Proportionate Progress: A Notion of Fairness in Resource Allocation*

S. K. Baruah *N. K. Cohen* *C. G. Plaxton* *D. A. Varvel*

Department of Computer Science

University of Texas at Austin

Abstract

We consider the problem of sharing multiple copies of some resource among various tasks. We introduce a new notion of temporal fairness based on proportionate progress, called P-fairness. We use this notion to design an algorithm that solves the multiple-resource periodic scheduling problem. We define a related problem and show how to solve it efficiently, which in turn leads to an efficient implementation of our scheduling algorithm. To our knowledge, this work represents the first non-trivial result concerning the multiple-resource scheduling of tasks that may make recurring requests.

1 Introduction

Scheduling is the act of assigning resources to activities or tasks. Scheduling problems typically involve a set of constraints (e.g., deadlines) that must be met by any acceptable schedule. Often these constraints are designed to enforce some notion of fairness; for example, a very weak fairness constraint might be that any task will eventually get to use the resource it has requested. For any particular set of constraints, there are two obvious problems to be addressed: (i) the "decision" problem (i.e., determining whether or not a given instance is schedulable), and (ii) the "scheduling" problem (i.e., actually constructing the schedule for a given schedulable instance). Many sets of constraints result in an intractable decision problem [4]. The concept of fairness has not generally been applied to such problems, although we will point out how certain kinds of deadlines amount to very strong fairness constraints.

We consider here a resource sharing problem defined as follows. There are m identical copies of some resource. Examples of such resources might be CPU cycles or university classrooms. There are n tasks numbered from 0 to $n-1$. A task x has a rational weight $x.w = x.e/x.p$ for positive integers $x.e$ and $x.p$, meaning that it needs a resource a fraction $x.w$ of the time. The sum of these weights must not exceed the number of resources, m. We wish to allocate resources as "fairly" as possible by keeping each task's resource usage close to $x.w$ times the elapsed time. One way to do this is by allocating a fraction $x.w$ of a resource to each task during each time unit, or *slot*. Each task's resource usage is thereby kept at exactly $x.w$ times the elapsed time. This is sometimes called the "time slicing" algorithm. It is easy to see that this definition of fairness, where each task gets its desired fraction of the resource in each slot, is incompatible with scheduling at integral boundaries. Such integral scheduling is necessary for many applications. Our first task is therefore to define a notion of fairness that is consistent with integral boundaries.

Given a set of n tasks, where each task x has rational weight $x.w = x.e/x.p$, $0 < x.w < 1$, a *periodic-fair* schedule is one that allocates a resource to a task x for exactly $x.e$ slots in each interval $[k \cdot x.p, (k+1) \cdot x.p)$ for all $k \in \mathbf{N}$. Under periodic fairness, the problem of fair resource sharing reduces to the periodic scheduling problem that was first discussed by Liu in 1969 [10].

There are several optimal single-resource scheduling algorithms for the periodic scheduling problem which respect integral boundaries. The Earliest Deadline algorithm is one example [11]. None of them extend directly to multiple resources. As Liu pointed out, "the simple fact that a task can use only one [resource] even when several [resources] are free at the same time adds a surprising amount of difficulty" to the scheduling of multiple resources [10].

Leung's application of the Least Slack algorithm to this problem represents a recent attempt at improvement on Earliest Deadline [9]. He was able to show that with multiple resources Least Slack dominates Earliest Deadline. However, the algorithm proposed does not respect integral boundaries and is not optimal.

The time-slicing algorithm can be used to show that

*This research was supported by NSF Research Initiation Award CCR–9111591, and Texas Advanced Research Program (TARP) Award #003658480.

non-integral schedules exist for the multiple-resource periodic scheduling problem. Baruah, Howell, and Rosier [1] used the network flow reduction of Horn [6] to show that the multiple-resource periodic scheduling problem does in fact have integral solutions. Thus, the decision problem for such a periodic task system reduces to checking that $\sum_{x=0}^{n-1} x.w \leq m$. Given any schedulable instance, it would be desirable to have an efficient *on-line scheduling algorithm*. Such an algorithm is executed at the beginning of each slot in order to determine which subset of the tasks should be scheduled (i.e., assigned a resource for the next slot). It is natural to define the complexity of a given on-line scheduler as the maximum number of steps it requires to schedule any single slot. Prior to this paper, no polynomial-time on-line scheduling algorithm was known for the multiple-resource periodic scheduling problem.

We solve the multiple-resource periodic scheduling problem by imposing an even stronger fairness constraint. Our approach is based on maintaining proportionate progress: each task is scheduled resources in proportion to its weight. We call this *proportionate fairness* or *P-fairness*. P-fairness is a strictly stronger form of fairness than periodic fairness, in that any P-fair schedule is periodic-fair while the converse is not generally true. P-fairness is a natural and desirable notion in certain practical applications. To the best of our knowledge, none of the scheduling algorithms currently known generate P-fair schedules even in the case of a single resource. We prove that any resource sharing problem instance for which $\sum_{x=0}^{n-1} x.w \leq m$ has a P-fair schedule. This proof makes use of certain results from network flow theory. We then describe and prove the correctness of a polynomial-time on-line scheduling algorithm that generates a P-fair schedule for any schedulable instance. Since every P-fair schedule is also periodic-fair, this algorithm solves the periodic scheduling problem, thus closing a problem that has remained open for over twenty years.

We consider the research described here to be significant for several reasons. First, we introduce a new and potentially important notion of fairness in resource sharing, prove that this notion of fairness is actually achievable, and demonstrate its applicability to certain real-life situations. Second, as a corollary to our main results, we solve the multiple-resource periodic scheduling problem. To our knowledge this work represents the first non-trivial complete result in on-line multiple-resource scheduling. We hope the lessons learned in solving this problem, and the techniques applied here, will generalize to other multiple-resource scheduling problems.

The remainder of this paper is organized as follows. In Section 2 we define P-fairness and some related concepts. We also give examples of real-life problems that might be approached using P-fairness. Section 3 established that P-fair schedules exist for the resource scheduling problem. The proof is non-constructive in the sense that the corresponding scheduling algorithm has exponential time complexity. In Section 4 we prove the correctness of a simple on-line algorithm for producing such P-fair schedules. The naive implementation of that algorithm schedules each slot in pseudo-polynomial time. A polynomial-time implementation is given in Section 5. Section 6 offers some concluding remarks.

2 P-Fairness

In this section we define P-fairness and some related concepts. We start with some conventions:

- Scheduling decisions occur at integral values of *time*, numbered from 0. The unit interval between time t and time $t+1$ will be referred to as slot t, $t \in \mathbf{N}$.
- For integers a and b, let $[a, b) = \{a, \ldots, b-1\}$. Furthermore, let $[a, b] = [a, b+1)$, $(a, b] = [a+1, b+1)$, $(a, b) = [a+1, b)$, and $[a] = [0, a)$.
- Our generic instance Φ of the fair resource sharing problem will involve m resources and n tasks. We assume that the tasks are numbered from 0 to $n-1$. Specific tasks will be denoted by identifiers x and y, which range over $[n]$.
- Each task x has an integer *period* $x.p$, $x.p > 1$, an integer *execution requirement* $x.e$, $x.e \in (0, x.p)$, and a rational *weight* $x.w = x.e/x.p$. Note that $0 < x.w < 1$. Without loss of generality we assume that $\sum_{x \in [n]} x.w = m$.
- Let σ_i denote the ith symbol of string σ.

