

Complexity and Approximation for the Precedence Constrained Scheduling Problem with Large Communication Delays

R. Giroudeau, J.C. König, F.K. Moulai, and J. Palaysi

LIRMM, 161 rue Ada, 34392 Montpellier Cedex 5, France, UMR 5056

Abstract. We investigate the problem of minimizing the makespan for the multiprocessor scheduling problem. We show that there is no hope of finding a ρ -approximation with $\rho < 1 + 1/(c+4)$ (unless $\mathcal{P} = \mathcal{NP}$) for the case where all the tasks of the precedence graph have unit execution times, where the multiprocessor is composed of an unrestricted number of machines, and where c denotes the communication delay between two tasks i and j submitted to a precedence constraint and to be processed by two different machines. The problem becomes polynomial whenever the makespan is at the most $(c+1)$. The $(c+2)$ case is still partially opened.

1 Introduction

Scheduling theory is concerned with the *optimal allocation of scarce resources to activities over time*. The *theory* of the design of algorithms for scheduling is younger, but still has a significant history.

In this article we adopt the *classical scheduling delay model* or *homogeneous model* in which an instance of a scheduling problem is specified by a set $J = \{j_1, \dots, j_n\}$ of n nonpreemptive tasks, a set of U of q precedence constraints (j_i, j_k) such that $G = (J, U)$ is a directed acyclic graphs (dag), the processing times $p_i, \forall j_i \in J$, and the communication times $c_{ik}, \forall (j_i, j_k) \in U$.

If the task j_i starts its execution at time t on processor π , and if task j_k is a successor of j_i in the dag, then either j_k starts its execution after the time $t + p_{j_i}$ on processor π , or after time $t + p_{j_k} + c_{j_i j_k}$ on some other processor. In the following we consider the case of $\forall j_k \in J, p_{j_k} = 1$ and $\forall (j_i, j_k) \in E, c_{j_i j_k} = c \geq 2$.

This model was first introduced by Rayward-Smith [13]. In this model we have a set of identical processors that are able to communicate in a uniform way. We want to use these processors in order to process a set of tasks that are subject to precedence constraints. The problem is to find a trade-off between the two extreme solutions, namely, execute all the tasks sequentially without communication, or try to use all the potential parallelism but at the cost of an increased communication overhead. This model has been extensively studied these last years both from the complexity and the (non)-approximability points of view [2].

Using the *three fields* notation scheme proposed by Graham et al. [6], the problem is denoted as $\bar{P}|prec, c_{ij} = c \geq 2; p_i = 1|C_{max}$ i.e. we have an unbounded number of identical processors in order to schedule a dag such that each task has the same execution time and each pair of tasks have the same communication time. The aim is to minimize the length of the schedule.

1.1 Complexity Results

The problems with unitary communication delay. If we consider the problem of scheduling a precedence graph with unitary communication delays and unit execution time (UET-UCT) on an unbounded number of processors, Hoogeveen et al. [7] proved that the decision problem associated to $\bar{P}|prec; c_{ij} = 1; p_i = 1|C_{max}$ becomes \mathcal{NP} -complete even for $C_{max} \geq 6$, and that it is polynomial for $C_{max} \leq 5$. Their proof is based on a reduction from the \mathcal{NP} -complete problem *3SAT* [3]. The \mathcal{NP} -completeness result for $C_{max} = 6$ implies that there is no polynomial time approximation algorithm with ratio guarantee better than $7/6$, unless $\mathcal{P} = \mathcal{NP}$.

Moreover, in the presence of a bounded number of processors, Hoogeveen et al. [7] establish that whether an instance of $P|prec; c_{ij} = 1; p_i = 1|C_{max}$ has a schedule of length of at the most 4 is \mathcal{NP} -complete (they use a reduction from the \mathcal{NP} -complete problem *Clique*), whereas Picouleau [11] develops a polynomial time algorithm for the $C_{max} = 3$. In the same way, the \mathcal{NP} -completeness result for $C_{max} = 4$ implies that there is no polynomial time approximation algorithm with ratio guarantee better than $5/4$, unless $\mathcal{P} = \mathcal{NP}$.

