New Algorithmic Results for Bin Packing and Scheduling

Klaus Jansen(✉)

Institut für Informatik, Christian-Albrechts-Universität zu Kiel, 24098 Kiel, Germany
kj@informatik.uni-kiel.de

Abstract. In this paper we present an overview about new results for bin packing and related scheduling problems. During the last years we have worked on the design of efficient exact and approximation algorithms for packing and scheduling problems. In order to obtain faster algorithms we studied integer linear programming (ILP) formulations for these problems and proved structural results for optimum solutions of the corresponding ILPs.

1 Introduction

In the first part of the paper we focus on the running times of approximation schemes for scheduling problems. A problem admits a polynomial-time approximation scheme (PTAS) for a minimization problem, if there is a family of algorithms $\{A_\varepsilon \mid \varepsilon > 0\}$ such that for any $\varepsilon > 0$ and any instance I, A_ε produces a $(1+\varepsilon)$-approximate solution in time polynomial in the size of the input. Two important restricted classes of approximation schemes were defined to distinguish the running times. An efficient polynomial-time approximation scheme (EPTAS) is a PTAS with running time of the form $f(1/\epsilon) \cdot poly(|I|)$, while a fully time polynomial time approximation scheme (FPTAS) runs in time $poly(|I|, 1/\epsilon)$.

In the second part we consider the classical bin packing problem and focus on the design of fixed parameter tractable (FPT) algorithms where the running time of the algorithm on an instance I is at most $f(k(I)) \cdot poly(|I|)$. Here f is a computable function and k the parameterization. A natural parameter k for bin packing is e.g. the optimum number $OPT(I)$ of bins used or the number d of different bin sizes.

2 Scheduling Problems

Minimum makespan scheduling is one of the fundamental problems in the literature on approximation algorithms [7,8]. In the *identical machine* setting the problem asks for an assignment of a set of n jobs $\mathcal{J}$ to a set of m identical

Research supported by the Deutsche Forschungsgemeinschaft (DFG), Project Ja 612/14-2, Entwicklung und Analyse von effizienten polynomiellen Approximationsschemata für Scheduling- und verwandte Optimierungsprobleme.

D. Fotakis et al. (Eds.): CIAC 2017, LNCS 10236, pp. 10–15, 2017.
DOI: 10.1007/978-3-319-57586-5_2

machines $\mathcal{M}$. Each job $j \in \mathcal{J}$ is characterized by a non-negative processing time $p_j \in \mathbb{Z}_{>0}$. The load of a machine is the total processing time of jobs assigned to it, and our objective is to minimize the *makespan*, that is, the maximum machine load. This problem is usually denoted $P||C_{\max}$.

For the setting with uniform machines the problem $Q||C_{max}$ is defined as follows. Suppose that we are given a set $\mathcal{J}$ of n independent jobs J_j with processing time p_j and a set $\mathcal{M}$ of m non-identical machines M_i that run at different speeds s_i. If job J_j is executed on machine M_i, the machine needs p_j/s_i time units to complete the job. The problem is to find an assignment $a : \mathcal{J} \rightarrow \mathcal{M}$ for the jobs to the machines that minimizes the total execution time $\max_{i=1,\ldots,m} \sum_{J_j:a(J_j)=P_i} p_j/s_i$. This is the minimum time needed to complete the execution of all jobs on the processors.

2.1 Known Results

It is well known that $P||C_{\max}$ admits a *polynomial time approximation scheme* (PTAS) [10], and there has been many subsequent works improving the running time or deriving PTAS's for more general settings. The first PTAS was found by Hochbaum and Shmoys [10] and had a running time of $(n/\varepsilon)^{O((1/\varepsilon)^2)} = n^{O((1/\varepsilon)^2 \log(1/\varepsilon))}$. This was improved to $n^{O((1/\varepsilon)\log^2(1/\varepsilon))}$ by Leung [17]. Subsequent articles improved further the running time. In particular Hochbaum and Shmoys (see [12]) and Alon et al. [1,2] obtained an *efficient PTAS* (EPTAS) with running time $2^{(1/\varepsilon)^{\mathrm{poly}(1/\varepsilon)}} + O(n \log n)$; doubly exponentional in $1/\epsilon$. The fastest previous known PTAS for $P||C_{\max}$ achieved a running time of $2^{O(1/\varepsilon^2)\log^3(1/\varepsilon))} + O(n \log n)$ for $(1+\varepsilon)$-approximate solutions [13].

