



HAL
open science

How the structure of precedence constraints may change the complexity class of scheduling problems

D Prot, Odile Bellenguez-Morineau

► To cite this version:

D Prot, Odile Bellenguez-Morineau. How the structure of precedence constraints may change the complexity class of scheduling problems. 2017. hal-01445300

HAL Id: hal-01445300

<https://hal.science/hal-01445300>

Preprint submitted on 24 Jan 2017

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

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

How the structure of precedence constraints may change the complexity class of scheduling problems.

D. Prot · O. Bellenguez-Morineau

Received: date / Accepted: date

Abstract This survey aims at demonstrating that the structure of precedence constraints plays a tremendous role on the complexity of scheduling problems. Indeed many problems can be \mathcal{NP} -hard when considering general precedence constraints, while they become polynomially solvable for particular precedence constraints. Add to this, the existence of many very exciting challenges in this research area is underlined.

Keywords Scheduling · Precedence constraints · Complexity

1 Introduction

Precedence constraints play an important role in many real-life scheduling problems. For example, when considering the scheduler of a computer, some operations have to be finished before some others begin. Other classical examples can be found in the book of Pinedo (Pinedo [2012]). In the most general cases, precedence constraints can be represented by an arbitrary directed acyclic graph. Nevertheless, in some cases, it is possible for precedence constraints to take a particular form. For example, if a problem includes only precedence constraints related to assembling steps, precedence constraints can be represented by a

D. Prot · O. Bellenguez-Morineau

Ecole des Mines de Nantes, 4 rue Alfred Kastler, B.P. 20722 F-44307 Nantes Cedex 3. France

E-mail: dprot@talend.com, odile.morineau@mines-nantes.fr

particular directed acyclic graph called intree. The fact that the precedence graph takes a particular form may transform the complexity of the problem. Most of the time, adding precedence constraints makes problem harder, since the empty graph is included in many graph classes. Yet, it may make the problem easier if it is not the case, for example for the class of graphs of bounded width. For this reason, the idea of this survey is to consider the complexity results in scheduling theory according to the structure of precedence constraints.

We assume that the reader is conversant with the theory of \mathcal{NP} -completeness, otherwise the book by Garey and Johnson [1979] is a good entry point. We will discuss, in the conclusion of this paper, the parameterized complexity. The reader can refer to the book of Downey and Fellows [2012] if needed. Lenstra and Rinnooy Kan [1978] offer a large set of complexity results for scheduling problems with precedence constraints. Graham et al. [1979] and Lawler et al. [1993] propose two surveys of complexity results for scheduling problems, but they are not necessarily focused on precedence constraints. For complexity results with the most classical precedence constraints (i.e., chains, trees, series-parallel and arbitrary precedence constraints), the reader can refer to the book by Brucker [2007], and/or to the websites Dürr [2016] and Brucker and Knust [2009]. Note that this website has been recently updated in order to include most of the results discussed in this article. Möhring [1989] proposes a very interesting survey dedicated to specific partial ordered sets, and studies their structure. Some applications to scheduling are also presented. Since this survey was conducted, many results arised in scheduling theory for specific precedence graphs, we hence believe that a new survey would be beneficial for the scheduling community. We restrict our field on purpose to complexity results and do not talk about approximations results (Williamson and Shmoys [2011]), despite the fact that many results arised in this field recently in scheduling theory, such as Svensson [2011] and Levey and Rothvoss [2016]. We believe that it is important to limit the scope of the survey, in order to be as complete as possible in a given field.

The paper is organized as follows: In Section 2, we introduce all the specific types of precedence constraints that will be studied in this paper, while scheduling notations are recalled in Section 3.

Section 4 is dedicated to single-machine scheduling problems, and Sections 5 and 6 are respectively dedicated to non-preemptive and preemptive parallel machines scheduling problems. Each of these three sections is based on the following structure: we first give the polynomial results, then the \mathcal{NP} -hard cases and last the most interesting open problems. Finally in Section 7 we give some concluding remarks.

2 Special types of precedence constraints

In this section we introduce, for the sake of completeness, the special types of precedence constraints that have already been studied in scheduling theory. Precedence constraints between jobs are easily modelled by a directed acyclic graph and we will use it as long as possible. Nevertheless, precedence relations may also be treated as partially ordered set (poset) in the terminology of order theory, and some definitions at the end of this section are much easier to understand from this point of view.

First, let us recall some basic graph theory definitions. Let $G = (X, A)$ be a directed acyclic graph, where X denotes the set of vertices and A , the set of arcs.

A directed acyclic graph is a collection of *chains* if each vertex has at most one successor and at most one predecessor. An *inforest* (resp. *outforest*) is a directed acyclic graph where each vertex has at most one successor (resp. predecessor). An *intree* (resp. *outtree*) is a connected inforest (resp. outforest). We call *forest* (resp. *tree*) a graph that is either an inforest or an outforest (resp. either an intree or an outtree). An *opposing forest* is a collection of inforests and outforests.

For each vertex $x \in X$, we can compute its *level* $h(x)$ which corresponds to the longest directed path starting from x in G . The height of a DAG, denoted $h(G)$, corresponds to the number of levels - 1 in this graph, as illustrated with Figure 1. Directed acyclic graphs of *bounded height* correspond to directed acyclic graphs where the height is bounded by a constant.

Definition 1 (Level order graph) *A directed acyclic graph is a level order graph if each vertex of a given level l is a predecessor of all the vertices of level $l - 1$ (see Figure 2).*

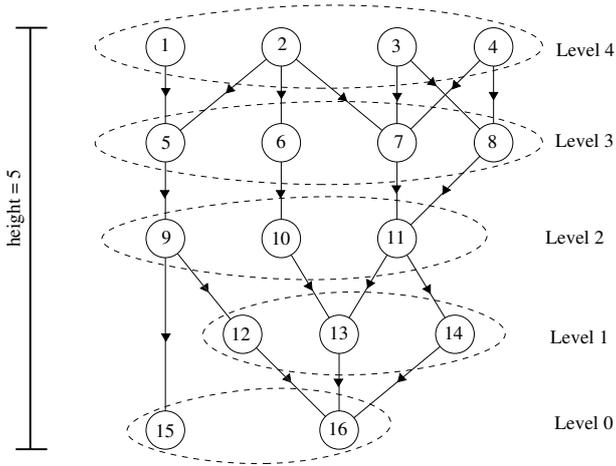


Fig. 1 The height of a DAG

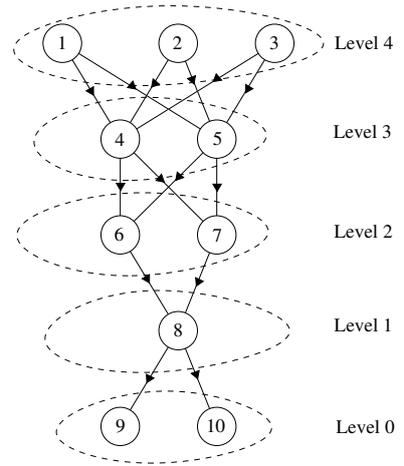


Fig. 2 A level order graph

Series-parallel graphs (sp-graph) are defined in many ways, we use the inductive definition of Lawler [1978].

Definition 2 (sp-graph) *A graph consisting of a single vertex is a sp-graph. Given two sp-graphs $G_1 = (X_1, A_1)$ and $G_2 = (X_2, A_2)$, the graph $G = (X_1 \cup X_2, A_1 \cup A_2)$ is a sp-graph (this is called parallel composition). Given two sp-graphs $G_1 = (X_1, A_1)$ and $G_2 = (X_2, A_2)$, the graph $G = (X_1 \cup X_2, A_1 \cup A_2 \cup X_1 \times X_2)$ is a sp-graph (this is called series composition).*

An example is given in Figure 3. Note that sp-graph can also be defined by the forbidden subgraph of Figure 4.

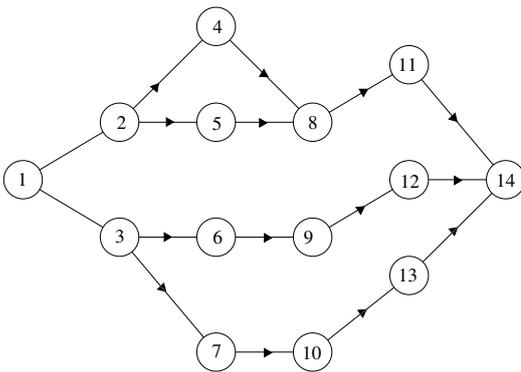


Fig. 3 A series-parallel graph

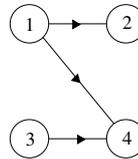


Fig. 4 The forbidden subgraph for sp-graphs

A divide-and-conquer-graph (DC-graph) is a special sp-graph built using symmetries. It can be used to model divide-and-conquer algorithms (for example binary search, merge sort, number multiplication...) and is formally defined as follows:

Definition 3 (DC-graph) *A single vertex is a DC-graph; given two vertices s and t and k DC-graphs $(X_1, A_1), \dots, (X_k, A_k)$, the graph $(\cup_{k=1}^n X_k \cup \{s\} \cup \{t\}, \cup_{k=1}^n (A_k \cup (s \times X_k) \cup (X_k \times t)))$ is a DC-graph.*

This definition is illustrated with Figure 5.