Now some definitions:

- A *schedule* S for instance Φ is a function from $[n] \times \mathbf{N}$ to $\{0, 1\}$, where $S(x, t) = 1$ if and only if task x is scheduled in slot t. We require that $\sum_{x \in [n]} S(x, t) \leq m$, $t \in \mathbf{N}$.
- A schedule S is *periodic* if and only if

$$\forall i, x : i \in \mathbf{N}, x \in [n] : \textstyle\sum_{t \in [i \cdot x.p]} S(x, t) = i \cdot x.e.$$

- The *lag* of a task x at time t with respect to schedule S, denoted $\mathsf{lag}(S, x, t)$, is defined as:

$$\mathsf{lag}(S, x, t) = x.w \cdot t - \textstyle\sum_{i \in [t]} S(x, i).$$

- A schedule S is *P-fair* if and only if

$$\forall x, t : x \in [n], t \in \mathbf{N} : -1 < \mathsf{lag}(S, x, t) < 1.$$

- A schedule S is *P-fair at time t* if and only if there exists a P-fair schedule S' such that

$$\forall x : x \in [n] : \mathsf{lag}(S, x, t) = \mathsf{lag}(S', x, t).$$

Informally speaking, $\mathsf{lag}(S, x, t)$ measures the difference between the number of resource allocations that task x "should" have received in the set of slots $[t]$ and the number that it actually received.

Periodic schedules can also be defined in terms of lag constraints. In particular, note that a schedule S is periodic if and only if

$$\forall i, x : i \in \mathbf{N}, x \in [n] : \mathsf{lag}(S, x, i \cdot x.p) = 0,$$

from which it follows that *every P-fair schedule is periodic*. (Note that in the definition of lag, the term $x.w \cdot t$ is independent of S, and the term $\sum_{i \in [t]} S(x, i)$ is an integer.)

P-fairness is a very strict requirement. It demands that the absolute value of the difference between the expected allocation and the actual allocation to every task always be strictly less than 1. In other words, a task never gets an entire slot ahead or behind. In general it is not possible to guarantee a smaller variation in lag. Consider n identical tasks sharing a single resource, where the weight of each task is $1/n$. For n sufficiently large, we can make the lag of the first (resp., last) task scheduled come arbitrarily close to -1 (resp., 1).

P-fairness is the natural notion of fairness for many resource allocation problems. Here are two examples:

Example 1 An airline has m airplanes and $n > m$ flight crews, all of which are based in the same city. Assume that exactly m flight crews are scheduled to work on any given day. Due to seniority, job performance, or other factors, it may be desirable to schedule some flight crews more often than others. For each flight crew x, set $x.w$ to the desired fraction of all days that x should work, while ensuring that $\sum_{x \in [n]} x.w = m$. A P-fair scheduler will produce a schedule in which every flight crew works at a steady rate: after t days, flight crew x will have worked either $\lfloor t \cdot x.w \rfloor$ or $\lceil t \cdot x.w \rceil$ days.

Example 2 Consider a node in a real-time communications network that needs to service a number of incoming and outgoing edges. The weight $x.w$ on an edge x corresponds to the relative amount of traffic expected on that edge. A P-fairness requirement may be necessary to maintain the real-time nature of the communications, and to prevent exceptionally long queueing delays from building up along certain edges.

3 Existence of a P-Fair Schedule

In Sections 4 and 5 we will develop a polynomial-time P-fair scheduling algorithm. The proof of correctness of that algorithm relies on the existence of a P-fair schedule for the resource sharing problem. In this section we use a network flow argument to prove the existence of such a schedule. In principle, the network reduction could itself serve as the basis for a P-fair scheduling algorithm. Unfortunately, the size of the network generated by our reduction is exponential in the size of the given scheduling instance, and so the network reduction argument does not by itself provide a polynomial-time algorithm.

With respect to instance Φ of the resource scheduling problem, let $\mathsf{earliest}(x, j)$ (resp., $\mathsf{latest}(x, j)$) denote the earliest (resp., latest) slot during which task x may be scheduled for the jth time, $j \in \mathbf{N}$, in any P-fair schedule. We can easily derive closed form expressions for $\mathsf{earliest}(x, j)$ and $\mathsf{latest}(x, j)$. Note that $\mathsf{earliest}(x, j) = \min t : t \in \mathbf{N} : x.w \cdot (t+1) - (j+1) > -1$ and $\mathsf{latest}(x, j) = \max t : t \in \mathbf{N} : x.w \cdot t - j < 1$. Hence,

$$\mathsf{earliest}(x, j) = \lfloor j / x.w \rfloor, \text{ and}$$
$$\mathsf{latest}(x, j) = \lceil (j+1) / x.w \rceil - 1.$$

Note that $\mathsf{earliest}(x, j) < \mathsf{latest}(x, j)$, $x \in [n]$, $j \in \mathbf{N}$. Furthermore, $\mathsf{earliest}(x, j+1) - \mathsf{latest}(x, j)$ is either 0 or 1. In words, there is at most one slot where either the jth or the $j+1$st scheduling of task x may occur.

The remainder of this section is devoted to proving the existence of a P-fair schedule for any instance of the resource sharing problem Φ. Our proof strategy is as follows: First, we will describe a reduction from instance Φ to a directed and weighted graph G with a designated source and sink, such that certain flows in G correspond exactly (in a manner that will be made precise) to a P-fair schedule for Φ. Then we will prove the existence of such a flow in G.

Throughout this section, let L denote the least common multiple of the task periods: $L = \mathrm{lcm}_{x \in [n]} x.p$.

Lemma 3.1 Instance Φ has a P-fair schedule if and only if there exists a schedule S such that

$$\forall x, t : x \in [n], t \in (0, L] : -1 < \mathsf{lag}(S, x, t) < 1.$$

Proof: An infinite P-fair schedule S' may be obtained from S by scheduling in slot t those tasks scheduled by S in slot $t \bmod L$. □

Theorem 1 Instance Φ has a P-fair schedule.

Before proving this theorem, we will present some definitions and an important lemma.

Recall that $x.w = x.e/x.p$. We describe below the construction of a weighted, directed graph G. The vertex set V of G is the union of 6 disjoint sets of vertices $V_0, \ldots, V_5$, and the edge set E of G is the union of 5 disjoint sets of edges $E_0, \ldots, E_4$, where E_i is a subset of $(V_i \times V_{i+1} \times \mathbf{N})$, $0 \leq i \leq 4$. That is, G is a "6-layered" graph, with all edges connecting vertices in adjacent layers. The sets of vertices are as follows:

$$\begin{aligned} V_0 &= \{\mathsf{source}\}, \\ V_1 &= \{\langle 1, x\rangle \mid x \in [n]\}, \\ V_2 &= \{\langle 2, x, j\rangle \mid x \in [n], j \in [x.w \cdot L]\}, \\ V_3 &= \{\langle 3, x, t\rangle \mid x \in [n], t \in [L]\}, \\ V_4 &= \{\langle 4, t\rangle \mid t \in [L]\}, \text{ and} \\ V_5 &= \{\mathsf{sink}\}. \end{aligned}$$

An edge is represented by a 3-tuple. For $u, v \in V$ and $w \in \mathbf{N}$, the 3-tuple $(u, v, w) \in E$ represents an edge from u to v of weight w. The sets of edges in G are as follows:

$$\begin{aligned} E_0 &= \{(\mathsf{source}, \langle 1, x\rangle, x.w \cdot L) \mid x \in [n]\}, \\ E_1 &= \{(\langle 1, x\rangle, \langle 2, x, j\rangle, 1) \mid x \in [n], j \in [x.w \cdot L]\}, \\ E_2 &= \{(\langle 2, x, j\rangle, \langle 3, x, t\rangle, 1) \mid x \in [n], \\ &\qquad j \in [x.w \cdot L], t \in [\mathsf{earliest}(x, j), \mathsf{latest}(x, j)]\}, \\ E_3 &= \{(\langle 3, x, t\rangle, \langle 4, t\rangle, 1) \mid x \in [n], t \in [L]\}, \text{ and} \\ E_4 &= \{(\langle 4, t\rangle, \mathsf{sink}, m) \mid t \in [L]\}. \end{aligned}$$

Lemma 3.2 If there is an integral flow of size $m \cdot L$ in G, then there exists a P-fair schedule for Φ.

