -normal and Lemma 13 implies the same abnormality point i for both \mathcal{P} and \mathcal{P}' . Finally, the A-maximality of \mathcal{P} implies




Fig. 10 Schedule transformations in the proof of Proposition 8

that a and b form an A-configuration in \mathcal{P}' . Therefore, if $h \leq e_{j'}$, then \mathcal{P}' is A-maximal and it can be ensured that $s(\mathcal{P}', b) > s(\mathcal{P}, b)$, which contradicts our assumption about \mathcal{P} . If $h > e_{j'}$, then \mathcal{P}' is A-maximal and satisfies (iii) as required. To simplify notation we set $\mathcal{P} := \mathcal{P}'$ in the remainder of the proof.

We now prove (i) and (ii). Observe that by (iii), $s(\mathcal{P}, b)$ is not an integer and thus

$$|\mathcal{J}(\xi_{j'-1})| \geq 2 \tag{10}$$

for otherwise \mathcal{P} would not be optimal—a contradiction.

Suppose first that $s(\mathcal{P}, a) = s(\mathcal{P}, b) = e_{j'}$. If a job a' in $\mathcal{J}(\xi_{j'-1})$ does not complete at $e_{j'}$, i.e., a' is preempted at $e_{j'}$, then $C(\mathcal{P}, a') > e_{\tau(b)}$ for otherwise, by Lemma 3, at most one of jobs $\{a, b\}$ can be spanning in block $\tau(a')$ and thus the other job in $\{a, b\}$ and a' would interlace, which contradicts Lemma 5. However, if $C(\mathcal{P}, a') > e_{\tau(b)}$, then job a' is spanning in block $\tau(b)$ for otherwise a' and b interlace, which again contradicts Lemma 5. Thus, $|\mathcal{J}(\xi_{\tau(b)})| = 2$ by Lemma 3. Therefore, a job in $\mathcal{J}(\xi_{j'-1}) \setminus \{a'\}$ completes at $e_{j'}$. The other conditions of the lemma trivially follow when $s(\mathcal{P}, a) = s(\mathcal{P}, b)$.

Now let $s(\mathcal{P}, a) \neq s(\mathcal{P}, b)$. By assumption, $s(\mathcal{P}, a) < s(\mathcal{P}, b)$. Informally, the proof is divided into two stages. In the first stage we consider block $j' - 1$ and we prove that $\mathcal{J}(\xi_{j'-1}) = \{a, c\}$ and that $\tau(c) = j' - 1$ —see Eqs. (11), (12) and (13) and Fig. 10b. In the second stage we prove that a starts at $e_{j'-1}$. The proof of the latter is done by contradiction, i.e., we suppose that a starts before $e_{j'-1}$. This assumption implies that \mathcal{P} looks as shown in Fig. 10c in the interval $[e_{j'-3}, e_{j'+1}]$, which allows us to get the desired contradiction thanks to Proposition 6.

First we prove by contradiction that

$$\xi_{j'-1}(a) = e_{j'} - e_{j'-1}. \tag{11}$$

By (10), $\mathcal{J}(\xi_{j'-1}) \setminus \{a\} \neq \emptyset$. Take any $c \in \mathcal{J}(\xi_{j'-1}) \setminus \{a\}$. By (7), $\xi_{j'}(c) = 0$. Since job a is non-spanning in block $j' - 1$, the conditions (i)–(iv) of Proposition 7 are all satisfied by jobs a and c , and $t = j'$. (Condition (iv) holds as j' is not

the abnormality point of \mathcal{P} by (7) and Lemma 12.) Therefore we get a contradictions, and (11) holds.

Next, we show that

$$\mathcal{J}(\xi_{j'-1}) = \{a, c\} \text{ for some } c \in \mathcal{J}. \tag{12}$$

If no job completes at $e_{j'}$, then Lemma 7 and (10) immediately imply (12). If some job, say c , completes at $e_{j'}$, then Lemma 3 implies that job c is spanning in block $j' - 1$. Since a completes after $e_{j'}$, $a \neq c$. This and (11) imply (12).

Finally to complete the first stage, we prove by contradiction that

$$C(\mathcal{P}, c) = e_{j'}. \tag{13}$$

To that end take $\varepsilon = \min\{\xi_{j'-1}(c), \xi_{j'}(b)\}$ and let \mathcal{P}' be a schedule with events e' and partition ξ' , where

$$(e', \xi') = \left\langle e, \xi, \varepsilon, (j' - 1 \overset{c}{\rightsquigarrow} j' \overset{b}{\rightsquigarrow} j' - 1) \right\rangle.$$

Note that

$$\begin{aligned} s(\mathcal{P}', b) &\geq e_{j'-1} = e_{j'+1} - (e_{j'+1} - e_{j'-1}) \\ &= e_{j'+1} - \xi_{j'-1}(a) - \xi_{j'}(a) \geq e_{j'+1} - 1, \end{aligned}$$

which, by (iii), implies that $s(\mathcal{P}', b) \geq r(b)$. Thus, \mathcal{P}' is feasible and, by assumption, optimal. Also, by (7), (12) and Lemma 12, we have $i \notin \{j' - 1, j'\}$, where i is the abnormality point of \mathcal{P} . Thus, as before, i is the abnormality point of \mathcal{P}' . Indeed, it follows from the fact that the two schedules are identical in $[0, e_{j'}]$ (which covers the case when $i < j'$), and from Lemma 13 (that covers the case when $i > j'$). Moreover, \mathcal{P}' contains a block that ends at $e_{j'+1}$ and contains the jobs a, b and c , none of which completes at $e_{j'+1}$ —a contradiction with Lemma 7. Therefore, (13) holds, and thus due to Eqs. (11), (12) and (13), the schedule in the interval $[e_{j'-1}, e_{j'+1}]$ looks like in Fig. 10b.

In the second stage we argue that

$$s(\mathcal{P}, a) = e_{j'-1}. \tag{14}$$

Suppose for a contradiction that this is not the case. By (11), c does not start at $e_{j'-1}$. Since a does not starts at $e_{j'-1}$ either, there is a job, say c' that ends at $e_{j'-1}$, otherwise $e_{j'-1}$ would not be an event. By Lemma 3,

$$\xi_{j'-2}(c') = e_{j'-1} - e_{j'-2}, \tag{15}$$

which implies

$$\xi_{j'-2}(a) = e_{j'-1} - e_{j'-2} \tag{16}$$

as follows: First we observe that there is no job $d \neq c'$ that completes at $e_{j'-1}$. Indeed, otherwise Lemma 8 applied to $c = d, c', d = c$, and $e_j = e_{j'-1}$ gives the required contradiction. Now, if $c \in \mathcal{J}(\xi_{j'-2})$, then the conditions (i)-(iv) of Proposition 7 are all satisfied by jobs a, b and c , and $t = j' - 1$ —a contradiction. (Condition (iv) holds as neither $j' - 1$ nor j' is the abnormality point of \mathcal{P} by (7), (12), and Lemma 12.) Therefore, $\xi_{j'-2}(c) = 0$. Thus, $\mathcal{J}(\xi_{j'-2}) \subseteq \{a, c'\}$, because if a job different than a and c' that does not complete at $e_{j'-1}$ is present in $\mathcal{J}(\xi_{j'-2})$ then, by (12) and (13), this job interlaces with c that contradicts Lemma 5. This implies (16) as required.

If job c' is non-spanning in block $j' - 3$, then by (12), (16) and $C(\mathcal{P}, c') = e_{j'-1}, s(\mathcal{P}, c') < e_{j'-2}$, which implies that c and c' form an A-configuration of length $e_{j'-1} - e_{j'-2}$ at $e_{j'-1}$, which leads to a contradiction with A-maximality of \mathcal{P} . Thus we have

$$\xi_{j'-3}(c') = e_{j'-2} - e_{j'-3}. \tag{17}$$

We prove that


$$\begin{aligned} \xi_{j'-3}(c') &= \xi_{j'-3}(c_1) = e_{j'-2} - e_{j'-3} \text{ and} \\ C(\mathcal{P}, c_1) &= e_{j'-2} \text{ for some } c_1 \in \mathcal{J} \setminus \{a, c\}, \end{aligned} \tag{18}$$

i.e., we prove that \mathcal{P} in the interval $[e_{j'-3}, e_{j'+1}]$ is as shown in Fig. 10c. First, we have $\xi_{j'-3}(c_1) > 0$ for some $c_1 \notin \{a, c, c'\}$. Otherwise $\mathcal{J}(\xi_{j'-3}) \subseteq \{a, c, c'\}$, and since $e_{j'-2}$ is an event, $s(\mathcal{P}, a) = e_{j'-2}$. Then, however, conditions (i)-(iv) of Proposition 7 are all satisfied by jobs a, b, c , and $t = j' - 2$ —a contradiction (observe that $h - s(\mathcal{P}, a) < 1$, thus (i) is satisfied; condition (iv) holds as none of $j' - 2, j' - 1, j'$ is the abnormality point of \mathcal{P} by (7), (12), (15), (16), and Lemma 12). Second, each such c_1 completes at $e_{j'-2}$ for otherwise, by (12), (16) and (17), c_1 and c interlace—a contradiction by Lemma 5. Thus, by Lemma 3, job c_1 is spanning in block $j' - 3$. This, and (17) imply (18). Thus, \mathcal{P} looks in the interval $[e_{j'-3}, e_{j'+1}]$ as shown in Fig. 10c. Finally, by Proposition 6 applied to $c = c', c_1, d = c, a$, and $e_j = e_{j'-2}$, we obtain that c and c' form an A-configuration at $e_{j'-1}$. Thus, again, we get a contradiction since \mathcal{P} is A-maximal. Hence, (14) holds. Therefore the lemma follows by (7), (12), and (14). \square

Given schedule $(\mathcal{P}, e, \xi), l \geq 1$ and $\{a_1, \dots, a_l\} \subseteq \mathcal{J}$, job sequence (a_1, \dots, a_l) is called a *sub-chain starting at t* in \mathcal{P} if:

- (S1) For each $j \in \{1, \dots, l - 1\}, a_j \preceq a_{j+1}$;
- (S2) For each $j \in \{1, \dots, l - 1\}, C(\mathcal{P}, a_j) = s(\mathcal{P}, a_{j+1})$;
- (S3) Job a_1 executes non-preemptively in $[t, C(\mathcal{P}, a_1)]$.

Fig. 11 (ε, d) -Exchanging of (a_1, \dots, a_l) and $(b_1, \dots, b_{l'})$ when $l' = l + 1, d = b_1$ and $s(\mathcal{P}, a) = t$



Moreover, job sequence (a_1, \dots, a_l) is a *chain* in \mathcal{P} if it satisfies conditions (S1), (S2) and additionally

- (S4) Time t is the earliest moment such that a_1 executes with no preemption in $[t, C(\mathcal{P}, a_1)]$;
- (S5) No predecessor of a_1 ends at t .

Suppose that jobs a and b form an A-configuration in \mathcal{P} with $C(\mathcal{P}, a) < C(\mathcal{P}, b)$. For $\varepsilon \geq 0$, we define an operation of ε -exchanging a and b in an interval $[e_k, C(\mathcal{P}, b)]$, $k < \tau(b)$, as follows. First, all pieces of a and b are removed from the blocks $k, \dots, \tau(b)$. Note that the total length of all removed pieces of a and b is $\sum_{t=k}^{\tau(b)} \xi_t(a)$ and $\sum_{t=k}^{\tau(b)} \xi_t(b)$, respectively. Then, the empty gaps are filled out with the total length $\sum_{t=k}^{\tau(b)} \xi_t(a) - \varepsilon$ of a and the total length of $\sum_{t=k}^{\tau(b)} \xi_t(b) + \varepsilon$ of b in such a way that b completes as early as possible. Note that the new schedule is feasible only if $\varepsilon = 0$. Whenever the transformation of ε -exchanging will be used with $\varepsilon > 0$, then some other appropriate changes in the schedule will be made to ensure feasibility.

For $\varepsilon > 0$, we extend the operation of ε -exchanging of two jobs to sub-chains as follows. Let $A = (a_1, \dots, a_l = a)$ and $B = (b_1, \dots, b_{l'} = b)$ be two sub-chains in \mathcal{P} that start at t , and such that a and b form an A-configuration in \mathcal{P} , where $s(\mathcal{P}, a) \leq s(\mathcal{P}, b)$. (Note that we either have $l = l'$ or $l = l' - 1$.) Let d be any job that executes in $[t - \varepsilon, t]$. The operation of (ε, d) -exchanging of A and B in \mathcal{P} leads to a schedule \mathcal{P}' obtained by making the following changes to \mathcal{P} :

- For each $j \in \{1, \dots, l\}$, a_j is executed in $[t_j - \varepsilon, t_{j+1} - \varepsilon]$ in \mathcal{P}' , where take $t_1 = t, t_j = s(\mathcal{P}, a_j)$ for $j \in \{2, \dots, l\}$ and $t_{l+1} = e_{j'+1}$ such that $e_{j'} = s(\mathcal{P}, b)$;
- For each $j \in \{1, \dots, l'\}$, b_j is executed in $[u_j + \varepsilon, u_{j+1} + \varepsilon]$ in \mathcal{P}' , where take $u_1 = t, u_j = s(\mathcal{P}, b_j)$ for $j \in \{2, \dots, l'\}$ and $u_{l'+1} = e_{j'+1}$;
- A piece of d executing in $[t - \varepsilon, t]$ is placed in $[t, t + \varepsilon]$ in \mathcal{P}' (the “room” for this job is made by postponing b_1);
- In the interval $[e_{j'+1}, C(\mathcal{P}, b)]$ ε -exchanging of a and b is made.

