



UNIVERSITÀ DEGLI STUDI DI ROMA TRE
Dipartimento di Informatica e Automazione
Via della Vasca Navale, 79 – 00146 Roma, Italy

Multi-agent single machine scheduling

A. AGNETIS¹, D. PACCIARELLI², A. PACIFICI³

RT-DIA-85-03

July 2003

- (1) Dipartimento di Ingegneria dell'Informazione, Università di Siena, Italy.
 - (2) Dipartimento di Informatica e Automazione, Università Roma Tre, Italy.
 - (3) Centro Vito Volterra, Università di Roma "Tor Vergata", Italy.
-

ABSTRACT

We consider the scheduling problems arising when several agents, each owning a set of nonpreemptive jobs, compete to perform their respective jobs on one shared processing resource. Each agent wants to minimize a certain cost function, which depends on the completion times of its jobs only. The cost functions we consider in this paper are maximum of regular functions (associated with each job), number of late jobs and total weighted completion time. The different combinations of the cost functions of each agent lead to various problems, whose computational complexity is analysed in this paper. In particular, we investigate the problem of finding schedules whose cost for each agent does not exceed a given bound for each agent.

1 Introduction and multi-agent literature

This paper addresses a class of scheduling problems in which multiple agents compete on the usage of a common machine. This kind of problems occurs in several different application environments and different methodological fields, such as artificial intelligence, decision theory, operations research etc. We focus on the following situation. There are k agents, each owning a set of jobs. The agents have to schedule their jobs on a common machine, and each agent wants to minimize a cost function which depends on its own jobs' completion times. The problem is how to compute schedules which account for each agent's cost function, and that can be used to support the negotiation among the agents.

Multi-agent scheduling problems occur in several application environments in which the need for negotiation/bidding procedures arises. Most of the papers on this subject investigate heuristic approaches for the construction of schedules that are acceptable to the agents, with no particular concern on optimality. For instance, Kim et al. [7] discuss complex negotiation procedures for project scheduling in a multi-agent environment, allowing the parties to come up with new schedules whenever unacceptable task timings occur. Other approaches are based on distributed artificial intelligence. Huang and Hallam [6] address a multi-agent scheduling problem in terms of a constraint satisfaction problem where a subset of constraints can be relaxed but are expected to be satisfied as well as possible. Chen et al. [4] propose a number of negotiation protocols for functional agent cooperation in a supply chain context. Brewer and Plott [3] devise a bidding mechanism for the problem of scheduling trains (agents) on a shared single railtrack. Schultz et al. [11] discuss the problem of the integration of multimedia telecommunication services for a Satellite-based Universal Mobile Telecommunication System (S-UMTS). The problem here is to fulfill the requirements of various integrated services (agents), such as Voice over IP, web browsing, file transfer via ftp, etc. Different agents have different objectives. For instance, the voice service may tolerate the loss of some packets, but under strict delay requirements. On the contrary, when transferring a data file, no packet can go lost, but some delay can be tolerated.

From the theoretical viewpoint, multi-agent scheduling problems are a special case of general multicriteria optimization problems. For instance, the two-agent problem in which the agent 1 wants to minimize the total completion time of the jobs in J^1 and the agent 2 wants to minimize the maximum lateness of the jobs in J^2 can be viewed as a bicriteria, single-agent problem with objective functions $\sum_{j \in J^1 \cup J^2} q_j^{(1)} C_j$ and $\max_{j \in J^1 \cup J^2} \{q_j^{(2)} T_j\}$, where $q_j^{(1)} = 1$ and $q_j^{(2)} = 0$ for $j \in J^1$, whereas $q_j^{(1)} = 0$ and $q_j^{(2)} = 1$ for $j \in J^2$. Hence, in principle, general methods for multicriteria optimization can be applied, although these may not exploit the peculiarity of the problem. However, in classical single-agent, multicriteria scheduling problems, all jobs contribute to all criteria, whereas, in a multi-agent situation, only the jobs belonging to an agent contribute to that agent's criterion. So, two-agent scheduling problems differ from the problems commonly referred to as *bicriteria scheduling problems* [12]. As a consequence, the complexity results known for a certain bicriteria scheduling problem in which there are two objectives f and g , in general do not imply similar complexity results for the corresponding two-agent problem in which the two agents have objectives f and g respectively. Moreover, note that some cases only make sense in the two-agent setting, such as, for instance, when both agents have the goal of minimizing the total (unweighted) completion time of their respective jobs.

In the multi-agent setting, a key issue is the determination of *nondominated* (or Pareto-

optimal) schedules, i.e., such that a better schedule for one agent necessarily results in a worse schedule for at least another agent.

For the two-agent and three-agent cases, Baker and Smith [2] analyzed the problem of minimizing a convex combination of the agents' objective functions. They characterize the complexity of the problem in some scenarios, when the objective functions are chosen among C_{max} , L_{max} , $\sum C_i$ and $\sum w_i C_i$. The optimal solutions of their model are nondominated schedules, but not all of them can be obtained in this way. Agnetis et al. [1] addressed the complexity of several two-agent, single-machine problems in which one agent computes the best solution for him/her, given that the other agent will not accept schedules of cost greater than a certain value. For various combinations of the two agents' objective functions the problem of generating nondominated schedules was addressed. In this paper we extend this approach to the general case of $k > 2$ agents.