Proof: By Lemma 3.1, it suffices to prove that the existence of an integral flow of size $m \cdot L$ in G implies the existence of a schedule S for Φ such that

$$\forall x, t : x \in [n], t \in [0, L] : -1 < \mathsf{lag}(S, x, t) < 1.$$

Suppose there is an integral flow of size $m \cdot L$ in G. The total capacity of E_0, the set of edges leading out of the source vertex, is equal to $\sum_{x \in [n]} x.w \cdot L = m \cdot L$. Hence, each edge in E_0 carries a unit flow, and each vertex $\langle 1, x\rangle$ receives exactly $x.w \cdot L$ units of flow. Since there are $x.w \cdot L$ vertices in V_2 each connected (by an edge of unit capacity) to vertex $\langle 1, x\rangle$, and no two vertices in V_1 are connected to the same vertex in V_2, it follows that each vertex in V_2 receives a unit flow. Accordingly, each vertex in V_2 must send a unit flow to some vertex in V_3.

We will construct the desired schedule S from the given flow according to the following rule: Allocate a resource to task x in slot t if and only if there is a unit flow from vertex $\langle 2, x, j\rangle$ to vertex $\langle 3, x, t\rangle$.

Because the total flow into the sink vertex is $m \cdot L$, each of the L edges of capacity m in E_4 must carry m units of flow. Hence, for all $t \in [L]$, vertex $\langle 4, t\rangle$ receives exactly m unit flows from vertices in V_3. Each vertex $\langle 3, x, t\rangle$ in V_3 is connected (by an edge of unit capacity) to vertex $\langle 4, t\rangle$, and is not connected to any other vertex in V_4. Thus, S schedules exactly m tasks in each time slot t, for all $t \in [L]$. To see that no lag constraints are violated by S, observe that for each task x and for all $j \in [x.w \cdot L]$, the jth scheduling of task x occurs at a slot in the interval $[\mathsf{earliest}(x, j), \mathsf{latest}(x, j)]$. (The jth scheduling corresponds to the unique unit flow out of vertex $\langle 2, x, j\rangle$.) □

We are now ready to proceed with the proof of Theorem 1.

Proof of Theorem 1: Since all edges of the graph have integral capacity, if there is a fractional flow of size $m \cdot L$ in the graph then there is an integral flow of that size [3]. It remains to be shown that such a fractional flow exists. We use the following flow assignments:

- Each edge $(\mathsf{source}, \langle 1, x\rangle, x.w \cdot L) \in E_0$ carries a flow of $x.w \cdot L$.
- Each edge $(\langle 1, x\rangle, \langle 2, x, j\rangle, 1) \in E_1$ carries a unit flow.
- Each edge $(\langle 3, x, t\rangle, \langle 4, t\rangle, 1) \in E_3$ carries a flow of size $x.w$.
- Each edge $(\langle 4, t\rangle, \mathsf{sink}, m) \in E_4$ carries a flow of size m.
- The flows through edges in E_2 are as follows:
  - Each edge $(\langle 2, x, j\rangle, \langle 3, x, \mathsf{earliest}(x, j)\rangle, 1)$ carries a flow of size $x.w - \left(j - x.w \cdot \left\lfloor \frac{j}{x.w} \right\rfloor\right)$.
  - Each edge $(\langle 2, x, j\rangle, \langle 3, x, \mathsf{latest}(x, j)\rangle, 1)$ such that $\mathsf{latest}(x, j) = \mathsf{earliest}(x, j + 1)$ carries a flow of size $(j + 1) - x.w \cdot \left\lfloor \frac{j+1}{x.w} \right\rfloor$.
  - Every other edge $(\langle 2, x, j\rangle, \langle 3, x, t\rangle, 1) \in E_2$ carries a flow of size $x.w$.

The flow out of the source vertex is $\sum_{x \in [n]} (x.w \cdot L) = m \cdot L$. The flow into the sink vertex is also $m \cdot L$. It remains to prove that flow is conserved at every interior vertex. The flow into each vertex in V_1 is $x.w \cdot L$, and there are $x.w \cdot L$ edges leaving, each carrying a unit flow. The flow into each vertex in V_2 is 1. Below we will prove that the flow out of each vertex in V_2 is 1, and that the flow into each vertex in V_3 is $x.w$. Each vertex in V_3 has only one outgoing edge carrying a flow of $x.w$. Each vertex in V_4 has n incoming edges each carrying a flow of size $x.w$; since $\sum_{x \in [n]} x.w = m$, the flow in is m, which equals the flow out on the one outgoing edge.

It remains to proved that: (i) the flow out of each vertex in V_2 is 1, and (ii) the flow into each vertex in V_3 is $x.w$.

Consider an arbitrary vertex $\langle 2, x, j\rangle$ in V_2. There are $\mathsf{latest}(x,j) - \mathsf{earliest}(x,j) + 1$, or $\lceil (j+1)/x.w\rceil - \lfloor j/x.w\rfloor$, outgoing edges from $\langle 2, x, j\rangle$. If $\mathsf{earliest}(x, j+1) = \mathsf{latest}(x,j)$ (equivalently, $\lceil (j+1)/x.w\rceil - 1 = \lfloor (j+1)/x.w\rfloor$), then the flow out of $\langle 2, x, j\rangle$ is $x.w - (j - x.w \cdot \lfloor j/x.w\rfloor) + x.w \cdot (\lceil (j+1)/x.w\rceil - \lfloor j/x.w\rfloor - 2) + (j+1) - x.w\lfloor (j+1)/x.w\rfloor$, which simplifies easily to 1. Otherwise, $\mathsf{earliest}(x, j+1) = \mathsf{latest}(x,j) + 1$ (equivalently, $\lceil (j+1)/x.w\rceil = \lfloor (j+1)/x.w\rfloor = (j+1)/x.w$), and the flow out of $\langle 2, x, j\rangle$ is $x.w - (j - x.w \cdot \lfloor j/x.w\rfloor) + x.w \cdot (\lceil (j+1)/x.w\rceil - \lfloor j/x.w\rfloor - 1)$, which also simplifies to 1.