For uniform processors, the decision problem for the scheduling problem with makespan at most T can be interpreted as a bin packing problem with different bin sizes. Using an ϵ-relaxed version of this bin packing problem, Hochbaum and Shmoys [11] were able to obtain a PTAS for scheduling jobs on uniform processors $Q||C_{max}$ with running time $(n/\epsilon)^{O(1/\epsilon^2)}$. The existence of an EPTAS for uniform processors was mentioned as an open problem by Epstein and Sgall [5]. Some years ago we found an EPTAS [13] with an improved running time for $Q||C_{max}$ based on an MILP formulation with a constant number of integral variables. For any $\epsilon > 0$ our algorithm A_ϵ produces a schedule for the jobs of length $A_\epsilon(I) \leq (1+\epsilon)OPT(I)$. The running time of A_ϵ is $2^{O(1/\epsilon^2 \log(1/\epsilon)^3)} + poly(n)$.

Very recently, Chen et al. [3] showed that, assuming the *exponential time hypothesis* (ETH), there is no PTAS that yields $(1+\varepsilon)$-approximate solutions for $\varepsilon > 0$ with running time $2^{(1/\varepsilon)^{1-\delta}} + poly(n)$ for any $\delta > 0$ [3].

2.2 New Results

We describe in this section the main new ideas for the identical machines setting; for the uniform setting we refer to [14]. Given a guess $T \in \mathbb{N}$ on the optimal makespan, which can be found with binary search, the problem reduces to deciding the existence of a packing of the jobs to m machines (or bins) of capacity T. If we aim for a $(1+\varepsilon)$-approximate solution, for some $\varepsilon > 0$, we can assume that all processing times are integral and T is a constant number, namely $T \in O(1/\varepsilon^2)$. This can be achieved with well known rounding and scaling techniques [1,2,12]. Let $\pi_1 < \pi_2 < \ldots < \pi_d$ be the job sizes appearing in the instance after rounding, and let b_k denote the number of jobs of size π_k. The mentioned rounding procedure implies that the number of different job sizes is $d = O((1/\varepsilon)\log(1/\varepsilon))$. Hence, for large n we obtain a highly symmetric problem where several jobs will have the same processing time. Consider the *knapsack polytope* $\mathcal{P} = \{c \in \mathbb{R}^d_{\geq 0} : \pi \cdot c \leq T\}$. A packing on one machine can be expressed as a vector $c \in Q = \mathbb{Z}^d \cap \mathcal{P}$, where c_k denotes the number of jobs of size π_k assigned to the machine. Elements in $Q = \mathbb{Z}^d \cap \mathcal{P}$ are called *configurations*. Considering a variable $x_c \in \mathbb{Z}_{\geq 0}$ that decides the multiplicity of configuration c in the solution, our problem reduces to solving the following linear integer program (ILP):

$$[\mathrm{conf}-\mathrm{ILP}] \quad \sum_{c \in Q} c \cdot x_c = b, \tag{1}$$

$$\sum_{c \in Q} x_c = m, \tag{2}$$

$$x_c \in \mathbb{Z}_{\geq 0} \qquad \text{for all } c \in Q. \tag{3}$$

In this paper we derive new insights on this ILP that help us to design faster algorithms for $P||C_{\max}$. We prove the following result.

Theorem 1 [14]**.** *The scheduling problem on identical machines admits an efficient polynomial time approximation scheme (EPTAS) with running time*

$$2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + O(n).$$

Hence, our algorithm is best possible up to polylogarithmic factors in the exponent assuming the ETH [3].

The support $supp(x)$ of a solution vector $x = (x_c)$ is defined as the set of non-negative components $x_c > 0$. Eisenbrand and Shmonin [4] proved that there always exists an optimum solution x of the configuration ILP with $|supp(x)| \leq 2(d+1)\log(4(d+1)T)$.