The transformation is illustrated in Fig. 11 for $d = b_1$. Note that in this particular case the total completion times of \mathcal{P} and \mathcal{P}' are equal.

The new schedule \mathcal{P}' is feasible under certain conditions. First, the value of ε must be selected in such a way that ε -exchanging of a and b is possible in the above-mentioned interval. Second, d should not be a predecessor of a_1 . Also, the release dates of jobs a_1, \dots, a_l need to be respected and a_1 must be non-preemptively executed in $[t, t + \varepsilon]$. We summarize those conditions in the following lemma.

Proposition 9 *Let $(a_1, \dots, a_l = a)$ and $(b_1, \dots, b_{l'} = b)$, starting at t , be two sub-chains in \mathcal{P} such that a and b form an A-configuration of length ℓ in \mathcal{P} . If $\varepsilon \leq \ell, r(a_1) \leq t - \varepsilon$ and $r(a_j) \leq s(\mathcal{P}, a_j) - \varepsilon$ for each $j \in \{2, \dots, l\}$, a_1 executes non-preemptively in $[t, t + \varepsilon]$, and some job d that is not a predecessor of a_1 executes non-preemptively in $[t - \varepsilon, t]$, then the schedule \mathcal{P}' obtained by (ε, d) -exchanging of the two sub-chains in \mathcal{P} is feasible. \square*

5.1 Proof of Theorem 2


Let (\mathcal{P}, e, ξ) be a maximal schedule that satisfies the properties in Proposition 8. Let $(a_1, \dots, a_l = a)$ be the chain in \mathcal{P} that starts at t_a and let $(b_1, \dots, b_{l'} = b)$ be the chain in \mathcal{P} that starts at t_b . By definition of chains and Proposition 8 we have

$$\begin{aligned}
 l = 1 &\Rightarrow t_a = s(\mathcal{P}, a), & (19) \\
 l \geq 2 &\Rightarrow s(\mathcal{P}, a) - (l - 1) \leq t_a < s(\mathcal{P}, a) - (l - 2), & (20) \\
 l' = 1 &\Rightarrow t_b = s(\mathcal{P}, b), \text{ and} & (21) \\
 l' \geq 2 &\Rightarrow s(\mathcal{P}, b) - (l' - 1) \leq t_b < s(\mathcal{P}, b) - (l' - 2). & (22)
 \end{aligned}$$

Case 1 $t_a \geq t_b$.

In this case we perform a transformation shown in Fig. 12 as described below. By Proposition 8, there exists an integer h such that both jobs a and b execute non-preemptively in

Fig. 12 Transformation from \mathcal{P} to \mathcal{P}' ($\varepsilon = t_a - h'$, $d = b'$) in Case 1 in proof of Theorem 2



$[s(\mathcal{P}, a), h]$ and $[s(\mathcal{P}, b), h]$, respectively. We have $e_p = t_a$ for some event e_p .

Let $e_{j'} = s(\mathcal{P}, b)$. By Proposition 8, $h - e_{j'} = \xi_{j'}(b)$. Let ℓ be the length of the A-configuration formed by a and b . Clearly $C(\mathcal{P}, a) - \ell > h$ by definition of A-configuration and Proposition 8.

Let $A = (a_1, \dots, a_l)$ and let B be the sub-chain of the chain $(b_1, \dots, b_{l'} = b)$ that starts at t_a with a job b' and ends with the job b . By definition of t_a , $\xi_{p-1}(a_1) < e_p - e_{p-1}$. Thus, $|\mathcal{J}(\xi_{p-1}) \setminus \{a_1\}| \geq 2$. Let $d \in \mathcal{J}(\xi_{p-1}) \setminus \{a_1\}$ be a job that does not complete at e_p (possibly b'), if such a job exists. Otherwise, let d be any job in $\mathcal{J}(\xi_{p-1}) \setminus \{a_1\}$. Take

$$\varepsilon = \min \{ \xi_{p-1}(d), e_p - e_{p-1} - \xi_{p-1}(a_1), y, \xi_{j'}(b), \ell, t_a - h' \},$$

where h' is the greatest integer smaller than t_a and

$$y = \begin{cases} h - s(\mathcal{P}, a), & \text{when } l = 1, \\ (C(\mathcal{P}, a_1) - e_p)/2, & \text{when } l > 1 \text{ and } C(\mathcal{P}, d) = t_a, \\ C(\mathcal{P}, a_1) - e_p, & \text{otherwise.} \end{cases}$$

The latter ensures that d , if it completes at t_a in \mathcal{P} , does not complete after $s(\mathcal{P}', a_2)$ in \mathcal{P}' . Note that, by definition of t_a , no predecessor of a_1 ends at t_a and $\xi_{p-1}(a_1) < e_p - e_{p-1}$. Hence, in particular, $\varepsilon > 0$. Let \mathcal{P}' be the schedule obtained by (ε, d) -exchanging of A and B in \mathcal{P} . By Proposition 9, \mathcal{P}' is feasible. If $\varepsilon = \ell$, then the total completion time of \mathcal{P}' is strictly smaller than that of \mathcal{P} and we get a contradiction since \mathcal{P} is optimal.

Thus, consider $\varepsilon < \ell$. Then, the total completion time of \mathcal{P}' does not exceed that of \mathcal{P} . To see that we observe that by (19) and (20) we have $s(\mathcal{P}, b) - t_a < \ell$. Also, if two jobs in $\mathcal{J}(\xi_{p-1}) \setminus \{a_1\}$ complete at t_a , then either at least one of them is a predecessor of b' , which implies that $s(\mathcal{P}, b') = t_a$, or otherwise we obtain from Lemma 8 that $s(\mathcal{P}, b') = t_a$. Therefore, no more than l jobs in $\{d, b_1, \dots, b_{l'}\}$ complete in $[t_a, h]$ in \mathcal{P} . Thus, no more than l jobs get delayed by ε each as a result of the exchange, however, each job in the

chain $(a_1, \dots, a_l = a)$ completes by ε earlier at the same time.

Finally, we show that \mathcal{P} and \mathcal{P}' have the same abnormality point. Clearly, this holds if $i < p - 1$. Also, if $i > j'$, then ε is p -normal. To see this we observe that $e_p, e_{p-1}, \xi_{p-1}(a_1)$ and t_a are clearly all p -normal. By Lemma 10, $C(\mathcal{P}, a_1) - t_a$ is p -normal. Also $e_{j'} = t_a + (C(\mathcal{P}, b') - t_a) + k - 2$, where k is the number of jobs in B , is p -normal. If $l = 1$, then $s(\mathcal{P}, a)$ and h are p -normal, which implies p -normality of y . For $l > 1$, we argue that y is also p -normal and for that we need only consider $y = (C(\mathcal{P}, a_1) - e_p)/2$. Then, b' is not present in the $(p - 1)$ -st block for otherwise b' would be selected as d . The length of $(p - 1)$ -st block, $e_p - e_{p-1}$, is by definition $(p - 1)$ -normal. By Lemma 3, $\xi_{p-1}(d) = e_p - e_{p-1}$. By Lemma 5, each job in $\mathcal{J}(\xi_{p-1}) \setminus \{a_1\}$ must complete at e_p . This proves, again by Lemma 3, that $|\mathcal{J}(\xi_{p-1})| = 2$. By Lemma 10, $\xi_{p-1}(a_1) + \xi_p(a_1) + \xi_{p+1}(a_1) = \xi_{p-1}(a_1) + C(\mathcal{P}, a_1) - t_a$ is $(p - 1)$ -normal. Since $\xi_{p-1}(a_1) \in \{0, e_p - e_{p-1}\}$, we obtain that $\xi_{p-1}(a_1)$ is $(p - 1)$ -normal. Thus, $(C(\mathcal{P}, a_1) - t_a)/2$ is p -normal as required. Therefore, ε is p -normal and, by Lemma 13, \mathcal{P} and \mathcal{P}' have the same abnormality point for $i > j'$. Also, by Proposition 8, and chain definition, we have $|\mathcal{J}(\xi_k)| = 2$ for each $k \in \{p, \dots, j'\}$. Thus, by Lemma 12, $i \notin \{p, \dots, j'\}$. Finally, consider $i = p - 1$. Then, if i is no longer the abnormality point i' of \mathcal{P}' , then $i' > i$ —a contradiction since \mathcal{P} is A-maximal. Therefore, i is the abnormality point of \mathcal{P}' , and hence we proved that \mathcal{P} and \mathcal{P}' have the same abnormality point. To complete the case proof we note that a and b form an A-configuration in \mathcal{P}' at $C(\mathcal{P}', a) = C(\mathcal{P}, a) - \varepsilon$, which contradicts the A-maximality of \mathcal{P} .

Case 2 $t_a < t_b$. We first define

$$t'_a = \max \{ t < t_b \mid t = t_a \text{ or } t \in \{s(\mathcal{P}, a_1), \dots, s(\mathcal{P}, a_l)\} \}$$

and a' to be the job from the chain (a_1, \dots, a_l) that starts or resumes at t'_a . By (19)–(22), it holds $a' = a_{l-l'+1}$ or $a' = a_{l-l'+2}$, and only one job, namely a' , from the chain




Fig. 13 **a** Illustration of Claim 1; **b** Proof of Claim 1

(a_1, \dots, a_l) is executed in (t'_a, t_b) . By definition $t'_a < t_b$, also we have $t'_a = e_p$ for some event e_p .

We first prove that exactly one job that is not in the chain (a_1, \dots, a_l) , call it d , executes in $[t'_a, t_b]$ and completes at t_b . Indeed, if $l' = 1$, then this follows from Proposition 8. If $l' > 1$, then any job not in the chain (a_1, \dots, a_l) that executes in $[t'_a, t_b]$ completes in $[t'_a, t_b]$, otherwise this job interlaces with b_1 —a contradiction with Lemma 5. Finally, we show that two or more jobs not in the chain (a_1, \dots, a_l) cannot complete in $[t'_a, t_b]$. If there are at least three such jobs, then the last two of them form an A-configuration, which contradicts the A-maximality of \mathcal{P} . For exactly two, c' and c completing in this order, by Claim 1, c and c' form an A-configuration at e_{p+1} —a contradiction since \mathcal{P} is A-maximal. Also, observe that for $l' = l + 1$, job b_1 resumes at t_b and thus b_1 and d form an A-configuration at e_{p+1} by Claim 1—a contradiction since \mathcal{P} is A-maximal. Thus, let $l \geq l'$ in the remainder of the lemma.

Now we prove that our schedule \mathcal{P} satisfies the following claim that we have used above (see Fig. 13a for illustration of Claim 1):

Claim 1 Suppose that $t'_a = e_p$ is an event in \mathcal{P} and that there exist jobs a' , c , and d such that

- (i) $C(\mathcal{P}, c) = e_{p+1}$, and $C(\mathcal{P}, d) = e_{p+2}$;
- (ii) $\mathcal{J}(\xi_p) = \{a', c\}$ and $\mathcal{J}(\xi_{p+1}) = \{a', d\}$;
- (iii) $s(\mathcal{P}, d) < e_p$;
- (iv) if $l' = l + 1$, then $d = b_1$; otherwise $C(\mathcal{P}, d) = t_b$.

Then jobs c and d form an A-configuration at e_{p+1} .

Proof If $\xi_{p-1}(c) < e_p - e_{p-1}$, then the jobs c and d form an A-configuration of length $e_{p+1} - e_p$ at e_{p+1} —the lemma holds. Thus,

$$\xi_{p-1}(c) = e_p - e_{p-1}. \quad (23)$$

We now prove that in interval $[e_{p-1}, e_{p+2}]$ the schedule \mathcal{P} is as in Fig. 13, i.e., there exists a job c_1 such that

$$\begin{aligned} \xi_{p-1}(c_1) &= e_p - e_{p-1} \wedge C(\mathcal{P}, c_1) \\ &= e_p \wedge s(\mathcal{P}, d) < e_{p-1}. \end{aligned} \quad (24)$$

First, we show that $\xi_{p-1}(c_1) > 0$ for some $c_1 \notin \{a', c, d\}$. Otherwise, by (23) and (ii), $s(\mathcal{P}, a') = e_p$ since e_p is an event. Thus, $\mathcal{J}(\xi_{p-1}) = \{c, d\}$. Now, take

$$\varepsilon = \min \{ \xi_{p-1}(d), \xi_{j'}(b), \ell, t'_a - h' \},$$

where h' is the greatest integer smaller than t'_a . Observe that $\varepsilon > 0$. Let $A = (a_1, \dots, a_l)$ and let $B = (d, b_1, \dots, b_{l'}) = b$ be the sub-chain that starts at e_{p+1} . Perform the (ε, d) -exchanging of A and B in \mathcal{P} as in Case 1 (the completion time of c does not change in this transformation when $d \neq b_1$ because a' from the $\tau(a')$ -th block is placed in the $(p-1)$ -st block and d from $(p-1)$ -st block is placed in the $(\tau(d)+1)$ -st block) to get a contradiction. Observe that, by (iv), $d = b_1$ for $l' = l + 1$ and thus the (ε, d) -exchanging of A and B indeed produces schedule \mathcal{P}' with total completion time that does not exceed that of schedule \mathcal{P} . Also, by Lemma 4, there is no idle time in the $(p-1)$ -st block of \mathcal{P} . This implies, by Lemma 10, that $\xi_{p-1}(d) = e_p - e_{p-1}$ is $(p-1)$ -normal, which by arguments in Case 1 implies that the abnormality points of \mathcal{P} and \mathcal{P}' are the same.

