

On the Fine-grained Parameterized Complexity of Partial Scheduling to Minimize the Makespan

Citation for published version (APA):

Nederlof, J., & Swennenhuis, C. M. F. (2022). On the Fine-grained Parameterized Complexity of Partial Scheduling to Minimize the Makespan. *Algorithmica*, 84(8), 2309-2334. [8]. <https://doi.org/10.1007/s00453-022-00970-8>

Document license:

CC BY

DOI:

[10.1007/s00453-022-00970-8](https://doi.org/10.1007/s00453-022-00970-8)

Document status and date:

Published: 01/08/2022

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



On the Fine-grained Parameterized Complexity of Partial Scheduling to Minimize the Makespan

Jesper Nederlof¹ · Céline M. F. Swennenhuis²

Received: 19 January 2021 / Accepted: 10 April 2022 / Published online: 7 May 2022
© The Author(s) 2022

Abstract

We study a natural variant of scheduling that we call *partial scheduling*: in this variant an instance of a scheduling problem along with an integer k is given and one seeks an optimal schedule where not all, but only k jobs, have to be processed. Specifically, we aim to determine the fine-grained parameterized complexity of partial scheduling problems parameterized by k for all variants of scheduling problems that minimize the makespan and involve unit/arbitrary processing times, identical/unrelated parallel machines, release/due dates, and precedence constraints. That is, we investigate whether algorithms with runtimes of the type $f(k)n^{\mathcal{O}(1)}$ or $n^{\mathcal{O}(f(k))}$ exist for a function f that is as small as possible. Our contribution is two-fold: First, we categorize each variant to be either in P, NP-complete and fixed-parameter tractable by k , or W[1]-hard parameterized by k . Second, for many interesting cases we further investigate the runtime on a finer scale and obtain run times that are (almost) optimal assuming the Exponential Time Hypothesis. As one of our main technical contributions, we give an $\mathcal{O}(8^k k(|V| + |E|))$ time algorithm to solve instances of partial scheduling problems minimizing the makespan with unit length jobs, precedence constraints and release dates, where $G = (V, E)$ is the graph with precedence constraints.

Keywords Fixed-parameter tractability · Scheduling · Precedence constraints

J. Nederlof ERC project No. 617951. and No. 853234. and NWO project No. 024.002.003. C.M.F. Swennenhuis NWO project No. 613.009.031b, ERC project No. 617951.

✉ Céline M. F. Swennenhuis
c.m.f.swennenhuis@tue.nl

Jesper Nederlof
j.nederlof@uu.nl

¹ Algorithms and Complexity Group, Utrecht University, Utrecht, Netherlands

² Combinatorial Optimization Group, Eindhoven University of Technology, Eindhoven, Netherlands

1 Introduction

Scheduling is one of the most central application domains of combinatorial optimization. In the last decades, huge combined effort of many researchers led to major progress on understanding the worst-case computational complexity of almost all natural variants of scheduling: By now, for most of these variants it is known whether they are NP-complete or not. Scheduling problems provide the context of some of the most classic approximation algorithms. For example, in the standard textbook by Shmoys and Williamson on approximation algorithms [29] a wide variety of techniques are illustrated by applications to scheduling problems. See also the standard textbook on scheduling by Pinedo [24] for more background.

Instead of studying approximation algorithms, another natural way to deal with NP-completeness is *Parameterized Complexity* (PC).

While the application of general PC theory to the area of scheduling has still received considerably less attention than the approximation point of view, recently its study has seen explosive growth, as witnessed by a plethora of publications (e.g. [2, 13, 18, 22, 27, 28]). Additionally, many recent results and open problems can be found in a survey by Mnich and van Bevern [21], and even an entire workshop on the subject was recently held [20].

In this paper we advance this vibrant research direction with a complete mapping of how several standard scheduling parameters influence the parameterized complexity of minimizing the makespan in a natural variant of scheduling problems that we call *partial scheduling*. Next to studying the classical question of whether parameterized problems are in P, in FPT parameterized by k , or W[1]-hard parameterized by k , we also follow the well-established modern perspective of ‘fine-grained’ PC and aim at run times of the type $f(k)n^{\mathcal{O}(1)}$ or $n^{f(k)}$ for the smallest function f of parameter k . *Partial Scheduling* In many scheduling problems arising in practice, the set of jobs to be scheduled is not predetermined. We refer to this as *partial scheduling*. Partial scheduling is well-motivated from practice, as it arises naturally for example in the following scenarios:

1. Due to uncertainties a *close-horizon approach* may be employed and only few jobs out of a big set of jobs will be scheduled in a short but fixed time-window,
2. In freelance markets typically a large database of jobs is available and a freelancer is interested in selecting only a few of the jobs to work on,
3. The selection of the jobs to process may resemble other choices the scheduler should make, such as to outsource non-processed jobs to various external parties.

Partial scheduling has been previously studied in the equivalent forms of *maximum throughput scheduling* [25] (motivated by the first example setting above), *job rejection* [26], *scheduling with outliers* [12], *job selection* [8, 16, 30] and its special case *interval selection* [5].

In this paper, we conduct a rigorous study of the parameterized complexity of partial scheduling, parameterized by *the number of jobs to be scheduled*. We denote this number by k . While several isolated results concerning the parameterized complexity of partial scheduling do exist, this parameterization has (somewhat surprisingly) not

been rigorously studied yet.¹ We address this and study the parameterized complexity of the (arguably) most natural variants of the problem. We fix as objective to minimize the makespan while scheduling at least k jobs, for a given integer k and study all variants with the following characteristics:

- 1 machine, identical parallel machines or unrelated parallel machines,
- release/due dates, unit/arbitrary processing times, and precedence constraints.

Note that a priori this amounts to $3 \times 2 \times 2 \times 2 \times 2 = 48$ variants.

1.1 Our Results

We give a classification of the parameterized complexity of these 48 variants. Additionally, for each variant that is not in P, we give algorithms solving them and lower bounds under ETH. To easily refer to a variant of the scheduling problem, we use the standard three-field notation by Graham et al. [11]. See Sect. 2 for an explanation of this notation. To accommodate our study of partial scheduling, we extend the $\alpha|\beta|\gamma$ notation as follows:

Definition 1 We let k -sched in the γ -field indicate that we only schedule k out of n jobs.

We study the fine-grained parameterized complexity of all problems $\alpha|\beta|\gamma$, where $\alpha \in \{1, P, R\}$, the options for β are all combinations for $r_j, d_j, p_j = 1, \text{prec}$, and γ is fixed to $\gamma = k$ -sched, C_{\max} . Our results are explicitly enumerated in Table 1.

The rows of Table 1 are lexicographically sorted on (i) precedence relations/no precedence relations, (ii) a single machine, identical machines or unrelated machines (iii) release dates and/or deadlines. Because their presence has a major influence on the character of the problem we stress the distinction between variants with and without *precedence constraints*.² On a high abstraction level, our contribution is two-fold:

1. We present a *classification* of the complexity of all aforementioned variants of partial scheduling with the objective of minimizing the makespan. Specifically, we classify all variants to be either solvable in polynomial time, to be fixed-parameter tractable in k and NP-hard, or to be W[1]-hard.
2. For most of the studied variants we present both an algorithm and a lower bound that shows that our algorithm cannot be significantly improved unless the Exponential Time Hypothesis (ETH) fails.

Thus, while we completely answer a classical type of question in the field of Parameterized Complexity, we pursue in our second contribution a more modern and fine-grained understanding of the best possible runtime with respect to the parameter k . For several of the studied variants, the lower bounds and algorithms listed in Table 1 follow relatively quickly. However, for many other cases we need substantial new insights to obtain (almost) matching upper and lower bounds on the runtime of the algorithms solving them. We have grouped the rows in *result types* [A]–[G] depending on our methods for determining their complexity.

¹ We compare the previous works and other relevant studied parameterization in the end of this section.

² A precedence constraint $a < b$ enforces that job a needs to be finished before job b can start.

Table 1 The fine-grained parameterized complexity of partial scheduling, where γ denotes k -sched, C_{\max} and S.I. abbreviates SUBGRAPH ISOMORPHISM (Color table online)

	Problem Description	Parameterized Complexity in k	Result Type	Lower Bound under ETH		Runtime
				Excluded Runtime	Reduction from	
Precedence Relations	1 $1 prec, p_j = 1 \gamma$	P	[A]			$n^{\mathcal{O}(1)}$
	2 $1 r_j, prec, p_j = 1 \gamma$	P	[A]			$n^{\mathcal{O}(1)}$
	3 $1 d_j, prec, p_j = 1 \gamma$	W[1]-hard	[B]	$n^{\mathcal{O}(k/\log k)}$	3-COLORING	$n^{\mathcal{O}(k)}$
	4 $1 r_j, d_j, prec, p_j = 1 \gamma$	W[1]-hard	[B]	$n^{\mathcal{O}(k/\log k)}$	3-COLORING	$n^{\mathcal{O}(k)}$
	5 $P prec, p_j = 1 \gamma$	FPT	[C]	$\mathcal{O}^*(2^{\mathcal{O}(\sqrt{k \log k})})$	$P prec, p_j = 1 C_{\max}$	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	6 $P r_j, prec, p_j = 1 \gamma$	FPT	[C]	$\mathcal{O}^*(2^{\mathcal{O}(\sqrt{k \log k})})$	$P prec, p_j = 1 C_{\max}$	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	7 $P d_j, prec, p_j = 1 \gamma$	W[1]-hard	[B]	$n^{\mathcal{O}(k/\log k)}$	3-COLORING	$n^{\mathcal{O}(k)}$
	8 $P r_j, d_j, prec, p_j = 1 \gamma$	W[1]-hard	[B]	$n^{\mathcal{O}(k/\log k)}$	3-COLORING	$n^{\mathcal{O}(k)}$
	9 $1 prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(\sqrt{k})}$	k -CLIQUE	$n^{\mathcal{O}(k)}$
	10 $1 r_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	11 $1 d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	12 $1 r_j, d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	13 $P prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	14 $P r_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	15 $P d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	16 $P r_j, d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	17 $R prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	18 $R r_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	19 $R d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
	20 $R r_j, d_j, prec \gamma$	W[1]-hard	[D]	$n^{\mathcal{O}(k/\log k)}$	PARTITIONED S.I.	$n^{\mathcal{O}(k)}$
No Precedence Relations	21 $1 p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	22 $1 r_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	23 $1 d_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	24 $1 r_j, d_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	25 $P p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	26 $P r_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	27 $P d_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	28 $P r_j, d_j, p_j = 1 \gamma$	P	[E]			$n^{\mathcal{O}(1)}$
	29 1γ	P	[F]			$n^{\mathcal{O}(1)}$
	30 $1 r_j \gamma$	P	[F]			$n^{\mathcal{O}(1)}$
	31 $1 d_j \gamma$	P	[F]			$n^{\mathcal{O}(1)}$
	32 $1 r_j, d_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	33 $P \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	34 $P r_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	35 $P d_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	36 $P r_j, d_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	37 $R \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	38 $R r_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	39 $R d_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$
	40 $R r_j, d_j \gamma$	FPT	[G]	$\mathcal{O}^*(2^{\mathcal{O}(k)})$	SUBSET SUM	$\mathcal{O}^*(2^{\mathcal{O}(k)})$

