Scheduling meets n-fold Integer Programming[⋆]

Dušan Knop and Martin Koutecký

Department of Applied Mathematics (KAM),
Charles University in Prague, Czech Republic.
{`knop, koutecky`}`@kam.mff.cuni.cz`

Abstract. Scheduling problems are fundamental in combinatorial optimization. Much work has been done on approximation algorithms for NP-hard cases, but relatively little is known about exact solutions when some part of the input is a fixed parameter. In 2014, Mnich and Wiese initiated a systematic study in this direction.

In this paper we continue this study and show that several additional cases of fundamental scheduling problems are fixed parameter tractable for some natural parameters. Our main tool is n-fold integer programming, a recent variable dimension technique which we believe to be highly relevant for the parameterized complexity community. This paper serves to showcase and highlight this technique.

Specifically, we show the following four scheduling problems to be fixed-parameter tractable, where p_{max} is the maximum processing time of a job and w_{max} is the maximum weight of a job:

- Makespan minimization on uniformly related machines ($Q||C_{max}$) parameterized by p_{max},
- Makespan minimization on unrelated machines ($R||C_{max}$) parameterized by p_{max} and the number of kinds of machines (defined later),
- Sum of weighted completion times minimization on unrelated machines ($R||\sum w_j C_j$) parameterized by $p_{max} + w_{max}$ and the number of kinds of machines,
- The same problem, $R||\sum w_j C_j$, parameterized by the number of distinct job times and the number of machines.

1 Introduction

Scheduling problems are one of the fundamental classes of problems in combinatorial optimization since 1960s [1,24,30] and many variants of scheduling turn out to be NP-hard. In response to this one can either look for an approximate solution, or restrict the input in a certain way. Approximation algorithms for scheduling have been an established area of research for a long time now [24]. On the other hand, parameterizing the input in order to obtain exact results has not been studied much before. We say that a problem P with input of size n is *fixed-parameter tractable* (FPT) with respect to parameter k if there exists a computable function f which does not depend on n and an algorithm solving P with running time $f(k) \operatorname{poly}(n)$; we call an algorithm with this running time FPT algorithm. Regarding scheduling, Mnich and Wiese [27] have recently initiated a systematic study of the relationship of various scheduling problems and their parameterizations, proving both positive and negative results. In this paper we continue in this direction, examining three additional fundamental scheduling problems and their parameterizations, and devising FPT algorithms for them.

However, our goal is not merely to prove new positive results. In their work, Mnich and Wiese rely on mathematical programming techniques in *fixed* dimension, which have been introduced in 1983 by Lenstra [25] and significantly extended in 2000 by Khachiyan and Porkolab [22]. These techniques are by now well established in the FPT community, even though the power of the latter extension of Khachiyan and Porkolab has not been fully utilized yet, as we will discuss further on. Independently of this, a new theory of *variable* dimension optimization has been developed in the past 15 years; see Onn's book [28]. A breakthrough result is an FPT algorithm for the so-called n-fold integer programming (n-fold IP) by Hemmecke, Onn and Romanchuk [16]. In contrast to the fixed dimension techniques, n-fold IP is not yet established as an indispensable part of an FPT researchers toolbox. In this paper we would like to help change that.

Let us now introduce n-fold IP. Given nt-dimensional integer vectors $\mathbf{b}, \mathbf{u}, \mathbf{l}, \mathbf{w}$, n-fold integer programming (n-fold IP) is the following problem in variable dimension nt:

$$\min\left\{\mathbf{w}\mathbf{x} : A^{(n)}\mathbf{x} = \mathbf{b},\ \mathbf{l} \leq \mathbf{x} \leq \mathbf{u},\ \mathbf{x} \in \mathbb{Z}^{nt}\right\}, \tag{1}$$

where

$$A^{(n)} \quad := \quad \begin{pmatrix} A_1 & A_1 & \cdots & A_1 \\ A_2 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_2 \end{pmatrix}$$

is an $(r + ns) \times nt$ matrix with A_1 an $r \times t$ matrix and A_2 an $s \times t$ matrix. Let a be the biggest absolute value of a number in $A^{(n)}$. The vector $\mathbf{x}$ is naturally partitioned into n *bricks* of size t, that is, we index it as $\mathbf{x} = (x_1^1, x_2^1, \ldots, x_t^1, \ldots, x_1^n, \ldots, x_t^n)$. As such, n-fold IP is best suited for *multi-index* problems whose IP formulation has variables indexed by $[n] \times [l_1] \times \cdots \times [l_k]$ for some integers $n, l_1, \ldots, l_k$ such that $l_1, \ldots, l_k$ are fixed parameters and only n is variable.

Hemmecke, Onn and Romanchuk [16] prove that there is an FPT algorithm solving problem (1) with parameters r, s, t and a. We will state and extend their result together with further observations in Section 2.1

1.1 Our contribution

We consider three non-preemptive scheduling models of increasing generality: parallel identical, uniformly related and unrelated machines (in the standard notation [24] denoted by P, Q and R, respectively), and the two most common objective functions: minimizing makespan and sum of weighted completion times (denoted C_{max} and $\sum w_j C_j$, respectively).

For *identical* machines, the problem consists of a set of n *jobs* $J = \{J_1, \ldots, J_n\}$ and m *machines* $M = \{M_1, \ldots, M_m\}$, and each job J_j has a *processing time* $p_j \in \mathbb{N}$. For *uniformly related* machines, we additionally have for each machine M_i its *speed* $s_i \in \mathbb{N}$, such that processing job J_j on machine M_i takes time p_j/s_i. For *unrelated* machines, we have for each job J_j an m-dimensional vector $\mathbf{p} = (p_j^1, \ldots, p_j^m)$, $p_j^i \in \mathbb{N} \cup \{\infty\}$ for all i, such that processing job J_j on machine M_i takes time p_j^i (in case $p_j^i = \infty$, J_j cannot be executed on M_i). We also consider a restricted variant of the unrelated machines model where there are K *kinds* of machines and the vector of processing times for a job J_j is given with respect to kinds of machines: $\mathbf{p} = (p_j^1, \ldots, p_j^K)$, such that processing J_j on machine M_i of kind k takes time p_j^k. Additionally, for the sum of weighted completion times objective, we are given for each job J_j its *weight* $w_j \in \mathbb{N}$.