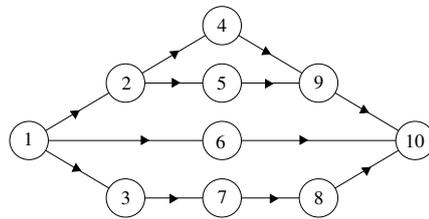


Fig. 5 A DC-graph

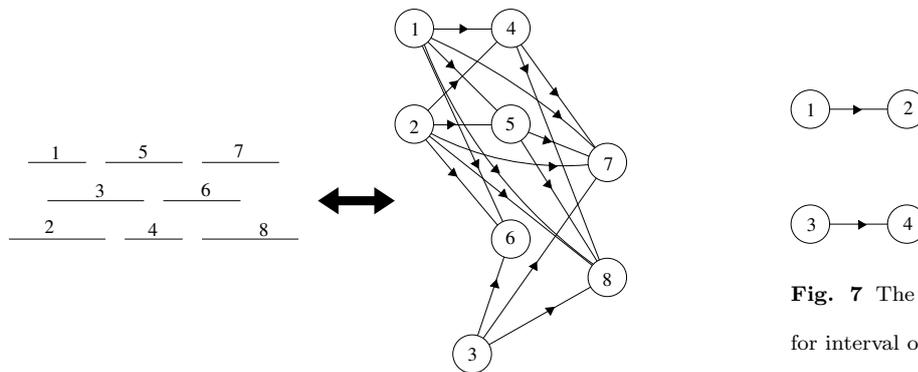


Fig. 6 A set of real intervals and the corresponding interval order graph

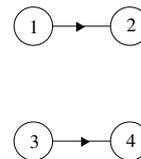


Fig. 7 The forbidden subgraph for interval order graphs

Definition 4 (Interval order graph) *The interval order for a collection of intervals on the real line is the partial order corresponding to their left-to-right precedence relation, so that the interval order graph is the Hasse diagram of an interval order. such that for any two vertices x and y , $(x, y) \in A \iff e_x \leq s_y$.*

An example is given in Figure 6. Papadimitriou and Yannakakis [1979] show that interval order graphs can also be defined by the forbidden induced subgraph presented in Figure 7. Larger classes of graphs were defined by forbidden subgraphs, such as *quasi-interval order graphs* and *over-interval order graphs* (respectively in Moukrim [1999] and Chardon and Moukrim [2005]), the quasi-interval order graphs being strictly included in over-interval order graphs. The corresponding forbidden subgraphs are drawn in Figures 8 and 9.

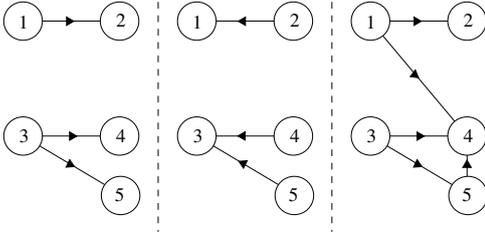


Fig. 8 The three forbidden subgraphs for quasi-interval order graphs

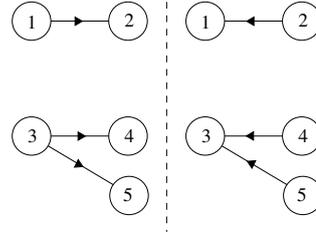


Fig. 9 The two forbidden subgraphs for over-interval order graphs

To ease the reading, the following definitions will be given within the order theory paradigm. We will hence talk about a partial order set $\mathcal{P} = (X, \preceq_{\mathcal{P}})$ rather than a directed acyclic graph $G = (X, A)$ to describe the precedence graph, and a partial order $\preceq_{\mathcal{P}}$ corresponds to the precedence constraints.

Definition 5 (Antichain) *Given a partial order set $\mathcal{P} = (X, \preceq_{\mathcal{P}})$, an antichain is a subset S of X such that any two elements of S are incomparable.*

Definition 6 (Width) *Given a poset $\mathcal{P} = (X, \preceq_{\mathcal{P}})$, the width of a poset is the size of a maximum antichain.*

By extension, for a given directed acyclic graph $G = (X, A)$ we define the width of the graph to be the width of the corresponding poset, denoted by $w(G)$.

The \mathcal{A}_m -order (first introduced in Moukrim and Quilliot [1997]) contains the over-interval order for any integer $m \geq 2$ and is defined in the following way:

Definition 7 (\mathcal{A}_m -order) *Let $\mathcal{P} = (X, \preceq_{\mathcal{P}})$ be a poset. For any two antichains A and B of size at most m , let us define the four sets : $\max(A, B) = \{x \in A \cup B | \exists y \in A \cup B, y \preceq_{\mathcal{P}} x\}$, $\min(A, B) = \{x \in A \cup B | \exists y \in A \cup B, x \preceq_{\mathcal{P}} y\}$, $\overline{\max}(A, B) = (A \cap B) \cup \max(A, B)$, and $\overline{\min}(A, B) = (A \cap B) \cup \min(A, B)$. $\preceq_{\mathcal{P}}$ is an \mathcal{A}_m order if and only if there do not exist two antichains A and B of size m at most, such that $|\overline{\max}(A, B)| \geq m + 1$ or $|\overline{\min}(A, B)| \geq m + 1$.*

We call \mathcal{A}_m -order graph a directed acyclic graph for which the set of arcs corresponds to an \mathcal{A}_m -order.

Definition 8 (Linear extension) *Given a partial order $\preceq_{\mathcal{P}}$ over a set X , a linear extension of $\preceq_{\mathcal{P}}$ over X is a total order respecting $\preceq_{\mathcal{P}}$.*

Definition 9 (Dimension) *The dimension of a poset $\mathcal{P} = (X, \preceq_{\mathcal{P}})$ is the minimum number t of linear extensions $\preceq_1, \dots, \preceq_t$ such that $x \preceq_{\mathcal{P}} y \iff \forall l \in 1..t, x \preceq_l y$. In other words, if $x \parallel y$ (x and y are incomparable in $\preceq_{\mathcal{P}}$), then there are at least two linear extensions, one with $x \preceq y$ and another one with $y \preceq x$.*

An interesting point is that series-parallel graphs are strictly included in directed acyclic graphs of dimension 2.

The fractional dimension of a poset is extending the notion of dimension (see Brightwell and Scheinerman [1992]).

Definition 10 (Fractional dimension) *For any integer k , $t(k)$ denotes the minimum number of linear extensions such that for any two incomparable elements x and y , there are at least k extensions with $x \preceq y$ and k with $y \preceq x$. The fractional dimension is the limit $\{t(k)/k\}$ as k tends to $+\infty$.*

The diagram in Figure 10 provides a better overview of the existing inclusion between the different classes.