Since $p_j = 1$ implies that the machines are identical, the mentioned number of 48 combinations reduces to 40 different scheduling problems. The \mathcal{O}^* notation omits factors polynomial in the input size. The highlighted table entries are new results from this paper

1.2 Our New Methods

We now describe some of our most significant technical contributions for obtaining the various types (listed as [A]–[G] in Table 1) of results. Note that we skip some less interesting cases in this introduction; for a complete argumentation of all results from Table 1 we refer to Sect. 6. The main building blocks and logical implications to obtain the results from Table 1 are depicted in Fig. 1. We now discuss these building blocks of Fig. 1 in detail.




Fig. 1 An illustration of the various result types as indicated in Table 1. Arrows indicate how a problem is generalized by another problem

1.2.1 Precedence Constraints

Our main technical contribution concerns result type [C]. The simplest of the two cases, $P|prec, p_j = 1|k\text{-sched}, C_{\max}$, cannot be solved in $\mathcal{O}^*(2^{o(\sqrt{k} \log k)})$ time assuming the Exponential Time Hypothesis and not in $2^{o(k)}$ unless sub-exponential time algorithms for the BICLIQUE problem exist, due to reductions by Jansen et al. [14]. Our contribution lies in the following theorem that gives an upper bound for the more general of the two problems that matches the latter lower bound:

Theorem 1 $P|r_j, prec, p_j = 1|k\text{-sched}, C_{\max}$ can be solved in $\mathcal{O}(8^k k(|V| + |E|))$ time,³ where $G = (V, E)$ is the precedence graph given as input.

Theorem 1 will be proved in Sect. 3. The first idea behind the proof is based on a natural⁴ dynamic programming algorithm indexed by anti-chains of the partial order naturally associated with the precedence constraints. However, evaluating this dynamic program naïvely would lead to an $n^{\mathcal{O}(k)}$ time algorithm, where n is the number of jobs.

³ We assume basic arithmetic operations with the release dates take constant time.

⁴ A similar dynamic programming approach was also present in, for example, [6].

Our key idea is to only compute a subset of the table entries of this dynamic programming algorithm, guided by a new parameter of an antichain called the *depth*. Intuitively, the depth of an antichain A indicates the number of jobs that can be scheduled after A in a feasible schedule without violating the precedence constraints.

We prove Theorem 1 by showing we may restrict attention in the dynamic programming algorithm to antichains of depth at most k , and by bounding the number of antichains of depth at most k indirectly by bounding the number of *maximal* antichains of depth at most k . We believe this methodology should have more applications for scheduling problems with precedence constraints.

Surprisingly, the positive result of Theorem 1 is in *stark contrast* with the seemingly symmetric case where only deadlines are present: Our next result, indicated as [B] in Fig. 1 shows it is much harder:

Theorem 2 *The problem $P|d_j, prec, p_j = 1|k$ -sched, C_{\max} is $W[1]$ -hard, and it cannot be solved in $n^{o(k/\log k)}$ time assuming the ETH.*

Theorem 2 is a consequence of a reduction outlined in Sect. 4. Note the $W[1]$ -hardness follows from a natural reduction from the k -CLIQUE problem (presented originally by Fellows and McCartin [9]), but this reduction increases the parameter k to $\Omega(k^2)$ and would only exclude $n^{o(\sqrt{k})}$ time algorithms assuming the ETH. To obtain the tighter bound from Theorem 2, we instead provide a non-trivial reduction from the 3-COLORING problem based on a new selection gadget.

For result type [D], we give a lower bound by a (relatively simple) reduction from PARTITIONED SUBGRAPH ISOMORPHISM in Theorem 6 and Corollary 4. Since it is conjectured that PARTITIONED SUBGRAPH ISOMORPHISM cannot be solved in $n^{o(k)}$ time assuming the ETH, our reduction is a strong indication that the simple $n^{O(k)}$ time algorithm (see Sect. 6) cannot be improved significantly in this case.

1.2.2 No Precedence Constraints

The second half of our classification concerns scheduling problems without precedence constraints, and is easier to obtain than the first half. Results [E], [F] are consequences of a greedy algorithm and Moore's algorithm [23] that solves the problem $1||\sum_j U_j$ in $O(n \log n)$ time. Notice that this also solves the problem $1|r_j|k$ -sched, C_{\max} , by reversing the schedule and viewing the release dates as the deadlines. For result type [G] we show that a standard technique in parameterized complexity, the color coding method, can be used to get a $2^{O(k)}$ time algorithm for the most general problem of the class, being $R|r_j, d_j|k$ -sched, C_{\max} . All lower bounds on the runtime of algorithms for problems of type [G] are by a reduction from SUBSET SUM, but for $1|r_j, d_j|k$ -sched, C_{\max} this reduction is slightly different.

1.3 Related Work

The interest in parameterized complexity of scheduling problems recently witnessed an explosive growth, resulting in e.g. a workshop [20] and a survey by Mnich and van Bevern [21] with a wide variety of open problems.

The parameterized complexity of partial scheduling parameterized by the number of processed jobs, or equivalently, the number of jobs ‘on time’ was studied before: Fellows et al. [9] studied a problem called k -TASKS ON TIME that is equivalent to $1|d_j, \text{prec}, p_j = 1|k\text{-sched}, C_{\max}$ and showed that it is $W[1]$ -hard when parameterized by k ,⁵ and in FPT parameterized by k and the width of the partially ordered set induced by the precedence constraints. Van Bevern et al. [27] showed that the JOB INTERVAL SELECTION problem, where each job is given a set of possible intervals to be processed on, is in FPT parameterized by k . Bessy et al. [2] consider partial scheduling with a restriction on the jobs called ‘Coupled-Task’, and also remarked the current parameterization is relatively understudied.

Another related parameter is the number of jobs that are *not scheduled*, that also has been studied in several previous works [4, 9, 22]. For example, Mnich and Wiese [22] studied the parameterized complexity of scheduling problems with respect to the number of rejected jobs in combination with other variables as parameter. If n denotes the number of given jobs, this parameter equals $n - k$. The two parameters are somewhat incomparable in terms of applications: In some settings only few jobs out of many alternatives need to be scheduled, but in other settings rejecting a job is very costly and thus will happen rarely. However, a strong advantage of using k as parameter is in terms of its computational complexity: If the version of the problem with all jobs mandatory is NP-complete it is trivially NP-complete for $n - k = 0$, but it may still be in FPT parameterized by k .

1.4 Organization of this Paper

This paper is organized as follows: We start with some preliminaries in Sect. 2. In Sect. 3 we present the proof of Theorem 1, and in Sect. 4 we describe the reductions for result types [B] and [D]. In Sect. 5 we give the algorithm for result type [G] and in Sect. 6 we motivate all cases from Table 1. Finally, in Sect. 7 we present a conclusion.

2 Preliminaries

2.1 The Three-Field Notation by Graham Et al

Throughout this paper we denote scheduling problems using the three-field notation by Graham et al. [11]. Problems are classified by parameters $\alpha|\beta|\gamma$. The α describes the machine environment. We use $\alpha \in \{1, P, R\}$, indicating whether there are one (1), identical (P) or unrelated (R) parallel machines available. Here identical refers to the fact that every job takes a fixed amount of time process independent of the machine, and unrelated means a job could take different time to process per machine. The β field describes the job characteristics, which in this paper can be a combination of the following values: prec (precedence constraints), r_j (release dates), d_j (deadlines) and $p_j = 1$ (all processing times are 1). We assume without loss of generality that all release dates and deadlines are integers.

⁵ Our results [C] and [D] build on and improve this result.

The γ field concerns the optimization criteria. A given schedule determines C_j , the completion time of job j , and U_j , the unit penalty which is 1 if $C_j > d_j$, and 0 if $C_j \leq d_j$. In this paper we use the following optimization criteria

- C_{\max} : minimize the makespan (i.e. the maximum completion time C_j of any job),
- $\sum_j U_j$: minimize the number of jobs that finish after their deadline,
- k -sched: maximize the number of processed jobs; in particular, process at least k jobs.

A schedule is said to be *feasible* if no constraints (deadlines, release dates, precedence constraints) are violated.

2.2 Notation for Posets

Any precedence graph G is a directed acyclic graph and therefore induces a partial order \preceq on $V(G)$. Indeed, if there is a path from x to y , we let $x \preceq y$. An *antichain* is a set $A \subseteq V(G)$ of mutually incomparable elements. We say that A is *maximal* if there is no antichain A' with $A \subset A'$, where ' \subset ' denotes strict inclusion. The set of *predecessors* of A is $\text{pred}(A) = \{x \in V(G) : \exists a \in A : x \preceq a\}$, and the set of *comparables* of A is $\text{comp}(A) = \{x \in V(G) : \exists a \in A : x \preceq a \text{ or } x \succeq a\}$. Note $\text{comp}(A) = V(G)$ if and only if A is maximal.