Consider an arbitrary vertex $\langle 3, x, t\rangle$ in V_3. If $t = \mathsf{latest}(x,j) = \mathsf{earliest}(x, j+1)$ for some $j \in \mathbf{N}$, then there are two incoming edges to $\langle 3, x, t\rangle$, namely $(\langle 2, x, j\rangle, \langle 3, x, t\rangle, 1)$ and $(\langle 2, x, j{+}1\rangle, \langle 3, x, t\rangle, 1)$. These edges carry flows of $(j+1) - x.w \cdot \lfloor (j+1)/x.w\rfloor$ and $x.w - ((j+1) - x.w \cdot \lfloor (j+1)/x.w\rfloor)$, respectively, for a total incoming flow of $x.w$. Otherwise, there is only one incoming edge to $\langle 3, x, t\rangle$, and it carries a flow of $x.w$. □

4 A P-Fair Scheduling Algorithm

We will now define a scheduling algorithm, and prove that it produces a P-fair schedule. Thus, we show that our algorithm is a solution for the multiple-resource periodic scheduling problem. It satisfies the integral boundary constraint and is suitable for use as an on-line scheduler.

First, some definitions:

- The *characteristic string* of task x, denoted $\alpha(x)$, is an infinite string over $\{-, 0, +\}$ with

  $$\alpha_t(x) = \text{sign}(x.w \cdot (t+1) - \lfloor x.w \cdot t\rfloor - 1),$$

  $t \in \mathbf{N}$.
- The *characteristic substring* of task x at time t is the finite string

  $$\alpha(x,t) \stackrel{\text{def}}{=} \alpha_{t+1}(x)\alpha_{t+2}(x)\cdots\alpha_{t'}(x),$$

  where $t' = \min i : i > t : \alpha_i(x) = 0$.
- With respect to P-fair schedule S at time t, we say that: task x is *ahead* if and only if $\mathsf{lag}(S,x,t) < 0$; task x is *behind* if and only if $\mathsf{lag}(S,x,t) > 0$; task x is *punctual* if and only if it is neither ahead nor behind.
- With respect to P-fair schedule S at time t, we say that: task x is *tnegru* if and only x is ahead and $\alpha_t(x) \neq +$; task x is *urgent* if and only if x is behind and $\alpha_t(x) \neq -$; task x is *contending* if and only if it is neither tnegru nor urgent.

Lemmas 4.1 to 4.5 provide the logical machinery that we will need in order to reason about the terms introduced above.

Lemma 4.1 If task x is ahead at time t under P-fair schedule S, then:

(a) If $\alpha_t(x) = -$, then $S(x,t) = 0$ and task x is ahead at time $t+1$.

(b) If $\alpha_t(x) = 0$, then $S(x,t) = 0$ and task x is punctual at time $t+1$.

(c) If $\alpha_t(x) = +$ and $S(x,t) = 1$, then task x is ahead at time $t+1$.

(d) If $\alpha_t(x) = +$ and $S(x,t) = 0$, then task x is behind at time $t+1$.

Proof: We will prove only Part (a), since the remaining parts are similar. Assuming that task x is ahead at time t under P-fair schedule S, we have $\sum_{i\in[t]} S(x,i) = \lceil x.w \cdot t\rceil$, where $x.w \cdot t \notin \mathbf{N}$. Under these conditions,

$$\begin{aligned}
& \alpha_t(x) = - \ \wedge \ \textstyle\sum_{i\in[t]} S(x,i) = \lceil x.w \cdot t\rceil \\
\Longrightarrow \ & x.w \cdot (t+1) - \lfloor x.w \cdot t\rfloor - 1 < 0 \\
& \wedge \ \textstyle\sum_{i\in[t]} S(x,i) = \lceil x.w \cdot t\rceil \\
\Longrightarrow \ & x.w \cdot (t+1) - \textstyle\sum_{i\in[t]} S(x,i) < 0 \\
\Longrightarrow \ & \mathsf{lag}(S,x,t+1) + S(x,t) < 0.
\end{aligned}$$

Note that $\mathsf{lag}(S,x,t+1) + S(x,t) < 0$ implies that $S(x,t) = 0$ and task x is ahead at time $t+1$, as required. □

The (straightforward) proofs of Lemmas 4.2 to 4.5 are similar to the proof of Lemma 4.1(a), and have been omitted.

Lemma 4.2 If task x is behind at time t under P-fair schedule S, then:

(a) If $\alpha_t(x) = -$ and $S(x,t) = 1$, then task x is ahead at time $t+1$.

(b) If $\alpha_t(x) = -$ and $S(x,t) = 0$, then task x is behind at time $t+1$.

(c) If $\alpha_t(x) = 0$, then $S(x,t) = 1$ and task x is punctual at time $t+1$.

(d) If $\alpha_t(x) = +$, then $S(x,t) = 1$ and task x is behind at time $t+1$.

Lemma 4.3 If task x is tnegru at time t under P-fair schedule S, then $S(x,t) = 0$.

Lemma 4.4 If task x is urgent at time t under P-fair schedule S, then $S(x,t) = 1$.

Lemma 4.5 If task x is contending at time t under P-fair schedule S, then:

(a) If $S(x,t) = 1$, then x is ahead at time $t+1$.

(b) If $S(x,t) = 0$, then x is behind at time $t+1$.

Given the preceding definitions and lemmas, it is now straightforward to present our on-line scheduling algorithm, which will be referred to as *Algorithm PF*. At any time t the task of Algorithm PF is to determine which m-subset of the n tasks to schedule. By Lemma 4.4, every urgent task must be scheduled in the current time slot in order to preserve P-fairness. Symmetrically, Lemma 4.3 implies that no tnegru task can be scheduled in the current time slot without violating P-fairness. Since our goal is to prove that Algorithm PF produces a P-fair schedule, it must be that Algorithm PF schedules all of the urgent tasks and none of the tnegru tasks. It remains to define the behavior of Algorithm PF on the set of contending tasks.

Before doing so, however, we should pause to address two possible pitfalls. Let n_0, n_1, and n_2 denote the number of tnegru, contending, and urgent tasks at time t, respectively. If $n_2 > m$, then it would be impossible for Algorithm PF to schedule all of the urgent tasks. Symmetrically, if $n_0 > n - m$, then Algorithm PF would be forced to schedule some tnegru task. (Because $\sum_{x \in [n]} x.w = m$, we cannot hope to schedule instance Φ correctly unless all m resources are allocated in every slot.) An immediate consequence of Theorem 2, stated below, is that neither of these pitfalls will ever arise under Algorithm PF. Thus, in defining the behavior of Algorithm PF on the set of contending tasks, we can assume that $n_0 \leq n - m$ and $n_2 \leq m$. The task of Algorithm PF is to determine which subset (of size $m - n_2 \leq n_1$) of the n_1 contending tasks to schedule.

We can define a total order $\succeq$ on the set of contending tasks as follows: $x \succeq y$ if and only if $\alpha(x,t) \geq \alpha(y,t)$, where the comparison between characteristic substrings $\alpha(x,t)$ and $\alpha(y,t)$ is resolved lexicographically with $+ > 0 > -$. Ties can be broken arbitrarily; for example, we could assume that ties are broken in favor of the higher-numbered task.

Algorithm PF schedules the $m - n_2$ highest-priority contending tasks according to this total order. Algorithm PF is summarized in its entirety below:

1. Schedule all urgent tasks.
2. Allocate the remaining resources to the highest-priority contending tasks according to the total order $\succeq$.

Throughout the remainder of this section, let S_{PF} denote the schedule produced by Algorithm PF on instance Φ. Lemma 4.6 below is proven in Appendix A.