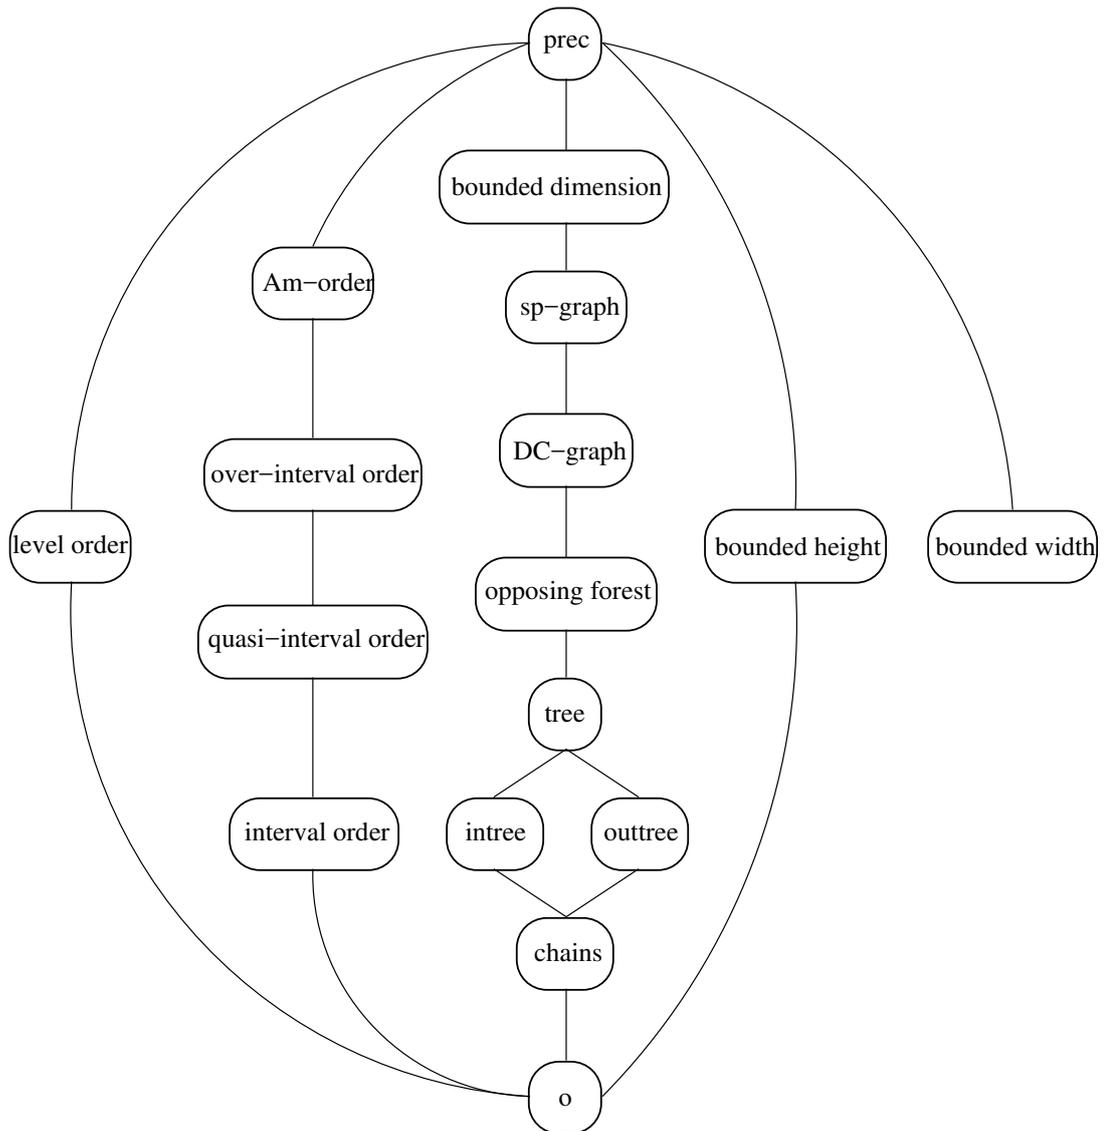


Fig. 10 Hasse diagram for different classes of directed acyclic graphs

3 Scheduling notation

In this paper, we will use the standard $\alpha|\beta|\gamma$ notation introduced in Graham et al. [1979], and updated in Brucker [2007]. We define below the different notations used all along the paper.

The α -field is used for the machine environment. $\alpha = 1$ corresponds to a single machine problem; if $\alpha = Pm$, there is a fixed number m of identical parallel machines. If this number is arbitrary, it is noted $\alpha = P$. Similarly, if $\alpha = Qm$ (resp. $\alpha = Q$), it corresponds to a fixed (resp. arbitrary) number of uniform parallel machines, i.e., each machine i has a speed s_i and the processing time of a job j on machine i is equal to p_j/s_i .

The field $\beta \subset \{\beta_1, \beta_2, \beta_3, \beta_4\}$ describes the jobs characteristics, the possible entries that we will deal with are the following ones:

- $\beta_1 \in \{pmtn, \circ\}$: *pmtn* means that preemption is allowed, i.e., a job may be interrupted and finished later. If $\beta_1 = \circ$, preemption is forbidden.

- β_2 describes the precedence constraints. If $\beta_2 = \circ$, there is no precedence constraint, whereas $\beta_2 = prec$ means that the precedence graph is a general directed acyclic graph. This field can take many values according to the structure of the directed acyclic graph, as presented in the previous section.

For sake of completeness, all the acronyms are recalled hereinafter: *chains*, *intree*, *outtree*, *opp.forest* (opposing forest), *io* (interval order graph), *qio* (quasi-interval order graph), *oio* (over-interval order graph), *sp-graph* (series-parallel graph), *DC-graph* (divide-and-conquer graph), *lo* (level order graph), $h(G) \leq k$ (directed acyclic graph with height bounded by k), $dim \leq k$ (directed acyclic graph with dimension bounded by k), $fdim < k$ (directed acyclic graph with fractional dimension bounded by k), $w(G) \leq k$ (directed acyclic graph with width bounded by k).

It appears there that classical literature often abuses of terms *intree* and *outtree* to handle in fact *inforest* and *outforest*, that can be mixed in *opp.forest* for some problems.

- $\beta_3 \in \{r_j, \circ\}$: if $\beta_3 = r_j$, each job j has a given release date. If $\beta_3 = \circ$, the release date is 0 for each job.

– β_4 represents the processing time of a job j . If $\beta_4 = \circ$, there is one processing time p_j for each job j . If $p_j = p$, all the jobs have the same processing time p . When $p_j = 1$, we use the acronym UET which stands for Unit Execution Time. In some cases, the processing time may increase or decrease with either the position of the job or the starting time of the job. If the job is in position r on a machine, the processing time will be noted $p_j^{[r]}$. If the processing time is depending of the starting time t of the job, it will be written $p_j^{(t)}$.

The γ -field is related to the objective function of the problem. Let C_j be the completion time of job j . The *makespan* is defined by $C_{\max} = \max_j C_j$ and the *total completion time* by $\sum_j C_j$. It is clear that, in general, makespan and total completion time are not equivalent. Nevertheless, for some problems, we can show that there exists an *ideal* schedule, which both minimizes makespan C_{\max} and total completion time $\sum C_j$. Due-date related objectives are also studied; if d_j is the due date of job j , then the *lateness* of job j is defined by $L_j = C_j - d_j$. The *tardiness* is $T_j = \max(0, C_j - d_j)$ and the *unit penalty* U_j is equal to one if $C_j > d_j$ and to zero otherwise. We then can define the *maximum lateness* $L_{\max} = \max L_j$, the *total tardiness* $\sum T_j$, the *total number of late jobs* $\sum U_j$. A weight w_j may also exist for each job j , leading to the corresponding objective functions: the *total weighted completion time* $\sum w_j C_j$, the *total weighted tardiness* $\sum w_j T_j$ and the *total weighted number of late jobs* $\sum w_j U_j$. All the functions presented so far are *regular* functions, i.e., they are non-decreasing with C_j .

4 Single machine problems

For single machine problems, most of the interesting results for this survey are related to the total weighted completion time $\sum w_j C_j$. It is mainly due to the fact that $1|chains, p_j = 1| \sum U_j$ (see Lenstra and Rinnooy Kan [1980]) and $1|chains, p_j = 1| \sum T_j$ (see Leung and Young [1990]) are already \mathcal{NP} -hard, while $1|prec, r_j|C_{\max}$ is solvable in polynomial time. Nevertheless, as we will see in Section 4.3, there are some interesting open problems for other criteria when considering preemption.

4.1 Polynomial cases

Lawler [1978] uses Sidney's theory (Sidney [1975]) to derive a polynomial-time algorithm to solve problem $1|sp-graph|\sum w_i C_i$. The most important results are based on the concept of a module in a precedence graph $G = (N, A)$: A non-empty subset $M \subset N$ is a module if, for each job $j \in N - M$, exactly one of the three conditions holds: 1. j must precede every job in M , 2. j must follow every job in M , 3. j is not constrained to any job in M . Using this concept leads to a very powerful theorem stating that there exists an optimal sequence consistent with any optimal sequence of any module.

A large improvement on this problem has been done recently, by using order theory and by proving that the problem is a special case of the vertex cover problem. More precisely, in Correa and Schulz [2005] the authors conjecture that $1|prec|\sum w_j C_j$ is a special case of vertex cover and prove that, under this conjecture, the problem $1|prec|\sum w_j C_j$ is polynomial if the precedence graph is of dimension 2. In Ambühl and Mastrolilli [2009], the authors prove the conjecture and hence the result provided in Lawler [1978] is considerably extended (since series-parallel graphs are strictly included in DAGs of dimension 2).

Using the same methodology as in Lawler [1978], Wang et al. [2008] extend these results to the case where jobs are deteriorating, i.e., processing times are an increasing function of their starting time, and they show that $1|sp-graph, p_j^{(t)} = p_j(1 + at)|\sum w_j C_j$ and $1|sp-graph, p_j^{(t)} = p_j + \alpha_j t|C_{\max}$ can be solved in a polynomial time, where p_j , a , and α_j are positive constants and t is the starting time of the job. The same framework is used in Wang and Wang [2013] for position-dependent processing times, and the authors show that $1|sp-graph, p_j^{[r]} = a_j + b_j r|C_{\max}$ and $1|sp-graph, p_j^{[r]} = a_j - b_j r|C_{\max}$, where r is the position of the job, a_j and b_j positive constants, are polynomially solvable. Gordon et al. [2008] propose a more general framework than the one introduced in Wang et al. [2008], that was first presented in Monma and Sidney [1979]. They extend it to different models with deterioration and learning, that were aimed at minimizing either the total weighted completion time or the makespan, with series-parallel precedence constraints.