An element $x \in V(G)$ is a *minimal* element if $x \preceq y$ for all $y \in \text{comp}(\{x\})$. An element $x \in V(G)$ is a *maximal* element if $x \succeq y$ for all $y \in \text{comp}(\{x\})$. Furthermore, $\min(G) = \{x \mid x \text{ is a minimal element in } G\}$ and $\max(G) = \{x \mid x \text{ is a maximal element in } G\}$.

Notice that $\max(G)$ is exactly the antichain A such that $\text{pred}(A) = V(G)$. We denote the subgraph of G induced by S with $G[S]$. We may assume that $r_j < r_{j'}$ if $j < j'$ since job j' will be processed later than r_j in any schedule. To handle release dates we use the following:

Definition 2 Let G be a precedence graph. Then G^t is the precedence graph restricted to all jobs that can be scheduled on or before time t , i.e. all jobs with release date at most t .

We assume $G = G^{C_{\max}}$, since all jobs with release date greater than C_{\max} can be ignored.

2.3 Parameterized Complexity

We say a problem is *Fixed-Parameter Tractable* (and in the complexity class FPT) parameterized by parameter k , if there exists an algorithm with runtime $\mathcal{O}(f(k) \cdot n^c)$, where n denotes the size of the instance, f is a computable function and c some constant. There also exist problems for which inclusion in FPT for some parameter is unlikely, such as k -CLIQUE. This is because k -CLIQUE is complete for the complexity class $W[1]$ and it is conjectured that $\text{FPT} \neq W[1]$. One could view FPT as the parameterized version of P and $W[1]$ of the parameterized version of NP. To prove a problem \mathcal{P} to be $W[1]$ -hard, one can use a *parameterized reduction* from another problem \mathcal{P}' that

is $W[1]$ -hard, where the reduction is a polynomial-time reduction with the following two additional restrictions: (1) the parameter k' of \mathcal{P}' is bounded by $g(k)$ for some function computable g and k the parameter of \mathcal{P} , (2) the runtime of the reduction is bounded by $f(k) \cdot n^c$ for f some computable function, n the size of the instance of \mathcal{P} and c a constant.

We exclude fixed-parameter tractable algorithms for problems that are $W[1]$ -hard. To exclude runtimes in a more *fine-grained* manner, we use the *Exponential Time Hypothesis* (ETH). Roughly speaking, the ETH conjectures that no $2^{o(n)}$ algorithm for 3-SAT exists, where n is the number of variables of the instance. As a consequence we can, for example, exclude algorithms with runtime $2^{o(n)}$ for SUBSET SUM where n is the number of input integers, and algorithms with runtime $n^{o(k)}$ for k -CLIQUE where n is the number of vertices of the input graph and k the size of the clique that we are after. The function $g(k)$ bounding the size of k' in the parameterized reductions plays an important role in these types of proofs, as for example a reduction with $g(k)$ from k -CLIQUE yields a lower bound under ETH of $n^{o(g^{-1}(k))}$.

3 Result Type C: Precedence Constraints, Release Dates and Unit Processing Times

In this section we provide a fast algorithm for partial scheduling with release dates and unit processing times parameterized by the number k of scheduled jobs (Theorem 1). There exists a simple, but slow, algorithm with runtime $\mathcal{O}^*(2^{k^2})$ that already proves that this problem is in FPT parameterized by k : This algorithm branches k times on jobs that can be processed next. If more than k jobs are available at a step, then processing these jobs greedily is optimal. Otherwise, we can recursively try to schedule all non-empty subsets of jobs to schedule next, and a $\mathcal{O}^*(2^{k^2})$ time algorithm is obtained via a standard (bounded search-tree) analysis. To improve on this algorithm, we present a dynamic programming algorithm based on table entries indexed by antichains in the precedence graph G describing the precedence relations. Such an antichain describes the maximal jobs already scheduled in a partial schedule. Our key idea is that, to find an optimal solution, it is sufficient to restrict our attention to a subset of all antichains. This subset will be defined in terms of the *depth* of an antichain. With this algorithm we improve the runtime to $\mathcal{O}(8^k k(|V| + |E|))$.

By binary search, we can restrict attention to a variant of the problem that asks whether there is a feasible schedule with makespan at most C_{\max} , for a fixed universal deadline C_{\max} .

3.1 The Algorithm

We start by introducing our dynamic programming algorithm for $P|r_j, prec, p_j = 1|k$ -sched, C_{\max} . Let m be the number of machines available. We start with defining

the table entries. For a given antichain $A \subseteq V(G)$ and integer t we define

$$S(A, t) = \begin{cases} 1, & \text{if there exists a feasible schedule of makespan } t \text{ that processes } \text{pred}(A), \\ 0, & \text{otherwise.} \end{cases}$$

Computing the values of $S(A, t)$ can be done by trying all combinations of scheduling at most m jobs of A at time t and then checking whether all remaining jobs of $\text{pred}(A)$ can be scheduled in makespan $t - 1$. To do so, we also verify that all the jobs in A actually have a release date at or before t . Formally, we have the following recurrence for $S(A, t)$:

Lemma 1

$$S(A, t) = (A \subseteq V(G^t)) \wedge \bigvee_{X \subseteq A: |X| \leq m} S(A', t - 1) : A' = \max(\text{pred}(A) \setminus X).$$

Proof If $A \not\subseteq V(G^t)$, then there is a job $j \in A$ with $r_j > t$. And thus $S(A, t) = 0$.

For any $X \subseteq A$, X is a set of maximal elements with respect to $G[\text{pred}(A)]$, and consists of pair-wise incomparable jobs, since A is an antichain. So, we can schedule all jobs from X at time t without violating any precedence constraints. Define $A' = \max(\text{pred}(A) \setminus X)$ as the unique antichain such that $\text{pred}(A) \setminus X = \text{pred}(A')$. If $S(A', t - 1) = 1$ and $|X| \leq m$, we can extend the schedule of $S(A', t - 1)$ by scheduling all X at time t . In this way we get a feasible schedule processing all jobs of $\text{pred}(A)$ before or at time t . So if we find such an X with $|X| \leq m$ and $S(A', t - 1) = 1$, we must have $S(A, t) = 1$.

For the other direction, if for all $X \subseteq A$ with $|X| \leq m$, $S(A', t - 1) = 0$, then no matter which set $X \subseteq A$ we try to schedule at time t , the remaining jobs cannot be scheduled before t . Note that only jobs from A can be scheduled at time t , since those are the maximal jobs. Hence, there is no feasible schedule and $S(A, t) = 0$. \square


The above recurrence cannot be directly evaluated, since the number of different antichains of a graph can be big: there can be as many as $\binom{n}{k}$ different antichains with $|\text{pred}(A)| \leq k$, for example in the extreme case of an independent set. Even when we restrict our precedence graph to have out degree k , there could be k^k different antichains, for example in k -ary trees. To circumvent this issue, we restrict our dynamic programming algorithm only to a specific subset of antichains. To do this, we use the following new notion of the *depth* of an antichain.

Definition 3 Let A be an antichain. Define the depth (with respect to t) of A as

$$d^t(A) = |\text{pred}(A)| + |\min(G^t - \text{comp}(A))|.$$

We also denote $d(A) = d^{C_{\max}}(A)$.

The intuition behind this definition is that it quantifies the number of jobs that can be scheduled before (and including) A without violating precedence constraints. See Fig. 2 for an example of an antichain and its depth. We restrict the dynamic



$$d(A) = |\text{pred}(A)| + |\min(G - \text{comp}(A))| = 2 + 2$$

Fig. 2 Example of an antichain and its depth in a perfect 3-ary tree. We see that $|\text{pred}(A)| = 2$, but $d(A) = 4$. If $k = 2$, the dynamic programming algorithm will not compute $S(A, t)$ since $d(A) > k$. The only antichains with depth ≤ 2 are the empty set and the root node r on its own as a set. Indeed, $d(\emptyset) = d(\{r\}) = 1$. Note that for instances with $k = 2$, a feasible schedule may exist. If so, we will find that $R(\{r\}, 1) = 1$, which will be defined later. In this way, we can still find the antichain A as a solution

programming algorithm to only compute $S(A, t)$ for A satisfying $d^t(A) \leq k$. This ensures that we do not go ‘too deep’ into the precedence graph unnecessarily at the cost of a slow runtime.

Because of this restriction in the depth, it could happen that we check no antichains with k or more predecessors, while there are corresponding feasible schedules. It is therefore possible that for some antichains A with $d^t(A) > k$, there is a feasible schedule for all $\geq k$ jobs in $\text{pred}(A)$ before time C_{\max} , but the value $S(A, C_{\max})$ will not be computed. To make sure we still find an optimal schedule, we also compute the following condition $R(A, t)$ for all $t \leq C_{\max}$ and antichains A with $d^t(A) \leq k$:

$$R(A, t) = \begin{cases} 1, & \text{if there exists a feasible schedule with makespan at most } C_{\max} \text{ that processes } \text{pred}(A) \text{ on or before } t \text{ and processes jobs from } \min(G - \text{pred}(A)) \text{ after } t, \text{ with a total of } k \text{ jobs processed,} \\ 0, & \text{otherwise.} \end{cases}$$

By definition of $R(A, t)$, if $R(A, t) = 1$ for any A and $t \leq C_{\max}$, then we find a feasible schedule that processes k jobs on time.⁶ We show that there is an algorithm, namely $\text{fill}(A, t)$, that quickly computes $R(A, t)$. The algorithm $\text{fill}(A, t)$ does the following: first it checks if $S(A, t) = 1$ and if so, greedily schedules jobs from $\min(G - \text{pred}(A))$ after t in order of smallest release date. If $k - |\text{pred}(A)|$ jobs can be scheduled before C_{\max} , it returns ‘true’ ($R(A, t) = 1$). Otherwise, it returns ‘false’ ($R(A, t) = 0$).