The paper is organized as follows. In Section 2 we formally introduce the notation and terminology used through the rest of the paper. In Section 3, we show that when one agent wants to minimize the weighted sum of completion times and another wants to minimize the makespan, even if $k = 2$, there is an exponential number of nondominated solutions (this issue was left open in [1]). In Section 4, a decomposition procedure is described, which proves useful to characterize the complexity of several special cases, as illustrated in Section 5. Conclusions follow in Section 6.

2 Problem Definition and notation

In this section we introduce the notation and terminology we use throughout the paper.

There are k competing agents $A = \{1, 2, \dots, k\}$. Each of them has a set of non-preemptive *jobs* to be processed on a single common machine. Every agent $h \in A$ has to execute the job set $J^h = \{J_1^h, J_2^h, \dots, J_{n(h)}^h\}$. We call h -jobs the jobs of J^h . The processing time of job J_i^h will be denoted by p_i^h . In some cases we will consider job due dates as well, d_i^h . We always assume zero release dates for all jobs. Each agent will have to schedule its jobs on the machine complying with the presence of the other agents' jobs.

Let σ indicate a feasible schedule of the $n = \sum_{h \in A} n(h)$ jobs, i.e., a feasible assignment of starting times to the jobs of the k agents. The completion times of job J_i^h in σ will be denoted as $C_i^h(\sigma)$ (or, simply, C_i^h if it does not generate confusion).

Each agent wants to minimize its own objective function, which depends on the completion times of its jobs only. We indicate the k objective functions by $f^h(\sigma) : \{C_i^h(\sigma) : 1 \leq i \leq n(h)\} \rightarrow \mathbb{R}$, $h \in 1, 2, \dots, k$. In this paper we consider the following objective functions for the generic agent h :

- *Maximum of regular functions* $f_{\max}^h(\sigma) = \max_{i=1, \dots, n(h)} \{f_i^h(C_i^h(\sigma))\}$, where each $f_i^h(\cdot)$ is a nondecreasing function of the completion time of job J_i^h . This includes the makespan C_{\max} and maximum tardiness T_{\max} as special cases.
- *Number of late jobs* $\sum U_i^h(\sigma) = \sum_{i=1}^{n(h)} U_i^h(\sigma)$, where $U_i^h(\sigma) = 1$ if job J_i^h is late in σ and zero otherwise.
- *Total weighted completion time* $\sum w_i C_i^h(\sigma) = \sum_{i=1}^{n(h)} w_i C_i^h(\sigma)$. When all the weights w_i are all equal, we obtain the special case of total completion time $\sum C_i$.

Since all these objective functions are regular (i.e., nondecreasing in the completion times) there is no convenience in keeping the machine idle, and therefore each job is started as soon as the previous job in the sequence is completed. We use $\sum C_i$ for $\sum w_i C_i$ when $w_i = 1$ for all i .

We say that a schedule σ is *nondominated* if there is no schedule $\bar{\sigma}$ such that $f^h(\bar{\sigma}) \leq f^h(\sigma)$, $h \in A$, and at least one of the k inequalities is strict, i.e., a schedule is nondominated if a better schedule for one of the k agents necessarily results in a worse schedule for another. Distinct nondominated schedules σ, σ', \dots may yield the same k -tuple of objective function values $(f^1(\sigma), \dots, f^k(\sigma)) = (f^1(\sigma'), \dots, f^k(\sigma')) = \dots = (y^1, \dots, y^k)$. We call (y^1, \dots, y^k) a *nondominated* k -tuple of objective function values. We say that σ, σ', \dots are *equivalent* schedules, and for each nondominated k -tuple we are interested in finding *one* of them, not all of them.

In this paper we refer to two types of problems, namely the decision problem of finding a single feasible solution, and the problem of finding all nondominated solutions. Following the usual classification scheme for scheduling problems, we denote a problem with three fields, $\psi_1|\psi_2|\psi_3$, where ψ_1 indicates the scheduling environment, ψ_2 describes the job characteristics or restrictive requirements, and ψ_3 defines the set of objective functions, one for each agent. In particular, in this paper we address only single machine problems, so we let $\psi_1 = 1$. Under ψ_3 , we will distinguish the two types of problems as follows.

DECISION PROBLEM ($1||f^1, f^2, \dots, f^k$). Given k integers Q^1, \dots, Q^k and, for each agent $h \in A$, the job set J^h and the objective function $f^h(\cdot)$ of the agent, find a schedule σ^* such that $f^h(\sigma^*) \leq Q^h$ for all $h \in A$.

PARETO-OPTIMIZATION PROBLEM ($1||f^1 \cdot f^2 \cdot \dots \cdot f^k$). Given the job sets J^h and the k objective functions $f^h(\cdot)$ of each agent $h \in A$, find the set of all nondominated k -tuples $(f^1(\cdot), \dots, f^k(\cdot))$ and a corresponding schedule of J^1, \dots, J^k for each k -tuple.

In some cases it will be necessary to specify the bounding integers in the decision problem, and so we will indicate the decision problem as $1||f^1 \leq Q^1, f^2 \leq Q^2 \dots, f^k \leq Q^k$. If some of the values Q^h , $h = 1, \dots, k$ is too small, an instance of $1||f^1, f^2, \dots, f^k$ may not have feasible solutions. If there is at least one feasible solution, we say that the instance is *feasible*.