A *schedule* is an assignment of jobs to machines and times, such that every machine is executing at most one job at any time. For a job J_j we denote by C_j its *completion time*, that is, the time J_j finishes. In makespan minimization, the goal is to minimize $C_{max} = \max_{J_j \in J} C_j$. When minimizing the sum of weighted completion times, the goal is to minimize $\sum_{J_j \in J} w_j C_j$. For example, the problem of minimizing makespan in the identical machines model is denoted $P||C_{max}$.

The parameters we consider are the following:

- p_{max}: the maximum processing time of any job,
- w_{max}: the maximum weight of any job,
- m: the number of machines,
- θ: the number of distinct job processing times and weights (in case of the $\sum w_j C_j$ objective); note that θ generalizes parameter p_{max},
- K: the number of kinds of machines (defined above).

In all of the cases we consider, we use such a combination of parameters that the number Θ of distinct job *types* is bounded, where jobs of a given type are indistinguishable from each other. This means that the set of jobs J on input can be given compactly by specifying integers $n_1, \ldots, n_\Theta$ with $n = n_1 + \cdots + n_\Theta$, such that

n_j denotes the number of jobs of type j that are to be scheduled. Modeling our terminology after Onn [29], we call a problem *huge* when the numbers n_j on input are given in binary. For example, the CUTTING STOCK problem can be seen as the huge variant of the BIN PACKING problem. All of our results work for the huge variant.

We show that:

Theorem 1. *The following scheduling problems are* FPT *with respect to parameter Θ as defined below and solvable in time $\Theta^{\mathcal{O}(\Theta^2)} n^{\mathcal{O}(1)}$, where*

1. $Q||C_{max}$: $\Theta = p_{max}$
2. $R||C_{max}$: $\Theta = p_{max}^K$
3. $R||\sum w_j C_j$: $\Theta = (\max\{p_{max}, w_{max}\})^K$

Theorem 2. $R||\sum w_j C_j$ *is* FPT *parameterized by $\Theta = m\theta^m$ and solvable in time $\Theta^{\mathcal{O}(\Theta)} n^{\mathcal{O}(1)}$.*

Note that part (1) of Theorem 1 for the easier $P||C_{max}$ problem was already shown by Mnich and Wiese [27]. However, our approach is substantially different and more straightforward, as demonstrated by the immediate extension to $Q||C_{max}$. Also, our result only has a single-exponential dependence on the parameter, unlike the double-exponential dependence of Mnich and Wiese. Thus, Theorem 1 serves to highlight the usefulness of n-fold integer programming in parameterized complexity. Theorem 2 differs as it is proved using the well known fixed dimension techniques of Khachiyan and Porkolab [22].

In order to prove part (3) of Theorem 1 we use an n-fold IP formulation and optimize a separable convex function over it. However, the algorithm of Hemmecke et al. [16] only works for linear and certain restricted separable convex objectives. Thus, using the ideas of Hemmecke, Köppe and Weismantel [15], we extend the previous result to optimizing any separable convex function.

We complement our positive findings by two hardness results:

Theorem 3. $P||C_{max}$ *and* $P||\sum w_j C_j$ *is* W*[1]-hard when parameterized by m, even when job processing times and weights are given in unary.*

1.2 Related work

We give a brief summary of known results related to scheduling and parameterized complexity. Many more results can be found in surveys on this topic (e.g. [24]). Note that these surveys focus on NP-hardness results and polynomial and approximation algorithms. There were several attempts to introduce scheduling problems to the FPT community. It started with the pioneering work of Bodlaender and Fellows [8] for the precedence constrained scheduling and continued (after nearly 10 years) with the first FPT algorithm of Fellows and McCartin [11]. A recent result of van Bevern et al. [3] resolves a question of Mnich and Wiese by showing that makespan minimization with precedence constraints $P|prec|C_{max}$ is W[2]-hard. Marx [10,26] also highlighted the importance of scheduling in the parameterized setting.

Many other settings are now popular in the scheduling community. These include *Open Shop*, where we are given for each job J a bunch of tasks $T_1^J, T_2^J, \ldots, T_k^J$ with associated machines $M_1^J, M_2^J, \ldots, M_k^J$ to be completed in the natural order (that is T_i^J has to finish before T_{i+1}^J starts to be processed) – see related work of van Bevern and Pyatkin [6] and Kononov et al. [23]. In *Two agent scheduling* there are two agents competing for one machine to schedule their respective jobs within each agents budget. Hermelin et al. [17] showed that if one agent has only k jobs and the other agent has jobs with unit processing times then the problem admits an FPT algorithm with parameter k. Halldórsson and Karlsson [14], followed by work of van Bevern et al. [4,5] investigate *interval scheduling* (or *job interval selection*) in which each job has several time slots in which it may be processed and the task is to schedule as many jobs as possible.

We now turn our attention towards more classical models of scheduling. A summary of what follows can be found in Table 1.

Makespan and identical machines – $P||C_{max}$ Most importantly, Mnich and Wiese [27] show that there is an FPT algorithm for this problem when parameterized by p_{max}. A remarkable result of Jansen et al. [21] shows that the UNARY BIN PACKING problem is W[1]-hard when parameterized by the number of bins. This immediately implies the W[1]-hardness of $P||C_{max}$ parameterized by m even when p_{max} is given in unary, and with some effort also implies the hardness of $P||\sum w_j C_j$ with parameter m when p_{max} and w_{max} are given in unary, as we show in Section 3.

Makespan and unrelated machines – $R||C_{max}$ Asahiro et al. [2] show that the problem $R||C_{max}$ is strongly NP-hard already for *restricted assignment* when there is a number p_j for each job such that for each machine $p_j^i \in \{p_j, \infty\}$ and all $p_j \in \{1, 2\}$ and for every job there are exactly two machines where it can run. Mnich and Wiese [27] proved that the problem is in FPT with parameters θ and m.

Sum of weighted completion times and unrelated machines – $R||\sum w_j C_j$ Surprisingly, in the unweighted case, $R||\sum C_j$ turns out to be solvable in polynomial time [9,20]. Preemption $(R|pmtn|\sum C_j)$ makes the problem strongly NP-hard [31]. The weighted case $R||\sum w_j C_j$ is strongly NP-hard [12, Problem SS13].