Second, by Lemma 5, c_1 and d cannot interlace, which implies $C(\mathcal{P}, c_1) = e_p$. By Lemma 3, job c_1 is spanning in block $p-1$. Thus, by (23) we have $\{c_1, c\} = \mathcal{J}(\xi_{p-1})$. Finally, $s(\mathcal{P}, d) < e_{p-1}$ is due to (iii) and $\mathcal{J}(\xi_{p-1}) = \{c_1, c\}$. This completes the proof of (24). Equation (24) allows us to apply Proposition 6 to c , c_1 , d , $a = a'$, and $e_j = e_p$ to conclude that c and d form an A-configuration at e_{p+1} . This contradicts the A-maximality of \mathcal{P} and completes the proof of Claim 1. \square

Let $e_{j'} = s(\mathcal{P}, b)$. Now, let $y = \sum_{t \geq p} \xi_t(b_1)$ and $z = \sum_{t \geq p} \xi_t(a_{l-l'+1})$. First we prove that $z \leq y$. This holds due to Proposition 8 when $l' = 1$ and hence let $l \geq l' > 1$. If $z > y$, then swap a and b and then do the $(\varepsilon, a_{l-l'+1})$ -exchanging of $(b_1, \dots, b_{l'})$ and $(a_{l-l'+1}, \dots, a_l)$ (note the order of the chain, which is important) both starting at t_b , where $\varepsilon = s(\mathcal{P}, b) - s(\mathcal{P}, a) + \lambda$ and $0 < \lambda < \min\{\ell, z - y, \xi_{j'}(b)\}$. This transformation is shown in Fig. 14a. Let the resulting schedule be \mathcal{P}' . The swapping increases the total completion time by $s(\mathcal{P}, b) - s(\mathcal{P}, a)$ and the $(\varepsilon, a_{l-l'+1})$ -exchanging decreases it by $((l' - 1) - l')\varepsilon$ —observe that after the swapping of a and b the completion time of job $a = a_l$ does not change in the exchange. Therefore, the overall change equals $-\lambda$ and thus to get a contradiction it suffices to prove that \mathcal{P}' is feasible.

Observe that $s(\mathcal{P}, b) - s(\mathcal{P}, a) = C(\mathcal{P}, b_1) - C(\mathcal{P}, a_{l-l'+1})$. By Proposition 8, $e_{j'+1}$ is an integer. Thus, $r(b) < e_{j'+1}$ implies $r(b) \leq e_{j'+1} - 1$. Moreover, $s(\mathcal{P}', b) \geq e_{j'+1} - 1$. Therefore, by the definition of a sub-chain, all jobs $b_2, \dots, b_{l'}$ respect their release dates in \mathcal{P}' . Since $z > y$, we have that $s(\mathcal{P}, b_1) < t'_a$ and hence b_1 respects its release date




Fig. 14 Transformations from \mathcal{P} to \mathcal{P}' in Case 2

in \mathcal{P}' . Thus, $z \leq y$ for the remainder of the proof. We consider the following three subcases.

Case 2a $t'_a = s(\mathcal{P}, a_{l-l'+1})$. (Schedule transformation performed in this case is shown in Fig. 14b.) Then, $z = 1$. Since $z \leq y$, we have $y = 1$. If some job in $\mathcal{J}(\xi_{p-1})$ does not complete at t'_a , then this case reduces to Case 1. Otherwise, two jobs in $\mathcal{J}(\xi_{p-1})$ complete at t'_a . Thus, by Lemma 7, for at least one job in $\mathcal{J}(\xi_{p-1})$, say job c' , we have $\xi_t(c') = 0$, where $e_t = s(\mathcal{P}, d)$. Therefore, c' and d interlace if $s(\mathcal{P}, c') < s(\mathcal{P}, d)$ —a contradiction by Lemma 5, or we can swap jobs c' and d if $s(\mathcal{P}, c') > s(\mathcal{P}, d)$. In the latter case the resulting schedule (see Fig. 14b) reduces the total completion time of \mathcal{P} by Lemma 2. This schedule is not feasible when $c' < a_{l-l'+1}$ and we restore feasibility by applying a 0-exchanging of b and a in $[e_{j'+1}, C(\mathcal{P}, b)]$ followed by $(t_b - t'_a, a_{l-l'+1})$ -exchanging of $(b_1, \dots, b_{l'})$ and $(a_{l-l'+1}, \dots, a_l)$ both starting at t_b . The new schedule \mathcal{P}' is feasible since $c' \not\prec b_1$ for in-trees, and since, by Proposition 8, $s(\mathcal{P}', b) \geq e_{j'+1} - 1 \geq r(b)$, which shows that all $b_1, \dots, b_{l'}$ respect their release dates in \mathcal{P}' . Thus, we get a contradiction since \mathcal{P} is optimal.

Case 2b $t'_a = s(\mathcal{P}, a_{l-l'+2})$. Since $l \geq l'$, $\xi_{p-1}(a_{l-l'+1}) = e_p - e_{p-1}$. Also, $t'_a < t_b$ implies that b_1 resumes at $t_b = e_{p+1}$. Thus, by Claim 1, d and b_1 form an A-configuration at e_{p+1} , which contradicts the A-maximality of \mathcal{P} .

Case 2c $t'_a \neq s(\mathcal{P}, a_{l-l'+1})$ and $t'_a \neq s(\mathcal{P}, a_{l-l'+2})$. By definition of t'_a , we have $t'_a = t_a$ and a_1 resumes at t_a . Since t_a is an event of \mathcal{P} some job, say c , completes at e_p . By Lemma 3, job c is spanning in block $p - 1$. If another job completes at e_p , then we get a contradiction by Lemma 8. Hence, by Lemma 5, $\mathcal{J}(\xi_{p-1}) \subseteq \{c, d, b_1, a_1\}$. Note that $z \leq y, l \geq l'$ and $t_a < t_b$ imply that $l' = l$. By definition of t_a , job a_1 is non-spanning in block $p - 1$. By Lemma 5, d and b_1 do not interlace, which implies $\xi_{p-1}(b_1) = 0$. Therefore, $\xi_{p-1}(d) > 0$. This allows us to obtain a contradiction by performing an analogous transformation as in Claim 1.

Observe that for the proof of Theorem 2 it is crucial to show that \mathcal{P} and \mathcal{P}' have the same abnormality point. This needs to be proven in Case 1, Claim 1, 5 and 5. In 5 and 5 the proof reduces to the proof for Case 1 and Claim 1. In Claim 1 the proof also reduces to the proof for Case 1 but the $\xi_{p-1}(d)$ is new in definition of ε as compared to Case 1 so we provide an appropriate comment about this $\xi_{p-1}(d)$ in Claim 1. Finally, in Case 1 we explicitly prove that \mathcal{P} and \mathcal{P}' have the same abnormality point.

6 Alternating chains

In this section we prove that there always exists an optimal normal schedule for $P2|pmtn, in-tree, r_j, p_j| \sum C_j$. Our proof is by contradiction. We show that an abnormality point $i \neq \infty$ in a maximal schedule implies that there is an alternating chain, see Sect. 6.1 for its definition, in the schedule. Each job in that chain completes at the moment which is not i -normal. This fact allows us to either make the alternating chain longer, which is shown in Sect. 6.3, or find an optimal schedule with an abnormality point higher than i . Thus in either case we get a contradiction, in the former, since the number of jobs is finite and we cannot extend the chain *ad infinitum*, in the latter since the initial schedule is maximal.

6.1 Basic definitions and properties

Given schedule (\mathcal{P}, e, ξ) , let $I = \{j, \dots, j'\} \subseteq \{1, \dots, q - 1\}$. For two jobs a and a' , We say that a covers a' in I if for each $t \in I$, $\xi_t(a') > 0$ implies $\xi_t(a) = e_{t+1} - e_t$. Let \mathcal{P} be a maximal schedule of abnormality point $i \neq \infty$. Job sequence (d_1, \dots, d_l) is called an *alternating chain* in \mathcal{P} if $d_1 \in A_i(\mathcal{P})$ and d_1 executes non-preemptively in $[e_{i+1}, C(\mathcal{P}, d_1)]$ and additionally, unless if $l = 1$, the following conditions are satisfied:

- (C1) $A_i(\mathcal{P}) = \{d_1, d_2\}$ and $\tau(d_1) = i + 1$,
- (C2) $C(\mathcal{P}, d_j) < C(\mathcal{P}, d_{j+1})$ for each $j \in \{1, \dots, l - 1\}$.
- (C3) the job d_j executes non-preemptively in the interval $[C(\mathcal{P}, d_{j-2}), C(\mathcal{P}, d_j)]$ for each $j \in \{2, \dots, l\}$, where $C(\mathcal{P}, d_0) = e_{i+1}$.

If (d_1, \dots, d_l) ($l \geq 2$) satisfies (C1), (C3) and

- (C2') $C(\mathcal{P}, d_j) < C(\mathcal{P}, d_{j+1})$ for each $j \in \{1, \dots, l - 2\}$ and $C(\mathcal{P}, d_{l-1}) = C(\mathcal{P}, d_l)$,

then (d_1, \dots, d_l) is said to be *almost alternating*.

The following lemma excludes almost alternating chains from maximal schedules. The lemma does not require that the maximal schedules are A-free.

Lemma 14 Let \mathcal{P} be a maximal schedule of abnormality point $i \neq \infty$. There exists no almost alternating chain in \mathcal{P} .

Proof Suppose for a contradiction that (d_1, \dots, d_l) is an alternating chain with $C(\mathcal{P}, d_{l-1}) = C(\mathcal{P}, d_l)$. Let first $l > 3$ and let U be the set of odd indices in $\{3, \dots, l\}$. Denote by ξ the partition of \mathcal{P} . We first prove that the total length of d_2 executed in $[C(\mathcal{P}, d_1), C(\mathcal{P}, d_l)]$, namely $\xi_{i+2}(d_2)$, is i -normal. From the definition of alternating chain we know that there is no idle time in this interval and only the jobs from the chain execute in it. Hence,

$$\begin{aligned} \xi_{i+2}(d_2) &= 2(C(\mathcal{P}, d_l) - C(\mathcal{P}, d_1)) \\ &\quad - \sum_{j=3}^l (C(\mathcal{P}, d_j) - C(\mathcal{P}, d_{j-2})) \\ &= 2 \sum_{j \in U} \sum_{j' \geq i} \xi_{j'}(d_j) - \sum_{j=3}^l \sum_{j' \geq i} \xi_{j'}(d_j). \end{aligned}$$

By Lemma 10, $\sum_{j'=1}^i \xi_{j'}(d_j)$ is i -normal for each $j \in \{3, \dots, l\}$. Therefore, $\xi_{i+2}(d_2)$ is i -normal. This implies, by Lemmas 10 and 11, that the three following numbers are i -normal:

$$\xi_i(d_1) + \xi_i(d_2), \quad \xi_i(d_1) + \xi_{i+1}(d_1), \quad \xi_i(d_2) + \xi_{i+1}(d_1)$$

since $\xi_{i+1}(d_1) = \xi_{i+1}(d_2)$. Therefore, $\xi_i(d_1)$ and $\xi_i(d_2)$ are $(i + 1)$ -normal, and since e_{i+1} is i -normal due to Lemma 3, i is not the abnormality point of \mathcal{P} —a contradiction.

Now consider $l = 2$. Let e be the events of \mathcal{P} . Denote $x = \xi_{i+1}(d_1) = \xi_{i+1}(d_2)$. By assumption and by definition of alternating chain, d_1 and d_2 complete at e_{i+2} and hence $\xi_i(d_1) + x = s_1/2^i$ and $\xi_i(d_2) + x = s_2/2^i$ for some integers s_1 and s_2 . By definition of alternating chain, $A_i(\mathcal{P}) = \{d_1, d_2\}$ and hence x is not $(i + 1)$ -normal. Thus, $x = s'/2^{i+1} + \varepsilon$ for some $0 < \varepsilon < 1/2^{i+1}$. Then, $\xi_i(d_1) = (s_1 - s')/2^i - \varepsilon$ and $\xi_i(d_2) = (s_2 - s')/2^i - \varepsilon$. By Lemma 11, $e_{i+1} - e_i$ is i -normal. By Lemma 12,

$$e_{i+1} - e_i = \xi_i(d_1) + \xi_i(d_2) = (s_1 + s_2 - 2s')/2^i - 2\varepsilon.$$

Thus, 2ε is i -normal, which implies that ε is $(i + 1)$ -normal—a contradiction with $0 < \varepsilon < 1/2^{i+1}$. \square

Lemma 15 Let \mathcal{P} be a maximal schedule of abnormality point $i \neq \infty$. If (d_1, \dots, d_l) ($l \geq 1$) is an alternating chain in \mathcal{P} , then $C(\mathcal{P}, d_j)$ is not $(i + 1)$ -normal for each $j \in \{1, \dots, l\}$.