Lemma 2 *There is an $\mathcal{O}(|V|k + |E|)$ time algorithm that, given an antichain A , integer t , and value $S(A, t)$, computes $R(A, t)$.*

⁶ The reverse direction is more difficult and postponed to Lemma 6.

Proof We show that $\text{fill}(A, t)$, defined above, fulfills all requirements. First we prove that if $\text{fill}(A, t)$ returns ‘true’, it follows that $R(A, t) = 1$. Since $S(A, t) = 1$, all jobs from $\text{pred}(A)$ can be finished at time t . Take that feasible schedule and process $k - |\text{pred}(A)|$ jobs from $\min(G - \text{pred}(A))$ between t and C_{\max} . This is possible because $\text{fill}(A, t)$ is true. All predecessors of jobs in $\min(G - \text{pred}(A))$ are in $\text{pred}(A)$ and therefore processed before t . Hence, no precedence constraints are violated and we find a feasible schedule with the requirements, i.e. $R(A, t) = 1$.

For the other direction, assume that $R(A, t) = 1$, i.e. we find a feasible schedule σ where exactly the jobs from $\text{pred}(A)$ are processed on or before t and only jobs from $\min(G - \text{pred}(A))$ are processed after t . Thus $S(A, t) = 1$. Define M as the set of jobs processed after t in σ . If M equals the set of jobs with the smallest release dates of $\min(G - \text{pred}(A))$, we can also process the jobs of M in order of increasing release dates. Then $\text{fill}(A, t)$ will be ‘true’, since M has size at least $k - |\text{pred}(A)|$. However, if M is not that set, we can replace a job which does not have one of the smallest $k - |\text{pred}(A)|$ release dates, by one which has and was not in M yet. This new set can then still be processed between $t + 1$ and C_{\max} because smaller release dates impose weaker constraints. We keep replacing until we end up with M being exactly the set of jobs with smallest release dates, which is then proved to be schedulable between t and C_{\max} . Hence, $\text{fill}(A, t)$ will return ‘true’.

Computing the set $\min(G - \text{pred}(A))$ can be done in $\mathcal{O}(|V| + |E|)$ time. Sorting them on release date can be done in $\mathcal{O}(|V|k)$ time, as there are at most k different release dates. Finally, greedily scheduling the jobs while checking feasibility can be done in $\mathcal{O}(|V|)$ time. Hence this algorithm runs in time $\mathcal{O}(|V|k + |E|)$. \square

Combining all steps gives us the algorithm as described in Algorithm 1. It remains to bound its runtime and argue its correctness.

```

1 foreach  $t = 1, \dots, C_{\max}$  do
2   Enumerate all antichains  $A$  in  $G^t$  with  $d^t(A) \leq k$  using Lemma 5
3   foreach antichain  $A$  in  $G^t$  with  $d^t(A) \leq k$  do
4     Compute  $S(A, t)$  using Lemma 1
5     if  $\text{fill}(S(A, t), A, t)$  then
6       return TRUE
7 return FALSE

```

Algorithm 1: Algorithm for $P|prec, p_j = 1|k$ -sched, C_{\max}

3.2 Runtime

To analyze the runtime of the dynamic programming algorithm, we need to bound the number of checked antichains. Recall that we only check antichains A with $d^t(A) \leq k$ for each time $t \leq C_{\max}$. We first analyze the number of antichains A with $d(A) \leq k$ in any graph and use this to upper bound the number of antichains checked at time t .

To analyze the number of antichains A with $d(A) \leq k$, we give an upper bound on this number via an upper bound on the number of maximal antichains. Recall from the

notations for posets, that for a maximal antichain A we have $\text{comp}(A) = V(G)$, and therefore $d(A) = |\text{pred}(A)|$. The following lemma connects the number of antichains and maximal antichains of bounded depth:

Lemma 3 *For any antichain A , there exists a maximal antichain A_{\max} such that $A \subseteq A_{\max}$ and $d(A) = d(A_{\max})$.*

Proof Let $A_{\max} = A \cup \min(G - \text{comp}(A))$. By definition, all elements in $\min(G - \text{comp}(A))$ are incomparable to each other and incomparable to any element of A . Hence A_{\max} is an antichain. Since $\text{comp}(A_{\max}) = V(G)$, A_{\max} is a maximal antichain. Moreover,

$$d(A) = |\text{pred}(A)| + |\min(G - \text{comp}(A))| = |\text{pred}(A_{\max})| = d(A_{\max}),$$

since the elements in $\min(G - \text{comp}(A))$ are minimal elements and all their predecessors are in $\text{pred}(A)$ besides themselves. □

For any (maximal) antichain A with $d(A) \leq k$, we derive that $|A| \leq k$ and so each maximal antichain of depth at most k has at most 2^k subsets. By Lemma 3, we see that each antichain is a subset of a maximal antichain with the same depth.

Corollary 1

$$|\{A : A \text{ antichain}, d(A) \leq k\}| \leq 2^k |\{A : A \text{ maximal antichain}, d(A) \leq k\}|.$$

This corollary allows us to restrict attention to only upper bounding the number of maximal antichains of bounded depth.

Lemma 4 *There are at most 2^k maximal antichains A with $d(A) \leq k$ in any precedence graph $G = (V, E)$, and they can be enumerated in $\mathcal{O}(2^k k (|V| + |E|))$ time.*

Proof Let $\mathcal{A}_k(G)$ be the set of maximal antichains in G with depth at most k . We prove that $|\mathcal{A}_k(G)| \leq 2^k$ for any graph G by induction on k . Clearly, $|\mathcal{A}_0(G)| \leq 1$ for any graph G , since the only antichain with $d(A) \leq 0$ is $A = \emptyset$ if $G = \emptyset$.

Let $k > 0$ and assume $|\mathcal{A}_j(G)| \leq 2^j$ for $j < k$ for any graph G . If we have a precedence graph G with minimal elements s_1, \dots, s_ℓ , we partition $\mathcal{A}_k(G)$ into $\ell + 1$ different sets $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{\ell+1}$. For $i = 1, \dots, \ell$, the set \mathcal{B}_i is defined as the set of maximal antichains A of depth at most k in which $\{s_{i'} : i' < i\} \subseteq A$, but $s_i \notin A$ (and no restrictions on elements in the set $\{s_{i'} : i' > i\}$). If $s_i \notin A$, then $s_i \in \text{pred}(A)$ since A is maximal, so any such maximal antichain has a successor of s_i in A . If we define S_j as the set of all successors of s_j (including s_j), we see that $\mathcal{B}_i = \mathcal{A}_{k-i} \left(G - \left(\bigcup_{i'=1}^{i-1} S_{i'} \cup \{s_i\} \right) \right)$. Indeed, if $A \in \mathcal{B}_i$, then $\{s_{i'} : i' < i\} \subseteq A$. Hence we can remove those elements and its successors from the graph, as they are comparable to any such antichain. Moreover, we can also remove s_i (but *not* its successors) from the graph, since it is in $\text{pred}(A)$. Thus \mathcal{B}_i is then exactly the set of maximal antichains with depth i less in the remaining graph. The set $\mathcal{B}_{\ell+1}$ is defined as all antichains not in some \mathcal{B}_i , which is all maximal antichains of A of depth at most

k for which $\{s_1, \dots, s_\ell\} \subseteq A$. Note that $\mathcal{B}_{\ell+1} = \{s_1, \dots, s_\ell\}$. We get the following recurrence relation:

$$|\mathcal{A}_k(G)| = \sum_{i=1}^{\ell} \left| \mathcal{A}_{k-i} \left(G - \left(\bigcup_{j=1}^{i-1} S_j \cup \{s_i\} \right) \right) \right| + 1, \tag{1}$$

since $|\mathcal{B}_{\ell+1}| = 1$. Notice that we may assume that $\ell \leq k$, because otherwise the depth of the antichain will be greater than k . Then if we use the induction hypothesis that $|\mathcal{A}_j(G)| \leq 2^j$ for $j < k$ for any graph G , we see by (1) that:

$$\begin{aligned} |\mathcal{A}_k(G)| &= \sum_{i=1}^{\ell} \left| \mathcal{A}_{k-i} \left(G - \left(\bigcup_{j=1}^{i-1} S_j \cup \{s_i\} \right) \right) \right| + 1, \\ &\leq 2^k \left(\sum_{i=1}^k \frac{1}{2^i} + \frac{1}{2^k} \right) \\ &= 2^k. \end{aligned}$$

The lemma follows since the above procedure can easily be modified in a recursive algorithm to enumerate the antichains, and by using a Breadth-First Search we can compute $G - \left(\bigcup_{j=1}^{i-1} S_j \cup \{s_i\} \right)$ in $\mathcal{O}(|V| + |E|)$ time. Thus, each recursion step takes $\mathcal{O}(k(|V| + |E|))$ time. □

Returning to (non-maximal) antichains, we see that we can enumerate all maximal antichains of depth at most k with Lemma 4 and by Corollary 1 we can find all antichains of depth at most k by taking all subsets of the found maximal antichains.

Corollary 2 *There are at most 4^k antichains A with $d(A) \leq k$ in any precedence graph $G = (V, E)$, and they can be enumerated within $\mathcal{O}(4^k(|V| + |E|))$ time.*

Notice that the runtime is indeed correct, as it dominates both the time needed for the construction of the set $\mathcal{A}_k(G)$ and the time needed for taking the subsets of $\mathcal{A}_k(G)$ (which is $2^k |\mathcal{A}_k(G)|$).

We now restrict the number of antichains A in G^t with $d^t(A) \leq k$. Take G^t to be the graph in Corollary 2 and notice that $d^t(A) = d(A)$ for any antichain A in G^t . By Corollary 2 we obtain Lemma 5.