Problem	n	m	p_{max}	θ	Complexity
	unary	unary	unary	—	NP-hard [21]
$P||C_{max}$	unary	binary	param.	—	FPT [27], $2^{2^{\mathcal{O}(p_{max}^2 \log p_{max})}}$
	unary	param.	unary	—	W[1]-hard [21]
$Q||C_{max}$	binary	unary	param.	—	FPT [*], $2^{\mathcal{O}(p_{max}^2 \log p_{max})}$
	binary	param.	—	param.	FPT [27]
$R||C_{max}$	unary	unary	—	constant	NP-hard [2]
	binary	constant	—	binary	NP-hard
	binary	unary	—	param.	FPT with K [*]
$R||\sum w_j C_j$	binary	unary	param.	—	FPT with w_{max}, K [*]
	binary	param.	—	param.	FPT [*]

Table 1: A summary of the complexity results we mention. The results contained in this paper are marked with [*].

2   Preliminaries

For a positive integer n, we denote $[n] = \{1, 2, \ldots n\}$ and $\langle n \rangle = \lceil \log_2(n) \rceil$ the length of binary encoding of n. We write vectors in bold such as $\mathbf{b} = (b_1, \ldots, b_n)$. By $\langle \mathbf{b} \rangle$ we denote the length of encoding of $\mathbf{b}$ which is $\sum_{i=1}^n \langle b_i \rangle$; similarly for matrices. Let $\mathbf{x} = (x_1, \ldots, x_n)$. A function $f : \mathbb{R}^n \to \mathbb{R}$ is *separable convex* if it can be written as $f(\mathbf{x}) = \sum_{i=1}^n f_i(x_i)$ such that f_i is univariate convex for all $1 \leq i \leq n$. Given a function f, a *comparison oracle*, queried on two vectors $\mathbf{x}, \mathbf{y}$, asserts whether or not $f(\mathbf{x}) \leq f(\mathbf{y})$. Time complexity measures the number of arithmetic operations and oracle queries.

2.1 N-fold Integer Programming

Recall the definition of the n-fold IP problem (1). Observe that in an n-fold IP, every equality $\mathbf{a}\mathbf{x} = b$ of $A^{(n)}\mathbf{x} = \mathbf{b}$ is of one of two kinds. Either $\mathbf{a} = (\boldsymbol{\alpha}, \boldsymbol{\alpha}, \ldots, \boldsymbol{\alpha})$ corresponds to a certain row $\boldsymbol{\alpha}$ of A_1 repeated n times, meaning that $a_i^j = \alpha_i$ for all $1 \leq j \leq n$; we call this equality *globally uniform*. Or, $\mathbf{a} = (\mathbf{0}, \ldots, \mathbf{0}, \boldsymbol{\alpha}, \mathbf{0}, \ldots, \mathbf{0})$, corresponding to kt zeros, a certain row $\boldsymbol{\alpha}$ of A_2, and $(n - k - 1)t$ zeros, such that there also must exist $n - 1$ other equalities of this form which have the same coefficients $\boldsymbol{\alpha}$ on the remaining $n - 1$ bricks and zeros elsewhere. We call this kind of constraint *locally uniform*. Thus, given an

IP in dimension nt, we can prove that it is an n-fold IP by showing that every equality is either globally uniform or locally uniform.

Let $L := \langle a, \mathbf{b}, \mathbf{l}, \mathbf{u}, \mathbf{w} \rangle$ be the length of the input. The first key result we use is the following:

Theorem 4. *([16, Theorem 6.1]) For any fixed r, s and t, there is an algorithm that, given n, a, $(r, s) \times t$, matrices A_1 and A_2 of appropriate dimensions with all entries bounded by a in absolute value, $\mathbf{b}, \mathbf{l}, \mathbf{u}$, and $\mathbf{w}$, solves problem (1) in time $\mathcal{O}(a^{3t(rs+st+r+s)}n^3 L)$.*

In other words, if r, s, t and a are parameters and $A^{(n)}, \mathbf{b}, \mathbf{l}, \mathbf{u}, \mathbf{w}$ are input, the n-Fold IP problem can be solved in FPT time.

2.2 Separable convex minimization

We also need a result regarding minimization of separable convex function. Given a separable convex function f, we consider integer programs of the form:

$$\min\left\{ f(\mathbf{x}) : A^{(n)}\mathbf{x} = \mathbf{b}, \; \mathbf{l} \le \mathbf{x} \le \mathbf{u}, \; \mathbf{x} \in \mathbb{Z}^{nt} \right\}. \tag{2}$$

Hemmecke, Onn and Romanchuk prove that:

Theorem 5. *([16, Theorem 4.1]) For matrices A_1 and A_2 of appropriate dimensions, there is an algorithm that, given n, $\mathbf{b}, \mathbf{l}, \mathbf{u}$, separable convex function f presented by a comparison oracle, and a feasible point $\mathbf{x}$ in the program (2), either asserts that $\mathbf{x}$ is optimal or finds an augmenting step $\mathbf{g}$ for $\mathbf{x}$ which satisfies $f(\mathbf{x} + \mathbf{g}) < f(\mathbf{x})$ in linear time $\mathcal{O}(n)$.*

In particular, if we are able to find an initial solution $\hat{\mathbf{x}}$ and guarantee that the optimum $\mathbf{x}^*$ is near, that is, $||\hat{\mathbf{x}} - \mathbf{x}^*||_\infty \le N$ for some $N \in \mathbb{N}$, applying Theorem 5 at most nN times will reach the optimum.