Proof Let first $1 \leq j \leq \min\{2, l\}$. Then, $C(\mathcal{P}, d_j) = e_{i+1} + 1 - \xi_i(d_j) - \sum_{t < i} \xi_t(d_j)$. By Lemma 10, the latter sum is i -normal. By Observations 1 and 2 and by Lemma 11, e_{i+1} is i -normal. However, $\xi_i(d_j)$ is not $(i + 1)$ -normal because

$d_j \in A_i(\mathcal{P})$ and hence, again by Observation 1, $C(\mathcal{P}, d_j)$ is not $(i + 1)$ -normal.

For $j > 2$, if j is even (respectively, odd), then let U be the set of even (respectively, odd) integers in $\{1, \dots, j - 1\}$. Let $u = 1$ if j is odd and let $u = 2$ if j is even. Then,

$$\begin{aligned} C(\mathcal{P}, d_j) &= e_{i+1} + \sum_{j' \in U} (1 - \xi_i(d_{j'})) - \sum_{t < i} \xi_t(d_{j'}) \\ &= e_{i+1} - \xi_i(d_u) + \sum_{j' \in U} (1 - \sum_{t < i} \xi_t(d_{j'})). \end{aligned}$$

Again, by Lemma 10, $\xi_i(d_u)$ is the only term in the above expression that is not $(i + 1)$ -normal. Thus, $C(\mathcal{P}, d_j)$ is not $(i + 1)$ -normal. \square

6.2 Transformations using alternating chains

Consider a schedule \mathcal{P} of abnormality point i and an alternating chain (d_1, \dots, d_l) ($l > 1$) and $\mathcal{J}(\xi_{\tau(d_l)}) = \{x, y, d_l\}$. Let $u = 2$ if l is odd, and $u = 1$ if l is even. Let $\varepsilon > 0$ be the largest ε such that

$$\begin{aligned} \varepsilon \leq \alpha &= \xi_i(d_u) \wedge \varepsilon \leq \beta \\ &= \frac{1}{2} \min \{e_{\tau(d_j)+1} - e_{\tau(d_j)} \mid j \in U\} \wedge \varepsilon \leq \gamma \\ &= \min \{e_{\tau(d_{j-1})} - r(d_j) \mid j \in U \setminus \{1, 2\}\}, \end{aligned} \tag{25}$$

where U is the set of the indices in $\{1, \dots, l\}$ having the same parity as l , and

$$\varepsilon \leq \min\{\xi_{\tau(d_l)}(x), \xi_{\tau(d_l)}(y)\}. \tag{26}$$

We define a transformation of ε -pushing of d_l that produces a schedule \mathcal{P}' as follows (see Fig. 15 for an illustration):

- the schedules \mathcal{P} and \mathcal{P}' are identical in time intervals $[0, e_j]$ and $[C(\mathcal{P}, d_l), \infty)$.
- To obtain the part of \mathcal{P}' in $[e_i, e_{i+1}]$, we increase (with respect to \mathcal{P}) the amount of d_{3-u} by ε and decrease the amount of d_u by ε . Then, a part of job d_{3-u} executes in $[e_{i+1}, C(\mathcal{P}, d_{3-u}) - \varepsilon]$ and a part of job d_u executes in $[e_{i+1}, C(\mathcal{P}, d_u) + \varepsilon]$. This in particular characterizes the execution of d_1 and d_2 in \mathcal{P}' .
- For each $j \in U \setminus \{1, 2\}$, the part of d_j that executes in $[C(\mathcal{P}, d_{j-2}), C(\mathcal{P}, d_j)]$ in \mathcal{P} is executed in $[C(\mathcal{P}, d_{j-2}) - \varepsilon, C(\mathcal{P}, d_j) - \varepsilon]$ in \mathcal{P}' . In this way we ensure that each job d_j , $j \in U$, completes ε earlier in \mathcal{P}' than in \mathcal{P} .
- For each $j \in \{3, \dots, l\} \setminus U$, the part of d_j that executes in $[C(\mathcal{P}, d_{j-2}), C(\mathcal{P}, d_j)]$ in \mathcal{P} is executed in $[C(\mathcal{P}, d_{j-2}) + \varepsilon, C(\mathcal{P}, d_j) + \varepsilon]$ in \mathcal{P}' . In this way we ensure that each job d_j , $j \notin U$, completes ε later in \mathcal{P}' than in \mathcal{P} .




Fig. 15 ε -Pushing of d_l when: **a** l is odd; **b** l is even

- Finally, the two jobs x and y are executed in the remaining free space in $[C(\mathcal{P}', d_{l-1}), C(\mathcal{P}, d_l)]$ on one machine and in $[C(\mathcal{P}', d_l), C(\mathcal{P}, d_l)]$ on the other machine.

The transformation of ε -pushing of d_l will be a key transformation used to extend an alternating chain of a maximal schedule \mathcal{P} in the proof of Theorem 3—the main result of the next section. The extension, as alluded earlier, requires that \mathcal{P} is A-free. Actually, it suffices that \mathcal{P} is A-free in the interval that starts with the completion of d_l , the last job of the chain. However, since the ε -pushing of d_l may change \mathcal{P} itself we need to prove that the resulting schedule is A-free in the interval that starts with the completion of d_l , which the transformation may have changed, in order to enable further extensions of the chain. Thus, we need the following lemma.

Lemma 16 Suppose that (d_1, \dots, d_l) (where $l > 1$) is an alternating chain in a maximal schedule \mathcal{P} that is A-free in $[C(\mathcal{P}, d_l), \infty)$, and $\mathcal{J}(\xi_{\tau(d_l)}) = \{x, y, d_l\}$. If ε is selected as in (25) and (26), then the schedule \mathcal{P}' obtained from \mathcal{P} by ε -pushing of d_l is maximal, A-free in $[C(\mathcal{P}', d_l), \infty)$, and (d_1, \dots, d_l) is an alternating chain in \mathcal{P}' .

Proof An ε -pushing of d_l results in a feasible schedule \mathcal{P}' (note that at most one of d_1 and d_2 can have release date in $[e_i, e_{i+1}]$) with the total completion time not exceeding that of \mathcal{P} . Thus, \mathcal{P}' is optimal. Note that an odd l would result in \mathcal{P}' having smaller total completion time than that of \mathcal{P} . Thus, l is even. If the ε makes at least one of the tree inequalities in (25) an equality, then the i -th block becomes i -normal and we get a contradiction in case of a maximal \mathcal{P} . On the other hand, if an ε makes all three inequalities in (25) holding strict, then $(d_1, \dots, d_l), l > 1$, is an alternating chain in \mathcal{P}' .

We prove, by contradiction, that the schedule \mathcal{P}' is A-free in $[C(\mathcal{P}', d_l), \infty)$. Suppose that some jobs a and b form an A-configuration at a point $t \geq C(\mathcal{P}', d_l)$ in \mathcal{P}' . Note that $t \geq C(\mathcal{P}, d_l)$ is not possible because \mathcal{P} and \mathcal{P}' are identical from $C(\mathcal{P}, d_l)$ on and \mathcal{P} is A-free in $[C(\mathcal{P}, d_l), \infty)$ by assumption. Thus, $t = C(\mathcal{P}', d_l)$. Therefore, $a = d_l$ and $b \in \{x, y\}$.

Let $\lambda \in [0, C(\mathcal{P}', b) - C(\mathcal{P}', d_l)]$ be a maximal real number such that for each $\varepsilon' \in [0, \lambda)$ there exists a feasible schedule $\mathcal{P}_{\varepsilon'}$ such that (d_1, \dots, d_l) is an alternating chain in $\mathcal{P}_{\varepsilon'}$ and ε' -pushing of d_l in $\mathcal{P}_{\varepsilon'}$ results in \mathcal{P} . Since l is even, the total completion time of $\mathcal{P}_{\varepsilon'}$ is the same as the total completion time of \mathcal{P}' for each $\varepsilon' \in [0, \lambda)$. Informally speaking, $\mathcal{P}_{\varepsilon'}$ is obtained by performing a modification that is ‘opposite’ to pushing of d_l . By definition of A-configuration, d_l is independent of any job that executes in the interval $(C(\mathcal{P}', d_l), C(\mathcal{P}', b))$ in \mathcal{P}' . Thus, the maximality of λ implies that taking $\varepsilon' = \lambda$ would result in a schedule $\mathcal{P}_{\varepsilon'}$ in which one of the following holds:

- Job sequence (d_1, \dots, d_l) is not an alternating chain in $\mathcal{P}_{\varepsilon'}$. Then we have two possibilities. The first possibility is that $t = C(\mathcal{P}_{\varepsilon'}, d_j) = C(\mathcal{P}_{\varepsilon'}, d_{j+1})$ for some $j \in \{1, \dots, l - 2\}$. Then, (d_1, \dots, d_{j+1}) is an almost alternating chain—a contradiction with Lemma 14. The second possibility is that either d_1 or d_2 is not present in the i -th block of $\mathcal{P}_{\varepsilon'}$. Then, the abnormality point of $\mathcal{P}_{\varepsilon'}$ is greater than i —a contradiction with the maximality of \mathcal{P} .
- $C(\mathcal{P}_{\varepsilon'}, d_l) = C(\mathcal{P}, b)$. This would imply that $C(\mathcal{P}_{\varepsilon'}, b) < C(\mathcal{P}, b)$ and this is not possible due to the optimality of \mathcal{P} .
- $s(\mathcal{P}_{\varepsilon'}, d_j) = r(d_j)$ for some $j \in \{3, \dots, l - 1\}$. In this case a contradiction follows from Lemma 15.
- $s(\mathcal{P}_{\varepsilon'}, b) = r(b)$. Since $C(\mathcal{P}_{\varepsilon'}, d_{l-1}) = s(\mathcal{P}_{\varepsilon'}, b)$, we again obtain a contradiction with Lemma 15.

Therefore, the lemma is proved. \square

Finally, we observe that the ε -pushing of d_l can readily be extended to the case where one of the jobs x or y starts in $(C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l))$ but neither of them completes in that interval.

6.3 Extending an alternating chain

We first prove that a single-job alternating chain is present in each maximal (and thus in A-maximal) schedule of abnormality point $i \neq \infty$.

Proposition 10 *If \mathcal{P} is a maximal schedule of abnormality point $i \neq \infty$, then a job in $A_i(\mathcal{P})$ with minimum completion time forms an alternating chain in \mathcal{P} .*

Proof Suppose that schedule (\mathcal{P}, e, ξ) is maximal. According to Observations 1 and 2 and Lemma 11, e_{i+1} is i -normal. By Lemma 12, we have $|A_i(\mathcal{P})| = 2$. Let $A_i(\mathcal{P}) = \{b, c\}$, where $C(\mathcal{P}, b) \leq C(\mathcal{P}, c)$. Denote $I = \{i + 1, \dots, \tau(b)\}$. Note that, by Lemma 10, $I \neq \emptyset$.

We prove the lemma by contradiction. More precisely, the assumption that b does not form an alternating chain in \mathcal{P} allows us to conclude that \mathcal{P} is not maximal. We may assume without loss of generality that c covers b in I . Indeed, if this is not the case, then we transform \mathcal{P} as follows. Let $t \in I$ be such that $\xi_t(b) > 0$ and job c is non-spanning in block t . Take $\varepsilon = \min\{\xi_t(b), \xi_t(c), e_{t+1} - e_t - \xi_t(c)\}$. Note that $\varepsilon > 0$ and, by Lemmas 4 and 7, $\xi_t(c) = e_{i+1} - e_i - \xi_t(b)$. The schedule obtained by a transformation

$$(e', \xi') = (e, \xi, \varepsilon, (t \overset{b}{\rightsquigarrow} i \overset{c}{\rightsquigarrow} t))$$

has the same total completion time and the same events as \mathcal{P} and either: $\xi'_t(b) = e_{i+1} - e_i$ (which happens when $\varepsilon = \xi_t(c)$); or $\xi'_t(b) = 0$ (which happens when $\varepsilon = \xi_t(b)$); or $\xi'_t(c) = e_{t+1} - e_t$ (which happens when $\varepsilon = e_{t+1} - e_t - \xi_t(c)$). In the former case we would obtain a schedule with abnormality point greater than i , which is not possible due to the maximality of \mathcal{P} . After repeating this transformation as long as c does not cover b in I we obtain the desired schedule.