The problems with large communication delay. If we consider the problem of scheduling a precedence graph with large communication delays and unit execution time (UET-LCT), on bounded number of processors, Bampis et al. in [1] proved that the decision problem denoted by $P|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$ for $C_{max} = c + 3$ is \mathcal{NP} -complete problem, and for $C_{max} = c + 2$ (for the special case $c = 2$), they develop a polynomial time algorithm. Their proof is based on a reduction from the \mathcal{NP} -complete problem *Balanced Bipartite Complete Graph*, *BBCG* [3]. Thus, Bampis et al. [1] proved that the $P|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$ problem does not possess a polynomial time approximation algorithm with ratio guarantee better than $(1 + \frac{1}{c+3})$, unless $\mathcal{P} = \mathcal{NP}$.

Remark: Notice that in the case of an unbounded number of processors ($\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$), the complexity to an associated decision problem is unknown.

1.2 Approximation Results

The problems with unitary communication delay. The best known approximation algorithm for $\bar{P}|prec; c_{ij} = 1; p_i = 1|C_{max}$ is due to Munier and König [10]. They presented a $(4/3)$ -approximation algorithm for this problem,

which is based on an integer linear programming formulation. The algorithm is based on the following procedure: an integrity constraint is relaxed, and feasible schedule is produced by rounding.

Munier and Hanen [9] proposed a $(\frac{7}{3} - \frac{4}{3m})$ -approximation algorithm for the problem $P|prec; c_{ij} = 1; p_i = 1|C_{max}$. They define and study a new list scheduling approximation algorithm based on the solution given on an unrestricted number of processors. They introduce the notion of *favourite successor* in order to define priorities between conflicting successors of a task. Note that, if we consider large communication delays, there is no ρ -polynomial time approximation algorithm known, except the trivial bound $(c + 1)$, one whose first step consists in executing the tasks and second step in initiating communication phasis and so on . . .

Concerning the case of a restricted number of processors, an only (as known) constant 2-approximation algorithm is given by Munier [8], for the special case where the precedence graph is tree in presence of large communication delays.

The problems with large communication delay. Contrary to the complexity results, as we know, an unique approximation algorithm is given by Rapine [12]. The author gives the lower bound $O(c)$ for the list scheduling in presence of large communication delays.

1.3 Presentation of the Paper

The challenge is to determinate a threshold for approximation algorithm for the problem $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$, to develop a non trivial approximation algorithm, and to improve, in the presence of a restricted number of processors, the bound given by Rapine [12].

This article is organized as follows: in the second section, we give a preliminary result. In the third section, we give the non-approximability result for the scheduling problem with the objective function of minimizing the length of the schedule. In the last section, we develop a $\frac{2(c+1)}{3}$ -approximation algorithm based on the notion of expansion of the makespan of a good feasible schedule.

2 Preliminary Result

In this part, we will define a variant of SAT problem [3], denoted in the following by Π_1 . The \mathcal{NP} -completeness of the scheduling problem $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$ (see section 3), is based on a reduction from this problem.

The problem Π_1 is a variant of the well known SAT problem [3]. We will call this variant the *One-in-(2, 3)SAT(2, $\bar{1}$)* problem. We denote by \mathcal{V} , the set of variables. Let n be a multiple of 3 and let \mathcal{C} be a set of clauses of cardinality 2 or 3. There are n clauses of cardinality 2 and $n/3$ clauses of cardinality 3 so that:

- each clause of cardinality 2 is equal to $(x \vee \bar{y})$ for some $x, y \in \mathcal{V}$ with $x \neq y$.
- each of the n literals x (resp. of the literals \bar{x}) for $x \in \mathcal{V}$ belongs to one of the n clauses of cardinality 2, thus to only one of them.

- each of the n literals x belongs to one of the $n/3$ clauses of cardinality 3, thus to only one of them.
- whenever $(x \vee \bar{y})$ is a clause of cardinality 2 for some $x, y \in \mathcal{V}$, then x and y belong to different clauses of cardinality 3.

Question: Is there a truth assignment $I : \mathcal{V} \rightarrow \{0, 1\}$ such that every clause in \mathcal{C} has exactly a true literal?