Under suitable assumptions on the function f, the continuous optimum $\hat{\mathbf{x}}$ of Problem 2 can be found in polynomial time using the ellipsoid method or an interior point method. The key insight of Hemmecke, Köppe and Weismantel [15] is adapting a proximity technique of Hochbaum and Shantikumar [19] for the context of Graver bases:

Theorem 6. *([15, Theorem 3.14]) Let $\hat{\mathbf{x}}$ be an optimal solution of the continuous relaxation of (2),*

$$\min\left\{ f(\mathbf{x}) : A^{(n)}\mathbf{x} = \mathbf{b}, \; \mathbf{l} \le \mathbf{x} \le \mathbf{u}, \; \mathbf{x} \in \mathbb{R}^{nt} \right\} \; .$$

Then there exists an optimal solution $\mathbf{x}^*$ of the integer optimization problem (2) with

$$||\hat{\mathbf{x}} - \mathbf{x}^*||_\infty \le n \cdot \max\{||v||_\infty \mid v \in \mathcal{G}(A)\} \; .$$

Here, $\mathcal{G}(A)$ is the Graver basis of the bimatrix A, and the quantity $\max\{||v||_\infty \mid v \in \mathcal{G}(A)\}$ is bounded by a $t \cdot a^{r+s}$ ([28, Lemma 3.20]). Hence, $n \cdot ||\hat{\mathbf{x}} - \mathbf{x}^*||_\infty \le n^2 \cdot t \cdot a^{r+s}$ applications of Theorem 5 (with $\mathbf{x} = \lfloor \hat{\mathbf{x}} \rfloor$) reach the integer optimum:

Theorem 7. *For any fixed r, s and t, there is an algorithm that, given n, a, $(r, s) \times t$, matrices A_1 and A_2 of appropriate dimensions with all entries bounded by a in absolute value, $\mathbf{b}, \mathbf{l}, \mathbf{u}$, and $\mathbf{w}$, solves problem (2) in time $\mathcal{O}(a^{3t(rs+st+r+s)}n^3 L)$.*
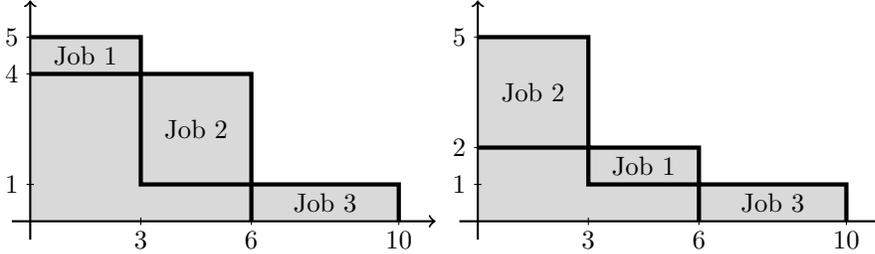
Fig. 1: An example of a 2D Gantt chart with three jobs of lengths 3, 3 and 4 and weights 1, 3 and 1, respectively. To the left is one possible ordering, to the right is the ordering given by Smith's rule, producing an optimal schedule (minimizing the gray area).

2.3 Convex minimization in fixed dimension

To prove Theorem 2 we formulate an integer linear program in fixed dimension and minimize a convex function over it:

Theorem 8. *(Khachiyan and Porkolab [22]) Minimizing a quasiconvex function over a semialgebraic convex set defined by k polynomials in dimension p is* FPT *with respect to k and p.*

In order not to delve into unnecessary details, let us say that a quasiconvex function is a generalization of a convex function, and semialgebraic convex sets are a generalization of convex sets, containing e.g. the feasible regions of semidefinite programs (SDPs); see the book by Blekherman et al. [7]. There are less general variants of this result which attain better running times, e.g. by Hildebrand and Köppe [18]. Since our aim is to simply prove fixed-parameter tractability, we choose to state the most general result.

Note that we are not aware of an application of Theorem 8 which would use the fact that one can optimize over a *region* more general than the integer hull of a polyhedron (i.e., a region given by non-linear convex constraints). There is a "linearization trick" which is widely used (including by Mnich and Wiese): a convex constraint $\mathbf{ax} \le f(x_i)$ whose domain is bounded by some number N given in unary can be rewritten as N linear constraints describing the piecewise linear approximation of $\mathbf{ax} \le f(x_i)$ which is exact on the N integers of its domain. Then, the feasible region is the integer hull of a polyhedron. To the best of our knowledge, there is no result whose feasible region is given by a set of constraints that cannot be "linearized" as described above. So both us and Mnich and Wiese [27] only need this result because of its generality in terms of the *objective function*, not the *feasible region*. It is an interesting open problem to find an application of Theorem 8 whose feasible region is given by, for example, a semidefinite program in fixed dimension.

2.4 Smith's rule – structure of solutions when minimizing $\sum w_j C_j$

Here we make a few basic observations about the structure of (optimal) solutions in the problem $R||\sum w_j C_j$. To do so, we utilize a useful way of visualizing the objective function $\sum w_j C_j$ called *two-dimensional Gantt charts* (2D Gantt charts), which was introduced by Goemans and Williamson [13].

Let us first introduce the following notation. Fix a machine M_i and assume that the set of jobs scheduled to run on M_i is given, denoted J^i. Whenever the index of a machine i is clear from context, we omit it. For any set of jobs $S \subseteq J$ let $w(S) = \sum_{J_j \in S} w_j$ and $p(S) = \sum_{J_j \in S} p_j$. A 2D Gantt chart starts at point $(0, w(J^i))$ and ends at $(p(J^i), 0)$. Each job $J_j \in J^i$ is represented by a rectangle of length p_j and height w_j whose position is defined by a startpoint and an endpoint. The startpoint (t, w) of a job is the endpoint of a previous job (or $(0, w(J^i))$ for the first job) while its endpoint is $(t + p_j, w - w_j)$. The value $\sum_{J_j \in J^i} w_j C_j$ is then simply the area beneath the upper sides of the rectangles; see Figure 1.

For $\sum w_j C_j$ minimization on one machine with no precedence constraints (no restrictions on the order of the jobs) there is a simple observation about the structure of an optimal schedule:
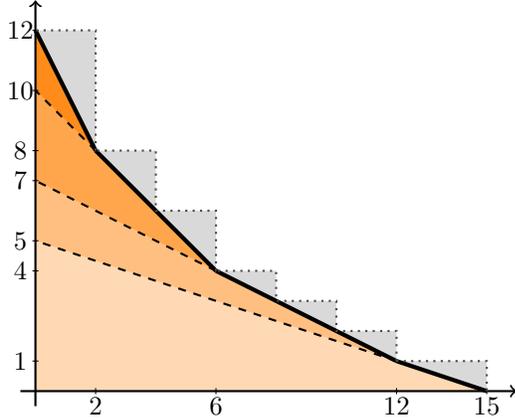
Fig. 2: An alternate way of computing the objective function $\sum_{J_j \in J^i} w_j C_j$. In this examples there are following jobs:

- 1 job with processing time 2 and weight 4,
- 2 jobs with processing time 2 and weight 2,
- 3 jobs with processing time 2 and weight 1,
- 1 job with processing time 3 and weight 1.