4.2 Minimum NP-hard cases

The problem $1|prec|\sum w_j C_j$ is known to be \mathcal{NP} -hard (see Lenstra and Rinnooy Kan [1978]). In order to identify for which precedence graph the problem may be polynomially solvable, we present here the minimal (with respect to precedence constraints) \mathcal{NP} -hard cases: the problem is still \mathcal{NP} -hard even if the precedence graph is:

- of indegree at most 2 (see Lenstra and Rinnooy Kan [1978]). Note that this result also stands with equal weights (i.e., $w_j = 1$)
- of bounded height (proof is straightforward by using the transformation proposed in Lawler [1978]).
- an interval order graph (see Ambühl et al. [2011]). This is a strong difference with parallel machine results for which interval order graphs provide often polynomial algorithms, as we will see in the dedicated section.
- of fractional dimension (see Definition 10) greater or equal to 3 (see Ambühl et al. [2011]).

4.3 Open problems

The problem $1|prec|\sum w_i C_i$ has been widely studied, and the boundary between polynomial and \mathcal{NP} -hard cases is globally well-defined. However, there are still boundaries to defined, as illustrated by the following examples. First, if the fractional dimension $fdim$ of the precedence graph lies in interval $]2, 3[$, the problem is open (it is polynomially solvable if $fdim \leq 2$ since the fractional dimension of a DAG is less than or equal to the dimension of this DAG and the problem is solvable in polynomial time if the precedence graph is of dimension 2). This is an interesting question but the gap is rather limited. A wider open question is when we consider the problem with equal weights, i.e., $1|prec|\sum C_i$. The problem is still \mathcal{NP} -hard, even if the precedence graph is of indegree at most 2 (see Lenstra and Rinnooy Kan [1978]). Nevertheless, we may hope that the problem becomes polynomial for precedence graphs with dimension larger than two. It is also possible that the problem is solvable in polynomial time if the

precedence graph is an interval order graph, and even for larger classes like quasi-interval order graphs and over-interval order graphs.

For the criteria related to tardiness T_j and unit penalty U_j , it is surprising to see that preemptive problems with precedence constraints have not been studied yet. More precisely, Tian et al. [2006] have shown that $1|pmtn, r_j, p_j = p|\sum T_j$ is solvable in polynomial time, but it is still open whether adding precedence constraints leads this problem to be \mathcal{NP} -hard or not; the problem remains also open when considering the more general criterion $\sum w_j T_j$. The same outline appears when considering unit penalty: $1|r_j; pmtn; p_j = p|\sum w_j U_j$ is solvable in polynomial time (with an algorithm in $O(n^{10})$ by Baptiste [1999], and in $O(n^4)$ by Baptiste et al. [2004b]); nevertheless, nothing has been shown when adding precedence constraints to the problem. The first research avenue is to study this problem including Chains as a first step and determine whether the problem is polynomial or not.

5 Parallel machines without preemption

When considering non-preemptive scheduling problems, whatever the structure of precedence graphs, we will mainly focus on problems with equal processing times, since problems $P2||C_{max}$ and $P2|chains|\sum C_j$ are already NP-hard (see Lenstra et al. [1977] and Du et al. [1991]).

5.1 Polynomial cases

Makespan criterion and arbitrary number of machines

Three seminal works on parallel machines with precedence constraints are the approaches of Hu [1961], Papadimitriou and Yannakakis [1979] and Möhring [1989] in which the authors are respectively dealing with trees, interval order graphs and graphs of bounded width.

In Hu [1961], the author proves that problem $P|tree, p_j = p|C_{max}$ is polynomially solvable by a list scheduling algorithm where the highest priority is given to the job with the highest level (this strategy is called HLF, for Highest Level First). It is unlikely to find a precedence graph that strictly includes

trees for which the problem is solvable in polynomial time, since it was proven in Garey et al. [1983] that scheduling opposing forest is \mathcal{NP} -hard. The hardness of the latter problem is mainly due to the arbitrary number of machines, since it is solvable in polynomial time for any fixed number of machines.

In Papadimitriou and Yannakakis [1979], the authors prove that problem $P|io, p_j = p|C_{\max}$ is polynomially solvable by a list scheduling algorithm where the highest priority is given to the job with the largest number of successors. This result has been improved twice; first Moukrim [1999] shows that the same algorithm gives an optimal solution if interval order graphs are replaced by quasi-interval order graphs, that properly contains the former. Chardon and Moukrim [2005] show that the same result does not stand for over-interval order graphs, but the Coffman-Graham algorithm (see Coffman Jr and Graham [1972]) can be applied to solve problem $P|oio, p_j = p|C_{\max}$ to optimality.

In Möhring [1989], the author studies the problem with bounded width, equal processing times and the makespan criterion. He shows that the problem can be solved in polynomial time by using dynamic programming on the digraph of order ideals. Middendorf and Timkovsky [1999] proposed their own approach to solve the problem $P|w(G) \leq k, r_j, p_j = 1|f$ for any regular function f . More precisely, their algorithm consists in searching a shortest path in the related transversal graph.

When adding release dates, the problem is already \mathcal{NP} -hard for intrees, yet it is solvable in polynomial time for outtrees (see Brucker et al. [1977], and Monma [1982] for a linear algorithm for the latter problem). Note that there is a strong relationship between scheduling with release dates and C_{\max} criterion and scheduling with L_{\max} , by simply looking at the schedule in the reverse way, and reversing the precedence constraints. Hence problems $P|intree, p_j = p|L_{\max}$ is polynomially solvable and $P|outtree, p_j = p|L_{\max}$ is \mathcal{NP} -hard.

Kubiak et al. [2009] recently opened new perspectives, since they show that $P|DC - graph, p_j = p|C_{\max}$ is solvable in polynomial time. They more precisely proved that the Highest Level First strategy (used in Hu [1961]) solves the problem to optimality when the precedence graph is a divide-and-conquer graph (see Definition 3).

Makespan criterion and fixed number of machines

Let us now focus on the case where the number of machines is fixed. Recall that the problem $Pm|prec, p_j = p|C_{\max}$ is still open (this problem is known as [OPEN8] in the book by Garey and Johnson [1979]) for $m \geq 3$, and it was solved for $m = 2$ in Coffman Jr and Graham [1972]. For opposing forests, Garey et al. [1983] propose an optimal polynomial algorithm (of complexity $O(n^{m^2+2m-5} \log n)$) that consists in a divide and conquer approach, that uses the HLF algorithm as a subroutine. A new algorithm with complexity $O(n^{2m-2} \log n)$ has been proposed by Dolev and Warmuth [1985], who also show that the problem is polynomially solvable for level order graphs (that are strictly included in series-parallel graphs). Dolev and Warmuth [1984] solve the case where the precedence graph is of bounded height by proposing an algorithm of time complexity $O(n^{h(m-1)+1})$. Recently, Aho and Mäkinen [2006] show that $Pm|prec, p_j = p|C_{\max}$ is solvable in polynomial time when the precedence graph is of bounded height and the maximum degree is bounded. This result is in fact a special case of the one proposed in Dolev and Warmuth [1984].

Other criteria and/or machine environment

The other well-studied criterion for parallel machine environment is the total completion time $\sum C_j$. One of the reasons for that is that, for some problems, it is equivalent to solve the total completion time and the makespan since they admit an ideal schedule. For example, ideal schedules exist when considering two machines, arbitrary precedence constraints and equal processing times, hence problem $P2|prec, p_j = p|\sum C_j$ is polynomially solvable with CG-algorithm (see Coffman Jr and Graham [1972]). When adding release dates, Baptiste and Timkovsky [2001] show that $P2|prec, r_j, p_j = 1|\sum C_j$ is solvable in polynomial time by reducing it to a shortest path problem.

For an arbitrary number of machines, if the precedence graph forms an outtree and the processing times are UET, the same result holds and hence $P|outtree, p_j = p|\sum C_j$ is solvable in polynomial time. Note that problem $P3|intree, p_j = p|\sum C_j$ is not ideal, see Huo and Leung [2006] for a counterexam-

ple. Nevertheless, for any fixed number of machines, problem $Pm|intree, p_j = p|\sum C_j$ is solvable in polynomial time (see Baptiste et al. [2004a]). Adding release dates maintains the same property: the algorithm proposed in Brucker et al. [2002] solves simultaneously problems $P|outtree, r_j, p_j = 1|C_{\max}$ and $P|outtree, r_j, p_j = 1|\sum C_j$. An improvement of this algorithm has been proposed in Huo and Leung [2005].

For interval order graphs, Möhring [1989] notices that $P|io, p_j = p|\sum C_j$ is solvable in polynomial time since the proof of the algorithm for $P|io, p_j = p|C_{\max}$ (in Papadimitriou and Yannakakis [1979]) only uses swaps between tasks, and this property is verified by the makespan and the total completion time. It has been noticed recently that the same result holds for overinterval orders (that properly contain interval orders), the problem admits also an ideal solution, so $P|oio, p_j = p|\sum C_j$ is solvable in polynomial time (see Wang [2015]).