Given an instance of $1||f^1, f^2, \dots, f^k$, the problem of finding, among feasible schedules, one which is also nondominated can always be addressed by *binary search*. Suppose first that a feasible solution σ to $1||f^1 \leq Q^1, f^2 \leq Q^2, \dots, f^k \leq Q^k$ is available. Now we want to find a schedule σ' such that $f^h(\sigma') \leq Q^h$ for all $h \in A \setminus \{1\}$ and $f^1(\sigma')$ is minimum. We may try solving $1||f^1 \leq \frac{1}{2}Q^1, f^2 \leq Q^2, \dots, f^k \leq Q^k$. If we get a feasible solution we may further decrease Q^1 to $\frac{1}{4}Q^1$, otherwise we try $\frac{3}{4}Q^1$. This goes on until we determine the smallest value Q^{1*} such that instance $1||f^1 \leq Q^{1*}, f^2 \leq Q^2, \dots, f^k \leq Q^k$ is feasible. Next, we apply the same procedure to find the smallest value Q^{2*} such that $1||f^1 \leq Q^{1*}, f^2 \leq Q^{2*}, f^3 \leq Q^3 \dots, f^k \leq Q^k$ is feasible and, iterating, we eventually get a solution σ^* such that $f^h(\sigma^*) \leq Q^{h*}$ for all $h \in A$. It is not hard to see that σ^* is nondominated. Moreover note that in general we may obtain different nondominated schedules if we perform the binary searches on the values Q^h according to different permutations of the agents' sequence.

In the Pareto-optimization problem $1||f^1 \cdot f^2 \cdot \dots \cdot f^k$, the agents want to list all possible nondominated k -tuples, in order to negotiate the most acceptable trade-off for them.

The main focus of the paper is to analyze the complexity of these problems and propose solution algorithms. Table 1 summarizes known results for the $k = 2$ case that will be used in the present work. The second column reports the computational complexity of the problem $1||f^1, f^2, \dots, f^k$ and the third column the number of distinct nondominated pairs. Except where indicated, the details concerning the results of the table are contained in [1].

<i>Scenario</i> ($k = 2$)	<i>Computational</i> <i>complexity of</i> $1 f^1, f^2$	<i>Number of nondominated</i> <i>pairs in</i> $1 f^1 \cdot f^2$
(f_{\max}^1, f_{\max}^2)	$O(n \log n)$ (Section 5.1)	$O(n(1)n(2))$
$(\sum w_i C_i^1, f_{\max}^2)$	binary NP-hard	exponential (Section 3)
$(\sum C_i^1, f_{\max}^2)$	$O(n \log n)$	$O(n(1)n(2))$
$(\sum U_i^1, f_{\max}^2)$	$O(n \log n)$	$O(n(1))$
$(\sum U_i^1, \sum U_i^2)$	$O(n^3)$	$\min\{n(1), n(2)\}$
$(\sum C_i^1, \sum U_i^2)$	Open	$O(n(2))$
$(\sum w_i C_i^1, \sum U_i^2)$	binary NP-hard	$O(n(2))$
$(\sum C_i^1, \sum C_i^2)$	binary NP-hard	exponential

Table 1: Summary of complexity results for decision and Pareto-optimization problems for the two-agent ($k = 2$) case.

3 Pareto-optimal solutions of $1|| \sum w_i C_i^1 \cdot f_{\max}^2$

In the following we show that $1|| \sum w_i C_i^1 \cdot f_{\max}^2$ may have an exponential number of Pareto optimal solution. This issue was left open in [1].

Consider the following instance of $1|| \sum w_i C_i^1 \cdot C_{\max}^2$. Agent 1 has $n(1)$ jobs. For each job $J_i^1 \in J^1$ the processing times and the weights are: $p_i^1 = w_i^1 = 2^{i-1}$, $i = 1, 2, \dots, n(1)$. Agent 2 has a single job of unit length.

THEOREM 3.1 *Consider an instance of $1|| \sum w_i C_i^1 \cdot C_{\max}^2$ in which agent 2 has a single job of unit length, while agent 1 has $n(1)$ jobs. For each job i ($i \in \{1, 2, \dots, n(1)\} = J^1$), $p_i^1 = w_i^1 = 2^{i-1}$. Then, for every active schedule, the quantity $C_1^2 + \sum_{i=1}^{n(1)} w_i C_i^1$ is constant and equal to $\frac{1}{3}(1 + 2^{2n(1)+1})$.*

Proof. Given any active schedule σ , consider two adjacent jobs j and k . Let t be the starting time of job j and $t + p_j$ the starting time of job k . The contribution to the objective function of the two jobs is then $w_j(t + p_j) + w_k(t + p_j + p_k)$. Consider now the schedule $\bar{\sigma}$ in which the two jobs are switched: the contribution of the two jobs to the objective function is now $w_k(t + p_k) + w_j(t + p_k + p_j)$. Observe now that $w_j p_k = w_k p_j$ for any pair of jobs in $J^1 \cup J^2$ (since $w_i = p_i$ for each job), thus proving that σ and $\bar{\sigma}$ has the same value of the objective function. This implies that any active schedule produces the same value of the quantity $C_1^2 + \sum_{i=1}^{n(1)} w_i C_i^1$. This value can be computed, for example, by considering the sequence: $J_1^2, J_1^1, J_2^1, \dots, J_n^1(1)$. We have:

$$C_1^2 + \sum_{i=1}^{n(1)} w_i C_i^1 = 1 + \sum_{i=1}^{n(1)} 2^{i-1} 2^i = 1 + \sum_{i=1}^{n(1)} 2^{2i-1} = 1 + \sum_{i=1}^{2n(1)} 2^i - \sum_{i=1}^{n(1)} 2^{2i}$$