Lemma 5 *For any t , there are at most 4^k antichains A with $d^t(A) \leq k$ in any precedence graph $G = (V, E)$, and they can be enumerated within $\mathcal{O}(4^k(|V| + |E|))$ time.*

To compute each $S(A, t)$, we look at a maximum of $\binom{k}{m} \leq 2^k$ different sets X . Computing the antichain A' such that $A' = \max(\text{pred}(A) \setminus X)$ takes $\mathcal{O}(|V| + |E|)$ time. After this computation, $R(A, t)$ is directly computed in $\mathcal{O}(|V|k + |E|)$ time. For each time $t \in \{1, \dots, C_{\max}\}$, there are at most 4^k different antichains A for which we compute $S(A, t)$ and $R(A, t)$. Since $C_{\max} \leq k$, we therefore have total runtime of $\mathcal{O}(4^k k (2^k (|V| + |E|) + (|V|k + |E|)))$. Hence, Algorithm 1 runs in time $\mathcal{O}(8^k k (|V| + |E|))$.

3.3 Correctness of Algorithm

To show that the algorithm described in Algorithm 1 indeed returns the correct answer, the following lemma is clearly sufficient:

Lemma 6 *A feasible schedule for k jobs with makespan at most C_{\max} exists if and only if $R(A, t) = 1$ for some $t \leq C_{\max}$ and antichain A with $d^t(A) \leq k$.*

Before we are able to prove Lemma 6, we need one more definition.

Definition 4 Let σ be a feasible schedule. Then $A(\sigma)$ is the antichain such that $\text{pred}(A(\sigma))$ is exactly the set of jobs that was scheduled in σ .

Equivalently, if X is the set of jobs processed by σ , then $A(\sigma) = \max(G[X])$.

Proof (Lemma 6) Clearly, if $R(A, t) = 1$ for some $t \leq C_{\max}$ and antichain A with $d^t(A) \leq k$, we have a feasible schedule with k jobs by definition of $R(A, t)$. Hence, it remains to prove that if a feasible schedule for k jobs exists, then $R(B, t) = 1$ for some $t \leq C_{\max}$ and antichain B with $d^t(B) \leq k$. Let

$$\Sigma^* = \{\sigma \mid \sigma \text{ is a feasible schedule that processes } k \text{ jobs and has a makespan of at most } C_{\max}\},$$

so Σ^* is the set of all possible solutions. Define

$$\sigma^* = \underset{\sigma}{\operatorname{argmin}}\{d(A(\sigma)) \mid \sigma \in \Sigma^*\},$$

i.e. σ^* is a schedule for which $A(\sigma^*)$ has minimal depth (with respect to C_{\max}). We now define t and B such that $R(B, t) = 1$.

- Let $t = \max\{t : \text{job not in } \max(G[\text{pred}(A(\sigma^*))]) \text{ was scheduled at time } t\}$, so from $t + 1$ and on, only maximal jobs (with respect to $G[\text{pred}(A(\sigma^*))]$) are scheduled.
- Let $M = \{x : \text{job } x \text{ was scheduled at } t + 1 \text{ or later in } \sigma^*\}$.
- Let $B = \max(\text{pred}(A(\sigma^*)) \setminus M)$, so $\text{pred}(B)$ is exactly the set of jobs scheduled on or before time t in σ^* .

See Fig. 3a for an illustration of these concepts. There are two cases to distinguish: $d^t(B) \leq k$. In this case we prove that $R(B, t) = 1$. The feasible schedule we are looking for in the definition of $R(B, t)$ is exactly σ^* . Indeed, all jobs from $\text{pred}(B)$ were finished at time t . Furthermore, all jobs in M are maximal, so all their predecessors are in $\text{pred}(B)$. Hence, $M \subseteq \min(G - \text{pred}(B))$. So, by definition $R(B, t) = 1$.

$d^t(B) > k$. In this case we prove that there is a schedule σ' such that $d(A(\sigma')) < d(A(\sigma^*))$, i.e. we find a contradiction to the fact that $d(A(\sigma^*))$ was minimal. This σ' can be found as follows: take schedule σ^* only up until time t . Let C be a subset of $\min(G^t - \text{comp}(B))$ such that $|C| = k - |\text{pred}(B)|$. This C can be found since $d^t(B) \geq k$. Process the jobs in C after time t in σ' . These can all be processed without precedence constraint or release date violations, since their predecessors were already scheduled and $C \subseteq G^t$. So, we find a feasible schedule that processes k jobs, called




Fig. 3 Visualization of the definitions of M and B and the schedule σ^* in the proof of Lemma 6 is shown in **a**. **b** Depicts the schedule σ' as chosen in the subcase $d(B) > k$. The grey boxes indicate which jobs are processed in the schedules. We will prove that $|D(A(\sigma'))| < |D(A(\sigma^*))|$

σ' . The choice of σ' is depicted in Fig. 3. Note that $C \subseteq \min(G^t - \text{comp}(B)) \subseteq \min(G - \text{comp}(B))$ and not all jobs of $\min(G - \text{comp}(B))$ are necessarily processed in σ' .

It remains to prove that $d(A(\sigma')) < d(A(\sigma^*))$. Define $D(A) = \text{pred}(A) \cup \min(G - \text{comp}(A))$ for any antichain A . So $D(A)$ is the set of jobs that contribute to $d(A)$ and so $|D(A)| = d(A)$. We will prove that $D(B) = D(A(\sigma')) \subset D(A(\sigma^*))$. This will be done in two steps, first we show that

$$D(B) = D(A(\sigma')) \subseteq D(A(\sigma^*)).$$

In the last step we prove $D(B) \neq D(A(\sigma^*))$, which gives us $d(A(\sigma')) < d(A(\sigma^*))$.

Notice that $C \subseteq D(B)$ since $C \subseteq \min(G - \text{comp}(B))$, hence $D(B) = D(B \cup C)$. Since $A(\sigma') = B \cup C$ it follows that $D(A(\sigma')) = D(B)$. Next we prove that $D(B) \subseteq D(A(\sigma^*))$. Clearly, if $x \in \text{pred}(B)$ then $x \in \text{pred}(A(\sigma^*))$. It remains to show that $x \in \min(G - \text{comp}(B))$ implies that $x \in D(A(\sigma^*))$. If $x \in \min(G - \text{comp}(B))$, then either $x \in M$ or $x \notin M$. If $x \in M$, then $x \in A(\sigma^*)$ so $x \in \text{pred}(A(\sigma^*))$. If $x \notin M$, then $x \notin \text{comp}(B \cup M)$ since x was a minimal element in $\min(G - \text{comp}(B))$. Since $A(\sigma^*) \subseteq B \cup M$, and thus $\text{comp}(A(\sigma^*)) \subseteq \text{comp}(B \cup M)$, we observe that $x \in \min(G - \text{comp}(A(\sigma^*)))$. We then conclude that $D(B) \subseteq D(A(\sigma^*))$.

We are left to show that $D(B) \neq D(A(\sigma^*))$. Remember that t was chosen such that there is a job processed at time t that was not in $\max(G[\text{pred}(A(\sigma^*))])$. In other words, there was a job $x \in B$ in σ^* at time t with $y \in M$ such that $y > x$. Note that $y \notin D(B)$, since $y \in M$, so y is not in $\text{pred}(B)$ and y is clearly comparable to x . However, $y \in D(A(\sigma^*))$, so we find that $d(A(\sigma')) = d(B) < d(A(\sigma^*))$. Hence, we found a schedule with smaller $d(A(\sigma'))$, which leads to a contradiction. \square

4 Result Types B and D: One Machine and Precedence Constraints

In this section we show that Algorithm 1 cannot be even slightly generalized further: if we allow job-dependent deadlines or non-unit processing times, the problem becomes $W[1]$ -hard parameterized by k and cannot be solved in $n^{o(k/\log k)}$ time unless the ETH fails. In the following reductions we reduce to a variant of the scheduling problems that asks whether there is a feasible schedule with makespan at most C_{\max} , where C_{\max} is given as input. If for a given instance such a schedule exists, we call the instance a yes instance, and a no instance otherwise. We may restrict ourselves to this variant because of binary search.

4.1 Job-Dependent Deadlines

The fact that combining precedence constraints with job-dependent deadlines makes the problem $W[1]$ -hard, is a direct consequence from the fact that $1|prec, p_j = 1|\sum_j U_j$ is $W[1]$ -hard, parameterized by $n - \sum_j U_j = k$ where n is the number of jobs [9]. It is important to notice that the notation of these problems implies that each job can have its own deadline. Hence, we conclude from this that $1|d_j, prec, p_j = 1|k$ -sched, C_{\max} is $W[1]$ -hard parameterized by k . This is a reduction from k -CLIQUE that yields a quadratic blow-up on the parameter, giving a lower bound on algorithms for the problem of $n^{\Omega(\sqrt{k})}$. Based on the Exponential Time Hypothesis, we now sharpen this lower bound with a reduction from 3-COLORING:

Theorem 3 $1|d_j, prec, p_j = 1|k$ -sched, C_{\max} is $W[1]$ -hard parameterized by k . Furthermore, there is no algorithm solving $1|d_j, prec, p_j = 1|k$ -sched, C_{\max} in $2^{o(n)}$ time where n is the number of jobs, assuming ETH.

Proof The proof will be a reduction from 3-COLORING, for which no $2^{o(|V|+|E|)}$ algorithm exists under the Exponential Time Hypothesis [7, pages 471–473]. Let the graph $G = (V, E)$ be the instance of 3-COLORING with $|V| = n'$ and $|E| = m'$. We label the vertices $v_1, \dots, v_{n'}$ and the edges $e_1, \dots, e_{m'}$. We then create the following instance for $1|d_j, prec, p_j = 1|k$ -sched, C_{\max} .

- For each vertex $v_i \in V$, create 6 jobs:

- v_i^1, v_i^2 and v_i^3 with deadline $d_{v_i} = i$,
- w_i^1, w_i^2 and w_i^3 with deadline $d_{w_i} = n' + 2m' + 1 - i$,

add precedence constraints $v_i^1 < w_i^1, v_i^2 < w_i^2$ and $v_i^3 < w_i^3$. These jobs represent which color for each vertex will be chosen (for instance if v_i^1 and w_i^1 are processed, vertex i gets color 1).