When considering uniform parallel machines, only few results are available; problem $Q2|chains, p_j = p|C_{\max}$ is solvable in polynomial time (see Brucker et al. [1999]). If one processor is going b times faster than the other (with b an integer), the problem $Q2|tree, p_j = p|C_{\max}$ is also polynomially solvable (see Kubiak [1989]); the problem is also ideal, and hence $Q2|tree, p_j = p|\sum C_j$ is also solvable in polynomial time.

5.2 Minimum NP-hard cases

The interesting results for this survey are the \mathcal{NP} -hardness of $P|prec, p_j = p|C_{\max}$ and $P|prec, p_j = p|\sum C_j$ (see Lenstra and Rinnooy Kan [1978]). Note that the proof for the two problems also holds when the precedence graph is of bounded height, and that the problem $P|opp. forest, p_j = p|C_{\max}$ is also \mathcal{NP} -hard (see Garey et al. [1983]). When adding release dates, the corresponding problem is already \mathcal{NP} -hard for intrees, for both the makespan and the total completion time (see Brucker et al. [1977]). Finally, what may not be obvious is that even if $P||C_{\max}$ is strongly \mathcal{NP} -hard, the problem $P|w(G) \leq K|C_{\max}$ is

solvable in pseudo-polynomial time (Middendorf and Timkovsky [1999]). This results is due to the fact that the class of empty graphs is not included in bounded-width graphs class.

5.3 Open problems

For an arbitrary number of machines and the makespan criterion, the boundary between polynomially solvable and \mathcal{NP} -hard problems seems very sharp, we believe that the efforts should not be focused on these problems. When the number of machines is fixed, this boundary is much larger. Surprisingly, to the best of our knowledge, no other structures of precedence graphs than the ones introduced in previous section have been studied for problem $P_m|prec, p_j = p|C_{\max}$. In our opinion, it could be a good opportunity to work on more general precedence graphs on this problem, to be able to arbitrate if $P_m|prec, p_j = p|C_{\max}$ is solvable in polynomial time or \mathcal{NP} -hard. The most natural extension in our opinion is to consider series-parallel graphs, since it is a generalization of opposing forests, level order graphs and DC-graphs, for which the problem is polynomially solvable.

For the total completion time criterion, the two most intriguing problems are $P|intree, p_j = p|\sum C_j$ and $P|outtree, r_j, p_j = p|\sum C_j$. For the former problem, the interest lies in the fact that there exists an ideal schedule for outtree precedences, but not for intree precedences. Nevertheless, we do believe that this problem admits an optimal polynomial algorithm. For the latter problem, an algorithm exists for $p_j = 1$ (i.e., release dates are multiple of the processing time, see Brucker et al. [2002]), and hence the gap to $p_j = p$ seems small.

For a fixed number of uniform parallel machines and the makespan criterion, the set of open problems is wide, since the only polynomial algorithm is for $Q2|chains, p_j = p|C_{\max}$ (Brucker et al. [1999]), and the problem $Qm|prec, r_j, p_j = p|L_{\max}$ is still open. The same behavior occurs for the total completion time: Dessouky et al. [1990] proved that $Qm|r_j, p_j = p|\sum C_j$ is solvable in polynomial time, and problem $Qm|prec, r_j, p_j = p|\sum C_j$ is still open. We hence believe that this set of problems deserves a deeper

study. A first approach may consist in trying to adapt the algorithms available for identical parallel machines.

6 Parallel machines with preemption

Timkovsky shows very strong links between preemption and chains, including the fact that a large set of scheduling problems with preemption can be reduced to problems without preemption, with UET tasks, and where each job is replaced by a chain of jobs (see Theorem 3.5 in Timkovsky [2003]). This interesting result can be applied in many cases, but does not hold for the total completion time criterion. Moreover, according to the structure of the precedence graph, the resulting graph may not have the same structure. For example, an intree where each job is replaced by a chain of jobs remains an intree, whereas it does not hold for interval order graphs (two independent jobs will be replaced by two chains of parallel jobs, which is not an interval order graph). That is why we will examine more precisely what happens in this section.

6.1 Polynomial cases

Makespan criterion

Since $P|tree, p_j = p|C_{\max}$ is polynomially solvable (see Hu [1961]), by Timkovsky's result, so is $P|pmtn, tree|C_{\max}$.

The first polynomial algorithm for this problem is proposed in Muntz and Coffman Jr [1970] with a running time of $O(n^2)$. An algorithm in $O(n \log m)$ was then proposed in Gonzalez and Johnson [1980]. Note that the latter algorithm produces at most $O(n)$ preemptions whereas the former may obtain a schedule with $O(nm)$ preemptions. Lawler [1982] studies the case with release dates and outtree, and shows that it can be solved in $O(n^2)$ with a dynamic priority list algorithm (i.e., priorities may change according to what has already been scheduled).

Timkovsky's result can not be applied to precedence graphs such as interval order graphs. Yet, it was proven that the problem is also solvable in polynomial time for this precedence structure; first, Sauer and

Stone [1989] show it for a fixed number of machines $Pm|pmtn,io|C_{\max}$ by using a linear programming approach based on the set of jobs scheduled at each instant. Later Djellab [1999] proposes another linear program that solves the problem for an arbitrary number of machines, i.e., $P|pmtn,io|C_{\max}$. For a fixed number m of machines, Moukrim and Quilliot [2005] extend the result of Sauer and Stone [1989], by proposing an linear programming approach for \mathcal{A}_m -order graphs (which properly contain interval order graphs).

Other criteria and/or machine environment

Du et al. [1991] explain that $P2|pmtn,chains|\sum C_j$ is strongly \mathcal{NP} -hard by showing that preemption is useless for this problem, that is why we will only focus on the UET case for the total completion time criterion. Chardon and Moukrim [2005] prove that $P2|pmtn,prec,p_j=p|\sum C_j$ is solvable in polynomial time by adapting the Coffman-Graham algorithm. Their algorithm finds an ideal schedule. By proving that preemption is redundant, Baptiste and Timkovsky [2001] prove that $P2|pmtn,outtree,r_j,p_j=1|\sum C_j$ is solvable in polynomial time. When the precedence graph is an intree, Coffman Jr et al. [2012] prove that the problem is not ideal, and Chen et al. [2015] provide a deep analysis of the structure of preemption. Using the same methodology than Baptiste and Timkovsky [2001], Brucker et al. [2002] show that the problem is solvable in polynomial time with an outtree and an arbitrary number of machines. Moreover, they provide a $O(n^2)$ algorithm, that admits a $O(n \log n)$ implementation according to Huo and Leung [2005]. Lushchakova [2006] slightly improves the result of Baptiste and Timkovsky [2001] and proposes an algorithm of complexity $O(n^2)$ for the problem $P2|pmtn,outtree,r_j,p_j=p|\sum C_j$.

6.2 Minimum NP-hard cases

For the makespan criterion, Ullman [1976] shows that $P|pmtn,prec,p_j=p|C_{\max}$ is \mathcal{NP} -hard. For the total completion time criterion, if we carefully look at the \mathcal{NP} -hardness proof of $P|prec,p_j=p|\sum C_j$ in Lenstra and Rinnooy Kan [1978], we can see that preemption is useless for the instance constructed

in the reduction and hence the preemptive version is still \mathcal{NP} -hard for precedence graphs of bounded height:

Theorem 1 $P|pmtn, prec, p_j = p| \sum C_j$ is \mathcal{NP} -hard even if the precedence graph is of bounded height.

Proof. We just need to modify slightly the proof in the reduction from CLIQUE of Lenstra and Rinnooy Kan [1978].

CLIQUE : $G = (V, E)$ is an undirected graph and k an integer. Does G have a clique on k vertices?

Let us recall this reduction; we denote by v (resp. e) the number of vertices (resp. edges) of G . We also define following parameters: $l = \frac{k(k-1)}{2}$, $k' = v - k$ and $l' = e - l$ (we use the notations of the original article). We construct an instance of $P|pmtn, prec, p_j = p| \sum C_j$ with $m = \max\{k, l + k', l'\} + 1$ machines and $n = 3m$ jobs:

- for each vertex $i \in V$, there is a job J_i .
- for each edge $[i, j] \in E$, there is a job $J_{[i,j]}$.
- dummy jobs $J_{h,t}$ with $h \in D_t, t \in \{1, 2, 3\}, D_1 = \{1, \dots, m - k\}, D_2 = \{1, \dots, m - l - k'\}, D_3 = \{1, \dots, m - l'\}$.

There are precedence constraints between any job $J_{g,t}$ and $J_{h,t+1}$, for any $g \in D_t, h \in D_{t+1}, t = 1, 2$. Moreover, for any edge $[i, j] \in E$, there is a precedence between J_i and $J_{[i,j]}$. Finally, the question is whether there is a schedule such that $\sum C_j \leq 6m$.

Clearly, if CLIQUE has a solution, so is the scheduling problem:

- dummy jobs $J_{h,t}$ are scheduled during the time-interval $[t - 1, t]$,
- the k jobs corresponding to the vertices of the clique are scheduled during the first period,
- the l jobs corresponding to the edges of the clique, and the k' jobs corresponding to the remaining vertices are scheduled at the second period,
- all the remaining jobs are scheduled at the third period.