Since $\sum_{i=1}^{n(1)} 2^{2i} = 2 \sum_{i=1}^{n(1)} 2^{2i-1}$, we can write: $\sum_{i=1}^{n(1)} 2^{2i-1} = \sum_{i=1}^{2n(1)} 2^i - 2 \sum_{i=1}^{n(1)} 2^{2i-1}$. Hence, we obtain: $\sum_{i=1}^{n(1)} 2^{2i-1} = \frac{1}{3}(\sum_{i=1}^{2n(1)} 2^i) = \frac{1}{3}(2^{2n(1)+1} - 2)$.

In conclusion, the quantity $C_1^2 + \sum_{i=1}^{n(1)} w_i C_i^1$ is equal to $\frac{1}{3}(1 + 2^{2n(1)+1})$, and the thesis follows. \square

In order to prove that the above instance has an exponential number of nondominated solutions, consider that, for any value $1 \leq Q^2 \leq 2^{n(1)}$, there is a subset of J^1 whose total length equals $Q^2 - 1$. This implies that there is a feasible solution to $1 || \sum w_i C_i^1, C_{\max}^2 \leq Q^2$ where $C_{\max}^2 = Q^2$ and $\sum w_i C_i^1 = \frac{1}{3}(1 + 2^{2n(1)+1}) - Q^2$. This is clearly a nondominated solution and therefore we have $2^{n(1)}$ nondominated pairs.

4 Separability property for f_{\max}

In this section, we describe a procedure for reducing a k -agent problem to a $(k-1)$ -agent problem, when the cost function of agent k is the maximum of regular functions. Under certain conditions on the other objective functions, this property may be exploited to devise efficient solution algorithms for the original problem.

The decomposition procedure operates in two phases, and can be outlined as follows. Consider a decision problem with k agents in which the cost function of the k -th agent is the maximum of regular functions:

$$1 || g^1, g^2, \dots, g^{k-1}, f_{\max}^k \quad (1)$$

The first phase consists in defining suitable *reserved time intervals* for the jobs of the k -th agent. Such intervals are *forbidden* to the remaining $k - 1$ agents, who have therefore to deal with a problem with forbidden intervals. As it will be clarified by Lemmas 4.1 and 4.2, due to the regularity of the all cost functions, given a preemptive solution to the problem with forbidden intervals, we may easily obtain a non-preemptive solution yielding the same cost values for all agents. The second phase consists in actually solving the problem with forbidden intervals. For certain cost functions g^h of the remaining $k - 1$ agents, this can be efficiently done by solving an equivalent instance of

$$1 || g^1, g^2, \dots, g^{k-1}. \quad (2)$$

The phases of the procedure are illustrated in the following sections.

4.1 Reserved intervals for agent k

In this section we always refer to the jobs of agent k , so we omit k in $D_i^k, C_i^k, J_i^k, Q^k \dots$.

Since $f_{\max}^k = \max_{i=1, \dots, n(k)} \{f_i^k(C_i^k(\sigma))\}$, for each job of agent k we can define a *deadline* D_i such that $f_i(C_i) \leq Q$ for $C_i \leq D_i$ and $f_i(C_i) > Q$ for $C_i > D_i$. In other words, job J_i must complete no later than D_i in a feasible solution.

In what follows, we call the *latest start time* (LS_i) of job J_i the maximum value the starting time of J_i can attain in a feasible schedule such that $C_i \leq D_i$ for all $J_i \in J^k$. The values LS_i can be computed as follows. Number the k -jobs in nondecreasing order of D_i . Schedule the last job $J_{n(k)}$ to start at time $D_{n(k)} - p_{n(k)}$. Continue backwards, letting

$LS_i := \min\{D_i, LS_{i+1}\} - p_i$, for all $i = n(k) - 1, \dots, 1$. Clearly, if job J_i starts after time LS_i , then at least one k -job (say, J_ℓ) attains $f_\ell(C_\ell) > Q$.

Consider now, for each k -job J_i , the latest processing interval $[LS_i, D_i]$. Let $I = \cup_{i=1}^{n(k)} [LS_i, D_i]$. Set I consists of a number $\beta \leq n(k)$ of intervals, $I_{1,h_1}, I_{h_1,h_2}, \dots, I_{h_{\beta-1},n(k)}$, call them *reserved intervals* for agent k . Each reserved interval $I_{u,v}$ ranges from LS_u to D_v . Note that, by construction, $|I_{u,v}| = D_v - LS_u = \sum_{i=u}^v p_i$. We say that jobs J_u, J_{u+1}, \dots, J_v are *associated with* $I_{u,v}$.

4.2 Preemptive problem with forbidden intervals

Consider now the variant of Problem (1) in which preemption is allowed:

$$1|pmtn|g^1, \dots, g^{k-1}, f_{max}^k. \quad (3)$$

The following lemmas extend known results on preemption redundancy for classical (single-agent) scheduling problems [5].