- For each edge $e_j \in E$, create 12 jobs:

- $e_j^{12}, e_j^{13}, e_j^{21}, e_j^{23}, e_j^{31}$ and e_j^{32} with deadline $d_{e_j} = n' + j$,
- $f_j^{12}, f_j^{13}, f_j^{21}, f_j^{23}, f_j^{31}$ and f_j^{32} with deadline $d_{f_j} = n' + m' + 1 - j$,

add precedence constraints $e_j^{ab} < f_j^{ab}$ for $a, b \in \{1, 2, 3\}$ with $a \neq b$. These jobs represent what the colors of the endpoints of an edge will be. So if the jobs e_j^{ab}

and f_j^{ab} are processed for $e = \{u, v\}$, then vertex u has color a and vertex v has color b . Since the endpoints should have different colors, the jobs e_j^{aa} and f_j^{aa} do not exist.

- For each e_j^{ab} with $e = \{v_i, v_{i'}\}$ and $a, b \in \{1, 2, 3\}$ with $a \neq b$, add the precedence constraints $v_i^a < e_j^{ab}$ and $v_{i'}^b < e_j^{ab}$.
- Set $C_{\max} = k = 2n' + 2m'$.

We now prove that the created instance is a yes instance if and only if the original 3-COLORING instance is a yes instance. Assume that there is a 3-coloring of the graph $G = (V, E)$. Then there is also a feasible schedule: For each vertex v_i with color a , process the jobs v_i^a and w_i^a at their respective deadlines. For each edge $e_j = \{u, v\}$ with u colored a and v colored b , process the jobs e_j^{ab} and f_j^{ab} exactly at their respective deadlines. Notice that because it is a 3-coloring, each edge has endpoints of different colors, so these jobs exist. Also note that no two jobs were processed at the same time. Exactly $2n' + 2m'$ jobs were processed before time $2n' + 2m'$. Furthermore, no precedence constraints were violated.

For the other direction, assume that we have a feasible schedule in our created instance of $1|d_j, \text{prec}, p_j = 1|k\text{-sched}, C_{\max}$. Let $\mathcal{V}_i = \{v_i^1, v_i^2, v_i^3\}$, $\mathcal{W}_i = \{w_i^1, w_i^2, w_i^3\}$ for all $i = 1, \dots, n'$, and let $\mathcal{E}_j = \{e_j^{12}, e_j^{13}, e_j^{21}, e_j^{23}, e_j^{31}, e_j^{32}\}$ and $\mathcal{F}_j = \{f_j^{12}, f_j^{13}, f_j^{21}, f_j^{23}, f_j^{31}, f_j^{32}\}$ for all $j = 1, \dots, m'$. We show by induction on i that out of each of the sets $\mathcal{V}_i, \mathcal{W}_i, \mathcal{E}_j$ and \mathcal{F}_j , exactly one job was scheduled at its deadline.

Since we have a feasible schedule, at time $2m' + 2n'$ one of the jobs of \mathcal{W}_1 must be scheduled, since they are the only jobs with a deadline greater than $2n + 2m - 1$. If w_1^a was scheduled at time $2m' + 2n'$, then the job v_1^a must be processed at time 1 because of precedence constraints and since its deadline is 1. No other jobs from \mathcal{V}_1 and \mathcal{W}_1 can be processed, due to their deadlines and precedence constraints.

Now assume that all sets $\mathcal{V}_1, \dots, \mathcal{V}_{i-1}, \mathcal{W}_1, \dots, \mathcal{W}_{i-1}$ have exactly one job scheduled at their respective deadline, and no more can be processed. Since we have a feasible schedule, one job should be scheduled at time $2n' + 2m' - (i - 1)$. However, since no more jobs from $\mathcal{W}_1, \dots, \mathcal{W}_{i-1}$ can be scheduled, the only possible jobs are from \mathcal{W}_i since they are the only other jobs with a deadline greater than $2n' + 2m' - i$. However, if w_i^a was scheduled at time $2n' + 2m' - (i - 1)$, then the job v_i^a must be processed at time i because of precedence constraints, its deadline at i and because at times $1, \dots, i - 1$ other jobs had to be processed. Also, no other job from \mathcal{V}_i can be processed in the schedule, since they all have deadline i . As a consequence, no other jobs from \mathcal{W}_i can be processed, as they are restricted to precedence constraints. So the statement holds for all sets \mathcal{V}_i and \mathcal{W}_i . In the exact same way, one can conclude the same about all sets \mathcal{E}_j and \mathcal{F}_j .

Because of this, we see that each job and each vertex have received a color from the schedule. They must form a 3-coloring, because a job from \mathcal{E}_j could only be processed if the two endpoints got two different colors. Hence, the 3-COLORING instance is a yes instance.

As $k = 2n' + 2m'$ we therefore conclude there is no $2^{o(n)}$ algorithm under the ETH. □

Note that this bound significantly improves the old lower bound of $2^{\Omega(\log n \sqrt{k})}$ implied by the the reduction from k -CLIQUE reduction. Since $k \leq n$ and $n/\log n$ is an increasing function, Theorem 3 implies that

Corollary 3 *Assuming ETH, there is no algorithm solving $1|d_j, prec, p_j = 1|k$ -sched, C_{\max} in $n^{o(k/\log(k))}$ where n is the number of jobs.*

4.2 Non-unit Processing Times

We show that having non-unit processing times combined with precedence constraints make the problem $W[1]$ -hard even on one machine. The proof of Theorem 4 heavily builds on the reduction from k -CLIQUE to k -TASKS ON TIME by Fellows and McCartin [9].

Theorem 4 $1|prec|k$ -sched, C_{\max} is $W[1]$ -hard, parameterized by k , even when $p_j \in \{1, 2\}$ for all jobs j .

Proof The proof is a reduction from k -CLIQUE. We start with $G = (V, E)$, an instance of k -CLIQUE. For each vertex $v \in V$, create a job J_v with processing time $p(J_v) = 2$. For each edge $e \in E$, create a job J_e with processing time $p(J_e) = 1$. Now for each edge (u, v) , add the following two precedence relations: $J_u < J_e$ and $J_v < J_e$, so before one can process a job associated with an edge, both jobs associated with the endpoints of that edge need to be finished. Now let $k' = k + \frac{1}{2}k(k - 1)$ and $C_{\max} = 2k + \frac{1}{2}k(k - 1)$. We will now prove that $1|prec|k'$ -sched, C_{\max} is a yes instance if and only if the k -CLIQUE instance is a yes instance.

Assume that the k -CLIQUE instance is a yes instance, then process first the k jobs associated with the vertices of the k -clique. Next process the $\frac{1}{2}k(k - 1)$ jobs associated with the edges of the k -clique. In total, $k + \frac{1}{2}k(k - 1) = k'$ jobs are now processed with a makespan of $2k + \frac{1}{2}k(k - 1)$. Hence, the instance of $1|prec|k'$ -sched, C_{\max} is a yes instance.

For the other direction, assume $1|prec|k'$ -sched, C_{\max} to be a yes instance, so there exists a feasible schedule. For any feasible schedule, if one schedules ℓ jobs associated with vertices, then at most $\frac{1}{2}\ell(\ell - 1)$ jobs associated with edges can be processed, because of the precedence constraints. However, because $k' = k + \frac{1}{2}k(k - 1)$ jobs were done in the feasible schedule before $C_{\max} = 2k + \frac{1}{2}k(k - 1)$, at most k jobs associated with vertices can be processed, because they have processing time of size 2. Hence, we can conclude that exactly k vertex-jobs and $\frac{1}{2}k(k - 1)$ edge-jobs were processed. Hence, there were k vertices connected through $\frac{1}{2}k(k - 1)$ edges, which is a k -clique. □

The proofs of Theorem 6 and Corollary 4 are reductions from PARTITIONED SUB-GRAPH ISOMORPHISM. Let $P = (V', E')$ be a ‘pattern’ graph, $G = (V, E)$ be a ‘target’ graph, and $\chi : V \rightarrow V'$ a ‘coloring’ of the vertices of G with elements from P . A χ -colorful P -subgraph of G is a mapping $\varphi : V' \rightarrow V$ such that (1) for each $\{u, v\} \in E'$ it holds that $\{\varphi(u), \varphi(v)\} \in E$ and (2) for each $u \in V'$ it holds that $\chi(\varphi(u)) = u$. If χ and G are clear from the context they may be omitted in this definition.

Definition 5 (PARTITIONED SUBGRAPH ISOMORPHISM) Given graphs $G = (V, E)$ and $P = (V', E')$, $\chi : V \rightarrow V'$, determine whether there is a χ -colorful P -subgraph of G .

Theorem 5 (Marx [19]) PARTITIONED SUBGRAPH ISOMORPHISM cannot be solved in $n^{o(|E'|/\log |E'|)}$ time assuming the Exponential Time Hypothesis (ETH), where n is the size of the input.

We will now reduce PARTITIONED SUBGRAPH ISOMORPHISM to $1|prec, r_j|k$ -sched, C_{\max} .

Theorem 6 $1|prec, r_j|k$ -sched, C_{\max} cannot be solved in $n^{o(k/\log k)}$ time assuming the Exponential Time Hypothesis (ETH).

Proof Let $G = (V, E)$, $P = (V', E')$ and $\chi : V \rightarrow V' = \{1, \dots, s\}$. Define for $i = 0, \dots, s$ the following important time stamps:

$$t_i := \sum_{j=1}^i 3^{s+1-j}.$$

Construct the following jobs for the instance of the $1|prec, r_j|k$ -sched, C_{\max} problem:

- For $i = 1, \dots, s$:
 - For each vertex $v \in V$ such that $\chi(v) = i$, create a job J_v with processing time $p(J_v) = 3^{s+1-i}$ and release date t_{i-1} .
- For each $\{v, w\} \in E$ such that $\{\chi(v), \chi(w)\} \in E'$, create a job $J_{\{v,w\}}$ with $p(J_{\{v,w\}}) = 1$ and release date t_s . Add precedence constraints $J_v \prec J_{\{v,w\}}$ and $J_w \prec J_{\{v,w\}}$.