This solution has a total completion time of exactly $6m$. Note that it does not use preemption, and is such that $C_{\max} = 3$.

Conversely, let us prove that if there is a schedule such that $\sum C_j \leq 6m$, then there is a clique of size k . First, we can easily prove that if there is a schedule such that $\sum C_j \leq 6m$, then there is a schedule such that $C_{\max} \leq 3$. Indeed, if there exists a job j^* such that $C_{j^*} > 3$, then in the best case (i.e., if there is no precedence constraint), we have $\sum C_j \geq 1 * m + 2 * m + 3 * (m - 1) + C_{j^*} > 6m$. We hence know that there is no idle time and that the dummy jobs $J_{h,t}$ are scheduled during interval $[t - 1, t]$. To conclude, we just need to see that if there is no clique of size k then, whatever the schedule on interval $[0, 1]$ without idle time, the number of eligible jobs at time 1 is strictly less than $k' + l$, which implies an idle time and hence no schedule such that $C_{\max} \leq 3$.

■

6.3 Open problems

Preemptive parallel machine scheduling problems did not receive as much attention as their non-preemptive counterpart, hence the set of open problems is wider.

For the makespan objective and an arbitrary number of machines, when the precedence graph is an interval order graph, it is known to be solvable in polynomial time, as in the UET non-preemptive case. Since for the UET non-preemptive problem, new classes strictly including interval order graphs (namely quasi-interval order graphs and over-interval order graphs) have led to polynomial algorithms, the same question arises for $P|pmtn, qio|C_{\max}$ and $P|pmtn, oio|C_{\max}$.

When the number of machines is fixed, there is a wide set of open problems, since $Qm|pmtn, prec, r_j|L_{\max}$ is the maximal open problem. A good challenge may be for example to try to fix the complexity of $Pm|pmtn, prec, p_j = p|C_{\max}$, or at least to try to find new precedence graphs for which the problem is polynomial, by taking advantage of the fact that tasks are UET (even if it may not always be helpful with preemption).

For the total completion time criterion, $P|pmtn, outtree, p_j = p|\sum C_j$ is maximal polynomially solvable, and we just show that $P|pmtn, prec, p_j = p|\sum C_j$ is \mathcal{NP} -hard even if the precedence graph

is of bounded height. It could be interesting to consider other structures of precedence graph to derive polynomial algorithms. In a similar way, Baptiste et al. [2004a] show that $P|pmtn, p_j = p|\sum T_j$ is solvable in polynomial time, and adding precedence constraints makes the problem \mathcal{NP} -hard, but there is no other result available in the literature, it hence would be interesting to search for new polynomial cases by testing different precedence graphs.

For the problem $P2|pmtn, prec, r_j, p_j = 1|\sum C_j$, the gap is even wider: it is polynomial if the precedence graph is an outtree (see Baptiste and Timkovsky [2001]) but all the other cases are open, from $P2|pmtn,intree, r_j, p_j = 1|\sum C_j$ to $P2|pmtn, prec, r_j, p_j = 1|\sum C_j$.

7 Conclusion

In this paper, we surveyed the complexity results for scheduling problems with precedence constraints, and we have seen that single machine scheduling problems have been much more studied than others. This looks quite normal since the single machine problem is the scheduling problem that is the closest to order theory. Nevertheless, we show that there still are a few open problems for the single machine case. We believe that the most interesting problems for which the complexity is open lie in the parallel machine case; more precisely, we do conjecture that $Pm|sp-graph, p_j = p|C_{\max}$ is solvable in polynomial time; this result will be a large breakthrough since series-parallel graphs are most of the time studied for single machine problems.

Another approach to understand the complexity of scheduling problems is to deal with the parameterized complexity (see Downey and Fellows [2012]). There are only very few results on parameterized complexity of scheduling problems. One can cite Fellows and McCartin [2003] who show that, if the precedence graph is of bounded width w (it is equal to the size of a maximum antichain), then problem $1|prec|\sum T_j \leq k$ is FPT when parameterized by (w, k) . The most recent result on the subject is that $P||C_{\max}$ is FPT for parameter $p_{max} = \max p_j$ (see Mnich and Wiese [2014]). A good graph measure is a powerful tool for parameterized complexity. For general (undirected) graphs, the creation of the

treewidth (see Robertson and Seymour [1986]) helped to discover many results in graph theory, including the Courcelle's theorem (Courcelle [1990]). For directed graphs, and more specifically DAGs, none of the existing measures (see Ganian et al. [2014]) is satisfactory. In our opinion, a major breakthrough will be achieved when one will be able to find a good measure on directed acyclic graphs.

References

- Isto Aho and Erkki Mäkinen. On a parallel machine scheduling problem with precedence constraints. *Journal of Scheduling*, 9(5):493–495, 2006.
- Christoph Ambühl and Monaldo Mastrolilli. Single machine precedence constrained scheduling is a vertex cover problem. *Algorithmica*, 53(4):488–503, 2009.
- Christoph Ambühl, Monaldo Mastrolilli, Nikolaus Mutsanas, and Ola Svensson. On the approximability of single-machine scheduling with precedence constraints. *Mathematics of Operations Research*, 36(4):653–669, 2011.
- P. Baptiste. Polynomial time algorithms for minimizing the weighted number of late jobs on a single machine with equal processing times. *Journal of Scheduling*, 2:245–252, 1999.
- Philippe Baptiste and Vadim G. Timkovsky. On preemption redundancy in scheduling unit processing time jobs on two parallel machines. *Operations Research Letters*, 28(5):205–212, 2001.
- Philippe Baptiste and Vadim G Timkovsky. Shortest path to nonpreemptive schedules of unit-time jobs on two identical parallel machines with minimum total completion time. *Mathematical Methods of Operations Research*, 60(1):145–153, 2004.
- Philippe Baptiste, Peter Brucker, Sigrid Knust, and Vadim G. Timkovsky. Ten notes on equal-execution-time scheduling. *4OR*, 2:111–127, 2004a.
- Philippe Baptiste, Marek Chrobak, Christoph Dürr, Wojciech Jawor, and Nodari Vakhania. Preemptive scheduling of equal-length jobs to maximize weighted throughput. *Operations Research Letters*, 32(3):258–264, 2004b.

-
- Graham R. Brightwell and Edward R. Scheinerman. Fractional dimension of partial orders. *Order*, 9(2): 139–158, 1992.
- Peter Brucker. *Scheduling algorithms*. Springer, 2007.
- Peter Brucker and Sigrid Knust. Complexity results for scheduling problems, 2009. URL <http://www2.informatik.uni-osnabrueck.de/knust/class/>.
- Peter Brucker, M.R. Garey, and D.S. Johnson. Scheduling equal-length tasks under treelike precedence constraints to minimize maximum lateness. *Mathematics of Operations Research*, 2(3):275–284, 1977.
- Peter Brucker, Johann Hurink, and Wiesław Kubiak. Scheduling identical jobs with chain precedence constraints on two uniform machines. *Mathematical Methods of Operations Research*, 49(2):211–219, 1999.
- Peter Brucker, Johann Hurink, and Sigrid Knust. A polynomial algorithm for $p\text{---}p_j=1, r_j, \text{outtree---}c_j$. *Mathematical Methods of Operations Research*, 56(3):407–412, 2002.
- Marc Chardon and Aziz Moukrim. The Coffman–Graham algorithm optimally solves UET task systems with overinterval orders. *SIAM Journal on Discrete Mathematics*, 19(1):109–121, 2005.
- Bo Chen, Ed Coffman, Dariusz Dereniowski, and Wiesław Kubiak. Normal-form preemption sequences for an open problem in scheduling theory. *Journal of Scheduling*, pages 1–28, 2015. ISSN 1099-1425. doi: 10.1007/s10951-015-0446-9. URL <http://dx.doi.org/10.1007/s10951-015-0446-9>.
- E. G. Coffman, J. Sethuraman, and V. G. Timkovsky. Ideal preemptive schedules on two processors. *Acta Informatica*, 39(8):597–612, 2003. ISSN 1432-0525. doi: 10.1007/s00236-003-0119-6. URL <http://dx.doi.org/10.1007/s00236-003-0119-6>.
- Edward G. Coffman Jr, Dariusz Dereniowski, and Wiesław Kubiak. An efficient algorithm for finding ideal schedules. *Acta informatica*, 49(1):1–14, 2012.
- E.G. Coffman Jr and R.L. Graham. Optimal scheduling for two-processor systems. *Acta informatica*, 1(3):200–213, 1972.
- José R. Correa and Andreas S. Schulz. Single-machine scheduling with precedence constraints. *Mathematics of Operations Research*, 30(4):1005–1021, 2005.