LEMMA 4.1 *If a feasible solution to $1|pmtn|g^1, \dots, g^{k-1}, f_{max}^k$ exists, then there is a feasible solution to $1||g^1, \dots, g^{k-1}, f_{max}^k$.*

Proof. Just observe that if in the feasible solution to $1|pmtn|g^1, \dots, g^{k-1}, f_{max}^k$ there is a job J_i (of any agent), ending at C_i , which is preempted at least once, we can always schedule the whole J_i in interval $[C_i - p_i, C_i]$, moving other (parts of) jobs backwards, without increasing the completion time of any job. Repeating this for each preempted job, we eventually obtain a nonpreemptive solution. \square

LEMMA 4.2 *If there exists a feasible solution to $1|pmtn|g^1, \dots, g^{k-1}, f_{max}^k$, there is one in which each k -job is nonpreemptively scheduled in the reserved interval it is associated with.*

Proof. In a feasible solution to $1|pmtn|g^1, \dots, g^{k-1}, f_{max}^k$, all the k -jobs associated with interval $I_{u,v}$ complete before D_v^k . Hence, if we move all the portions of each such job to exactly fit the interval $I_{u,v}$, we obtain a solution in which the completion time of no h -job ($h \in A \setminus \{k\}$) has increased, since we only moved pieces of h -jobs backwards. \square

Lemma 4.2 allows us to fix the position of the k -jobs in a feasible solution to Problem (3). We have therefore showed that, in order to solve the original Problem (1) with k agents we may equivalently solve an instance of a preemptive problem involving only agents $1, \dots, k-1$, in which intervals $I = \{I_{1,h_1}, I_{h_1,h_2}, \dots, I_{h_{\beta-1},n(k)}\}$ cannot be used. We indicate such problem as:

$$1|pmtn, I|g^1, \dots, g^{k-1}. \quad (4)$$

and call it *problem with forbidden intervals*.

4.3 Solving problem $1|pmtn, I|g^1, \dots, g^{k-1}$.

In some cases the preemptive problem with forbidden intervals, and consequently the original instance, can be solved efficiently by a simple sequencing rule. In other cases, it is necessary to define an auxiliary problem, as it is shown in the following paragraphs.

4.3.1 SPT sequencing for $1|pmtn, I|\sum C_i^1$

Consider the single-agent problem $1|pmtn, I|\sum C_i^1$. The following proposition holds.

PROPOSITION 4.3 *Problem $1|pmtn, I|\sum C_i^1$, with set of forbidden intervals I , is solved by sequencing the 1-jobs in shortest-processing-time order.*

Proof. The result follows straightforwardly by a pairwise interchange argument. \square

Referring to the case $k = 2$, Lemma 4.1, Lemma 4.2 and Proposition 4.3 ensure that a feasible solution to $1||\sum C_i^1, f_{\max}^2$ can be derived from a feasible solution to $1|pmtn, I|\sum C_i^1$, where I is the set of intervals reserved to agent 2. Observe that the same result does not hold if f^1 is $\sum w_i C_i^1$. In fact, sequencing the 1-jobs according to non-increasing w_i/p_i order (Smith's rule), does not ensure optimality when forbidden intervals exist. The problem $1||\sum w_i C_i^1, f_{\max}^2$ is in fact NP-hard (Table 1).

4.3.2 Auxiliary problem with modified deadlines

In other cases, for different combinations of the objective functions g^1, \dots, g^{k-1} , the problem with forbidden intervals (4) may be solved by resorting to an auxiliary instance of a problem *without* forbidden intervals. In particular if, for all $h \in A \setminus \{k\}$, the cost functions are of the type:

1. $g^h = f_{\max}^h$, or
2. $g^h = \sum U_i^h$,

we can define an auxiliary problem by Procedure 4.4 below, where we use the symbol d_i^h indifferently to denote a due date, if $f^h = \sum U_i$, or a deadline, computed as in Section 4.1, if $f^h = f_{\max}$.

PROCEDURE 4.4

Input: Problem with forbidden intervals I : $1|pmtn, I|g^1, g^2, \dots, g^{k-1}$ and values d_i^h for $h \in A \setminus \{k\}, i \in J^h$.

Output: Auxiliary instance of the problem $1|pmtn|g^1, g^2, \dots, g^{k-1}$ with modified due dates (or deadlines) \mathcal{D}_i^h , for $h \in A \setminus \{k\}, i \in J^h$.

Begin

For each job J_i^h of each agent $h \in A \setminus \{k\}$:

1. if d_i^h falls outside of any forbidden interval, subtract from d_i^h the total length of all the forbidden intervals preceding d_i^h , i.e., define the modified due date

$$\mathcal{D}_i^h := d_i^h - \sum_{u,v: D_u^k \leq d_i^h} \|I_{u,v}\|$$

2. **else** if d_i^h falls within the forbidden interval $I_{p,q}$, do the same, but instead of d_i^h use the left extreme of $I_{p,q}$, i.e.,

$$\mathcal{D}_i^h := LS_p - \sum_{u,v:D_v^k < d_i^h} \|I_{u,v}\|$$

Return \mathcal{D}_i^h for all $h \in A \setminus \{k\}$, $i \in J^h$.

End

The following proposition summarizes the results obtained in the previous paragraphs.