Then ask whether there exists a solution to the scheduling problem for $k = s + |E'|$ with makespan $C_{\max} = t_s + |E'|$.

Let the PARTITIONED SUBGRAPH ISOMORPHISM instance be a yes-instance and let $\varphi : V' \rightarrow V$ be a colorful P -subgraph. We claim the following schedule is feasible:

- For $i = 1, \dots, s$:
 - Process $J_{\varphi(i)}$ at its release date t_{i-1} .
- Process for each $\{i, i'\} \in E'$ the job $J_{\{\varphi(i), \varphi(i')\}}$ somewhere in the interval $[t_s, t_s + |E'|]$.

Notice that all jobs are indeed processed after their release date and that in total there are $k = s + |E'|$ jobs processed before $C_{\max} = t_s + |E'|$. Furthermore, all precedence constraints are respected as any edge job is processed after both its predecessors. Also, the edge jobs $J_{\{\varphi(i), \varphi(i')\}}$ must exist, as φ is a properly colored P -subgraph. Therefore, we can conclude that indeed this schedule is feasible.

For the other direction, assume that there is a solution to the created instance of $1|prec, r_j|k$ -sched, C_{\max} . Define $J_i = \{J_v : \chi(v) = i\}$. We will first prove that at most one job from each set J_i can be processed in a feasible schedule. To do this, we first

prove that at most one job from each set J_i can be processed before t_s . Any job in J_i has release date $t_{i-1} = \sum_{j=1}^{i-1} 3^{s+1-j}$. Therefore, there is only $t_s - t_{i-1} = \sum_{j=i}^s 3^{s+1-j}$ time left to process the jobs from J_i before time t_s . However, the processing time of any job in J_i is 3^{s+1-i} , and since $2 \cdot 3^{s+1-i} > \sum_{j=i}^s 3^{s+1-j}$, at most one job from J_i can be processed before t_s . Since all jobs not in some J_i have their release date at t_s , at most s jobs are processed at time t_s . Thus, at time t_s , there are $|E'|$ time units left to process $|E'|$ jobs, because of the choice of k and makespan. Hence, the only way to get a feasible schedule is to process exactly one job from each set J_i at its respective release date and process exactly $|E'|$ edge jobs after t_s .

Let v^i be the vertex, such that J_v was processed in the feasible schedule with color i . We will show that $\varphi : V' \rightarrow V$, defined as $\varphi(i) = v^i$, is a properly colored P -subgraph of G . Hence, we are left to prove that for each $\{i, i'\} \in E'$, the edge $\{\varphi(i), \varphi(i')\} \in E$, i.e. that for each $\{i, i'\} \in E'$, the job $J_{\{\varphi(i), \varphi(i')\}}$ was processed. Because only the vertex jobs $J_{\varphi(1)}, J_{\varphi(2)}, \dots, J_{\varphi(s)}$ were processed, the precedence constraints only allow for edge jobs $J_{\{\varphi(i), \varphi(i')\}}$ to be processed. We created edge job $J_{\{v, w\}}$ if and only if $\{v, w\} \in E$ and $\{\chi(v), \chi(w)\} \in E'$, hence the $|E'|$ edge jobs have to be exactly the edge jobs $J_{\{\varphi(i), \varphi(i')\}}$ for $\{i, i'\} \in E'$. Therefore, we proved indeed that φ is a colorful P -subgraph of G .

Notice that $k = s + |E'| \leq 3|E'|$ as we may assume the number of vertices in P is at most $2|E'|$. The given bound follows. □

Corollary 4 $2|prec|k$ -sched, C_{\max} cannot be solved in $n^{o(k/\log k)}$ time assuming the Exponential Time Hypothesis (ETH).

Proof We can use the same idea for the reduction from PARTITIONED SUBGRAPH ISOMORPHISM as in the proof of Theorem 6, except for the release dates, as they are not allowed in this type of scheduling problem. To simulate the release dates, we use the second machine as a release date machine, meaning that we will create a job for each upcoming release date and will require these new jobs to be processed. More formally: For $i = 1, \dots, s$, create a job J_{r_i} with processing time 3^{s+1-i} and precedence constraints $J_{r_i} < J$ for any job J that had release date t_i in the original reduction. Furthermore, let $J_{r_i} < J_{r_{i+1}}$. Then we add $|E'|$ jobs J' with processing time 1 and with precedence relations $J_{r_s} < J'$. We then ask whether there exists a feasible schedule with $k = 2s + 2|E'|$ and with makespan $t_s + |E'|$. All newly added jobs are required in any feasible schedule and therefore, all other arguments from the previous reduction also hold. Finally, note that k is again linear in $|E'|$. □

5 Result Type G: k -scheduling without Precedence Constraints

The problem $P||k$ -sched, C_{\max} cannot be solved in $\mathcal{O}^*(2^{o(k)})$ time assuming the ETH, since there is a reduction to SUBSET SUM for which $2^{o(n)}$ algorithms were excluded by Jansen et al. [15].

We show that the problem is fixed-parameter tractable with a matching runtime in k , even in the case of unrelated machines, release dates and deadlines, denoted by $R|r_j, d_j|k$ -sched, C_{\max} .

Theorem 7 $R|r_j, d_j|k$ -sched, C_{\max} is fixed-parameter tractable in k and can be solved in $\mathcal{O}^*((2e)^k k^{\mathcal{O}(\log k)})$ time.

Proof We give an algorithm that solves any instance of $R|r_j, d_j|k$ -sched, C_{\max} within $\mathcal{O}^*((2e)^k k^{\mathcal{O}(\log k)})$ time. The algorithm is a randomized algorithm that uses the color coding method; it can be derandomized as described by Alon et al. [1]. The algorithm first (randomly) picks a coloring $c : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$, so each job is given one of the k available colors. We then compute whether there is a feasible colorful schedule, i.e. a feasible schedule that processes exactly one job of each color. If this colorful schedule can be found, then it is possible to schedule at least k jobs before C_{\max} . □

Given a coloring c , we compute whether there exists a colorful schedule in the following way. Define for $1 \leq i \leq m$ and $X \subseteq \{1, \dots, k\}$:

$$B_i(X) = \text{minimum makespan of all schedules on machine } i \text{ processing } |X| \text{ jobs, each from a different color in } X.$$

Clearly $B_i(\emptyset) = 0$, and all values $B_i(X)$ can be computed in $\mathcal{O}(2^k n)$ time using the following:

Lemma 7 Let $\min\{\emptyset\} = \infty$. Then

$$B_i(X) = \min_{\ell \in X} \min_{j:c(j)=\ell} \{C_j = \max\{r_j, B_i(X \setminus \{\ell\})\} + p_{ij} : C_j \leq d_j\}.$$

Proof In a schedule on one machine with $|X|$ jobs using all colors from X , one job should be scheduled as last, defining the makespan. So for all possible jobs j , we compute what the minimal end time would be if j was scheduled at the end of the schedule. This j cannot start before its release date or before all other colors are scheduled. □

Next, define for $1 \leq i \leq m$ and $X \subseteq [k]$, $A_i(X)$ to be 1 if $B_i(X) \leq C_{\max}$, and to be 0 otherwise. So $A_i(X) = 1$ if and only if $|X|$ jobs, each from a different color of X , can be scheduled on machine i before C_{\max} . A colorful feasible schedule exists if and only if there is some partition X_1, \dots, X_m of $\{1, \dots, k\}$ such that $\prod_{i=1}^m A_i(X_i) = 1$. The subset convolution of two functions is defined as $(A_i * A_{i'}) (X) = \sum_{Y \subseteq X} A_i(Y) A_{i'}(X \setminus Y)$. Then there is some partition X_1, \dots, X_m of $\{1, \dots, k\}$ such that $\prod_{i=1}^m A_i(X_i) = 1$ if and only if $(A_1 * \dots * A_m)(\{1, \dots, k\}) > 0$. The value of $(A_1 * \dots * A_m)(\{1, \dots, k\}) > 0$ can be computed in $2^k k^{\mathcal{O}(1)}$ time using fast subset convolution [3].

An overview of the randomized algorithm is given in Algorithm 2. If the k jobs that are processed in an optimal solution are all in different colors, the algorithm outputs true. By standard analysis, k jobs are all assigned different colors with probability at least $1/e^k$, and thus e^k independent trials to boost the error probability of the algorithm to at most $1/2$.

By using the standard methods by Alon et al. [1], Algorithm 2 can be derandomized. □

```

1 For a given coloring  $c$ :
2 foreach  $i = 1, \dots, m$  do
3   foreach  $X \subseteq \{1, \dots, k\}$  in order of increasing size do
4     Compute  $B_i(X)$  using Lemma 7.
5     Set  $A_i(X) = 1$  if  $B_i(X) \leq C_{\max}$ , set  $A_i(X) = 0$  otherwise.
6 Compute  $(A_1 * \dots * A_m)(\{1, \dots, k\})$  using fast subset convolution [3].
7 if  $(A_1 * \dots * A_m)(\{1, \dots, k\}) > 0$  then
8   return True
9 return False
    
```

Algorithm 2: Algorithm for solving $R|r_j, d_j|k$ -sched, C_{\max}

6 Argumentation of the Results in Table 1

For completeness and the readers convenience, we explain in this section for each row of Table 1 how the upper and lower bounds are obtained.

First notice that the most general variant $R|r_j, d_j, \text{prec}|k$ -sched, C_{\max} can be solved in $n^{\mathcal{O}(k)}$ time as follows: Guess for each machine the set of jobs that are scheduled on it, and guess how they are ordered in an optimal solution, to get sequences $\sigma_1, \dots, \sigma_m$ with a joint length equal to k . For each such $(\sigma_1, \dots, \sigma_m)$, run the following simple greedy algorithm to determine whether the minimum makespan achieved by a feasible schedule that schedules for each machine i the jobs as described in σ_i : Iterate $t = 1, \dots, n$ and schedule the job $\sigma_i(t)$ at machine i as early as possible without violating release dates/deadline and precedence constraints (if this is not possible, return NO). Since each optimal schedule can be assumed to be normalized in the sense that no single job can be executed earlier, it is easy to see that this algorithm always returns an optimal schedule for some choice of $\sigma_1, \dots, \sigma_m$. Since there are only $n^{\mathcal{O}(k)}$ different sequences $\sigma_1, \dots, \sigma_m$ of combined length k , the runtime follows.