Now we prove that $\tau(b) = i + 1$. Suppose for a contradiction that $\tau(b) > i + 1$. Thus, since c covers b in I , and $e_{\tau(b)}$ is an event of \mathcal{P} , we have $\xi_{\tau(b)-1}(b) = \xi_{\tau(b)-1}(c) = 0$ and there exists $a \in \mathcal{J} \setminus \{b, c\}$ such that $C(\mathcal{P}, a) = e_{\tau(b)}$. Find the maximum j , $j < \tau(b) - 1$, such that $\xi_j(b) \neq 0$. Note that $j \geq i$. Since job c covers b in I , c is spanning in block j . By Lemma 7, $a \notin \mathcal{J}(\xi_j)$. Lemma 9 applied to $a = b$, $a' = a$, j and $j' = \tau(b)$ gives $\tau(b) > j' = \tau(b)$ —a contradiction. This proves that (b) is an alternating chain in \mathcal{P} . \square

Proposition 11 *Let (d_1, \dots, d_l) be an alternating chain in a maximal schedule \mathcal{P} . Then, there is no idle time in block $\tau(d_{l-1}) + 1$ of \mathcal{P} .*

Proof Let e be the q events of \mathcal{P} and let ξ be its partition. Let i be the abnormality point of \mathcal{P} . Since \mathcal{P} has an alternating chain, $i \neq \infty$. Suppose for a contradiction that there is idle time in the $(\tau(d_{l-1}) + 1)$ -st block of \mathcal{P} . By Lemma 3, at most one job completes in the $(\tau(d_{l-1}) + 1)$ -st block of \mathcal{P} . By

Lemma 4, no job that does not complete in the $(\tau(d_{l-1}) + 1)$ -st block can be present in this block. Thus, d_l is the only job in block $(\tau(d_{l-1}) + 1)$ and the total length of the idle time is $x = e_{\tau(d_{l-1})+2} - e_{\tau(d_{l-1})+1}$. Construct a schedule \mathcal{P}' by performing an ε -pushing of d_l in \mathcal{P} with

$$\varepsilon = \min\{\alpha, \beta, \gamma, x/2\}.$$

Denote the resulting schedule by \mathcal{P}' . To complete the proof we observe that $C(\mathcal{P}', d_j) = r(a)$ for some job a and some $j \in \{3, \dots, l\}$ (when $\varepsilon = \gamma$) or $C(\mathcal{P}', d_{j-1}) = C(\mathcal{P}', d_j)$ for some $j \in \{2, \dots, l\}$ (when $\varepsilon \in \{\beta, x/2\}$) or there is no d_j in the i -th block of \mathcal{P}' for some $j \in \{1, 2\}$ (when $\varepsilon = \alpha$). Therefore, the choice of ε always results in \mathcal{P}' that has all blocks j , $j \in \{1, \dots, i\}$, being j -normal—a contradiction with the lemma assumption that \mathcal{P} is maximal. \square

The next theorem states that if a maximal schedule \mathcal{P} with an alternating chain (d_1, \dots, d_l) has no A-configuration at $t \geq C(\mathcal{P}, d_{l-1})$, then there exists another maximal schedule \mathcal{P}' with longer alternating chain $(d_1, \dots, d_l, d_{l+1})$ with no A-configuration at $t \geq C(\mathcal{P}', d_l)$.

Theorem 3 *Let \mathcal{P} be a maximal schedule of abnormality point $i \neq \infty$. If (d_1, \dots, d_l) , $l \geq 1$, is an alternating chain in \mathcal{P} and \mathcal{P} is A-free in $[C(\mathcal{P}, d_{l-1}), \infty)$, where $C(\mathcal{P}, d_0)$ is the $(i + 1)$ -st event of \mathcal{P} , then there exists a job d_{l+1} such that $(d_1, \dots, d_l, d_{l+1})$ is an alternating chain in some maximal schedule \mathcal{P}' that is A-free in $[C(\mathcal{P}', d_l), \infty)$.*

We leave the proof of the theorem to the end of the section. Theorem 2 guarantees that a maximal A-free schedule exists for in-trees. If this schedule is not normal, then we have a single-job alternating chain in it by Proposition 10. Theorem 3 guarantees that the alternating chain can be always extended by one job. However, the process of extending the alternating chain may result in a schedule which is not A-free in general—the source of this lies in Lemma 16. Luckily, we do not need the resulting schedule to be A-free—it suffices that it has no A-configuration at a completion of d_l , the last job from the alternating chain, or later—see the assumptions of Theorem 3. The above gives a sketch of the proof of the following theorem.

Theorem 4 *There exists a normal optimal schedule for each instance of problem P2|pmtn, in-tree, r_j, p_j | $\sum C_j$.*

Proof Take a maximal schedule \mathcal{P} and suppose for a contradiction that \mathcal{P} is not normal. Let $i \neq \infty$ be its abnormality point. By Proposition 10, $A_i(\mathcal{P}) = \{d_1, d_2\}$, and (d_1) is an alternating chain in \mathcal{P} . Next, by Theorem 2, there is a maximal schedule \mathcal{P}^1 with its abnormality point i and alternating chain (d_1) which is also A-free. Thus, in particular, \mathcal{P}^1 is A-free in $[x, \infty)$, where x is the $(i + 1)$ -st event of \mathcal{P}^1 . Finally, by Theorem 3 and a simple inductive argument, there exists

a schedule \mathcal{P}^{n+1} with an alternating chain of $l = n + 1$ jobs, contradicting the fact that the number of jobs equals n . \square

Corollary 1 *For the given set of n jobs, there exists an optimal schedule for $P2|pmtn, in-tree, r_j, p_j|\sum C_j$ such that each job start, preemption, resumption, or completion occurs at a time point that is a multiple of $1/2^{2^n}$.*

6.3.1 Proof of Theorem 3

By Lemma 12, $|A_i(\mathcal{P})| = 2$. If $l = 1$, then take d_{l+1} to be the job in $A_i(\mathcal{P}) \setminus \{d_1\}$. If $l > 1$, then by Proposition 11, there is no idle time in $(\tau(d_{l-1}) + 1)$ -st block of \mathcal{P} . Thus, $|\mathcal{J}(\xi_{\tau(d_{l-1})+1})| > 1$. Take d_{l+1} to be a job in $\mathcal{J}(\xi_{\tau(d_{l-1})+1}) \setminus \{d_l\}$ that starts or resumes at $C(\mathcal{P}, d_{l-1})$.

We will perform several schedule modifications leading to some maximal schedule \mathcal{P}' that is A-free in $[C(\mathcal{P}', d_l), \infty)$ and has an alternating chain (d_1, \dots, d_{l+1}) . We point out that some steps of the proof redefine the job d_{l+1} selected above. By assumption, $d_1 \in A_i(\mathcal{P})$, and by the choice of $d_2, d_2 \in A_i(\mathcal{P})$. Thus, (C1) follows for (d_1, \dots, d_{l+1}) .

We now prove that

$$C(\mathcal{P}, d_{l+1}) > C(\mathcal{P}, d_l). \tag{27}$$

Note that if $l = 1$, then (27) follows from Proposition 10 and from the choice of d_1 and d_2 . Thus, $l \geq 2$ from now on.

We begin by proving, by contradiction, that d_{l+1} executes non-preemptively in $[C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$. First, we observe that no job, except for d_l , completes in the interval $(C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$ for otherwise job d_{l+1} must complete in $(C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$ and thus jobs d_{l-1} and d_{l+1} form an A-configuration in \mathcal{P} —a contradiction since \mathcal{P} is A-free in $[C(\mathcal{P}, d_{l-1}), \infty)$.

Second, $|J \setminus \{d_l, d_{l+1}\}| \leq 1$, where J is the set of jobs executed in $(C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$. Otherwise, there is a pair of jobs in $\{x, y, d_l\}$, where $\{x, y\} \subseteq J \setminus \{d_l, d_{l+1}\}$, that interlace—contradiction by Lemma 5 and by the fact that d_l is preempted in $(C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$. Indeed, this pair of job consists of a job $z \in \{x, y, d_l\}$ with minimum completion time among those three jobs, and a job in $\{x, y, d_l\} \setminus \{z\}$ that is non-spanning in block $\tau(z)$.

Finally, we show that without loss of generality $J = \{d_l, d_{l+1}\}$. Suppose otherwise, i.e., $J = \{x, d_l, d_{l+1}\}$. The ε -pushing of d_l with

$$\varepsilon = \min \{ \alpha, \beta, \gamma, C(\mathcal{P}, d_l) - C(\mathcal{P}, d_{l-1}) - \xi_{\tau(d_{l-1})+1}(d_{l+1}), \xi_{\tau(d_{l-1})+1}(d_{l+1}) \}$$

results in a schedule \mathcal{P}' that either has all blocks $j, j \in \{1, \dots, i\}$, being j -normal (this happens when $\varepsilon \in \{\alpha, \beta, \gamma\}$)—a contradiction with the lemma assumption that \mathcal{P} is maximal; or it has d_l and d_{l+1} as the only two




Fig. 16 Schedule transformations in the proof of Theorem 3

jobs executed in interval $[C(\mathcal{P}', d_{l-1}), C(\mathcal{P}', d_l)]$; or it has d_l and x as the only two jobs executed in interval $[C(\mathcal{P}', d_{l-1}), C(\mathcal{P}', d_l)]$. (See Fig. 16a for this transformation when $\varepsilon = C(\mathcal{P}, d_l) - C(\mathcal{P}, d_{l-1}) - \xi_{\tau(d_{l-1})+1}(d_{l+1})$.) In the latter case, i.e., when $\varepsilon = \xi_{\tau(d_{l-1})+1}(d_{l+1})$, we take x as d_{l+1} from now on. The schedule \mathcal{P}' is maximal, by Lemma 16 it is A-free in $[C(\mathcal{P}, d_l), \infty)$, and (d_1, \dots, d_l) is an alternating chain in \mathcal{P}' . Thus, without loss of generality we can take \mathcal{P} as being \mathcal{P}' from now on. Then, we have that d_{l+1} executes non-preemptively in $[C(\mathcal{P}, d_{l-1}), C(\mathcal{P}, d_l)]$ and, by Lemma 14, (27) holds. Thus, (C2) holds for (d_1, \dots, d_{l+1}) in \mathcal{P} .


We argue that there is no idle time in the $(\tau(d_l) + 1)$ -st block of \mathcal{P} . Our argument is by contradiction. If there is idle time in the block, then by (27), d_{l+1} is the only job there. Thus, d_{l+1} completes at $e_{\tau(d_l)+2}$ or there is a release date pinned job a that starts at $e_{\tau(d_l)+2}$, i.e., $r(a) = e_{\tau(d_l)+2}$. In the former case we have an alternating chain (d_1, \dots, d_{l+1}) with idle time in block $\tau(d_l) + 1$ which contradicts Proposition 11 and completes the proof. In the latter case we use an extended ε -pushing of d_{l+1} (the operation of the ε -pushing can be generalized in a straightforward way to the case when d_{l+1} is preempted in interval $[C(\mathcal{P}, d_l), C(\mathcal{P}, d_{l+1})]$, see also Fig. 16b—we omit a formal definition of the extended pushing) with

$$\varepsilon = \min \{ \alpha, \beta, \gamma, r(a) - C(\mathcal{P}, d_l) \},$$

which results in a schedule \mathcal{P}' that has all blocks $j, j \in \{1, \dots, i\}$, being j -normal—a contradiction with the lemma assumption that \mathcal{P} is maximal. (See Fig. 16b for an illustration of this schedule transformation when $\varepsilon = r(a) - C(\mathcal{P}, d_l)$.) Therefore, without loss of generality we may assume there is no idle time block $\tau(d_l) + 1$, and thus some job $a \neq d_{l+1}$ starts or resumes at $e_{\tau(d_l)+1}$.

We next describe a finite iterative process that starts with \mathcal{P} , produces a schedule \mathcal{P}_u in its u -th iteration, $u \geq 1$, and stops after $T \geq 1$ iterations. We then show that if $T = 1$,

Fig. 17 The transformation from \mathcal{P}_{u-1} to \mathcal{P}'_u



then the schedule \mathcal{P}_1 satisfies the conditions of the lemma. However, if $T > 1$, then we prove that there is a pair of jobs x and y that allows the iterative process to construct maximal schedules $\mathcal{P}_1, \dots, \mathcal{P}_{T-1}$ each with alternating chain (d_1, \dots, d_l) , yet at the same time the pair prevents d_{l+1} from satisfying (C3) in $\mathcal{P}_1, \dots, \mathcal{P}_{T-1}$. However, an exit after $T > 1$ iterations is only possible through a schedule \mathcal{P}_T such that the total completion time of \mathcal{P}_T is smaller than that of \mathcal{P} or both schedules have the same total completion times but the abnormality point of \mathcal{P}_T is greater than i , which contradicts the maximality of \mathcal{P} . Therefore the iterative process must exit after exactly one iteration producing the desired schedule \mathcal{P}_1 —more than one iteration leads to a contradiction.