PROPOSITION 4.5 *Problem (1)*

$$1||g^1, g^2, \dots, g^{k-1}, f_{max}^k$$

with $g^h = f_{max}^h$, or $g^h = \sum U_i^h$ for all $h \in A \setminus \{k\}$, is feasible if and only if the instance of Problem (2)

$$1||g^1, g^2, \dots, g^{k-1}$$

obtained modifying the due dates and deadlines according to Procedure 4.4 is feasible.

Proof. Lemma 4.1 ensures that Problem (1)

$$1||g^1, g^2, \dots, g^{k-1}, f_{max}^k$$

and Problem (3)

$$1|pmtn|g^1, g^2, \dots, g^{k-1}, f_{max}^k$$

are equivalent from the feasibility viewpoint. On the other hand, Lemma 4.2 shows that we may always non-preemptively schedule the k -jobs in the associated reserved intervals, and therefore Problem (3) can be solved by solving an instance of Problem (4)

$$1|pmtn, I|g^1, g^2, \dots, g^{k-1}.$$

When the cost functions of the first $k-1$ agents consist of minimizing the maximum of regular functions or the number of tardy jobs ($g^h = f_{max}^h$, or $g^h = \sum U_i^h$ for all $h \in A \setminus \{k\}$), we may define an auxiliary instance of the Problem

$$1|pmtn|g^1, g^2, \dots, g^{k-1} \tag{5}$$

in which the due dates and deadlines of the jobs have been modified according to Procedure 4.4. Such modified due dates and deadlines take into account the room needed by the reserved intervals associated with the k -jobs: since preemption is allowed, it is not hard to see that if and only if the auxiliary instance of (2) is feasible then so does the instance of Problem (4) with forbidden intervals. Invoking again Lemma 4.1, Problem (5) is in turn equivalent to the nonpreemptive Problem (2)

$$1||g^1, g^2, \dots, g^{k-1}.$$

This proves that, when $g^h = f_{max}^h$, or $g^h = \sum U_i^h$ for all $h \in A \setminus \{k\}$, Problems (2) and (1) are indeed equivalent. \square

The previous proposition emphasizes that it is always possible to reduce an instance of Problem (1), when $g^h = f_{max}^h$, or $g^h = \sum U_i^h$ for all $h \in A \setminus \{k\}$, to an equivalent instance of Problem (2) in polynomial time.

5 Polynomial special cases

In this section, we apply the concepts introduced in Section 4 to provide polynomial algorithms for certain scenarios.

5.1 $1||\sum C_i^1, f_{\max}^2, \dots, f_{\max}^k$

We next show polynomial solvability of problem $1||\sum C_i^1, f_{\max}^2, \dots, f_{\max}^k$.

THEOREM 5.1 $1||\sum C_i^1, f_{\max}^2, \dots, f_{\max}^k$ can be solved in time $O(n(1) \log n(1) + \tilde{n} \log \tilde{n})$, where $\tilde{n} = \sum_{h \in A \setminus \{1\}} n(h)$.

Proof. If $k = 2$, the problem is $1||\sum C_i^1, f_{\max}^2$. The procedure can be applied and, from Proposition 4.3, we may easily solve the problem with forbidden intervals by sequencing the 1-jobs in SPT order. Consider now the general case of k agents. We may proceed exactly in the same way that in the $k = 2$ case and consider all the jobs of agents $2, 3, \dots, k$ as if they were jobs of a single agent. Thus we order all the jobs in $\bigcup_{h \in A \setminus \{1\}} J^h$ in nondecreasing order of the deadlines D_i and, starting from one with the largest deadline, compute backward the $LS_i := \min\{D_i, LS_{i+1}\} - p_i$ in order to define proper reserved intervals for those jobs. After this, the problem is indeed equivalent to $1||\sum C_i^1, f_{\max}^2$.

Let us turn to complexity issues. In the $k = 2$ case, the complexity can be measured for the basic steps of the procedure as follows. The definition of reserved intervals for the 2-jobs implies that the 2-jobs are ordered first: the complexity for this is $O(n(2) \log n(2))$. Moreover, the computation of the reserved intervals takes time $O(n(2))$. The definition and the solution of the auxiliary preemptive instance requires $O(n(1) \log n(1))$ time for the SPT ordering of the 1-jobs and $O(n(1) + n(2))$ time to reconstruct the preemptive schedule. The solution of the original instance of $1||\sum C_i^1, f_{\max}^2$ is obtained from the preemptive schedule in time $O(n(1) + n(2))$.

The overall complexity is thus dominated by the ordering steps and it is therefore $O(n(1) \log n(1) + n(2) \log n(2))$.

For the general case with k agents, the computational complexity can be immediately derived from the previous considerations: the cardinality of the job set $\bigcup_{h \in A \setminus \{1\}} J^h$ is equal to $\sum_{h \in A \setminus \{1\}} n(h)$ and the thesis follows. \square

(In [1] a different procedure was proposed for solving $1||\sum C_i^1, f_{\max}^2$, having the same complexity, based on a simple adaptation of the well-known Lawler's algorithm for the single-agent problem $1|prec|f_{\max}$ [8].)

In the above proof we have not applied directly the procedure illustrated in Section 4, and based our result on the observation that, basically, all the jobs of agents with f_{\max} objective functions may be regarded as the jobs of a single agent so that we reduce the problem to a special instance of $1||\sum C_i^1, f_{\max}^2$. Still, if we applied the general procedure of Section 4, by iteratively reducing an instance of $1||\sum C_i^1, f_{\max}^2, \dots, f_{\max}^h$ to a special instance of $1||\sum C_i^1, f_{\max}^2, \dots, f_{\max}^{h-1}$, for $h = k$ down to $h = 3$, we would obtain the same complexity, as $O(n \log n) = O(\sum_h n(h) \log(\sum_h n(h)))$.