Our main technical contribution is a new structural result on the configuration ILP. More precisely, we show the existence of a highly symmetric and sparse optimal solution, in which all but a constant number of machines are assigned a configuration with small support. We say that a configuration c is simple if its support (or number of non-negative components $c_k > 0$) is of size at most $\log(T+1)$, otherwise it is complex. Then we can prove the following structural result:

Theorem 2 [14] **.** *Suppose that the* [conf-ILP] *is feasible. Then there exists a feasible solution x to* [conf-ILP] *such that*

(1) *if $x_c > 1$ then the configuration c is simple,*
(2) *the support of x satisfies $|supp(x)| \leq 4(d+1)\log(4(d+1)T)$, and*
(3) $\sum_{c \in Q_c} x_c \leq 2(d+1)\log(4(d+1)T)$*, where Q_c denotes the set of complex configurations.*

This structure can then be exploited by integer programming techniques [16,18] and dynamic programming. Interestingly, the result can be generalized also to the uniform machine setting with the same running time $2^{O((1/\varepsilon)\log^4(1/\varepsilon))} + O(n)$ using an MILP formulation and the same structural result. We believe that our structural result is of independent interest and should find applications to other settings.

3 Bin Packing

We consider the classical bin packing problem with d different item sizes $s_1, \ldots, s_d$ and build upon the results by Goemans and Rothvoß [6] to obtain a new polynomial time algorithm for the bin packing problem when d is constant [15]. Therefore, we present new techniques on how solutions of an instance can be modified and we give a new structural theorem that relies on the set of vertices of the underlying integer polytope.

3.1 Known Results

Given a polytope $\mathcal{P} = \{x \in \mathbb{R}^d \mid Ax \leq c\}$ for some matrix $A \in \mathbb{Z}^{m \times d}$ and a vector $c \in \mathbb{Z}^d$. We consider the integer cone

$$int.cone(\mathcal{P} \cap \mathbb{Z}^d) = \{ \sum_{p \in \mathcal{P} \cap \mathbb{Z}^d} \lambda_p p \mid \lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d} \}$$

of integral points inside the polytope $\mathcal{P}$. When we choose $\mathcal{P}$ to be the knapsack polytope, i.e. $\mathcal{P} = \{x \in \mathbb{Z}^d_{\geq 0} \mid s^T x \leq 1\}$, then each integral point of the polytope represents one possibility of packing a single bin with items from $s_1, \ldots, s_d$. Hence a vector $\lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d}$ of $int.cone(\mathcal{P} \cap \mathbb{Z}^d)$ represents a packing of the bin packing problem.

A long standing open question was, whether the bin packing problem can be solved in polynomial time when the number of different item sizes d is constant. This problem was recently solved by Goemans and Rothvoß [6] using structural properties of the integer cone. Essentially, they proved the existence of a distinguished set $X \subseteq \mathcal{P} \cap \mathbb{Z}^d$ of bounded size $|X| \leq m^d d^{O(d)} (\log \Delta)^d$ such that for every vector $b \in int.cone(\mathcal{P} \cap \mathbb{Z}^d)$ there exists an integral vector $\lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d}$ where most of the weight lies in X.

3.2 New Results

In this section we show that a similar structural theorem holds for a rather natural choice of the distinguished set X. Therefore, we consider the so called integer polytope $\mathcal{P}_I$. It is defined by the convex hull of all integer points inside $\mathcal{P}$, i.e. $\mathcal{P}_I = Conv(\mathcal{P} \cap \mathbb{Z}^d)$. Let V_I be the vertices of the integer polytope $\mathcal{P}_I$ i.e. $\mathcal{P}_I = Conv(V_I)$. Based on the set V_I, we can show the following structural result for solutions λ of $int.cone(\mathcal{P} \cap \mathbb{Z}^d)$:

Theorem 3 [15]**.** *Let $\mathcal{P} = \{x \in \mathbb{R}^d \mid Ax \leq c\}$ be a polytope with $A \in \mathbb{Z}^{m \times d}, c \in \mathbb{Z}^d_{\geq 0}$ and let $supp(\lambda)$ be the set of non-zero components of λ. Then for any vector $b \in int.cone(\mathcal{P} \cap \mathbb{Z}^d)$, there exists an integral vector $\lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d}$ such that $b = \sum_{p \in \mathcal{P} \cap \mathbb{Z}^d} \lambda_p p$ and*

(1) $\lambda_p \leq 2^{2^{O(d)}} \quad \forall p \in (\mathcal{P} \cap \mathbb{Z}^d) \setminus V_I$,
(2) $|supp(\lambda) \cap V_I| \leq d \cdot 2^d$,
(3) $|supp(\lambda) \setminus V_I| \leq 2^{2d}$.