Lemma 4.6 If schedule S_{PF} is P-fair at time t, then it is P-fair at time $t+1$, $t \in \mathbf{N}$.

Theorem 2 Schedule S_{PF} is P-fair.

Proof: By Theorem 1, schedule S_{PF} is P-fair at time 0. Hence, Lemma 4.6 implies that schedule S_{PF} is P-fair at time t, $t \in \mathbf{N}$. □

5 The Comparison Algorithm

We will now present two implementations of the characteristic substring comparison function required by Algorithm PF. The first, which we call NaiveCompare, we prove correct. The second, Compare, we prove equivalent to the first and show that it runs in polynomial time. Both subroutines use only integer variables, and the integer operations $\{-, +, \cdot, \bmod\}$. We will prove that the number of integer operations performed by Compare on tasks x and y is at most linear in the size of the binary representation of $\min\{x.p, y.p\}$. (Furthermore, all intermediate values can be represented in $\lceil \lg(\max\{x.p, y.p\}) \rceil$ bits.)

Subroutine Compare can be used as the basis for an implementation of Algorithm PF that requires at most linear time (in the size of instance Φ) to decide which m-subset of the n tasks to schedule in a given slot. (The idea is to compare all strings at once, rather than two at a time.)

5.1 A Naive Implementation

This subsection presents a naive implementation of the characteristic substring comparison algorithm. Given contending tasks x and y at time t, our goal is to determine whether: (i) $\alpha(x,t) < \alpha(y,t)$, (ii) $\alpha(x,t) > \alpha(y,t)$, or (iii) $\alpha(x,t) = \alpha(y,t)$. The naive approach is to compare the two substrings one symbol at a time. Note that for any P-fair schedule S and $i \in [|\alpha(x,t)|]$:

$$\begin{aligned}
& \alpha_i(x,t) \\
= \; & \alpha_{t+i+1}(x) \\
= \; & \mathrm{sign}(x.w \cdot (t+i+2) - \lfloor x.w \cdot (t+i+1) \rfloor - 1) \\
= \; & \mathrm{sign}(\mathsf{lag}(S,x,t) + x.w \cdot (i+2) \\
& \qquad - \lfloor \mathsf{lag}(S,x,t) + x.w \cdot (i+1) \rfloor - 1) \\
= \; & \mathrm{sign}(x.e - x.p \\
& \qquad + (x.p \cdot \mathsf{lag}(S,x,t) + x.e \cdot (i+1)) \bmod x.p).
\end{aligned}$$

If task x is contending at time t under P-fair schedule S, we have $(x.p \cdot \mathsf{lag}(S,x,t) + x.e) \in (0, x.p)$. Hence

$$\begin{aligned}
\alpha_i(x,t) \; = \; & \mathrm{sign}(x.p \cdot \mathsf{lag}(S,x,t) + 2 \cdot x.e \\
& \qquad - x.p + (x.e \cdot i) \bmod x.p).
\end{aligned} \tag{1}$$

Let

$$\begin{aligned} a_0 &\stackrel{\text{def}}{=} x.p - x.e, \\ b_0 &\stackrel{\text{def}}{=} x.e, \text{ and} \\ c_0 &\stackrel{\text{def}}{=} x.p \cdot \mathsf{lag}(S, x, t) + 2 \cdot x.e - x.p. \end{aligned}$$

Note that $a_0 \in (0, x.p)$, $b_0 \in (0, x.p)$, and $c_0 \in (-a_0, b_0)$. Define a_1, b_1, and c_1 similarly with respect to task y. Given a_0, b_0, and c_0, it is straightforward to compute $\alpha(x, t)$ one symbol at a time, using a constant number of integer operations per symbol. Of course, $\alpha(y, t)$ can be computed in a similar fashion. This is the approach taken in subroutine NaiveCompare below.

```
(1)  NaiveCompare(a_0, b_0, c_0, a_1, b_1, c_1)
(2)  int a_0, b_0, c_0, a_1, b_1, c_1;
(3)  {
(4)    do TRUE ⟶
(5)      if c_0 = 0 ∧ c_1 = 0 ⟶ return TIE fi;
(6)      if c_0 ≥ 0 ∧ c_1 ≤ 0 ⟶ return 0
(7)       [] c_0 ≤ 0 ∧ c_1 ≥ 0 ⟶ return 1
(8)      fi;
(9)      if c_0 < 0 ⟶ c_0 := c_0 + b_0
(10)      [] c_0 > 0 ⟶ c_0 := c_0 − a_0
(11)     fi;
(12)     if c_1 < 0 ⟶ c_1 := c_1 + b_1
(13)      [] c_1 > 0 ⟶ c_1 := c_1 − a_1
(14)     fi
(15)   od
(16) }
```

The return values of NaiveCompare are 0, 1, and TIE. The return value 0 indicates that the task corresponding to the triple (a_0, b_0, c_0) should be given priority over the one corresponding to the triple (a_1, b_1, c_1). Conversely, the return value 1 indicates that the triple (a_1, b_1, c_1) should have priority. The return value TIE indicates that either can be scheduled ahead of the other. As mentioned in Section 4, such a tie could be broken using the task numbers.

Definition 5.1 A triple (a, b, c) is *admissible* if and only if: (i) a and b are positive integers, and (ii) c is an integer in the interval $(-a, b)$ such that

$$c \equiv i \cdot a \pmod{a + b}$$

for some integer i. We will say that a 6-tuple $(a_0, b_0, c_0, a_1, b_1, c_1)$ is *admissible* if and only if (a_0, b_0, c_0) and (a_1, b_1, c_1) are admissible triples.

It is immediate from the foregoing discussion that every input 6-tuple passed to NaiveCompare by our scheduling algorithm is admissible. Condition (ii) implies that NaiveCompare will eventually terminate. Unfortunately, the running time of NaiveCompare is not very good; it is pseudo-polynomial in the input size. This deficiency will be addressed in the next section.

5.2 An Efficient Implementation

In this section, we present a polynomial-time subroutine Compare with the same input-output behavior as the NaiveCompare subroutine of Section 5.1. The algorithm is recursive. As argued in Section 5.1, we can assume that any 6-tuple of arguments passed to the NaiveCompare subroutine is admissible. Correspondingly, the arguments of any top-level call to Compare may be assumed to be admissible. Lemma 5.1 below proves that this assumption can be extended to any non-trivial depth of recursion.

Lemma 5.1 If algorithm Compare is called with an admissible 6-tuple, then every resulting recursive call will also involve an admissible 6-tuple.

Proof: Assume that subroutine Compare is called with admissible 6-tuple $(a_0, b_0, c_0, a_1, b_1, c_1)$. It is sufficient to prove that, for $0 \leq i \leq 1$: (i) if the recursive call in Line 5 of Compare is executed then $(b_i, a_i, -c_i)$ is an admissible triple, and (ii) if the recursive call in Line 18 is executed then (a_i', b_i', c_i'), defined as

$$(a_i - (b_i \bmod a_i),\ b_i \bmod a_i,\ c_i + (b_i \bmod a_i)),$$

is an admissible triple. These facts are straightforward to prove. (Note that b_i cannot be a multiple of a_i because then c_i would be 0, and Line 18 would not have been reached.) □

Theorem 3 Let d denote the minimum number of bits in the binary representations of a_0, b_0, a_1, and b_1. Then algorithm Compare performs $O(d)$ integer operations.