Lemma 1. *(Smith's rule [13]) Given a set of jobs J^i, a schedule minimizing $\sum_{J_j \in J^i} w_j C_j$ is given by ordering the jobs by non-increasing $\rho_i(j) = w_j / p^i_j$.*

Since the ratios $\rho_i(j)$ correspond to slopes of the rectangles in a 2D Gantt chart, Smith's rule implies that the chart of an optimal schedule will have slopes which form a piecewise linear convex function. Goemans and Williams then go on to observe that for such a chart there is an alternate way of computing its area based on splitting it into triangles; see Figure 2. That leads us to this lemma:

Lemma 2. *Given jobs $J^i = \{J_1, \ldots, J_l\}$ scheduled to run on machine M_i such that $\rho(j) \geq \rho(j+1)$ for all $1 \leq i \leq l - 1$, the optimal schedule has value*

$$\sum_{j=1}^{l} \left(\frac{1}{2} p(\{J_1, \ldots, J_j\})^2 (\rho(j) - \rho(j+1)) + \frac{1}{2} w_j p_j \right).$$

Proof. Given the set J^i, the optimal schedule on the machine M_i is determined according to Smith's rule. It is possible to divide the area as can be seen on Figure 2. Note that the gray area is determined by the set of jobs J^i and in fact can be computed just from the knowledge of J^i as $\frac{1}{2} w_j p_j$ for each job $J_j \in J^i$. This results in the linear term in the statement.

It remains to compute the area under the bold line (orange area in color printing) of Figure 2 – i.e. the area under the piecewise linear function (again determined by the observed structure of the solution). We divide the area into triangles and compute the total area as a sum of those.

For a job J_j we compute the contribution of the job as the associated area. The total area can be expressed as a difference of area of two impedance triangles – see Figure 3 for illustration. The length of the common ordinate parallel to processing time axis is $p(\{J_1, \ldots, J_j\})$, so it remains to establish the height b of the two triangles. We express it with the help of a tangent rule as $b = a \tan \varphi$ (and $b' = a \tan \varphi$)

Fig. 3: It is possible to compute the area of the lighter gray (green in color printing) rectangle as a difference between the area of an auxiliary triangle ($\frac{1}{2}a \cdot b'$ in the figure) and the dark gray (red in color printing) triangle. Note that it is possible to compute b, b' from the value of a and the tangent of φ, φ'.

It is straightforward to express the total area contribution of the job J_j as

$$
\begin{aligned}
J_j &= \frac{1}{2} p\big(\{J_1, \ldots, J_j\}\big) \cdot (b' - b) \\
&= \frac{1}{2} p\big(\{J_1, \ldots, J_j\}\big) \cdot p\big(\{J_1, \ldots, J_j\}\big) \big(\rho(j) - \rho(j+1)\big) \\
&= \frac{1}{2} p\big(\{J_1, \ldots, J_j\}\big)^2 \big(\rho(j) - \rho(j+1)\big).
\end{aligned}
\tag{3}
$$

Summing over all jobs in J^i finishes the proof.

$\square$

In our setup the set J^i will be given by integers $x_1^i, \ldots, x_\Theta^i$ representing how many jobs of each type are scheduled to run on machine M_i. Observe that for each type $1 \le j \le \Theta$, jobs of type j have identical slope $\rho(j)$ and thus correspond to a single triangle in the chart. We have the following corollary:

Corollary 1. *Given integers $x_1^i, \ldots, x_\Theta^i$ representing numbers of jobs of each type scheduled to run on machine M_i and a permutation $\pi_i : [\Theta] \to [\Theta]$ such that $\rho_i(\pi_i(j)) \ge \rho_i(\pi_i(j+1))$ for all $1 \le j \le \Theta - 1$, the optimal schedule has value $\sum_{i=1}^{\Theta}(\frac{1}{2}(z_j^i)^2 + \frac{1}{2}x_j^i p_j^i w_j)$, where $z_j^i = \sum_{l=1}^{j} p_l^i x_l^i$.*

3   W[1]-hardness of unary $P||C_{max}$ and $P||\sum w_j C_j$

The BIN PACKING problem asks whether k bins of size B suffice to contain a set of n items represented by integers $o_1, \ldots, o_n$. It is straightforward from the 2-PARTITION problem that this is NP-complete when $o_1, \ldots, o_n$ are large numbers given in binary. Remarkably, Jansen et al. [21] prove that even when the numbers are given in unary (i.e., assume $\max_i o_i \le n$), there is likely no $f(k)n^{\mathcal{O}(1)}$ algorithm as the problem is W[1]-hard parameterized by k. This hardness obviously translates to makespan minimization: deciding whether there is a schedule of jobs $o_1, \ldots, o_n$ on k machines with $C_{max} = B$ is equivalent to deciding the aforementioned UNARY BIN PACKING problem.

To the best of our knowledge, the analogous question regarding the complexity of $P||\sum w_j C_j$ parameterized by the number of machines m was not yet considered. We prove that it is W[1]-hard by once again reducing UNARY BIN PACKING to it.

Jansen et al. note that their hardness result stands even for *tight* instances of UNARY BIN PACKING, that is, instances where $\sum_i o_i = kB$. Given a tight instance of UNARY BIN PACKING, construct an instance of $P||\sum w_j C_j$ consisting of k machines and n jobs with $p_j = w_j = o_j$ for $j = 1, \ldots, n$. Let J^ℓ and $\hat{C}_\ell$ for $\ell = 1, \ldots, k$ denote the set of jobs scheduled on machine ℓ and the completion time of machine ℓ, respectively.