In order to describe the iterative process formally, we introduce a key definition and related notation. We say that a schedule $(\mathcal{P}_u, \xi^u, e^u)$ with $\tau_u = \tau_{\mathcal{P}_u}$ is d_{l+1} -preempted if there exists a pair of jobs x and y such that the following conditions are satisfied:

- (11) \mathcal{P}_u is maximal, (d_1, \dots, d_l) is an alternating chain in \mathcal{P}_u , i is the abnormality point of \mathcal{P}_u , and d_l and d_{l+1} are the only two jobs executed in $[C(\mathcal{P}_u, d_{l-1}), C(\mathcal{P}_u, d_l)]$, and $C(\mathcal{P}_u, d_{l+1}) > C(\mathcal{P}_u, d_l)$;
- (12) Some job a_u starts or resumes at $C(\mathcal{P}_u, d_l)$;
- (13) d_{l+1} covers each job in $\{a_1, \dots, a_u\}$ in $I_u = \{\tau_u(d_l) + 1, \dots, k_u\}$, where $k_u = \min\{\tau_u(a_u), \tau_u(d_{l+1})\}$;
- (14) $C(\mathcal{P}_u, a_u) > C(\mathcal{P}_u, d_{l+1})$;
- (15) There exists $j_u < \tau_u(d_{l+1})$ such that $\mathcal{J}(\xi_{j_u}^u) = \{x, y\}$ and $\min\{r(x), r(y)\} > C(\mathcal{P}_u, d_l)$.

In the following we prove that all schedules $\mathcal{P}_1, \dots, \mathcal{P}_{T-1}$ are d_{l+1} -preempted when $T > 1$. Let initially $u = 1$, and the iterative process is as follows. (\mathcal{P}_0 refers to \mathcal{P} .)

Step 1: Moving a_{u-1} from block $\tau_{u-1}(d_l) + 1$ to block k_{u-1} .

If $u = 1$, then let $\mathcal{P}'_u = \mathcal{P}$ and go to Step 2. If $u > 1$, then we construct \mathcal{P}'_u by an extended ε -pushing of d_{l+1} in \mathcal{P}_{u-1} and moving a piece of a_{u-1} of length ε from block

$\tau_{u-1}(d_l) + 1$ to block k_{u-1} , where

$$\varepsilon = \min \{ \alpha, \beta, \gamma, \lceil C(\mathcal{P}_{u-1}, d_l) \rceil - C(\mathcal{P}_{u-1}, d_l), \xi_{\tau_{u-1}(d_l)+1}^{u-1}(a_{u-1}), e_{k_{u-1}+1}^{u-1} - e_{k_{u-1}}^{u-1} - \xi_{k_{u-1}}^{u-1}(a_{u-1}) \}.$$

Denote by e' and ξ' the events and the partition of \mathcal{P}'_u , respectively. Let for brevity $\tau_{\mathcal{P}'_u} = \tau'$. Figure 17 depicts the transition from \mathcal{P}_{u-1} to \mathcal{P}'_u for $u > 1$. Note that \mathcal{P}'_u and \mathcal{P}_{u-1} have the same total completion times.

If $\varepsilon \in \{ \alpha, \beta, \gamma, \lceil C(\mathcal{P}'_u, d_l) \rceil - C(\mathcal{P}'_u, d_l) \}$, then the abnormality point of \mathcal{P}'_u is greater than i and, having the required contradiction, we stop the iterative process with $T = u$.

For the two remaining values we have that the number of blocks in $[C(\mathcal{P}'_u, d_l), C(\mathcal{P}'_u, d_{l+1})]$ in \mathcal{P}'_u is one less than the number of blocks in $[C(\mathcal{P}_{u-1}, d_l), C(\mathcal{P}_{u-1}, d_{l+1})]$ in \mathcal{P}_{u-1} . (We give an appropriate argument at the end of the proof of the lemma.)

Also, if $\varepsilon = e_{k_{u-1}+1}^{u-1} - e_{k_{u-1}}^{u-1} - \xi_{k_{u-1}}^{u-1}(a_{u-1})$, then the total completion time of \mathcal{P}'_u is smaller than the total completion time of \mathcal{P}_{u-1} provided that $(\xi_{k_{u-1}}^{u-1}(a_{u-1}) = 0$ and $\xi_{k_{u-1}-1}^{u-1}(d_{l+1}) < e_{k_{u-1}}^{u-1} - e_{k_{u-1}-1}^{u-1}$) because d_{l+1} completes in \mathcal{P}'_u strictly prior to $e_{\tau_u(d_{l+1})} = C(\mathcal{P}_u, d_{l+1}) - \varepsilon$. Then, we stop the iterative process with $T = u$. Otherwise, set $a_u := a_{u-1}$ and go to Step 2.

If $\varepsilon = \xi_{\tau_{u-1}(d_l)+1}^{u-1}(a_{u-1})$, then there is a job a_u that starts or resumes at $C(\mathcal{P}'_u, d_l)$. Go to Step 2.

Step 2: Making d_{l+1} cover a_u .

Denote for brevity $\tau' = \tau_{\mathcal{P}'_u}$ and let e' and ξ' be the events and the partition of \mathcal{P}'_u , respectively. If d_{l+1} covers a_u in

$$I'_u = \{ \tau'(d_l) + 1, \dots, \min\{ \tau'(d_{l+1}), \tau'(a_u) \} \},$$

then set $\mathcal{P}_u = \mathcal{P}'_u$ and go to Step 3. Otherwise we obtain \mathcal{P}_u from \mathcal{P}'_u as follows. Find $t \in I'_u$ such that $\xi'_t(a_u) > 0$ and $\xi'_t(d_{l+1}) < e'_{t+1} - e'_t$. By Lemma 5, $t < \tau'(a_u)$, and by

Lemma 3, $t < \tau'(d_{l+1})$. Let

$$\varepsilon' = \min \left\{ e'_{\tau'(d_{l+1})+1} - r(a_u) - \xi'_{\tau'(d_l)}(a_u), \right. \\ \xi'_t(a_u), e'_{t+1} - e'_t - \xi'_t(d_{l+1}), \\ \left. (e'_{\tau'(d_{l+1})+1} - e'_{\tau'(d_l)})/2 - \xi'_{\tau'(d_l)}(a_u) \right\}.$$

We construct \mathcal{P}_u with events e^u and partition ξ^u , where

$$(e^u, \xi^u) = (e', \xi', \varepsilon', (t \xrightarrow{a_u} \tau'(d_l) \xrightarrow{\sim} t)),$$

and then:

- If $\varepsilon' \in \{e'_{\tau'(d_{l+1})+1} - r(a_u) - \xi'_{\tau'(d_l)}(a_u), (e'_{\tau'(d_{l+1})+1} - e'_{\tau'(d_l)})/2 - \xi'_{\tau'(d_l)}(a_u)\}$, then we do ε'' -pushing of d_l in \mathcal{P}'_u with

$$\varepsilon'' = \min \{\alpha, \beta, \gamma, \varepsilon'\}$$

to get a schedule \mathcal{P}_u that has all blocks $j, j \in \{1, \dots, i\}$, being j -normal—in such case we stop the iterative process with $T = u$; (This ε'' -pushing is shown in Fig. 16a with $x = a_u$ and $\varepsilon' = \varepsilon''$.)

- If $\varepsilon' = \xi'_t(a_u)$, then $\xi''_t(a_u) = 0$, i.e., a_u is no longer in block t as required;
- If $\varepsilon' = e'_{t+1} - e'_t - \xi'_t(d_{l+1})$, then $\xi''_t(d_{l+1}) = e''_{t+1} - e''_t$ as required.

If d_{l+1} does not cover a_u in I_u and $\varepsilon' > 0$, then repeat Step 2 for \mathcal{P}_u . Thus, either in the resulting schedule \mathcal{P}_u, d_{l+1} covers a_u in I_u , in which case go to Step 3, or d_{l+1} does not cover a_u in I_u and $\varepsilon' = 0$ (then $\xi''_{\tau_u(d_l)}(a_u) = \xi''_{\tau_u(d_l)}(d_{l+1})$ or $e''_{\tau_u(d_{l+1})+1} = r(a_u)$), in which case the abnormality point of \mathcal{P}_u is greater than i and, having the required contradiction, we stop the iterative process with $T = u$.

Step 3: Pushing a_u out of $[C(\mathcal{P}_u, d_{l-1}), C(\mathcal{P}_u, d_l)]$. If a part of a_u executes in $[C(\mathcal{P}_u, d_{l-1}), C(\mathcal{P}_u, d_l)]$ in \mathcal{P}_u , then perform an ε'' -pushing of d_l in \mathcal{P}_u as in Fig. 16a with $x = a_u$, where

$$\varepsilon'' = \min \left\{ \alpha, \beta, \gamma, [C(\mathcal{P}_u, d_l)] - C(\mathcal{P}_u, d_l), \xi''_{\tau_u(d_l)}(a_u) \right\}.$$

If $\varepsilon'' = \min \{\alpha, \beta, \gamma, [C(\mathcal{P}_u, d_l)] - C(\mathcal{P}_u, d_l)\}$, then the abnormality point of the resulting schedule is greater than i and, having the required contradiction, we stop the iterative process with $T = u$. If (d_1, \dots, d_{l+1}) is an alternating chain in \mathcal{P}_u , then also stop with $T = u$. Otherwise, go to Step 4.

Step 4: Moving to the next iteration. Set $u := u + 1$ and return to Step 1.

We now briefly sketch the reminder of the proof. For the time being let us assume that the iteration process ends after

$T \geq 1$ iterations, and that $\varepsilon > 0$ in Step 1 for $u > 1$. We prove these two assumptions at the end of the proof. In the following $\mathcal{P}'_u, \mathcal{P}''_u$, and \mathcal{P}_u refer to the schedules obtained at the end of Steps 1, 2 and 3, respectively, $u \geq 1$.

Let $u = 1$. Note that $\mathcal{P}'_1 = \mathcal{P}$. Then either d_{l+1} executes without preemption in $[C(\mathcal{P}'_1, d_l), C(\mathcal{P}'_1, d_{l+1})]$ or not. In the former case, Step 3 ensures that d_{l+1} executes in \mathcal{P}_1 without preemption in $[C(\mathcal{P}_1, d_{l-1}), C(\mathcal{P}_1, d_{l+1})]$. The latter implies that (C3) holds for \mathcal{P}_1 and by Claim 4 below that \mathcal{P}_1 is A-free, which proves the lemma. We then have $T = 1$. If d_{l+1} is preempted in $[C(\mathcal{P}'_1, d_l), C(\mathcal{P}'_1, d_{l+1})]$, then $T > 1$ and it suffices to show that in this case we get a contradiction. To that end we show that, if $T > 1$, then \mathcal{P}_1 is d_{l+1} -preempted (see Claims 2, 3, and 5 below), and if \mathcal{P}_{u-1} is d_{l+1} -preempted, then \mathcal{P}_u is d_{l+1} -preempted as well, $u \in \{2, \dots, T - 1\}$ (see Claims 2, 3, and 6 below). This process of generating d_{l+1} -preempted schedules cannot continue *ad infinitum* since the process exits after T iterations. However, any exit schedule \mathcal{P}_T certifies that \mathcal{P} is not maximal which gives the required contradiction.

We now proceed with details. Note that, for each $u \in \{1, \dots, T\}$, (C1) and (C2) hold for (d_1, \dots, d_{l+1}) in \mathcal{P}_u and

$$\xi''_{\tau_u(d_l)}(a_u) = 0 \wedge \xi''_{\tau_u(d_{l+1})}(a_u) > 0, \tag{28}$$

and that (I1), (I2) and (I3) follow directly from the definition of the iterative process above:

Claim 2 *Let $T > 1$. If $u = 1$, or $u \in \{2, \dots, T - 1\}$ and \mathcal{P}_{u-1} is d_{l+1} -preempted, then \mathcal{P}_u satisfies conditions (I1), (I2) and (I3).*

Proof Note that, by construction, the total completion times and abnormality points of the schedules $\mathcal{P}, \mathcal{P}_1, \dots, \mathcal{P}_{T-1}$ are the same. Thanks to Step 3 and the fact that $u < T, d_l$ and d_{l+1} are the only two jobs executed in $[C(\mathcal{P}_u, d_{l-1}), C(\mathcal{P}_u, d_l)]$ for each $u \in \{1, \dots, T - 1\}$. Finally, by Lemma 14, we have $C(\mathcal{P}_u, d_{l+1}) > C(\mathcal{P}_u, d_l)$ for $u = 1$, and by (I5) for schedule \mathcal{P}_{u-1} we have $C(\mathcal{P}_{u-1}, d_l) < \min\{r(x), r(y)\} < C(\mathcal{P}_{u-1}, d_{l+1})$ for $u > 1$. By construction, $C(\mathcal{P}_u, d_l) < [C(\mathcal{P}_{u-1}, d_l)]$ and $j_u < \tau_u(d_{l+1})$ such that $\mathcal{J}(\xi''_{j_u}) = \{x, y\}$. Thus, $C(\mathcal{P}_u, d_l) < \min\{r(x), r(y)\} < C(\mathcal{P}_u, d_{l+1})$. Therefore, \mathcal{P}_u satisfies (I1).