Example The following logic formula is a valid instance of Π_1 :

$$(x_0 \vee x_1 \vee x_2) \wedge (x_3 \vee x_4 \vee x_5) \wedge (\bar{x}_0 \vee x_3) \wedge (\bar{x}_3 \vee x_0) \wedge (\bar{x}_4 \vee x_2) \wedge (\bar{x}_1 \vee x_4) \wedge (\bar{x}_5 \vee x_1) \wedge (\bar{x}_2 \vee x_5).$$

The answer to Π_1 is *yes*. It suffices to choose $x_0 = 1, x_3 = 1$ and $x_i = 0$ for $i = \{1, 2, 4, 5\}$. This yields a truth assignment satisfying the formula, and there is exactly one true literal in every clause. For the proof of the \mathcal{NP} -completeness see [4].

3 Non-approximability Results

In this section, we show in the first part, that the problem denoted by $\bar{P}|prec; c_{ij} = c \geq 3; p_i = 1|C_{max}$ cannot be approximated by a polynomial time approximation algorithm with ratio guarantee better than $1 + \frac{1}{c+4}$ for the minimization of the length of the schedule.

3.1 The Minimization of Length of the Schedule

Theorem 1. *The problem of deciding whether an instance of $\bar{P}|prec; c_{ij} = c; p_i = 1|C_{max}$ has a schedule of length at most $(c + 4)$ is \mathcal{NP} -complete with $c \geq 3$.*

Proof. It is easy to see that $\bar{P}|prec; c_{ij} = c; p_i = 1|C_{max} = c + 4 \in \mathcal{NP}$.

Our proof is based on a reduction from Π_1 . Given an instance π^* of Π_1 , we construct an instance π of the problem $\bar{P}|prec; c_{ij} = c; p_i = 1|C_{max} = c + 4$, in the following way:

Remark: n designs the number of variables of π^* .

1. For all $x \in \mathcal{V}$, we introduce $(c + 6)$ variables-tasks: $\alpha_{x'\bar{x}'}, x', \bar{x}', \hat{x}', \beta_j^x$ with $j \in \{1, 2, \dots, c + 2\}$. We add the precedence constraints: $\alpha_{x'\bar{x}'} \rightarrow x', \alpha_{x'\bar{x}'} \rightarrow \bar{x}', \beta_1^x \rightarrow \hat{x}', \beta_1^x \rightarrow \bar{x}', \beta_j^x \rightarrow \beta_{j+1}^x$ with $j \in \{1, 2, \dots, c + 1\}$.
2. For all clauses of length three denoted by $C_i = (y \vee z \vee t)$, we introduce $2 \times (2 + c)$ clauses-tasks C_j^i and $A_j^i, j \in \{1, 2, \dots, c + 2\}$, with precedence constraints: $C_j^i \rightarrow C_{j+1}^i$ and $A_j^i \rightarrow A_{j+1}^i, j \in \{1, 2, \dots, c + 1\}$. We add the constraints $C_1^i \rightarrow l$ with $l \in \{y', z', t'\}$ and $l \rightarrow A_{c+2}^i$ with $l \in \{\hat{y}', \hat{z}', \hat{t}'\}$.
3. For all clauses of length two denoted by $C_i = (x \vee \bar{y})$, we introduce $2(c + 3)$ clauses-tasks D_j^i (resp. $D_j'^i$), $j \in \{1, 2, \dots, c + 3\}$ with precedence constraints: $D_j^i \rightarrow D_{j+1}^i$ (resp. $D_j'^i \rightarrow D_{j+1}'^i$) with $j \in \{1, 2, \dots, c + 2\}$ and $x' \rightarrow D_{c+3}^i$ (resp. $\bar{y}' \rightarrow D_{c+3}'^i$).

- Each variable-task l' is executed on a processor of P_C at slot 3 or on a processor of P_D at slot $(c + 2)$ or $(c + 3)$,
- Each variable-task \bar{l}' is executed on a processor of P_β at slot 3 or on a processor of P_D at slot $(c + 2)$ or $(c + 3)$,
- Each variable-task \hat{l}' is executed on a processor of P_β at slot 2 or 3 or on a processor of P_A at slot $(c + 2)$ or $(c + 3)$,
- The variables-tasks \bar{l}' and \hat{l}' cannot be executed together on a processor of P_β (they have a common predecessor).