Because the ratio $\frac{w_j}{p_j}$ is identical for all jobs, the ordering of jobs on a each machine is irrelevant. Thus, the contribution of a machine ℓ to the objective function is

$$\frac{1}{2}\hat{C}_\ell^2 + \sum_{j \in J^\ell} \frac{p_j^2}{2} \ .$$

Summing over all machines, we get

$$\Big(\sum_{j \in J^\ell} \frac{p_j^2}{2} \Big) + \Big(\sum_{\ell=1}^k \frac{1}{2}\hat{C}_\ell^2 \Big) \ .$$

We argue that a schedule with cost $\frac{kB^2}{2} + kB$ exists if and only if the original UNARY BIN PACKING instance is a "yes" instance. Observe that there is only one schedule of this cost, namely one where $\hat{C}_\ell = B$ for all $\ell = 1, \ldots, k$. Let $\Delta_\ell = \hat{C}_\ell - B$ for all $\ell = 1, \ldots, k$. Disregarding the term $\sum_{j \in J^\ell} \frac{p_j^2}{2}$ which is independent of the schedule, the contribution becomes

$$\sum_{\ell=1}^k \frac{1}{2}(B + \Delta_\ell)^2 = \sum_{\ell=1}^k \frac{1}{2}(B^2 + 2\Delta_\ell B + \Delta_\ell^2) \ .$$

Since $\sum_\ell \Delta_\ell = 0$, $\sum_\ell 2\Delta_\ell B = 0$. Thus $\sum_\ell \Delta_\ell^2 = 0$ exactly when $\Delta_\ell = 0$ for all ℓ, that is, when $\hat{C}_\ell = B$ for all ℓ. That concludes the proof of Theorem 3.

4 FPT results

4.1 Warmup: $P||C_{max}$ and $Q||C_{max}$ parameterized by p_{max}

In this section we show that $P||C_{max}$ and its generalization $Q||C_{max}$ is FPT parameterized by p_{max}.

Denote $\Theta := p_{max}$. We say that a job of length j is of *type j*. On input we have n jobs of at most Θ types, given as numbers $n_1, \ldots n_\Theta$ encoding the number of jobs of given type (in binary). For every machine i we have variables $x_1^i, \ldots, x_\Theta^i$; in the solution the interpretation of x_j^i is "how many jobs of type j are scheduled on machine i". Let us fix a time $T \in \mathbb{N}$; the IP we will formulate will be feasible if there is a schedule with $C_{max} \leq T$.

To assure that each job is scheduled on some machine, we state these globally uniform constraints:

$$\sum_{i=1}^m x_j^i = n_j \quad \forall 1 \leq j \leq \Theta. \tag{4}$$

To assure that, for every machine M_i, $1 \leq i \leq m$, the lengths of jobs scheduled on M_i sum up to at most T, we state a locally uniform constraint:

$$\sum_{j=1}^\Theta j x_j^i \leq T \tag{5}$$

(Note here that using inequalities instead of equations does not cause problems: we can simply introduce a slack variable for every inequality: $\sum_{j=1}^\Theta j x_j^i + x_s^i = T, x_s^i \geq 0$; it is also possible to add a suitable number of unit-processing time jobs and work with equalities directly; similarly in the following.)

Clearly if this program is feasible, then there exists a schedule with $C_{max} \leq T$. Finding minimum T can be then done in polynomially many steps by binary search. Thus we want to show that checking feasibility is FPT by applying Theorem 4.

To apply Theorem 4, we need to bound the values r, s, t and a. Clearly the brick size is $t = \Theta$, the number of globally uniform constraints is $r = \Theta$ and the number of locally uniform constraints per brick is $s = 1$. Finally, the largest coefficient is $a = \Theta$.

9

Q||C$_{\textbf{max}}$ Now we are given speeds s_i for every machine, such that executing a job with processing time j takes time j/s_i. The globally uniform constraints (4) are the same, but the locally uniform constraints (5) now become:

$$\sum_{j=1}^{\Theta} j x_j^i \leq s_i T. \tag{6}$$

Observe that only the right hand side differs for every machine. This finishes the proof of part (1) of Theorem 1.

4.2 $R||C_{max}$ parameterized by p_{max} and K

Now we turn to the unrelated machines model. Observe that with parameters p_{max} and the number of kinds of machines K, there are at most $\Theta = (p_{max} + 1)^K$ possible vectors $\mathbf{p}_i$ of processing times with respect to kinds of machines, which we call *types*, and each job is of a certain type. The input is then again given (in binary) by Θ integers $n_1, \ldots, n_\Theta$ specifying the number of jobs of each type.

As before, we will describe an n-fold IP solving the problem. We have $n\Theta$ variables x_j^i with the same interpretation as above. The globally uniform constraints are the same as before, (4).

In the previous examples the locally uniform constraints were used to specify that the jobs assigned to each machine finish by time T. However, now we need to specify a different constraint for each *kind* of machine, which might seem hard to do "uniformly". Fortunately, because the number of kinds of machines is bounded, we can actually specify all constraints *simultaneously* and make all but "the right one" irrelevant by differing right hand side.

Formally, let B be some number bigger than np_{max}, then for machine M_i which is of kind k we have locally uniform constraints

$$\sum_{j=1}^{\Theta} p_j^{k'} x_j^{k'} \leq B \qquad \forall 1 \leq k' \neq k \leq K,$$

$$\sum_{j=1}^{\Theta} p_j^k x_j^k \leq T.$$

In the above constraints whenever $p_j^k = \infty$ we replace it by zero and forbid the job to be run on this kind of machine by specifying appropriate upper bounds:

$$x_j^i \leq n_j \qquad \forall 1 \leq k \leq K, p_j^k \in [p_{max}] \tag{7}$$

$$x_j^i \leq 0 \qquad \forall 1 \leq k \leq K, p_j^k = \infty, \tag{8}$$

$$x_j^i \geq 0 \qquad \forall i, j. \tag{9}$$

Observe that the constraints above are indeed locally uniform – they are identical for every machine up to the right hand side. Thus we have defined an n-fold IP, which is feasible if there is a schedule with $C_{max} \leq T$. It remains to observe that all of r, s, t and a are bounded by our choice of parameters p_{max} and K: clearly $t = \Theta$, $r = \Theta$, $s = K$ and $a = p_{max}$. Using Theorem 4 concludes the proof of part (2) of Theorem 1.

4.3 $R||\sum w_j C_j$ parameterized by $p_{max} + w_{max}$ and K

Here we turn our attention to the $\sum w_j C_j$ objective.