By Proposition 11, there is no idle time in block $\tau_{\mathcal{P}}(d_l) + 1$ in \mathcal{P} . Thus the choice of a_1 ensures that it starts or resumes at $C(\mathcal{P}_1, d_l)$ for $u = 1$. For $u > 1$, the job a_u always exists because there is no idle time in block $\tau_u(d_l) + 1$ in \mathcal{P}_u . This follows from the fact that otherwise, by construction, there would be idle time in $[C(\mathcal{P}_{u-1}, d_l), C(\mathcal{P}_{u-1}, d_{l+1})]$ in \mathcal{P}_{u-1} , which, since (I3) and (I4) hold for \mathcal{P}_{u-1} , implies that a_{u-1} can be completed earlier in \mathcal{P}_{u-1} , which contradicts its optimality. Hence, \mathcal{P}_u satisfies (I2).

Finally, Steps 2 and 3 and $u < T$ ensure that (I3) holds for \mathcal{P}_u . □

Claim 3 Let $T > 1$. If $u = 1$, or $u \in \{2, \dots, T - 1\}$ and \mathcal{P}_{u-1} is d_{l+1} -preempted, then \mathcal{P}_u satisfies condition (I4).

Proof It suffices to argue that

$$k_u < \tau_u(a_u). \tag{29}$$

Denote for brevity $k = k_u$. Suppose for a contradiction that $k = \tau_u(a_u)$. By Lemma 3, $\xi_k^u(a_u) = e_{k+1}^u - e_k^u$. By Claim 2, (I3) holds for \mathcal{P}_u . Thus, d_{l+1} covers a_u in I_u and, by Lemma 5, $\xi_k^u(d_{l+1}) = e_{k+1}^u - e_k^u$. Hence, there exists $a' \in \mathcal{J}$ such that $C(\mathcal{P}_u, a') = e_k^u$ because e_k^u is an event in \mathcal{P}_u .

If $a' \neq d_l$, then $\tau_u(d_l) < k$ and $\xi_{k-1}^u(a_u) = 0$ because, again, d_{l+1} covers a_u in I_u . Then, let $j < k - 1$ be such that $\xi_j^u(a_u) > 0$ and $\xi_{j'}^u(a_u) = 0$ for each $j' \in \{j + 1, \dots, k - 1\}$. By (28) such a j exists. By Lemma 7 and the fact that d_{l+1} covers a_u in I_u , $\xi_j^u(a_u) = \xi_j^u(d_{l+1}) = e_{j+1}^u - e_j^u$ which implies $\xi_j^u(a') = 0$. Lemma 9 applied to $a = a_u, a', j$ and $j' = k$, leads to a contradiction.

It remains to consider the case when $a' = d_l$. We have that $u > 1$ because otherwise the jobs d_l and a_1 form an A-configuration in \mathcal{P} —a contradiction since \mathcal{P} is A-free. We have

$$\begin{aligned} C(\mathcal{P}_{u-1}, a_{u-1}) &> C(\mathcal{P}_{u-1}, d_{l+1}) \\ &\geq C(\mathcal{P}_u, d_{l+1}) \geq C(\mathcal{P}_u, a_u). \end{aligned} \tag{30}$$

The first inequality follows from (I4) for \mathcal{P}_{u-1} , which is d_{l+1} -preempted; the second inequality follows by construction of \mathcal{P}_u , while the last one holds by assumption that $k = \tau_u(a_u)$. Moreover, the construction ensures that the inequality $C(\mathcal{P}_u, d_{l+1}) \geq C(\mathcal{P}_u, a_u)$ implies

$$C(\mathcal{P}_{u-1}, d_{l+1}) \geq C(\mathcal{P}_{u-1}, a_u). \tag{31}$$

Thus, by (30), we have $C(\mathcal{P}_{u-1}, a_{u-1}) > C(\mathcal{P}_{u-1}, a_u)$. Therefore, (31) and (I2), (I3) for \mathcal{P}_{u-1} imply that a_u and a_{u-1} interlace in \mathcal{P}_{u-1} . Finally, by (I1) applied to \mathcal{P}_{u-1} , \mathcal{P}_{u-1} is optimal and hence we arrive at a contradiction with Lemma 5. Hence, (29) follows, and the proof of the lemma is completed. \square

Claim 4 If the job d_{l+1} executes with no preemption in interval $[C(\mathcal{P}_1, d_{l-1}), C(\mathcal{P}_1, d_{l+1})]$ in \mathcal{P}_1 , then \mathcal{P}_1 is A-free in $[C(\mathcal{P}_1, d_l), \infty)$.

Proof In view of Lemma 16, it suffices to prove that \mathcal{P}_1'' at the end of Step 2 is A-free in $[C(\mathcal{P}', d_l), \infty)$, if the step is not vacuous. Let $\tau(a_1) \equiv \tau_{\mathcal{P}_1''}$ and $k = \min\{\tau(a_1), \tau(d_{l+1})\}$ for convenience. Let e'' and ξ'' be the events and the partition of \mathcal{P}_1'' , respectively. We start with an observation concerning the construction of \mathcal{P}_1'' , namely, the sequence of events, that is the start and completion times of jobs, is the same in $\mathcal{P} = \mathcal{P}'_1$ and \mathcal{P}_1'' . More precisely,

(B1) $\mathcal{J}(\xi_j) = \mathcal{J}(\xi_j'')$ for $j \in \{1, \dots, \tau(d_l) - 1\}$ and $j \geq k + 1$;

(B2) $\mathcal{J}(\xi_{\tau_{\mathcal{P}}(d_l)}) = \{d_l, d_{l+1}\}$ and $\mathcal{J}(\xi_{\tau(d_l)}'') = \{d_l, d_{l+1}, a_1\}$;

(B3) $\xi_j''(x) = \xi_j(x)$ for $x \notin \{a_1, d_{l+1}\}$, $\xi_j''(a_1) \leq \xi_j(a_1)$, $\xi_j''(d_{l+1}) \geq \xi_j(d_{l+1})$, and $\xi_j''(a_1) + \xi_j''(d_{l+1}) = \xi_j(a_1) + \xi_j(d_{l+1})$ for each $j \in \{\tau(d_l) + 1, \dots, k\}$.

Suppose for a contradiction that schedule \mathcal{P}_1'' is not A-free in interval $[C(\mathcal{P}_1'', d_l), \infty)$. Then, by (B1)–(B3) and since the schedule \mathcal{P} is A-free in $[C(\mathcal{P}, d_{l-1}), \infty)$, $a = a_1$ must be one of the two jobs that form an A-configuration in \mathcal{P}_1'' .

Let a and a job x form an A-configuration at $e_{\tau(a)}''$ in \mathcal{P}_1'' . By definition of A-configuration, x is not preempted in $(e_j'', e_{\tau(a)}'')$ for some $j < \tau(a)$ in \mathcal{P}_1'' and $\xi_{j-1}''(x) < e_{j-1}'' - e_j''$. By (B2), $\tau(d_l) + 1 \leq j$. Thus, $\tau(d_l) + 1 \leq j \leq k$ for otherwise, by (B1), a and x or form an A-configuration in \mathcal{P} —a contradiction. Moreover, there is a block $t \in \{j, \dots, k\}$ in \mathcal{P} with $\{a, x\} \subseteq \mathcal{J}(\xi_t)$ for otherwise again a and x form an A-configuration in \mathcal{P} —a contradiction. Therefore, if $x = d_{l+1}$, then $\xi_t(a) > 0$ and $\xi_t(d_{l+1}) = e_{t+1} - e_t$ in \mathcal{P} and it remains so in \mathcal{P}_1'' which follows from the transformation in Step 1. Thus, a and d_{l+1} do not form an A-configuration in \mathcal{P}_1'' which contradicts our assumption that $x = d_{l+1}$. If $x \neq d_{l+1}$, then $\xi_t(d_{l+1}) < e_{t+1} - e_t$ and hence a job $a' \in \{a, x\}$ with $\xi_{\tau_{\mathcal{P}}(d_{l+1})}(d_{l+1}) < e_{\tau_{\mathcal{P}}(d_{l+1})+1} - e_{\tau_{\mathcal{P}}(d_{l+1})}$ and d_{l+1} interlace in \mathcal{P} when $\tau(a) > \tau(d_{l+1}) = k$ —a contradiction with Lemma 5. Thus it remains to consider $\tau(a) \leq \tau(d_{l+1})$. By Step 3, no preemption of d_{l+1} in $[C(\mathcal{P}_1, d_{l-1}), C(\mathcal{P}_1, d_{l+1})]$ in \mathcal{P}_1 implies no preemption of d_{l+1} in $[C(\mathcal{P}_1'', d_l), C(\mathcal{P}_1'', d_{l+1})]$ in \mathcal{P}_1'' . Therefore, the jobs a and x interlace in \mathcal{P}_1'' if $\tau(d_l) \geq \tau(a)$ —a contradiction by Lemma 5.

Now, suppose that a and a job x such that a completes at $e_{\tau(x)}''$ form an A-configuration in \mathcal{P}_1'' . Since d_{l+1} covers a in I'_1 in \mathcal{P}_1'' , we have $x \neq d_{l+1}$. By definition of A-configuration, a is not preempted in $(e_j'', e_{\tau(x)}'')$ for some $j \leq \tau(a)$ in \mathcal{P}_1'' and $\xi_{j-1}''(a) < e_{j-1}'' - e_j''$. Also, $\tau(d_l) + 1 \leq j \leq k$ for otherwise, by (B1), a and x form an A-configuration in \mathcal{P} —a contradiction. Moreover, there is a block $t \in \{j, \dots, k - 1\}$ in \mathcal{P} with $\{a, x\} \subseteq \mathcal{J}(\xi_t)$ for otherwise again a and x form an A-configuration in \mathcal{P} —a contradiction. Therefore, $\{d_{l+1}, x\} \subseteq \mathcal{J}(\xi_t'')$, and a and x interlace in \mathcal{P}_1'' —a contradiction by Lemma 5. \square

Claim 5 If $T = 1$, then \mathcal{P}_1 is maximal, A-free in interval $[C(\mathcal{P}, d_l), \infty)$ and (d_1, \dots, d_{l+1}) is an alternating chain in \mathcal{P}_1 . If $T > 1$, then \mathcal{P}_1 satisfies condition (I5).

Proof If $T = 1$, then by the maximality of \mathcal{P} we have that \mathcal{P}_1 is maximal and (d_1, \dots, d_{l+1}) is an alternating chain in \mathcal{P}_1 . By Claim 4, \mathcal{P}_1 is (A, $[C(\mathcal{P}, d_l), \infty)$)-free, which completes the proof in this case.

Suppose now that $T > 1$. Let for brevity $k = k_1$ and $\tau_1 = \tau_{\mathcal{P}_1}$ in the proof of Claim 5. By Claim 3, \mathcal{P}_1 satisfies (I4).

Hence, $k = \tau_1(d_{l+1})$, and $\xi_k^1(d_{l+1}) = e_{k+1}^1 - e_k^1$ by Lemma 3. Note that $\xi_j^1(d_{l+1}) = e_{j+1}^1 - e_j^1$ for each $j \in \{\tau_1(d_l) + 1, \dots, k\}$ is not possible because then (d_1, \dots, d_{l+1}) is an alternating chain in \mathcal{P}_1 and hence the iterative process would stop with $T = 1$ in Step 3. We show that otherwise we can find the desired jobs x and y in (15). The key to finding the jobs is the existence of a block $j \in \{\tau_1(d_l) + 1, \dots, k - 1\}$ such that $\xi_j^1(d_{l+1}) < e_{j+1}^1 - e_j^1$. Take the smallest such j . Let $\{x, y\} \subseteq \mathcal{J}(\xi_j^1) \setminus \{a_1, d_{l+1}\}$. Such two jobs exist since there is no idle time in block j and $a_1 \notin \mathcal{J}(\xi_j^1)$ because, by Claim 2, d_{l+1} covers a_1 in I_1 . We now prove that

$$r(x) > C(\mathcal{P}_1, d_l) \quad \text{and} \quad r(y) > C(\mathcal{P}_1, d_l). \tag{32}$$

To that end we first argue that no predecessor of x or y is executed in $(C(\mathcal{P}_1, d_l), e_j^1]$. By contradiction, suppose z is a predecessor of x or y that completes in $(C(\mathcal{P}_1, d_l), e_j^1]$. Then, z must also start in $[C(\mathcal{P}_1, d_l), e_j^1]$ for otherwise by (28) and the fact that d_{l+1} covers a_1 in I_1 (by Claim 2) we get that z interlaces either a_1 or d_{l+1} in \mathcal{P}_1 —a contradiction by Lemma 5. Therefore, z starts in $[C(\mathcal{P}_1, d_l), e_j^1]$ and thus there is a block $j' \in \{\tau_1(d_l) + 1, \dots, j - 1\}$ such that $\xi_{j'}^1(z) > \xi_{j'}^1(d_{l+1})$. The latter is guaranteed by the fact that z executes in $[C(\mathcal{P}_1, d_l), e_j^1]$. Since, by Claim 2, d_{l+1} covers a_1 in I_1 , we have $\xi_{j'}^1(a_1) = 0$. Therefore, there is a job $w, w \notin \{a_1, d_{l+1}\}$, such that $\xi_{j'}^1(w) > 0$. Thus, we get a contradiction by our definition of j since $j' < j$.