Another relevant case is when *all* the agents have to minimize the maximum of regular functions. Similarly to Theorem 5.1, it is easy to prove that the following proposition holds.

PROPOSITION 5.2 $1||f_{\max}^1, \dots, f_{\max}^k$ can be solved in time $O(n \log n)$ where $n = \sum_{h \in A} n(h)$.

Note that this result improves that presented in [1] for $k = 2$, having complexity of $O(n(1)^2 + n(2) \log n(2))$.

5.2 $1||\sum U_i^1, \dots, \sum U_i^p, f_{\max}^{p+1}, \dots, f_{\max}^k$

In view of the decomposition procedure described in Section 4, this problem reduces to solving an auxiliary instance of $1||\sum U_i^1, \dots, \sum U_i^p$. Hence, for $p = 1$ the problem is solvable by Moore's algorithm [9], while for $p = 2$ the problem can be solved by dynamic programming with an algorithm described in [1]. We next show that it is indeed possible to deal with a generic number $p \geq 1$ of agents. The following lemma relates to the structure of a feasible schedule.

LEMMA 5.3 *For any feasible instance, there is a feasible schedule σ^* of $1||\sum U_i^1, \dots, \sum U_i^p$ in which all the late jobs are scheduled consecutively at the end of the schedule, and all the early jobs are scheduled consecutively in Earliest Due Date (EDD) order at the beginning of the schedule.*

Proof. Consider an optimal schedule σ^* and move all the late jobs to the end of the schedule, thus obtaining a new schedule σ' . Clearly, $\sum U_i^h(\sigma') \leq \sum U_i^h(\sigma^*)$, for $h = 1, \dots, p$, since we are moving backward the early jobs. Consider now all the early jobs in σ' , that are sequenced consecutively at the beginning of the schedule, and resequence them in EDD order. This does not increase the number of late jobs, thus completing the proof. \square

In the remaining part of this section we assume that the $n = \sum_{h=1}^p n(h)$ jobs in $\bigcup_{h=1}^p J^h$ are numbered from J_1 to J_n according to the EDD order.

We next illustrate a recursion relation that can be exploited to design a polynomial dynamic programming algorithm for σ^* of $1||\sum U_i^1, \dots, \sum U_i^p$.

Let $C(i, l_1, \dots, l_p)$ be the minimum completion time of the last early job in a partial schedule of the job set $\{J_1, \dots, J_p\}$ in which there are at most l_h late h -jobs for each set J^h , $h = 1, \dots, p$. By definition, we set $C(i, l_1, \dots, l_p) = +\infty$ if no such schedule exists.

The following relations hold:

Boundary Conditions:

$$\begin{aligned} C(0, 0, \dots, 0) &= 0 \\ C(i, l_1, \dots, l_p) &= +\infty \quad \text{if } i < 0 \text{ or } l_h < 0, h = 1, \dots, p. \end{aligned}$$

Recursion Relation:

$$f(i, l_1, \dots, l_p) = \begin{cases} +\infty & \text{if } C(i-1, l_1, \dots, l_p) + p_i > d_i \\ 0 & \text{otherwise.} \end{cases}$$

$$C(i, l_1, \dots, l_p) = \min \left\{ \begin{array}{l} C(i-1, l_1, \dots, l_p) + p_i + f(i, l_1, \dots, l_p); \\ C(i-1, l_1, \dots, l_{h-1}, l_h - 1, l_{h+1}, \dots, l_p), \quad \text{where } h : J_i \in J^h. \end{array} \right\}$$

LEMMA 5.4 *If $C(i, l_1, \dots, l_p)$ is finite, then it is the minimum completion time of the last early job over all feasible schedules for the job set $\{J_1, \dots, J_i\}$, with at most l_h late h -jobs, $h = 1, \dots, p$. If $C(i, l_1, \dots, l_p)$ is infinite, then there is no such feasible schedule.*

Proof. The proof is by induction on i . Clearly the property holds for $i = 1$, for any $l_h = 1, \dots, n(h)$, with $h = 1, \dots, p$. Now, assume that the property holds until $(i-1)$. We will show that the property holds also for i and for any $l_h = 1, \dots, n(h)$, and $h = 1, \dots, p$. Let σ be a feasible schedule for the job set $\{J_1, \dots, J_i\}$, such that the completion time τ of the last early job in σ is minimum among all feasible schedules with at most l_h late h -jobs, $h = 1, \dots, p$. Without loss of generality, assume that $J_i \in J^{\hat{h}}$. If J_i is late in σ , then, from the inductive hypothesis, $\tau = C(i-1, l_1, \dots, l_{\hat{h}-1}, l_{\hat{h}}-1, l_{\hat{h}+1}, \dots, l_p)$. If J_i is early in σ then, again from the inductive hypothesis, $\tau = C(i-1, l_1, \dots, l_p) + p_i$. Hence, the above recursion relation correctly chooses the smallest between the two quantities. Note that the schedule attaining $C(i, l_1, \dots, l_p)$ is feasible if either $C(i-1, l_1, \dots, l_p) + p_i < d_i$ or $C(i-1, l_1, \dots, l_{\hat{h}-1}, l_{\hat{h}}-1, l_{\hat{h}+1}, \dots, l_p) < +\infty$. If none of the two holds, there can be no feasible schedule of $\{J_1, \dots, J_i\}$ with at most l_h late h -jobs, $h = 1, \dots, p$, and the algorithm sets $C(i, l_1, \dots, l_p) = +\infty$. \square