-
- Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Information and computation*, 85(1):12–75, 1990.
- M.I. Dessouky, B.J. Lageweg, J.K. Lenstra, and S.L. van de Velde. Scheduling identical jobs on uniform parallel machines. *Statist. Neerlandica*, 44(3):115–123, 1990.
- Khaled Djellab. Scheduling preemptive jobs with precedence constraints on parallel machines. *European journal of operational research*, 117(2):355–367, 1999.
- Danny Dolev and Manfred K. Warmuth. Scheduling precedence graphs of bounded height. *Journal of Algorithms*, 5(1):48–59, 1984.
- Danny Dolev and Manfred K. Warmuth. Profile scheduling of opposing forests and level orders. *SIAM Journal on Algebraic Discrete Methods*, 6(4):665–687, 1985.
- Rodney G. Downey and Michael R. Fellows. *Parameterized complexity*. Springer Science & Business Media, 2012.
- Jianzhong Du, Joseph Y.T. Leung, and Gilbert H. Young. Scheduling chain-structured tasks to minimize makespan and mean flow time. *Information and Computation*, 92(2):219–236, 1991.
- Christoph Dürr. The scheduling zoo, 2016. URL <http://http://schedulingzoo.lip6.fr>.
- Michael R. Fellows and Catherine McCartin. On the parametric complexity of schedules to minimize tardy tasks. *Theoretical computer science*, 298(2):317–324, 2003.
- Robert Ganian, Petr Hliněný, Joachim Kneis, Alexander Langer, Jan Obdržálek, and Peter Rossmanith. Digraph width measures in parameterized algorithmics. *Discrete Applied Mathematics*, 168:88–107, 2014.
- Michael R. Garey and David S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. WH Freeman New York, 1979.
- M.R. Garey, D.S. Johnson, R.E. Tarjan, and M. Yannakakis. Scheduling opposing forests. *SIAM Journal on Algebraic Discrete Methods*, 4(1):72–93, 1983.
- Teofilo F. Gonzalez and Donald B. Johnson. A new algorithm for preemptive scheduling of trees. *Journal of the ACM (JACM)*, 27(2):287–312, 1980.

-
- Valery S. Gordon, Chris N. Potts, Vitaly A. Strusevich, and J. Douglass Whitehead. Single machine scheduling models with deterioration and learning: handling precedence constraints via priority generation. *Journal of Scheduling*, 11(5):357–370, 2008.
- Ronald L. Graham, Eugene L. Lawler, Jan Karel Lenstra, and A.H.G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of discrete mathematics*, 5:287–326, 1979.
- T.C. Hu. Parallel sequencing and assembly line problems. *Operations research*, 9(6):841–848, 1961.
- Yumei Huo and Joseph Y.-T. Leung. Minimizing total completion time for uet tasks with release time and outtree precedence constraints. *Mathematical Methods of Operations Research*, 62(2):275–279, 2005.
- Yumei Huo and Joseph Y-T Leung. Minimizing mean flow time for uet tasks. *ACM Transactions on Algorithms (TALG)*, 2(2):244–262, 2006.
- W. Kubiak. Optimal scheduling of unit-time tasks on two uniform processors under tree-like precedence constraints. *Zeitschrift für Operations Research*, 33(6):423–437, 1989.
- Wieslaw Kubiak, Djamel Rebaine, and Chris Potts. Optimality of HLF for scheduling divide-and-conquer UET task graphs on identical parallel processors. *Discrete Optimization*, 6:79–91, 2009.
- E.L. Lawler. Preemptive scheduling of precedence-constrained jobs on parallel machines. In M.A.H. Dempster, J.K. Lenstra, and A.H.G. Rinnooy Kan, editors, *Deterministic and Stochastic Scheduling*, volume 84 of *NATO Advanced Study Institutes Series*, pages 101–123. Springer Netherlands, 1982.
- Eugene L. Lawler. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Annals of Discrete Mathematics*, 2:75–90, 1978.
- Eugene L. Lawler, Jan Karel Lenstra, Alexander H.G. Rinnooy Kan, and David B. Shmoys. Sequencing and scheduling: Algorithms and complexity. *Handbooks in operations research and management science*, 4:445–522, 1993.
- Jan Karel Lenstra and A.H.G. Rinnooy Kan. Complexity of scheduling under precedence constraints. *Operations Research*, 26(1):22–35, 1978.

-
- Jan Karel Lenstra and A.H.G. Rinnooy Kan. Complexity results for scheduling chains on a single machine. *European Journal of Operational Research*, 4(4):270–275, 1980.
- Jan Karel Lenstra, A.H.G. Rinnooy Kan, and Peter Brucker. Complexity of machine scheduling problems. *Annals of discrete mathematics*, 1:343–362, 1977.
- J.K. Lenstra. Not published.
- Joseph Y.-T. Leung and Gilbert H. Young. Minimizing total tardiness on a single machine with precedence constraints. *ORSA Journal on Computing*, 2(4):346–352, 1990.
- Elaine Levey and Thomas Rothvoss. A Lasserre-based $(1+\epsilon)$ -approximation for $Pm|p_j = 1, prec|C_{max}$. To appear in *STOC'16*, 2016. URL <http://arxiv.org/abs/1509.07808>.
- Irene N. Lushchakova. Two machine preemptive scheduling problem with release dates, equal processing times and precedence constraints. *European Journal of Operational Research*, 171(1):107–122, 2006.
- Martin Middendorf and Vadim G. Timkovsky. Transversal graphs for partially ordered sets: Sequencing, merging and scheduling problems. *Journal of Combinatorial Optimization*, 3(4):417–435, 1999.
- Matthias Mnich and Andreas Wiese. Scheduling and fixed-parameter tractability. *Mathematical Programming*, pages 1–30, 2014. doi: 10.1007/s10107-014-0830-9.
- Rolf H. Möhring. Computationally tractable classes of ordered sets. In Ivan Rival, editor, *Algorithms and Order*, volume 255 of *NATO ASI Series*, pages 105–193. Springer Netherlands, 1989.
- Clyde L. Monma. Linear-time algorithms for scheduling on parallel processors. *Operations Research*, 30(1):116–124, 1982.
- Clyde L. Monma and Jeffrey B. Sidney. Sequencing with series-parallel precedence constraints. *Mathematics of Operations Research*, 4(3):215–224, 1979.
- Aziz Moukrim. Optimal scheduling on parallel machines for a new order class. *Operations research letters*, 24(1):91–95, 1999.
- Aziz Moukrim and Alain Quilliot. A relation between multiprocessor scheduling and linear programming. *Order*, 14(3):269–278, 1997.

- Aziz Moukrim and Alain Quilliot. Optimal preemptive scheduling on a fixed number of identical parallel machines. *Operations research letters*, 33(2):143–150, 2005.
- Richard R. Muntz and Edward G. Coffman Jr. Preemptive scheduling of real-time tasks on multiprocessor systems. *Journal of the ACM (JACM)*, 17(2):324–338, 1970.
- Christos H. Papadimitriou and Mihalis Yannakakis. Scheduling interval-ordered tasks. *SIAM Journal on Computing*, 8(3):405–409, 1979.
- Michael L. Pinedo. *Scheduling: theory, algorithms, and systems*. Springer Verlag New York, 2012.
- Neil Robertson and Paul D. Seymour. Graph minors. ii. algorithmic aspects of tree-width. *Journal of algorithms*, 7(3):309–322, 1986.
- N.W. Sauer and M.G. Stone. Preemptive scheduling of interval orders is polynomial. *Order*, 5(4):345–348, 1989.
- Jeffrey B. Sidney. Decomposition algorithms for single-machine sequencing with precedence relations and deferral costs. *Operations Research*, 23(2):283–298, 1975.
- Ola Svensson. Hardness of precedence constrained scheduling on identical machines. *SIAM Journal on Computing*, 40(5):1258–1274, 2011.
- Z. Tian, C.T. Ng, and T.C.E. Cheng. An $O(n^2)$ algorithm for scheduling equal-length preemptive jobs on a single machine to minimize total tardiness. *Journal of Scheduling*, 9(4):343–364, 2006.
- Vadim G. Timkovsky. Identical parallel machines vs. unit-time shops and preemptions vs. chains in scheduling complexity. *European Journal of Operational Research*, 149(2):355–376, 2003.
- J.D. Ullman. Complexity of sequencing problems. In J.L. Bruno, E.G. Coffman, Jr., R.L. Graham, W.H. Kohler, R. Sethi, K. Steiglitz, and J.D. Ullman, editors, *Computer and Job/Shop Scheduling Theory*. John Wiley & Sons Inc., 1976.
- Ji-Bo Wang and Jian-Jun Wang. Single-machine scheduling with precedence constraints and position-dependent processing times. *Applied Mathematical Modelling*, 37(3):649–658, 2013.
- Ji-Bo Wang, C.T. Ng, and T.C.E. Cheng. Single-machine scheduling with deteriorating jobs under a series-parallel graph constraint. *Computers & Operations Research*, 35(8):2684–2693, 2008.

Tianyu Wang. Ordonnancement parallèle non-préemptif avec contraintes de précédence. Technical report, Ecole Centrale Nantes, 2015. URL <https://hal.archives-ouvertes.fr/hal-01202104>.

David P. Williamson and David B. Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.