Second, arguing by contradiction, suppose without loss of generality that $r(x) \leq C(\mathcal{P}_1, d_l)$. If $\xi_k^1(a_1) < e_{k+1}^1 - e_k^1$, then take

$$\begin{aligned} \varepsilon = \min \{ & \xi_{\tau_1(d_l)+1}^1(a_1), e_{k+1}^1 - e_k^1 - \xi_k^1(a_1), e_{j+1}^1 - e_j^1 \\ & - \xi_j^1(d_{l+1}), \xi_j^1(x), e_{\tau_1(d_l)+2}^1 - e_{\tau_1(d_l)+1}^1 \\ & - \xi_{\tau_1(d_l)+1}^1(x) \} \end{aligned}$$

and let \mathcal{P}'' be a schedule with events e'' and partition ξ'' , where

$$(e'', \xi'') = (e', \xi', \varepsilon, (\tau_1(d_l) + 1 \overset{a_1}{\rightsquigarrow} k \overset{d_{l+1}}{\rightsquigarrow} j \overset{x}{\rightsquigarrow} \tau_1(d_l) + 1)).$$

Note that $j > \tau_1(d_l) + 1$ because $\xi_j^1(a_1) = 0$. The assumption $r(x) \leq C(\mathcal{P}_1, d_l)$ implies that \mathcal{P}'' is feasible. Thus, we get a contradiction since the total completion time of \mathcal{P}'' is smaller than that of \mathcal{P}_1 .

On the other hand, if $\xi_k^1(a_1) = e_{k+1}^1 - e_k^1$, then $x \notin \mathcal{J}(\xi_k^1)$, and (28) and the fact that d_{l+1} covers a_1 in I_1 (by Claim 2) imply that $\mathcal{J}(\xi_{\tau_1(d_l)+1}^1) \cap \{x, y\} = \emptyset$. Thus, $C(\mathcal{P}_1, x) \leq e_k^1$ or $C(\mathcal{P}_1, x) > e_{k+1}^1$. In the former case a_1 and x interlace, and in the latter case x and d_{l+1} interlace—contradiction

with Lemma 5. Therefore, (32) holds and \mathcal{P}_1 satisfies (15) as required. \square

Claim 6 *Let $T > 1$. Then, \mathcal{P}_u satisfies condition (15) for $u \in \{1, \dots, T - 1\}$.*

Proof By Claim 5, \mathcal{P}_1 satisfies (15). By induction on $u = 1, \dots, T - 1$, we have $\min\{r(x), r(y)\} > C(\mathcal{P}_u, d_l)$. Let j_u be the earliest j such that $\mathcal{J}(\xi_j^u) = \{x, y\}$ in \mathcal{P}_u . Again by induction on $u = 1, \dots, T - 1$ we have $j_u < \tau_u(d_{l+1})$. Therefore, \mathcal{P}_u satisfies condition (15) for $u \in \{1, \dots, T - 1\}$. \square

Claim 7 *For each $u \in \{1, \dots, T\}$, if \mathcal{P}_u is d_{l+1} -preempted, then $\xi_{k_u}^u(a_u) < e_{k_u+1}^u - e_{k_u}^u$.*

Proof Suppose that $\xi_{k_u}^u(a_u) = e_{k_u+1}^u - e_{k_u}^u$ and \mathcal{P}_u is d_{l+1} -preempted. By (I4), it holds $k_u = \tau_u(d_{l+1})$ and by Lemma 3, $\xi_{k_u}^u(d_{l+1}) = e_{k_u+1}^u - e_{k_u}^u$. Then, there is a job a' that completes at e_{k_u} because e_{k_u} is an event in \mathcal{P}_u . The job a' starts after $C(\mathcal{P}_u, d_l)$ for otherwise a' interlaces with either a_u or d_{l+1} in an optimal \mathcal{P}_u —contradiction with Lemma 5. Perform swapping of a' and d_{l+1} . By Lemma 2, this leads to a feasible schedule \mathcal{P}' . If $\xi_{k_u-1}(d_{l+1}) < e_{k_u} - e_{k_u-1}$, then the total completion time of the new schedule is smaller than that of \mathcal{P}_u —a contradiction again. If $\xi_{k_u-1}(d_{l+1}) = e_{k_u} - e_{k_u-1}$, then there is a job a'' that completes at e_{k_u-1} in \mathcal{P}_u because e_{k_u-1} is an event in \mathcal{P}_u . Again, the job a'' starts after $C(\mathcal{P}_u, d_l)$ for otherwise a'' interlaces with either a_u or d_{l+1} in an optimal \mathcal{P}_u —contradiction with Lemma 5. Clearly, $C(\mathcal{P}_u, a'') = C(\mathcal{P}', a'')$ and $s(\mathcal{P}_u, a'') = s(\mathcal{P}', a'')$. Take $\mathcal{P}_u := \mathcal{P}'$ and $a' := a''$ and repeat the above swapping.

After a finite number of the above ‘swappings’ we obtain a feasible schedule with lower total completion time than the initial one—a contradiction. This completes the proof of Claim 7. \square

Now we return to the proof of Theorem 3. If $T = 1$, then the lemma holds by Claim 5. To complete the proof we show that $T > 1$ leads to a contradiction. First, for $T > 1$, Claims 2, 3, and 5 imply that the schedule \mathcal{P}_1 is d_{l+1} -preempted. Thus, Claims 2, 3 and 6 and an induction on $u \in \{1, \dots, T - 1\}$ give that each of the schedules $\mathcal{P}_1, \dots, \mathcal{P}_{T-1}$ is d_{l+1} -preempted. Since the iterative process exits in iteration $u = T$, producing either \mathcal{P}'_T or \mathcal{P}''_T , we get a contradiction since either these two exit schedules have smaller total completion times than \mathcal{P}_{T-1} or they have the same total completion times but their abnormality points are greater than i . Hence, \mathcal{P}_{T-1} is not maximal, contradicting the fact that it is d_{l+1} -preempted.

It remains to show that T exists, i.e., that the number of iterations is finite. To that end let c_u and C_u be the numbers of jobs executed in \mathcal{P}_u which are not covered by d_{l+1} , and the number of blocks, respectively, in $[C(\mathcal{P}_u, d_l), C(\mathcal{P}_u, d_{l+1})]$.

By Claim 7, $\varepsilon > 0$ in Step 1 of the construction of \mathcal{P}'_u for each $u \in \{2, \dots, T\}$. If $\varepsilon = \xi_{\tau_{u-1}(d_l)+1}^{u-1}(a_{u-1})$ in Step 1, then the block $\tau_{u-1}(d_l) + 1$ disappears. Otherwise, $\varepsilon = e_{k_{u-1}+1}^{u-1} - e_{k_{u-1}}^{u-1} - \xi_{k_{u-1}}^{u-1}(a_{u-1})$ in Step 1. If $\xi_{k_{u-1}}^{u-1}(a_{u-1}) > 0$, then $a_u = a_{u-1}$ at the end of Step 1 and $\xi_{k_u}^u(a_u) = e_{k_u+1}^u - e_{k_u}^u$. Therefore, both Steps 2 and 3 are vacuous and thus $\mathcal{P}'_u = \mathcal{P}_u$ is d_{l+1} -preempted. However, $\xi_{k_u}^u(a_u) = e_{k_u+1}^u - e_{k_u}^u$ contradicts Claim 7. Thus, we have $\xi_{k_{u-1}}^{u-1}(a_{u-1}) = 0$ in Step 1, and therefore the block k_{u-1} disappears.

Next, Step 2 does not change the number of blocks. Finally, Step 3, if not vacuous, may increase the number of blocks by at most one. Thus, we have $C_u = C_{u-1} - 1$ and $c_u = c_{u-1}$, if Step 3 is vacuous, and $C_u \leq C_{u-1}$ and $c_u = c_{u-1} - 1$, if Step 3 is not vacuous. Consequently $C_u + c_u < C_{u-1} + c_{u-1}$ and $T \leq C_1 + c_1 \leq 3n$. Thus, the iterative process indeed stops with some schedule \mathcal{P}'_T or \mathcal{P}''_T . This, by Claim 5, completes the proof of the theorem.

7 Summary

In this paper we have provided some structural characterization of the preemptions in optimal schedules for the problem $P2|pmtn, in-tree, r_j, p_j | \sum C_j$. The advantages of our characterization are as follows:

- It narrows down a search space of optimal solutions from an infinite one to a finite one.
- The understanding of the possible structure of preemptions is a step towards determining the complexity of the problem $P2|pmtn, in-tree, r_j, p_j | \sum C_j$. On the one hand, the normality of an optimal schedule may lead us to a polynomial-time algorithm. On the other hand, the fact that a single job may need many (of the order of $\log n$) preemptions as stated in Theorem 1 could be useful in proving NP-completeness, although the complexity of the problem $P2|pmtn, in-tree, r_j, p_j | \sum C_j$ is left as an interesting and challenging open problem.
- It significantly improves the lower bound on the resolution for the problem $P2|pmtn, in-tree, r_j, p_j | \sum C_j$.

Note that we rely on the in-trees precedence constraints between the jobs in our proof. This assumption is crucial when proving in Sect. 5 that a maximal A-free schedule exists. The generalization of our result to arbitrary precedence constraints or providing an example that an analogous statement as the one in Corollary 1 is false for more general precedence constraints is left as an open problem.

Acknowledgements This research has been supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Grant OPG0105675. Dariusz Dereniowski was partially supported by a schol-

arship for outstanding young researchers founded by the Polish Ministry of Science and Higher Education, and Polish National Science Center under Contract DEC-2011/02/A/ST6/00201.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Baptiste, P., Brucker, P., Knust, S., & Timkovsky, V. (2004). Ten notes on equal-processing-time scheduling. *4OR*, 2(2), 111–127.
- Baptiste, P., Carlier, J., Kononov, A., Queyranne, M., Sevastyanov, S., & Sviridenko, M. (2011). Properties of optimal schedules in preemptive shop scheduling. *Discrete Applied Mathematics*, 159(5), 272–280.
- Baptiste, P., & Timkovsky, V. (2001). On preemption redundancy in scheduling unit processing time jobs on two parallel machines. *Operations Research Letters*, 28(5), 205–212.
- Baptiste, P., & Timkovsky, V. (2004). Shortest path to nonpreemptive schedules of unit-time jobs on two identical parallel machines with minimum total completion time. *Mathematical Methods of Operations Research*, 60(1), 145–153.
- Brucker, P., Hurink, J., & Knust, S. (2003). A polynomial algorithm for $P | p_j = 1, r_j, \text{outtree} | \sum c_j$. *Mathematical Methods of Operations Research*, 56(3), 407–412.
- Carlier, A., Hanen, C., & Munier-Kordon, A. (2014). Equivalence of two classical list scheduling algorithms for dependent typed tasks with release dates, due dates and precedence delays. New Challenges in Scheduling Theory, March 31–April 4, Aussois, France.
- Coffman, E. Jr, Dereniowski, D., & Kubiak, W. (2012). An efficient algorithm for an ideal scheduling problem. *Acta Informatica*, 6, 1–14.
- Coffman, E. Jr, & Graham, R. (1972). Optimal scheduling for two-processor systems. *Acta Informatica*, 1, 200–213.
- Coffman, E. Jr, Ng, C., & Timkovsky, V. (2015). How small are shifts required in optimal preemptive schedules? *Journal of Scheduling*, 18(2), 155–163.
- Coffman, E. Jr, Sethuraman, J., & Timkovsky, V. (2003). Ideal preemptive schedules on two processors. *Acta Informatica*, 39(8), 597–612.
- Fujii, M., Kasami, T., & Ninomiya, K. (1971). Optimal sequencing of two equivalent processors. *SIAM Journal on Applied Mathematics*, 17, 784–789. Erratum 20 (1971) 141.
- Gabow, H. (1982). An almost-linear algorithm for two-processor scheduling. *Journal of ACM*, 29, 766–780.
- Garey, M., & Johnson, D. (1976). Scheduling tasks with nonuniform deadlines on two processors. *Journal of ACM*, 23(3), 461–467.
- Garey, M., & Johnson, D. (1977). Two-processor scheduling with start-times and deadlines. *SIAM Journal on Computing*, 6(3), 416–426.
- Herrbach, L. A., & Leung, J. Y. T. (1990). Preemptive scheduling of equal length jobs on two machines to minimize mean flow time. *Operations Research*, 38(3), 487–494.
- Huo, Y., & Leung, J. Y. (2005). Minimizing total completion time for UET tasks with release time and outtree precedence constraints. *Mathematical Methods of Operations Research*, 62(2), 275–279.
- Leung, A., Palem, K., & Pnueli, A. (2001). Scheduling time-constrained instructions on pipelined processors. *ACM Transactions on Programming Languages and Systems*, 23(1), 73–103.
- Sauer, N. W., & Stone, M. G. (1987). Rational preemptive scheduling. *Order*, 4, 195–206.