THEOREM 5.5 *An instance of $1||\sum U_i^1 \leq Q_1, \dots, \sum U_i^p \leq Q_p$ is feasible if and only if $C(n, Q_1, \dots, Q_p) < +\infty$. This value can be computed in time $O(\sum_{h=1}^p n(h) \prod_{j=1}^p n(j))$.*

Proof. Suppose first that $C(n, Q_1, \dots, Q_p) < +\infty$. Lemma 5.4 implies that there exists a feasible solution to $1||\sum U_i^1 \leq Q_1, \dots, \sum U_i^p \leq Q_p$. Conversely, if $C(n, Q_1, \dots, Q_p) = +\infty$, again Lemma 5.4 implies that there is no feasible solution to $1||\sum U_i^1 \leq Q_1, \dots, \sum U_i^p \leq Q_p$. The only if part also follows from Lemma 5.4, since the cases in which the instance is feasible and $C(n, Q_1, \dots, Q_p) = +\infty$, or the instance is infeasible and $C(n, Q_1, \dots, Q_p) < +\infty$ cannot occur.

Let us now turn to complexity. Computing each $C(i, l_1, \dots, l_p)$ requires constant time, and therefore computing all of them requires $O(\sum_{h=1}^p n(h) \prod_{j=1}^p n(j))$ time. \square

6 Conclusions

In this paper we have investigated the complexity of some scheduling problems in which several agents have to negotiate the usage of a common processing resource. The focus of this paper was the problem of finding a schedule attaining minimum quality requirements specified by each single agent. The model presented in this paper can be regarded as a tool for complex negotiation processes. A lot has to be done in the area of modeling multi-agent scheduling situations.

- From the modeling viewpoint, connections with concepts of bargaining theory (e.g. Nash solution or other solution concepts in the case of $k = 2$, see for instance [10]) when specialized to the scheduling context might deserve further investigation.
- From the theoretical viewpoint, the problem $1||\sum U_i^1, \sum C_i^2$ is minimally open. Depending on its status, other multi-agent problems might deserve further investigation.

- From the algorithmic viewpoint, there is a need for effective enumeration algorithms as well as approximation algorithms and polynomial approximation schemes for NP-hard cases.

Acknowledgements

This paper is the most recent result of a stream of work that was started together with Prof. Mario Lucertini. Here we wish to recall his generosity and enthusiasm during those fruitful discussions that we miss so much.

References

- [1] Agnetis, A., Mirchandani, P.B., Pacciarelli, D., Pacifici, A. 2003. Scheduling Problems with Two Competing Users, *Operations Research* to appear.
- [2] Baker, K.R., Smith, J.C. 2003. A multiple-criterion model for machine scheduling. *Journal of Scheduling*, 6(1), 7–16.
- [3] Brewer, P.J., C.R. Plott. 1996. A Binary Conflict Ascending Price (BICAP) Mechanism for the Decentralized Allocation of the Right to Use Railroad Tracks. *International Journal of Industrial Organization*. **14** 857-886.
- [4] Chen, Y., Y. Peng, T. Finin, Y. Labrou, S. Cost, B. Chu, J. Yao', R. Sun, and B. Wilhelm. 1999. A negotiation-based multi-agent system for supply chain management. In *Proc. of the Agents'99 Workshop "Agent-Based Decision-Support for Managing the Internet-Enabled Supply-Chain"*, Seattle, WA. 15–20.
- [5] Conway, R.W., W.L. Maxwell, L.W. Miller (1967), *Theory of Scheduling*, Addison Wesley, Reading, MA.
- [6] Huang, X., J.C. Hallam. 1995. Spring-Based Negotiation for Conflict Resolution in AGV Scheduling. *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*.
- [7] Kim, K., B.C. Poulon, C.J. Petrie, and V.R. Lesser. 2000. Compensatory Negotiation for Agent-Based Project Schedule Coordination, CIFE Working Paper #55. Stanford University.
- [8] Lawler, E.L. 1973. Optimal sequencing of a single machine subject to precedence constraints. *Management Science*. **19** 544–546.
- [9] Moore, J.M. 1968. An n Job, One Machine Sequencing Algorithm for Minimizing the Number of Late Jobs, *Management Science*. **15** 102–109.
- [10] Osborne, M.J., Rubinstein, A., 1994, *A Course in Game Theory*, MIT Press, Cambridge.

- [11] Schultz, D., Seoung-Hoon Oh, C. F. Grecas, M. Albani, J. Sanchez, C. Arbib, V. Arvia, M. Servilio, F. Del Sorbo, A. Giralda, G. Lombardi. 2002. *A QoS Concept for Packet Oriented S-UMTS Services*. Proceedings of the IST MOBILE SUMMIT 2002, Thessaloniki (Greece).
- [12] T'kindt, V., Billaut, J.-C. 2002. *Multicriteria Scheduling*, Springer-Verlag, Berlin.