- Cases 1–2** The polynomial time algorithms behind result [A] are obtained by a straightforward greedy algorithm: For $1|r_j, \text{prec}, p_j = 1|k$ -sched, C_{\max} , build the schedule from beginning to end, and schedule an arbitrary job if any is available; otherwise wait until one becomes available.
- Cases 3–4, 7–8** The given lower bound is by Corollary 3.
- Cases 5–6** The upper bound is by the algorithm of Theorem 1. The lower bound is due to reduction by Jansen et al. [14]. In particular, if no subexponential time algorithm for the BICLIQUE problem exists, there exist no algorithms in $n^{o(k)}$ time for these problems.
- Case 9** The lower bound is by Theorem 4, which is a reduction from k - CLIQUE and heavily builds on the reduction from k - CLIQUE to k -TASKS ON TIME by Fellows and McCartin [9]. This reduction increases the parameter k to $\Omega(k^2)$, hence the lower bound of $n^{o(\sqrt{k})}$.
- Cases 10–20** The given lower bound is by Theorem 6, which is a reduction from PARTITIONED SUBGRAPH ISOMORPHISM. It is conjectured that there exist no algorithms solving PARTITIONED SUBGRAPH ISOMORPHISM in $n^{o(k)}$ time assuming ETH, which would imply that the $n^{\mathcal{O}(k)}$ algorithm for these problems cannot be improved significantly.

- Cases 21–28** Result [E] is established by a simple greedy algorithm that always schedules an available job with the earliest deadline.
- Cases 29–31** Result [F] is a consequence of Moore’s algorithm [23] that solves the problem $1||\sum_j U_j$ in $\mathcal{O}(n \log n)$ time. The algorithm creates a sequence j_1, \dots, j_n of all jobs in earliest due date order. It then repeats the following steps: It tries to process the sequence (in the given order) on one machine. Let j_i be the first job in the sequence that is late. Then a job from j_1, \dots, j_i with maximal processing time is removed from the sequence. If all jobs are on time, it returns the sequence followed by the jobs that have been removed from the sequence. Notice that this also solves the problem $1|r_j|k$ -sched, C_{\max} , by reversing the schedule and viewing the release dates as the deadlines.
- Cases 32** The lower bound for this problem is a direct consequence of the reduction from KNAPSACK to $1|r_j|\sum_j U_j$ by Lenstra et al. [17], which is a linear reduction. Jansen et al. [15] showed that SUBSET SUM (and thus also KNAPSACK) cannot be solved in $2^{o(n)}$ time assuming ETH.
- Cases 33–40** Since $2||C_{\max}$ is equivalent to SUBSET SUM and can therefore not be solved in $2^{o(n)}$ time assuming ETH, as shown by Jansen et al. [15]. Therefore, its generalizations, in particular those mentioned in cases 33–40, have the same lower bound on run times assuming ETH. The upper bound is by the algorithm of Theorem 7.

7 Concluding Remarks

We classify all studied variants of partial scheduling parameterized by the number of jobs to be scheduled to be either in P, NP-complete and fixed-parameter tractable by k , or W[1]-hard parameterized by k . Our main technical contribution is an $\mathcal{O}(8^k k(|V| + |E|))$ time algorithm for $P|r_j, \text{prec}, p_j = 1|k$ -sched, C_{\max} .

In a fine-grained sense, the cases we left open are cases 3–20 from Table 1. We believe in fact algorithms in rows 5–6 and 10–20 are optimal: An $n^{o(k)}$ time algorithm for any case from result type [C] or [D] would imply either a $2^{o(n)}$ time algorithm for BICLIQUE or an $n^{o(k)}$ time algorithm for PARTITIONED SUBGRAPH ISOMORPHISM, which both would be surprising. It would be interesting to see whether for any of the remaining cases with precedence constraints and unit processing times a ‘subexponential’ time algorithm exists.

A related case is $P3|\text{prec}, p_j = 1|C_{\max}$ (where $P3$ denotes three machines). It is a famously hard open question (see e.g. [10]) whether this can be solved in polynomial time, but maybe it is doable to try to solve this question in subexponential time, e.g. $2^{o(n)}$?

Acknowledgements The authors would like to thank the anonymous referees for useful comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included

in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alon, N., Yuster, R., Zwick, U.: Color-coding. *J. ACM* **42**(4), 844–856 (1995)
- Bessy, S., Giroudeau, R.: Parameterized complexity of a coupled-task scheduling problem. *J. Sched.* **22**(3), 305–313 (2019)
- Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: Fourier meets Möbius: fast subset convolution. In: *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, pp. 67–74. ACM (2007)
- Bodlaender, H.L., Fellows, M.R.: W[2]-hardness of precedence constrained k -processor scheduling. *Oper. Res. Lett.* **18**(2), 93–97 (1995)
- Chuzhoy, J., Ostrovsky, R., Rabani, Y.: Approximation algorithms for the job interval selection problem and related scheduling problems. *Math. Oper. Res.* **31**(4), 730–738 (2006)
- Cygan, M., Pilipczuk, M., Pilipczuk, M., Wojtaszczyk, J.O.: Scheduling partially ordered jobs faster than 2^n . *Algorithmica* **68**(3), 692–714 (2014)
- Cygan, M., Fomin, F.V., Kowalik, Ł., Lokshantov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: *Parameterized Algorithms*, vol. 3. Springer, Berlin (2015)
- Eun, J., Sung, C.S., Kim, E.S.: Maximizing total job value on a single machine with job selection. *J. Oper. Res. Soc.* **68**(9), 998–1005 (2017)
- Fellows, M.R., McCartin, C.: On the parametric complexity of schedules to minimize tardy tasks. *Theoret. Comput. Sci.* **298**(2), 317–324 (2003)
- Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York (1979)
- Graham, R.L., Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A.H.G.: Optimization and approximation in deterministic sequencing and scheduling: a survey. *Ann. Discret. Math.* **5**(2), 287–326 (1979)
- Gupta, A., Krishnaswamy, R., Kumar, A., Segev, D.: Scheduling with outliers. In: Dinur, I., Jansen, K., Naor, J., Rolim, J.D.P. (eds.) *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21–23, 2009. *Proceedings, Lecture Notes in Computer Science*, vol. 5687, pp. 149–162. Springer, Berlin (2009). https://doi.org/10.1007/978-3-642-03685-9_12
- Hermelin, D., Mnich, M., Omlor, S.: Single machine batch scheduling to minimize the weighted number of tardy jobs. *CoRR* (2019)
- Jansen, K., Land, F., Kaluza, M.: Precedence scheduling with unit execution time is equivalent to parametrized biclique. In: *International Conference on Current Trends in Theory and Practice of Informatics*, pp. 329–343. Springer, Berlin (2016)
- Jansen, K., Land, F., Land, K.: Bounding the running time of algorithms for scheduling and packing problems. *SIAM J. Discret. Math.* **30**(1), 343–366 (2016)
- Koullamas, C., Panwalkar, S.S.: A note on combined job selection and sequencing problems. *Nav. Res. Logist.* **60**(6), 449–453 (2013)
- Lenstra, J.K., Rinnooy Kan, A.H.G., Brucker, P.: Complexity of machine scheduling problems. *Ann. Discret. Math.* **1**, 343–362 (1977)
- Lenté, C., Liedloff, M., Soukhal, A., T'Kindt, V.: Exponential algorithms for scheduling problems. *Tech. rep.* (2014). <https://hal.archives-ouvertes.fr/hal-00944382>
- Marx, D.: Can you beat treewidth? *Theory Comput.* **6**, 85–112 (2010)
- Megow, N., Mnich, M., Woeginger, G.: *Lorentz Workshop 'Scheduling Meets Fixed-Parameter Tractability'* (2019)
- Mnich, M., van Bevern, R.: Parameterized complexity of machine scheduling: 15 open problems. *Comput. Oper. Res.* (2018)
- Mnich, M., Wiese, A.: Scheduling and fixed-parameter tractability. *Math. Program.* **154**(1–2), 533–562 (2015)

23. Moore, J.M.: An n job, one machine sequencing algorithm for minimizing the number of late jobs. *Manag. Sci.* **15**(1), 102–109 (1968)
24. Pinedo, M.L.: *Scheduling: Theory, Algorithms, and Systems*, 3rd edn. Springer Publishing Company, Incorporated (2008)
25. Sgall, J.: Open problems in throughput scheduling. In: Epstein, L., Ferragina, P. (eds.) *Algorithms—ESA 2012—20th Annual European Symposium*, Ljubljana, Slovenia, September 10–12, 2012. *Proceedings, Lecture Notes in Computer Science*, vol. 7501, pp. 2–11. Springer, Berlin (2012). https://doi.org/10.1007/978-3-642-33090-2_2
26. Shabtay, D., Gaspar, N., Kaspi, M.: A survey on offline scheduling with rejection. *J. Sched.* **16**(1), 3–28 (2013)
27. van Bevern, R., Mnich, M., Niedermeier, R., Weller, M.: Interval scheduling and colorful independent sets. *J. Sched.* **18**(5), 449–469 (2015)
28. van Bevern, R., Bredereck, R., Bulteau, L., Komusiewicz, C., Talmon, N., Woeginger, G.J.: Precedence-constrained scheduling problems parameterized by partial order width. In: *Discrete Optimization and Operations Research—9th International Conference, DOOR 2016, Vladivostok, Russia, September 19–23, 2016*, *Proceedings*, pp. 105–120 (2016)
29. Williamson, D.P., Shmoys, D.B.: *The Design of Approximation Algorithms*. Cambridge University Press, Cambridge (2011)
30. Yang, B., Geunes, J.: A single resource scheduling problem with job-selection flexibility, tardiness costs and controllable processing times. *Comput. Ind. Eng.* **53**(3), 420–432 (2007)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.