Recall what follows from Corollary 1: we have shown that, given integers $x_1^i, \ldots, x_\Theta^i$ representing numbers of jobs of each type to be scheduled on machine M_i, the contribution of M_i to the objective function is $f^i(\mathbf{x}^i, \mathbf{z}^i) = \sum_{j=1}^{\Theta}(\frac{1}{2}(z_j^i)^2(\rho_i(j) - \rho_i(j+1)) + \frac{1}{2}x_j^i p_j^i w_j)$ where $z_j^i = \sum_{l=1}^{j} p_l^i x_l^i$ and $\pi_i : [\Theta] \to [\Theta]$ is a permutation such that $\rho_i(\pi_i(j)) \geq \rho_i(\pi_i(j+1))$ for all $j \in [\Theta - 1]$. Observe that f^i is separable convex and thus also $f = \sum_i f^i$. Our goal now is to once again formulate an n-fold IP, however this time we need

to introduce new variables $z^i_{j,k}$, for $j \in [\Theta], k \in [K]$ and $i \in [m]$. For a machine M_i of kind k, we want $z^i_j = z^i_{j,k}$, so that we can use the formulation of f which we just stated. Notice that we are introducing many "unnecessary" variables $z^i_{j,k}$ for kinds $k' \neq k$. This is in order to have a uniform set of local constraints.

The globally uniform constraints (4) stay the same. The locally uniform constraints serve to project the brick $\mathbf{x}^i$ to the variables $z^i_{j,k}$, with permutations π_k as defined above:

$$\sum_{l=1}^{j} x^i_l p^i_{\pi_k(l)} = z^i_{j,k} \qquad \forall j \in [\Theta], \forall k \in [K], \forall i \in [m].$$

It is in the objective function where we distinguish which $z^i_{j,k}$ are relevant for which machine. In the following, let $z^i_j = z^i_{j,k}$ if machine M_i is of kind k:

$$f(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^{m} \sum_{j=1}^{\Theta} \left(\frac{1}{2}(z^i_j)^2 (\rho_i(j) - \rho_i(j+1)) + \frac{1}{2} x^i_j p^i_j w_j \right).$$

Lower and upper bounds (7)-(9) stay as before. Applying Theorem 7 concludes the proof of part (3) of Theorem 1.

4.4 $R || \sum w_j C_j$ parameterized by m and θ

Finally, we examine the same scenario as before, but this time we restrict the number of machines m to be a parameter, but, in turn, relax the restriction from p_{max}, w_{max} to θ. We use the same ILP formulation which Mnich and Wiese used to show that $R || C_{max}$ is FPT with respect to m and θ. However, the careful analysis in Corollary 1 was needed to show that the objective function is convex in order to apply Theorem 8.

Let $\Theta \leq \theta^m$ be the number of distinct types of jobs. We have variables x^i_j, $j \in [\Theta], i \in [m]$ and permutations π_i with the same meaning as above. Notice that the following can be seen as a subset of the previous n-fold IP:

minimize

$$f(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^{m} \sum_{j=1}^{\Theta} \left(\frac{1}{2}(z^i_j)^2 (\rho_i(j) - \rho_i(j+1)) + \frac{1}{2} x^i_j p^i_j w_j \right)$$

subject to

$$\sum_{i=1}^{m} x^i_j = n_j \qquad\qquad \forall j \in [\Theta]$$

$$\sum_{l=1}^{j} x^i_l p^i_{\pi_i(l)} = z^i_j \qquad\qquad \forall j \in [\Theta], \forall i \in [m]$$

In order to apply Theorem 8, we observe that both the number of variables $2\Theta m$ and the number of constraints $\Theta + \Theta m$ are fixed parameters, and that the objective function f is convex as shown in the previous subsection. This concludes the proof of Theorem 2.

5 Conclusions

Although much is known about approximating scheduling problem, little is known from the parameterized complexity point of view about the most basic problems. The purpose of this paper is twofold. The first is to show new FPT algorithms for some scheduling problems. The second is to demonstrate the use of n-fold integer programming, a recent and powerful variable dimension technique. We hope to encourage research in both directions. To facilitate this research, we point out the following open problems:

- Minimizing weighted flow time $P|r_j|\sum w_j F_j$ parameterized by $p_{max} + w_{max}$.
- $P||C_{max}$ parameterized by θ instead of p_{max}.
- $R|pmtn|\sum C_j$ parameterized by m and p_{max}; this is justified by the problem being strongly NP-hard [31], so parameterizing by p_{max} is not enough.
- $P||C_{max}$ parameterized by p_{max} with both m and n given in binary; this might be possible using the recently developed result for huge n-fold IPs due to Onn and Sarrabezolles [29].
- Multi-agent scheduling was studied by Hermelin et al. [17]; what results can be obtained by applying n-fold IP?
- Turning our attention to developing the techniques we use, we ask if 4-block n-fold IP (a generalization of n-fold IP) is FPT or W[1]-hard; only an XP algorithm is known so far [15].
- We are also interested in further applications of n-fold IP and quasiconvex minimization over convex sets in fixed dimension.