Appendix A: List of results

For an easier reading of all the complexity results that are reviewed in this survey, we proposed a synthesis in the following tables. In each table, we write the polynomial cases, some open cases (the ones that seem the most promising in our opinion) and the \mathcal{NP} -hard problems.

Problem	Complexity	Reference
$1 sp - graph \sum w_j C_j$	\mathcal{P}	Lawler [1978]
$1 dim \leq 2 \sum w_j C_j$	\mathcal{P}	Ambühl and Mastrolilli [2009]
$1 sp - graph, p_j^{(t)} = p_j(1 + at) \sum w_j C_j$	\mathcal{P}	Wang et al. [2008]
$1 sp - graph, p_j(t) = p_j + \alpha_j t C_{\max}$	\mathcal{P}	Wang et al. [2008]
$1 sp - graph, p_j^{[r]} = a_j + b_j r C_{\max}$	\mathcal{P}	Wang and Wang [2013]
$1 sp - graph, p_j^{[r]} = a_j - b_j r C_{\max}$	\mathcal{P}	Wang and Wang [2013]
$1 sp - graph, p_j^{[r]} = p_j r C_{\max}$	\mathcal{P}	Gordon et al. [2008]
$1 sp - graph, p_j^{[r]} = p_j \gamma^{r-1} \sum C_j$ with $\gamma \geq 2$ or $0 < \gamma < 1$	\mathcal{P}	Gordon et al. [2008]
$1 sp - graph, p_j^{(t)} = p_j(1 - at) \sum w_j C_j$	\mathcal{P}	Gordon et al. [2008]
$1 sp - graph, p_j^{(t)} = p_j + at \sum C_j$	\mathcal{P}	Gordon et al. [2008]
$1 sp - graph, p_j^{(t)} = p_j - at \sum C_j$	\mathcal{P}	Gordon et al. [2008]
$1 pmtn, r_j, p_j = p \sum T_j$	\mathcal{P}	Tian et al. [2006]
$1 pmtn, r_j, p_j = p \sum w_j U_j$	\mathcal{P}	Baptiste [1999]
$1 2 < fdim < 3 \sum w_j C_j$	Open	
$1 io \sum C_j$	Open	
$1 pmtn, chains, r_j, p_j = p \sum T_j$	Open	
$1 pmtn, prec, r_j, p_j = p \sum w_j T_j$	Open	
$1 pmtn, chains, r_j, p_j = p \sum U_j$	Open	
$1 pmtn, prec, r_j, p_j = p \sum w_j U_j$	Open	
$1 prec, p_j = p \sum w_j C_j$	\mathcal{NP} -hard	Lawler [1978]
$1 prec \sum C_j$	\mathcal{NP} -hard	Lawler [1978]
$1 indegree \leq 2 \sum C_j$	\mathcal{NP} -hard	Lenstra and Rinnooy Kan [1978]
$1 h(G) \leq k \sum w_j C_j$	\mathcal{NP} -hard	Lawler [1978]
$1 io \sum w_j C_j$	\mathcal{NP} -hard	Ambühl et al. [2011]
$1 fdim \geq 3 \sum w_j C_j$	\mathcal{NP} -hard	Ambühl et al. [2011]
$1 chains, p_j = 1 \sum U_j$	\mathcal{NP} -hard	Lenstra and Rinnooy Kan [1980]
$1 chains, p_j = 1 \sum T_j$	\mathcal{NP} -hard	Leung and Young [1990]

Table 1 Complexity results for single machine problems

Problem	Complexity	Reference
$P tree, p_j = p C_{\max}$	\mathcal{P}	Hu [1961]
$P outtree, p_j = p \sum C_j$	\mathcal{P}	Hu [1961]
$Pm opp. forest, p_j = p C_{\max}$	\mathcal{P}	Garey et al. [1983]
$Pm lo, p_j = p C_{\max}$	\mathcal{P}	Dolev and Warmuth [1985]
$P io, p_j = p C_{\max}$	\mathcal{P}	Papadimitriou and Yannakakis [1979]
$P io, p_j = p \sum C_j$	\mathcal{P}	Möhring [1989]
$P qio, p_j = p C_{\max}$	\mathcal{P}	Moukrim [1999]
$P oio, p_j = p C_{\max}$	\mathcal{P}	Chardon and Moukrim [2005]
$P oio, p_j = p \sum C_j$	\mathcal{P}	Wang [2015]
$P outtree, r_j, p_j = p C_{\max}$	\mathcal{P}	Brucker et al. [1977]
$P DC - graph, p_j = p C_{\max}$	\mathcal{P}	Kubiak et al. [2009]
$Pm h(G) \leq k, p_j = p C_{\max}$	\mathcal{P}	Dolev and Warmuth [1984]
$P w(G) \leq k, r_j, p_j = 1 f$	\mathcal{P}	Middendorf and Timkovsky [1999]
$P2 prec, p_j = p C_{\max}$	\mathcal{P}	Coffman Jr and Graham [1972]
$P2 prec, p_j = p \sum C_j$	\mathcal{P}	Coffman Jr and Graham [1972]
$P2 prec, r_j, p_j = 1 \sum C_j$	\mathcal{P}	Baptiste and Timkovsky [2004]
$Q2 chains, p_j = p C_{\max}$	\mathcal{P}	Brucker et al. [1999]
$Pm intree, p_j = p \sum C_j$	\mathcal{P}	Baptiste et al. [2004a]
$P outtree, r_j, p_j = 1 \sum C_j$	\mathcal{P}	Brucker et al. [2002]
$Qm r_j, p_j = p \sum C_j$	\mathcal{P}	Dessouky et al. [1990]
$Pm sp - graph, p_j = p C_{\max}$	Open	
$Pm prec, p_j = p C_{\max}$	Open	
$P intree, p_j = p \sum C_j$	Open	
$P outtree, r_j, p_j = p \sum C_j$	Open	
$Qm prec, r_j, p_j = p L_{\max}$	Open	
$Qm prec, r_j, p_j = p \sum C_j$	Open	
$P2 C_{\max}$	\mathcal{NP} -hard	Lenstra et al. [1977]
$P2 chains \sum C_j$	\mathcal{NP} -hard	Du et al. [1991]
$P opp. forest, p_j = p C_{\max}$	\mathcal{NP} -hard	Garey et al. [1983]
$P h(G) \leq k, p_j = p C_{\max}$	\mathcal{NP} -hard	Lenstra and Rinnooy Kan [1978]
$P intree, r_j, p_j = p C_{\max}$	\mathcal{NP} -hard	Brucker et al. [1977]
$P h(G) \leq k, p_j = p \sum C_j$	\mathcal{NP} -hard	Lenstra and Rinnooy Kan [1978]
$P intree, r_j, p_j = 1 \sum C_j$	\mathcal{NP} -hard	Lenstra

Table 2 Complexity results for parallel machine problems without preemption

Problem	Complexity	Reference
$P pmtn, tree C_{\max}$	\mathcal{P}	Muntz and Coffman Jr [1970]
$P pmtn, outtree, r_j C_{\max}$	\mathcal{P}	Lawler [1982]
$P pmtn, io C_{\max}$	\mathcal{P}	Djellab [1999]
$Pm pmtn, \mathcal{A}_m C_{\max}$	\mathcal{P}	Moukrim and Quilliot [2005]
$P2 pmtn, outtree, r_j, p_j = 1 \sum C_j$	\mathcal{P}	Baptiste and Timkovsky [2001]
$P2 pmtn, outtree, r_j, p_j = p \sum C_j$	\mathcal{P}	Lushchakova [2006]
$P pmtn, outtree, r_j, p_j = 1 \sum C_j$	\mathcal{P}	Brucker et al. [2002]
$P2 pmtn, prec, p_j = 1 \sum C_j$	\mathcal{P}	Coffman et al. [2003]
$Pm pmtn, qio, p_j = p C_{\max}$	Open	
$Pm pmtn, oio, p_j = p C_{\max}$	Open	
$Pm pmtn, lo, p_j = p C_{\max}$	Open	
$Pm pmtn, sp - graph, p_j = p C_{\max}$	Open	
$Pm pmtn, prec, p_j = p C_{\max}$	Open	
$Qm pmtn, prec, r_j L_{\max}$	Open	
$P2 pmtn, intree, r_j, p_j = 1 \sum C_j$	Open	
$P2 pmtn, prec, r_j, p_j = 1 \sum C_j$	Open	
$P pmtn, prec, p_j = p C_{\max}$	\mathcal{NP} -hard	Ullman [1976]
$P2 pmtn, chains \sum C_j$	\mathcal{NP} -hard	Du et al. [1991]
$P pmtn, h(G) \leq k, p_j = p \sum C_j$	\mathcal{NP} -hard	[this paper]
$P2 pmtn, chains, p_j = 1 \sum w_j C_j$	\mathcal{NP} -hard	Du et al. [1991]
$P2 pmtn, chains, p_j = 1 \sum U_j$	\mathcal{NP} -hard	Baptiste et al. [2004a]

Table 3 Complexity results for parallel machine problems with preemption