Proof: Since algorithm Compare does not contain any loops and uses only tail recursion, it is sufficient to prove that the maximum depth of recursion is $O(d)$. More precisely, we will prove by induction that the maximum depth of recursion is $2d-2$ if $\min\{a_0, a_1\} \leq \min\{b_0, b_1\}$, and $2d-1$ otherwise.

For the base case, note that if (a, b, c) is an admissible triple, then $a + b \geq 2$. Thus, using Lemma 5.1, we have $a_i + b_i \geq 2$, $0 \leq i \leq 1$. Hence $d = 0$ is impossible and if $d = 1$ we must have $a_0 = a_1 = b_0 = b_1 = 1$. This implies $c_0 = c_1 = 0$. Thus the depth of recursion is $2d - 2 = 0$, as claimed.

For the induction step, assume that $d \geq 2$ and that the claim holds for smaller values of d. We consider two cases:

1. If $\min\{a_0, a_1\} \leq \min\{b_0, b_1\}$ (and the depth of recursion is greater than 0), Line 18 of Compare must be executed. Let a_i', b_i', and c_i' be defined as in the proof of Lemma 5.1, and assume without loss of generality that $a_0 \leq a_1$. Thus, a_0 has exactly $d \geq 2$ bits in its binary representation. Since

$a'_0 + b'_0 = a_0$, one of a'_0, b'_0 must have strictly fewer than d bits in its binary representation. The claim then follows by the induction hypothesis.

2. If $\min\{a_0, a_1\} > \min\{b_0, b_1\}$ then the recursive call in Line 5 of Compare will be executed. That call will terminate within at most $2d-2$ additional levels of recursion by the argument of the preceding case. Thus, the maximum depth of recursion is at most $2d - 1$, as claimed.

□

```
(1)  Compare(a_0, b_0, c_0, a_1, b_1, c_1)
(2)  int a_0, b_0, c_0, a_1, b_1, c_1;
(3)  {
(4)     if min{a_0, a_1} > min{b_0, b_1} ⟶
(5)        return Compare(b_1, a_1, −c_1, b_0, a_0, −c_0)
(6)     fi;
(7)     if ⌈c_0/a_0⌉ > ⌈c_1/a_1⌉ ⟶ return 0
(8)      [] ⌈c_0/a_0⌉ < ⌈c_1/a_1⌉ ⟶ return 1
(9)     fi;
(10)    c_0, c_1 := c_0 − a_0 · ⌈c_0/a_0⌉, c_1 − a_1 · ⌈c_1/a_1⌉;
(11)    if c_0 = 0 ∧ c_1 = 0 ⟶ return TIE
(12)     [] c_0 ≠ 0 ∧ c_1 = 0 ⟶ return 0
(13)     [] c_0 = 0 ∧ c_1 ≠ 0 ⟶ return 1
(14)    fi;
(15)    if ⌊b_0/a_0⌋ > ⌊b_1/a_1⌋ ⟶ return 0
(16)     [] ⌊b_0/a_0⌋ < ⌊b_1/a_1⌋ ⟶ return 1
(17)    fi;
(18)    return Compare(a_0 − (b_0 mod a_0),
(19)                   b_0 mod a_0, c_0 + (b_0 mod a_0),
(20)                   a_1 − (b_1 mod a_1),
(21)                   b_1 mod a_1, c_1 + (b_1 mod a_1))
(22) }
```

It remains to argue that: (i) Compare never executes a division by 0, and (ii) Compare always returns the correct value. Claim (i) is easy to justify: all divisions are by a_0 or a_1, which are strictly positive. Claim (ii) is addressed by the following theorem, which is proven in Appendix A.

Theorem 4 On any admissible input 6-tuple, algorithms NaiveCompare and Compare return the same value.

6 Conclusions

We have defined a new notion of fairness, called P-fairness, which we believe to be quite useful in a variety of resource allocation problems. We have shown that P-fair schedules exist for the resource sharing problem, which is a slight generalization of the periodic scheduling problem. Furthermore, we have provided an efficient algorithm that produces a P-fair schedule on-line.

The swapping argument of Lemma 4.6 captures the essence of P-fairness by modeling exchanges that are permissible in P-fair schedules. An interesting problem for future research is to identify generalizations of the resource sharing problem that can be handled within the same framework.

The Compare subroutine appears to be closely related to Euclid's GCD algorithm, as well as to various algorithms that have been proposed for 2-ILP, that is, integer linear programming with two variables [5, 7, 12, 13]. (ILP is NP-complete in general, but can be solved in polynomial-time for any fixed number of variables [8].) Deng has extensively studied the relationship between GCD and 2-ILP [2].

7 Acknowledgments

We would like to thank Al Mok and C. L. Liu for their encouragement and interest in this problem.

References

[1] S. K. Baruah, R. R. Howell, and L. E. Rosier. Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor. *Real-Time Systems*, 2:301–324, 1990.

[2] X. Deng. *Mathematical Programming: Complexity and Applications*. PhD thesis, Department of Operations Research, Stanford University, September 1989.

[3] L. R. Ford, Jr. and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, NJ, 1962.

[4] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, NY, 1979.

[5] D. S. Hirschberg and C. K. Wong. A polynomial-time algorithm for the knapsack problem with two variables. *Journal of the ACM*, 23:147–154, 1976.

[6] W. A. Horn. Some simple scheduling algorithms. *Naval Research Logistics Quarterly*, 21:177–185, 1974.

[7] R. Kannan. A polynomial algorithm for the two-variable integer programming problem. *Journal of the ACM*, 27:118–122, 1980.

[8] H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.

[9] J. Y.-T. Leung. A new algorithm for scheduling periodic, real-time tasks. *Algorithmica*, 4:209–219, 1989.

[10] C. L. Liu. Scheduling algorithms for multiprocessors in a hard-real-time environment. JPL Space Programs Summary 37–60, vol. II, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, pages 28–37, November 1969.

[11] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM*, 20:46–61, 1973.

[12] H. E. Scarf. Production sets with indivisibilities, Part I: Generalities. *Econometrica*, 49:1–32, 1981.

[13] H. E. Scarf. Production sets with indivisibilities, Part II: The case of two activities. *Econometrica*, 49:395–423, 1981.

A Proofs

Proof of Lemma 4.6: Assume that schedule S_{PF} is P-fair at time t for some $t \in \mathbf{N}$. Hence, there exists a P-fair schedule S such that $\mathsf{lag}(S_{PF}, x, t) = \mathsf{lag}(S, x, t)$, $x \in [n]$. Let X (resp., Y) denote the m-subset of tasks scheduled by S (resp., S_{PF}) in slot t. If $X = Y$ then S_{PF} is P-fair at time $t+1$ because S is P-fair at time $t+1$. If $X \neq Y$, there exist tasks $x \in X$ and $y \in Y$ such that $x \in X \setminus Y$ and $y \in Y \setminus X$. In the argument that follows we will demonstrate the existence of a P-fair schedule S' such that: (i) $\mathsf{lag}(S_{PF}, x, t) = \mathsf{lag}(S', x, t)$, $x \in [n]$, and (ii) S' schedules the m-subset $X \setminus \{x\} \cup \{y\}$ in slot t. By repeating this argument $|X \setminus Y|$ times, we can obtain a sequence of P-fair schedules such that the last schedule in the sequence, S^*, satisfies $\mathsf{lag}(S_{PF}, x, t+1) = \mathsf{lag}(S^*, x, t+1)$. Hence, schedule S_{PF} is P-fair at time $t+1$, proving the lemma.