Notation and property: For each $l \in \mathcal{V}$, we can associate the three tasks l' , \bar{l}' , \hat{l}' . We denote by $X = \{l' | l \in \mathcal{V}\}$, $\bar{X} = \{\bar{l}' | l \in \mathcal{V}\}$ and $\hat{X} = \{\hat{l}' | l \in \mathcal{V}\}$ three sets of tasks. For each subset A of \bar{X} (resp. \hat{X}), we can associate a subset B of X in the following way: $l' \in B$ if and only if $\bar{l}' \in A$ (resp. $\hat{l}' \in A$).

Let be the following sets: $X_1 = \{l' \setminus \pi(l') = \pi(P_C)\}$ where $\pi(l')$ (resp. $\pi(P_C)$) designs the processor on which the task l' is scheduled, $X_2 = \{l' \setminus \pi(l') = \pi(P_D)\}$, $X_3 = \{\bar{l}' \setminus \pi(\bar{l}') = \pi(P_\beta)\}$, $X_4 = \{\bar{l}' \setminus \pi(\bar{l}') = \pi(P_D)\}$, $X_5 = \{\hat{l}' \setminus \pi(\hat{l}') = \pi(P_\beta)\}$, $X_6 = \{\hat{l}' \setminus \pi(\hat{l}') = \pi(P_A)\}$. Let be $x_i = |X_i|$ for $i \in \{1, \dots, 6\}$.

We can stem from the construction of an instance of the scheduling problem the following table,

		P_C	P_β	P_A	P_D
x'		X_1			X_2
\bar{x}'			X_3		X_4
\hat{x}'			X_5	X_6	

From the previous table, using the variable x_i , we obtain the following in-equations system: $x_1 + x_2 = n(1)$, $x_3 + x_4 = n(2)$, $x_5 + x_6 = n(3)$, $x_1 \leq \frac{n}{3}(4)$, $x_6 \leq \frac{2n}{3}(5)$, $x_3 + x_5 \leq n(6)$, $x_2 + x_4 \leq n(7)$.

We will give some details about the previous system:

- For the equations (1), (2) and (3): We must execute all the tasks of the sets X , \bar{X} and \hat{X} .
- For the equation (4), on the processor which executes the path C_j^i of the clause $C_i = (y \vee z \vee t)$, we can execute at most one of the three variables-tasks y' , z' , t' . Indeed, all variables-tasks l' as a successor which is executed on a processor of P_D . If it is executed on the processor which scheduled the tasks from the path P_C it cannot be executed before the slot 3 and so, the variable-task $\alpha_{l'\bar{l}'}$ must be executed on the same processor which becomes saturated. So, we have $|X_3| < |P_C|$.
- For the equation (5), each processor of the paths P_A has two free slots and $|P_A| = \frac{n}{3}$.
- For the equation (6), all the variables-tasks \bar{l}' or \hat{l}' which are executed on a processor of the path P_β must be finished before slot 3 (it has a successor executed on another processor). So the variable-task $\alpha_{l'\bar{l}'}$ must be executed on the same processor which becomes saturated. Therefore, at the most one task between the variables-tasks \bar{l}' and \hat{l}' can be executed on a processor of the path P_β and so, $|X_3| + |X_5| \leq |P_\beta|$.

- For the equation (7), it is clear that, $|P_D| = n$ and there is at the most one free slot on each processor of P_D .

On the one hand, we have $x_3 + x_5 = n$ (indeed, we have $x_3 + x_4 + x_5 + x_6 = 2n$ and $x_6 \leq \frac{2n}{3}$, $x_4 \leq \frac{n}{3}$, so $x_3 + x_5 \geq n$) and on the other hand, $\forall l'$ only one variable-task between the variables-tasks \bar{l}' and \hat{l}' can be executed on a processor of P_β , thus we obtain $X_3 \cap X_5 = \emptyset$. Consequently, we have $X_3 \cup X_5 = X$. As the set X_4 (resp. X_6) is the complementary of the set X_3 (resp. X_5) we have $X_4 \cup X_6 = X$. Moreover, if the variable-task l' is executed on a processor of P_C then the variable-task $\alpha_{l'\bar{l}'}$ is executed on the same processor. Thus, the variable-task \bar{x}' cannot be executed before the slot $(c + 2)$, thus it is executed on a processor of P_D . We can deduce that $X_1 = X_4$ (the two sets are the same cardinality). Finally, we have $X_1 \cup X_2 = X$, $X_3 \cup X_4 = X$, $X_5 \cup X_6 = X$, $X_4 \cup X_6 = X$, $X_3 \cup X_5 = X$, $X_1 = X_4$ and therefore $X_1 = X_4 = X_5$ and $X_2 = X_3 = X_6$.