References

1. Allahverdi, A.: The third comprehensive survey on scheduling problems with setup times/costs. European Journal of Operational Research **246**(2), 345–378 (2015). DOI 10.1016/j.ejor.2015.04.004
2. Asahiro, Y., Jansson, J., Miyano, E., Ono, H., Zenmyo, K.: Approximation algorithms for the graph orientation minimizing the maximum weighted outdegree. In: M.Y. Kao, X.Y. Li (eds.) AAIM, *LNCS*, vol. 4508, pp. 167–177. Springer (2007). URL `http://dx.doi.org/10.1007/978-3-540-72870-2_16`
3. van Bevern, R., Bredereck, R., Bulteau, L., Komusiewicz, C., Talmon, N., Woeginger, G.J.: Precedence-constrained scheduling problems parameterized by partial order width. CoRR **abs/1605.00901** (2016). URL `http://arxiv.org/abs/1605.00901`
4. van Bevern, R., Mnich, M., Niedermeier, R., Weller, M.: Interval scheduling and colorful independent sets. J. Scheduling **18**(5), 449–469 (2015). DOI 10.1007/s10951-014-0398-5. URL `http://dx.doi.org/10.1007/s10951-014-0398-5`
5. van Bevern, R., Niedermeier, R., Suchý, O.: A parameterized complexity view on non-preemptively scheduling interval-constrained jobs: few machines, small looseness, and small slack. CoRR **abs/1508.01657** (2015). URL `http://arxiv.org/abs/1508.01657`
6. van Bevern, R., Pyatkin, A.V.: Completing partial schedules for open shop with unit processing times and routing. In: CSR 2016, pp. 73–87 (2016). DOI doi:10.1007/978-3-319-34171-2_6. URL `http://dx.doi.org/10.1007/978-3-319-34171-2_6`
7. Blekherman, G., Parrilo, P.A., Thomas, R.R.: Semidefinite Optimization and Convex Algebraic Geometry. SIAM (2012)
8. Bodlaender, H.L., Fellows, M.R.: W[2]-hardness of precedence constrained k-processor scheduling. Oper. Res. Lett. **18**(2), 93–97 (1995). DOI 10.1016/0167-6377(95)00031-9. URL `http://dx.doi.org/10.1016/0167-6377(95)00031-9`
9. Bruno, J., Coffman Jr., E.G., Sethi, R.: Scheduling independent tasks to reduce mean finishing time. Commun. ACM **17**(7), 382–387 (1974). DOI 10.1145/361011.361064. URL `http://doi.acm.org/10.1145/361011.361064`
10. Demaine, E.D., Hajiaghayi, M., Marx, D. (eds.): Parameterized complexity and approximation algorithms, 13.12. - 17.12.2009, *Dagstuhl Seminar Proceedings*, vol. 09511. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany (2009). URL `http://drops.dagstuhl.de/portals/09511/`
11. Fellows, M.R., McCartin, C.: On the parametric complexity of schedules to minimize tardy tasks. Theor. Comput. Sci. **2**(298), 317–324 (2003). DOI 10.1016/S0304-3975(02)00811-3. URL `http://dx.doi.org/10.1016/S0304-3975(02)00811-3`
12. Garey, M.R., Johnson, D.S.: Computers and intractability: A guide to the theory of np-completeness (1979)
13. Goemans, M., Williamson, D.P.: Two-dimensional Gantt charts and a scheduling algorithm of Lawler. SIAM Journal on Disc. Mat. **13**(3), 281–294 (2000). DOI http://dx.doi.org/10.1137/S0895480197330254. URL `http://epubs.siam.org/sam-bin/dbq/article/33025`
14. Halldórsson, M.M., Karlsson, R.K.: Strip graphs: Recognition and scheduling. In: WG 2006, pp. 137–146 (2006). DOI 10.1007/11917496_13. URL `http://dx.doi.org/10.1007/11917496_13`

15. Hemmecke, R., Köppe, M., Weismantel, R.: Graver basis and proximity techniques for block-structured separable convex integer minimization problems. Math. Program **145**(1-2), 1–18 (2014). URL `http://dx.doi.org/10.1007/s10107-013-0638-z`

16. Hemmecke, R., Onn, S., Romanchuk, L.: n-fold integer programming in cubic time. Math. Program **137**(1-2), 325–341 (2013). URL `http://dx.doi.org/10.1007/s10107-011-0490-y`

17. Hermelin, D., Kubitza, J., Shabtay, D., Talmon, N., Woeginger, G.J.: Scheduling two competing agents when one agent has significantly fewer jobs. In: IPEC 2015, pp. 55–65 (2015). DOI 10.4230/LIPIcs.IPEC.2015.55. URL `http://dx.doi.org/10.4230/LIPIcs.IPEC.2015.55`

18. Hildebrand, R., Köppe, M.: A new Lenstra-type algorithm for quasiconvex polynomial integer minimization with complexity $2^{O(n\log n)}$. Discrete Optimization **10**(1), 69–84 (2013)

19. Hochbaum, Shanthikumar: Convex separable optimization is not much harder than linear optimization. JACM: Journal of the ACM **37** (1990)

20. Horn, W.A.: Technical Note—Minimizing Average Flow Time with Parallel Machines. Operations Research **21**(3), 846–847 (1973). DOI 10.1287/opre.21.3.846. URL `http://dx.doi.org/10.1287/opre.21.3.846`

21. Jansen, K., Kratsch, S., Marx, D., Schlotter, I.: Bin packing with fixed number of bins revisited. Journal of Computer and System Sciences **79**(1), 39–49 (2013). DOI 10.1016/j.jcss.2012.04.004

22. Khachiyan, L., Porkolab, L.: Integer Optimization on Convex Semialgebraic Sets. Discrete & Computational Geometry **23**(2), 207–224 (2000)

23. Kononov, A.V., Sevastyanov, S., Sviridenko, M.: A complete 4-parametric complexity classification of short shop scheduling problems. J. Scheduling **15**(4), 427–446 (2012). DOI 10.1007/s10951-011-0243-z. URL `http://dx.doi.org/10.1007/s10951-011-0243-z`

24. Lawler, E.L., Lenstra, J.K., Kan, A.H.R., Shmoys, D.B.: Sequencing and scheduling: Algorithms and complexity. Handbooks in operations research and management science **4**, 445–522 (1993)

25. Lenstra Jr., H.W.: Integer programming with a fixed number of variables. Mathematics of Operations Research **8**(4), 538–548 (1983)

26. Marx, D.: Packing and scheduling algorithms for information and communication services (dagstuhl seminar 11091). Dagstuhl Reports **1**(2), 67–93 (2011). DOI 10.4230/DagRep.1.2.67. URL `http://dx.doi.org/10.4230/DagRep.1.2.67`

27. Mnich, M., Wiese, A.: Scheduling and fixed-parameter tractability. Mathematical Programming **154**(1), 533–562 (2014). DOI doi:10.1007/s10107-014-0830-9. URL `http://dx.doi.org/10.1007/s10107-014-0830-9`

28. Onn, S.: Nonlinear discrete optimization. Zurich Lectures in Advanced Mathematics, European Mathematical Society (2010)

29. Onn, S., Sarrabezolles, P.: Huge unimodular n-fold programs. SIAM Journal on Discrete Mathematics **29**(4), 2277–2283 (2015). DOI http://dx.doi.org/10.1137/151004227

30. Potts, C.N., Strusevich, V.A.: Fifty years of scheduling: a survey of milestones. JORS **60**(S1) (2009). DOI 10.1057/jors.2009.2. URL `http://dx.doi.org/10.1057/jors.2009.2`

31. Sitters, R.: Complexity of preemptive minsum scheduling on unrelated parallel machines. J. Algorithms **57**(1), 37–48 (2005). URL `http://dx.doi.org/10.1016/j.jalgor.2004.06.011`