Accordingly, it is sufficent to prove existence of a P-fair schedule S' as defined above. We begin by claiming that $S(x,i) \neq S(y,i)$ for some $i > t$. (If not, it follows easily that $x.w = y.w$, $\mathsf{lag}(S, x, t) = \mathsf{lag}(S, y, t)+1$, and hence that Algorithm PF would have given priority to task x over task y at time t, a contradiction.) We transform schedule S into S' as follows. Let

$$t' \stackrel{\text{def}}{=} \min i : i > t : S(x,i) \neq S(y,i).$$

We prove below that in fact $S(x,t') = 0$ and $S(y,t') = 1$. Schedule S' is defined to be identical to S except that we "swap" the allocations to tasks x and y at slots t and t', setting $S'(x,t) = 0$, $S'(y,t) = 1$, $S'(x,t') = 1$, and $S'(y,t') = 0$.

In the arguments that follow, all statements "categorizing" tasks x and y (e.g., "task x is not urgent at time t") are being made with respect to the P-fair schedule S. (Note that at time t, it makes no difference whether our claims are made with respect to S_{PF} or S, since $\mathsf{lag}(S_{PF}, x, t) = \mathsf{lag}(S, x, t)$, $x \in [n]$.)

Consider the following predicates:

$$\begin{aligned} P_0(i) &\stackrel{\text{def}}{=} \text{task } x \text{ is ahead at time } i, \\ P_1(i) &\stackrel{\text{def}}{=} \text{task } y \text{ is behind at time } i, \\ P_2(i) &\stackrel{\text{def}}{=} \alpha_j(x) = \alpha_j(y) \neq 0,\ j \in (t,i), \text{ and} \\ P(i) &\stackrel{\text{def}}{=} P_0(i) \wedge P_1(i) \wedge P_2(i). \end{aligned}$$

We will prove by induction on i that $P(i)$ holds, $i \in (t, t']$. For the base case, set $i = t+1$. Since $S_{PF}(x,t) = 0$ (resp., $S(y,t) = 0$), task x (resp., y) is not urgent at time t. Similarly, since $S(x,t) = 1$ (resp., $S_{PF}(y,t) = 1$), task x (resp., y) is not tnegru at time t. Hence, tasks x and y are both contending at time t. We can now use Lemma 4.5(a) to establish $P_0(t+1)$. Similarly, Lemma 4.5(b) implies $P_1(t+1)$. Note that $P_2(t+1)$ is vacuously true. This completes the base case of the induction.

For the induction step, we assume that $P(i)$ holds over the interval $(t, i]$, and prove that it holds over $(t, i+1]$, where $i \in (t, t')$. By the definition of t', we have

$$S(x,i) = S(y,i). \tag{2}$$

Given that $P_2(i)$ is part of our induction hypothesis, $P_2(i+1)$ will follow if we can establish that

$$\alpha_i(x) = \alpha_i(y) \neq 0. \tag{3}$$

Assuming that Equation (3) fails to hold, there are four cases to consider: (i) $\alpha_i(x) = -$ and $\alpha_i(y) = 0$, (ii) $\alpha_i(x) = -$ and $\alpha_i(y) = +$, (iii) $\alpha_i(x) = 0$ and $\alpha_i(y) = +$, and (iv) $\alpha_i(x) = 0$ and $\alpha_i(y) = 0$. (The symmetric versions of the first three cases are impossible since $y \succeq x$ at time t and $P_2(i)$ holds.) Assume that Case (i) holds. Lemma 4.1(a) and $P_0(i)$ imply that $S(x,i) = 0$. Lemma 4.2(c) and $P_1(i)$ imply that $S(y,i) = 1$, contradicting Equation (2). Similar reasoning applies to Cases (ii) through (iv). Hence, Equation (3) holds and $P_2(i+1)$ holds. It remains to establish $P_0(i+1)$ and $P_1(i+1)$.

By Lemma 4.1 and $P_0(i)$, if $S(x,i) = 1$ then $\alpha_i(x) = +$. Assume the converse fails to hold, i.e., that $\alpha_i(x) = +$ and yet $S(x,i) = 0$. By Equations (2) and (3), $\alpha_i(y) = +$ and $S(y,i) = 0$. However, Lemma 4.2(d) and $P_1(i)$ imply that $S(y,i) = 1$, a contradiction. Hence, we have proven that

$$S(x,i) = 1 \iff \alpha_i(x) = +. \tag{4}$$

By Equations (2), (3), and (4), at time i we either had: (i) $S(x,i) = S(y,i) = 0$ and $\alpha_i(x) = \alpha_i(y) = -$, or (ii) $S(x,i) = S(y,i) = 1$ and $\alpha_i(x) = \alpha_i(y) = +$. Consider Case (i). By Lemma 4.1(a) and $P_0(i)$, $P_0(i+1)$ holds. By Lemma 4.2(b) and $P_1(i)$, $P_1(i+1)$ holds. Similar reasoning applies to Case (ii). Hence, $P_0(i+1)$ and $P_1(i+1)$ hold. This completes our proof by induction.

Given that $P(i)$ holds, $i \in (t, t']$, it is now quite easy to prove the two remaining claims that we need, namely: (i) $S(x,t') = 0$ and $S(y,t') = 1$, and (ii) S' is P-fair. Suppose that Claim (i) fails to hold; hence, $S(x,t') = 1$ and $S(y,t') = 0$. By Lemma 4.1 and $P_0(t')$, $\alpha_{t'}(x) = +$. By Lemma 4.2 and $P_1(t')$, $\alpha_{t'}(y) = -$. But $P_2(t')$ then implies that Algorithm PF would have given a strictly higher priority to task x than to task y at time t, a contradiction. Hence, Claim (i) holds.

For Claim (ii), it is sufficient to prove that

$$\forall i : i \in (t, t'] : -1 < \mathsf{lag}(S', x, i) < 1 \tag{5}$$

and

$$\forall i : i \in (t, t'] : -1 < \mathsf{lag}(S', y, i) < 1, \tag{6}$$

since all other lags are the same as under schedule S. Note that $\mathsf{lag}(S', x, i) = \mathsf{lag}(S, x, i) + 1$, $i \in (t, t']$. Since S is P-fair, Equation (5) will hold if we can show that $\mathsf{lag}(S, x, i) < 0$, $i \in (t, t']$. Since $P_0(i)$ holds, $i \in (t, t']$, it remains only to prove that $\mathsf{lag}(S, x, i) \neq 0$, $i \in (t, t']$. Now suppose that for some $i \in (t, t']$, $\mathsf{lag}(S, x, i) = 0$, that is, task x is punctual at time i. Hence $\alpha_{i-1}(x) = 0$. If $i = t+1$, Lemma 4.5 leads to a contradiction since task x is contending at time t. Otherwise, $i-1$ belongs to the interval (t, t'), and we obtain a contradiction to $P_2(t')$. Thus, Equation (5) holds. A symmetric argument proves that Equation (6) holds. Hence, Claim (ii) holds, and our proof is complete. □

Proof of Theorem 4: In the following, let σ_i denote the characteristic substring associated with the admissible triple (a_i, b_i, c_i), $0 \leq i \leq 1$.