We can deduce from the previous equations that $x_1 = x_4 = x_5 = \frac{n}{3}$ and $x_2 = x_3 = x_6 = \frac{2n}{3}$.

So, if we affect the value “true” to the variable l iff the variable-task l' is executed on a processor of P_C it is trivial to see that in the clause of length 3 we have one and only one literal equal to “true”.

Let be $c = (x \vee \bar{y})$, a clause of length 2.

- If $x' \in X_1 \implies y' \in X_4 \implies y' \in X_1$. The first implication (resp. the second) is due to the fact that each processor of the path P_D must be saturated ($x_2 + x_4 = n$) (resp. $X_1 = X_4$). Only the literal x is “true” between the variables x and \bar{y} .
- If $x' \in X_2 \implies y' \in X_3 \implies y' \in X_2$. The first (resp. the second) implication is due to the fact that there is only one free slot on each processor executing the path P_D (resp. $X_3 = X_2$). Only the literal \bar{y} is “true” between the variables x and \bar{y} .

In conclusion, there is only one true literal per clause. *This concludes the proof of Lemma 1.*

- Conversely, we suppose that there is a truth assignment $I : \mathcal{V} \rightarrow \{0, 1\}$, such that each clause in \mathcal{C} has exactly one true literal.

Suppose that the true literal in the clause $C_i = (y \vee z \vee t)$ is t . Therefore, the variable-task t' (resp. y' and z') is processed at the slot 2 (resp. at the slot $(c+2)$) on the same processor as the path P_{C_i} (resp. as the path P_D and $P_{D'}$, where D and D' indicates a clause of length two where the variables y and z occurred). The $\frac{2n}{3}$ other variables-tasks y' not yet scheduled are executed at slot 3 on processor P_β as the variable-task $\alpha_{y'\bar{y}'}$. The variable-task \hat{t}' (resp. \hat{y}' and \hat{z}') is executed at the slot 2 (resp. $c + 2$ and $c + 3$) on a processor of the path P_β (resp. P_A). *This concludes the proof of Theorem 1.*

In the full version of this paper [5], we proved the following results:

Corollary 1. *There is no polynomial-time algorithm for the problem $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$ with performance bound smaller than $1 + \frac{1}{c+4}$ unless $\mathcal{P} \neq \mathcal{NP}$.*

Theorem 2. *There is no polynomial-time algorithm for the problem $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1 | \sum_j C_j$ with performance bound smaller than $1 + \frac{1}{2c+5}$ unless $\mathcal{P} \neq \mathcal{NP}$.*

Theorem 3. *The problem of deciding whether an instance of $\bar{P}|prec; c_{ij} = c; p_i = 1 | C_{max}$ with $c \in \{2, 3\}$ has a schedule of length at most $(c + 2)$ is solvable in polynomial time.*

4 Approximation by Expansion

4.1 Introduction, Notation and Description of the Method

Notation: We denote by σ^∞ , the UET-UCT schedule, and by σ_c^∞ the UET-LCT schedule. Moreover, we denote by t_i (resp. t_i^c) the starting time of the task i in the schedule σ^∞ (resp. in the schedule σ_c^∞).

Principle: We keep an assignment for the tasks given by a “good” feasible schedule on an unbounded number of processors σ^∞ . We proceed to an expansion of the makespan, while preserving communication delays ($t_j^c \geq t_i^c + 1 + c$) for two tasks, i and j with $(i, j) \in E$, processing on two different processors.

Let be a precedence graph $G = (V, E)$, we determinate a feasible schedule σ^∞ , for the model UET-UCT, using an $(4/3)$ -approximation algorithm proposed by Munier and König [10]. This algorithm gives a couple $\forall i \in V, (t_i, \pi)$ on the schedule σ^∞ corresponding to: t_i the starting time of the task i for the schedule σ^∞ and π the processor on which the task i is processed at t_i .

Now, we determinate a couple $\forall i \in V, (t_i^c, \pi')$ on the schedule σ_c^∞ in the following ways: The starting time $t_i^c = d \times t_i - i = \frac{(c+1)}{2} t_i$ and, $\pi = \pi'$. The justification of the expansion coefficient is given below. An illustration of the expansion is given by Figure 2.

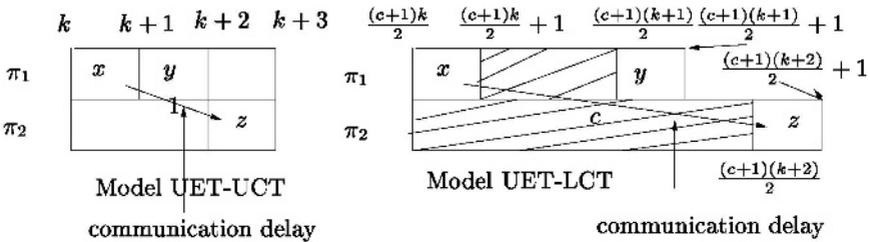


Fig. 2. Illustration of notion of an expansion

4.2 Analysis of the Method

Lemma 2. *The coefficient of an expansion is $d = \frac{(c+1)}{2}$.*

Proof. Let be two tasks i and j such that $(i, j) \in E$, which are processed on two different processors in the feasible schedule σ^∞ . We are interested in having a coefficient d such that $t_i^c = d \times t_i$ and $t_j^c = d \times t_j$. After an expansion, in order to

respect the precedence constraints and the communication delays we must have $t_j^c \geq t_i^c + 1 + c$, and so $d \times t_i - d \times t_j \geq c + 1$, $d \geq \frac{c+1}{t_i-t_j}$, $d \geq \frac{c+1}{2}$. It is sufficient to choose $d = \frac{(c+1)}{2}$.

Lemma 3. *An expansion algorithm gives a feasible schedule for the problem denoted by $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$.*

Proof. It sufficient to check that the solution given by an expansion algorithm produces a feasible schedule for the model UET-LCT. Let be two tasks i and j such that $(i, j) \in E$. We denote by π_i (resp. π_j) the processor on which the task i (resp. the task j) is executed in the schedule σ^∞ . Moreover, we denote by π'_i (resp. π'_j) the processor on which the task i (resp. the task j) is executed in the schedule σ_c^∞ . Thus,

- If $\pi_i = \pi_j$ then $\pi'_i = \pi'_j$. Since the solution given by Munier and König [10] gives a feasible schedule on the model UET-UCT, then we have $t_i + 1 \leq t_j$, $\frac{2}{c+1}t_i^c + 1 \leq \frac{2}{c+1}t_j^c$; $t_i^c + 1 \leq t_i^c + \frac{c+1}{2} \leq t_j^c$.
- If $\pi_i \neq \pi_j$ then $\pi'_i \neq \pi'_j$. We have $t_i+1+1 \leq t_j$, $\frac{2}{c+1}t_i^c+2 \leq \frac{2}{c+1}t_j^c$; $t_i^c+(c+1) \leq t_j^c$.

Theorem 4. *An expansion algorithm gives a $\frac{2(c+1)}{3}$ -approximation algorithm for the problem $\bar{P}|prec; c_{ij} = c \geq 2; p_i = 1|C_{max}$.*

Proof. We denote by C_{max}^h (resp. C_{max}^{opt}) the makespan of the schedule computed by the Munier and König (resp. the optimal value of a schedule σ^∞). In the same way we denote by C_{max}^{h*} (resp. $C_{max}^{opt,c}$) the makespan of the schedule computed by our algorithm (resp. the optimal value of a schedule σ_c^∞).