We will prove the theorem by induction on the depth of recursion used by algorithm Compare. By Theorem 3, this depth is finite. For the base case, assume that Compare does not call itself recursively, i.e., that the maximum depth of recursion is 0. Thus, one of the non-recursive **return** statements must be executed (the two recursive **return** statements are in Lines 5 and 18). In the argument that follows we will deal with each of the non-recursive **return** statements in turn.

Since Line 5 is not executed, we can assume that $\min\{a_0, a_1\} \leq \min\{b_0, b_1\}$. Now consider the two quantities, $\lceil c_0/a_0 \rceil$ and $\lceil c_1/a_1 \rceil$, being compared in Lines 7 and 8. Note that the string σ_i must begin with $\lceil c_i/a_i \rceil$ +'s, followed by either a $-$ or a 0. Thus, the **return** statements of Lines 7 and 8 correctly handle any case where $\lceil c_0/a_0 \rceil \neq \lceil c_1/a_1 \rceil$.

If execution proceeds beyond Line 9, we must have $t \stackrel{\text{def}}{=} \lceil c_0/a_0 \rceil = \lceil c_1/a_1 \rceil$. Note that Line 10 then sets c_0 and c_1 to the values these variables would have attained in NaiveCompare after processing the common prefix of t +'s in σ_0 and σ_1. Let σ'_i denote the string σ_i with this common prefix removed, $0 \leq i \leq 1$. It remains to compare strings σ'_0 and σ'_1.

Note that after executing Line 10, we have $c_i \in (-a_i, 0]$, $0 \leq i \leq 1$. If either c_0 or c_1 is equal to 0, we can immediately determine the outcome of the comparison between strings σ'_0 and σ'_1. For example, if $c_0 = 0$ and $c_1 \neq 0$ then $\sigma'_0 > \sigma'_1$) because $\sigma_0 = +^t 0$, whereas the first $t+1$ symbols of the string σ_1 are $+^t -$. Reasoning in this manner, we can see that the three **return** statements of Lines 11, 12, and 13 correctly handle any case where either c_0 or c_1 is equal to 0.

If execution proceeds beyond Line 14, we must have $c_i \in (-a_i, 0)$, $0 \leq i \leq 1$. For each i, $0 \leq i \leq 1$, we now consider three cases:

- Case 1: $c_i = -(b_i \bmod a_i)$. In this case, it is easy to verify that $\sigma'_i = -+^{\lfloor b_i/a_i \rfloor} 0$. In what follows, let $\odot_i$ denote the string $-+^{\lfloor b_i/a_i \rfloor} 0$.
- Case 2: $c_i \in (-a_i, -(b_i \bmod a_i))$. In this case, the first $\lfloor b_i/a_i \rfloor + 1$ symbols of σ'_i are $\ominus_i \stackrel{\text{def}}{=} -+^{\lfloor b_i/a_i \rfloor}$. Let c'_i denote the new value of c_i after processing these symbols as in NaiveCompare. Then

$$\begin{aligned} c'_i &= c_i + b_i - a_i \cdot \lfloor b_i/a_i \rfloor \\ &= c_i + (b_i \bmod a_i). \end{aligned}$$

  Note that $c'_i \in (-a_i, 0)$.
- Case 3: $c_i \in (-(b_i \bmod a_i), 0)$. In this case, the first $\lceil b_i/a_i \rceil + 1$ symbols of σ'_i are $\oplus_i \stackrel{\text{def}}{=} -+^{\lceil b_i/a_i \rceil}$. Let c'_i denote the new value of c_i after processing these symbols as in NaiveCompare. Then

$$\begin{aligned} c'_i &= c_i + b_i - a_i \cdot \lceil b_i/a_i \rceil \\ &= c_i - (a_i - (b_i \bmod a_i)). \end{aligned}$$

  Note that $c'_i \in (-a_i, 0)$.

In Case 1 above, we completely characterize the string σ'_i. In Cases 2 and 3, we identify a prefix of σ'_i and find that after processing that prefix, the new value of c_i remains in the interval $(-a_i, 0)$, meaning that the preceding case analysis can be repeated on the remaining suffix of σ'_i. In other words, the string σ'_i may be viewed as a sequence of $\ominus_i$'s and $\oplus_i$'s, followed by a single occurrence of $\odot_i$. Given this, it is not difficult to check that whenever $\lfloor b_0/a_0 \rfloor \neq \lfloor b_1/a_1 \rfloor$, we can immediately determine which of the strings σ'_0 and σ'_1 is lexicographically greater. In particular, the **return** statements of Lines 15 and 16 correctly handle any case in which $\lfloor b_0/a_0 \rfloor \neq \lfloor b_1/a_1 \rfloor$.

We have now completed the base case of the induction, that is, we have proven that Compare works correctly (i.e., returns the same value as NaiveCompare) on any admissible input for which no recursive call is generated. It remains to consider the induction step. Accordingly, let us assume that algorithm Compare works correctly on any admissible input leading to a maximum depth of recursion strictly less than d, $d > 0$. It remains to prove that Compare works correctly on any admissible input $(a_0, b_0, c_0, a_1, b_1, c_1)$ with associated maximum depth of recursion $d > 0$. There are two cases to be considered: (i) the top-level recursive call is made in Line 5, and (ii) the top-level recursive call is made in Line 18.

The case in which the top-level recursive call occurs in Line 5 is quite easy to handle. Let τ_i denote the characteristic substring associated with the admissible triple $(b_i, a_i, -c_i)$, $0 \leq i \leq 1$. Note that the strings σ_i and τ_i are closely related. In particular, they are "complementary" strings in the sense that one can be obtained from the other by changing $-$'s to +'s, +'s to $-$'s, and leaving the 0 symbol unchanged. Once we have made this observation, it is easy to see that NaiveCompare will return the same result on $(b_1, a_1, -c_1, b_0, a_0, -c_0)$ as it would on $(a_0, b_0, c_0, a_1, b_1, c_1)$. By the induction hypothesis, the recursive call of Line 5 will function correctly, completing the analysis of this case.

It remains to consider the case in which the top-level recursive call occurs in Line 18. Our base case analysis implies that when Line 18 is executed: (i) $a_0, a_1, b_0, b_1 > 0$, (ii) $c_0 \in (-a_0, 0)$, (iii) $c_1 \in (-a_1, 0)$, (iv) $\ominus \stackrel{\text{def}}{=} \ominus_0 = \ominus_1$, (v) $\oplus \stackrel{\text{def}}{=} \oplus_0 = \oplus_1$, (vi) $\odot \stackrel{\text{def}}{=} \odot_0 = \odot_1$, (vii) strings σ'_0 and σ'_1 (as defined in the base case analysis) can be viewed as strings of $\ominus$'s and $\oplus$'s terminated by a $\odot$.

Furthermore, it is straightforward to prove that σ'_i, viewed as a string over $\{\ominus, \odot, \oplus\}$, corresponds to the characteristic substring of the admissible triple

$$(a_i - (b_i \bmod a_i),\ b_i \bmod a_i,\ c_i + (b_i \bmod a_i)),$$

$0 \leq i \leq 1$. (To make the correspondence, replace $\ominus$ by $-$, $\odot$ by 0, and $\oplus$ by +.) Thus, the induction hypothesis implies that the recursive call of Line 18 correctly compares strings σ'_0 and σ'_1. □