We know that $C_{max}^h \leq \frac{4}{3}C_{max}^{opt}$. Thus, we obtain $\frac{C_{max}^{h*}}{C_{max}^{opt,c}} = \frac{(c+1)}{2} \frac{C_{max}^h}{C_{max}^{opt,c}} \leq \frac{(c+1)}{2} \frac{C_{max}^h}{C_{max}^{opt,c}} \leq \frac{(c+1)}{2} \frac{4}{3} \frac{C_{max}^{opt}}{C_{max}^{opt,c}} \leq \frac{2(c+1)}{3}$.

Remark: this expansion method can be used for another problems.

5 Conclusion

In this paper, we first proved the problem of deciding whether an instance of $\bar{P}|prec; c_{ij} = c \geq 3; p_i = 1|C_{max}$ has a schedule of length at most $(c + 4)$ is \mathcal{NP} -complete. This result is to be compared with the result of [7] (resp. [1]), which states that $\bar{P}|prec; c_{ij} = 1; p_i = 1|C_{max} = 6$ (resp. $P|prec; c_{ij} = c \geq 3; p_i = 1|C_{max} = c + 3$) is \mathcal{NP} -complete. Our result implies that there is no ρ -approximation algorithm with $\rho < 1 + \frac{1}{c+4}$, unless $\mathcal{P} = \mathcal{NP}$. Secondly, we also propose a $\frac{2(c+1)}{3}$ -approximation algorithm based on the notion of expansion. In the full version [5], we show that there is no hope of finding a ρ -approximation algorithm with ρ strictly less than $\rho < 1 + \frac{1}{2c+5}$ for the problem of the minimization of the sum of the completion time. We established that the problem of

deciding whether an instance of $\bar{P}|prec; c_{ij} = c; p_i = 1|C_{max}$ with $c \in \{2, 3\}$ has a schedule of length at most $(c + 2)$ is solvable in polynomial time.

Remark: We conjecture that the problem of deciding whether an instance of $\bar{P}|prec; c_{ij} = c; p_i = 1|C_{max}$ with $c \geq 2$ has a schedule of length at most $(c + 3)$ is solvable in polynomial time.

References

1. E. Bampis, A. Giannakos, and J.C. König. On the complexity of scheduling with large communication delays. *European Journal of Operation Research*, 94:252–260, 1996.
2. B. Chen, C.N. Potts, and G.J. Woeginger. A review of machine scheduling: complexity, algorithms and approximability. Technical Report Woe-29, TU Graz, 1998.
3. M.R. Garey and D.S. Johnson. *Computers and Intractability, a Guide to the Theory of \mathcal{NP} -Completeness*. Freeman, 1979.
4. R. Giroudeau. *L'impact des délais de communications hiérarchiques sur la complexité et l'approximation des problèmes d'ordonnancement*. PhD thesis, Université d'Évry Val d'Essonne, 2000.
5. R. Giroudeau, J.C. König, F.K. Moulai, and J. Palaysi. Complexity and approximation for the precedence constrained scheduling problem with large communications delays. Technical Report 11903, Laboratoire d'Informatique, de Robotique et Microélectronique de Montpellier, 2005.
6. R.L. Graham, E.L. Lawler, J.K. Lenstra, and A.H.G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling theory: a survey. *Ann. Discrete Math.*, 5:287–326, 1979.
7. J.A. Hoogeveen, J.K. Lenstra, and B. Veltman. Three, four, five, six, or the complexity of scheduling with communication delays. *O. R. Lett.*, 16(3):129–137, 1994.
8. A. Munier. Approximation algorithms for scheduling trees with general communication delays. *Parallel Computing*, 25(1):41–48, January 1999.
9. A. Munier and C. Hanen. An approximation algorithm for scheduling unitary tasks on m processors with communication delays. Non publié, 1996.
10. A. Munier and J.C. König. A heuristic for a scheduling problem with communication delays. *Operations Research*, 45(1):145–148, 1997.
11. C. Picouleau. New complexity results on scheduling with small communication delays. *Discrete Applied Mathematics*, 60:331–342, 1995.
12. C. Rapine. *Algorithmes d'approximation garantie pour l'ordonnancement de tâches, Application au domaine du calcul parallèle*. PhD thesis, Institut National Polytechnique de Grenoble, 1999.
13. V.J. Rayward-Smith. UET scheduling with unit interprocessor communication delays. *Discr. App. Math.*, 18:55–71, 1987.