As a consequence of our structural result, we obtain an algorithm for the bin packing problem with a running time of $|V_I|^{2^{O(d)}} \cdot \log(\Delta)^{O(1)}$, where Δ is the maximum over all multiplicities b and denominators in s. Since $|V_I| \geq d + 1$ this is an FPT-algorithm parameterized by the number of vertices of the integer knapsack polytope V_I.

Theorem 4 [15]**.** *The bin packing problem can be solved in time $|V_I|^{2^{O(d)}} \cdot (\log \Delta)^{O(1)}$ and hence in* FPT*-time, parameterized by the number of vertices V_I.*

This algorithmic result shows that the bin packing problem can be solved efficiently when the underlying knapsack polytope has an easy structure, i.e. has not too many vertices. However, since the total number of vertices is bounded by $O(\log \Delta)^d$ (see also [9]) the algorithm has a worst case running time of $(\log \Delta)^{2^{O(d)}}$, which is identical to the running time of the algorithm by Goemans and Rothvoß [6].

Furthermore, we were able to complement this result by giving a matching lower bound. We prove that the double exponential bound of the structure theorem is actually tight, even in the mentioned special case of bin packing, when all items sizes $s_1, \ldots, s_d$ are of the form $s_i = \frac{1}{a_i}$ for some $a_i \in \mathbb{Z}_{\geq 1}$.

References

1. Alon, N., Azar, Y., Woeginger, G., Yadid, T.: Approximation schemes for scheduling. In: 8th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 493–500 (1997)
2. Alon, N., Azar, Y., Woeginger, G., Yadid, T.: Approximation schemes for scheduling on parallel machines. J. Sched. **1**, 55–66 (1998)
3. Chen, L., Jansen, K., Zhang, G.: On the optimality of approximation schemes for the classical scheduling problem. In: 25th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 657–668 (2014)
4. Eisenbrand, F., Shmonin, G.: Carathéodory bounds for integer cones. Oper. Res. Lett. **34**, 564–568 (2006)

5. Epstein, L., Sgall, J.: Approximation schemes for scheduling on uniformly related and identical parallel machines. Algorithmica **39**, 43–57 (2004)
6. Goemans, M.X., Rothvoß, T.: Polynomiality for bin packing with a constant number of item types. In: 25th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 830–839 (2014)
7. Graham, R.L.: Bounds for certain multiprocessing anomalies. Bell Syst. Tech. J. **45**, 1563–1581 (1966)
8. Graham, R.L.: Bounds on multiprocessing timing anomalies. SIAM J. Appl. Math. **17**, 416–429 (1969)
9. Hayes, A.C., Larman, D.G.: The vertices of the knapsack polytope. Discret. Appl. Math. **6**(2), 135–138 (1983)
10. Hochbaum, D.S., Shmoys, D.B.: Using dual approximation algorithms for scheduling problems: theoretical and practical results. J. ACM **34**, 144–162 (1987)
11. Hochbaum, D.S., Shmoys, D.B.: A polynomial approximation scheme for scheduling on uniform processors: using the dual approximation approach. SIAM J. Comput. **17**, 539–551 (1988)
12. Hochbaum, D.S.: Approximation Algorithms for NP-Hard Problems. PWS Publishing Company, Boston (1997)
13. Jansen, K.: An EPTAS for scheduling jobs on uniform processors: using an MILP relaxation with a constant number of integral variables. SIAM J. Discret. Math. **24**, 457–485 (2010)
14. Jansen, K., Klein, K.M., Verschae, J.: Closing the gap for makespan scheduling via sparsification techniques. In: 43rd International Colloquium on Automata, Languages, Programming, (ICALP), pp. 72:1–72:13 (2016). arXiv:1604.07153
15. Jansen, K., Klein, K.M.: About the structure of the integer cone, its application to bin packing. In: 28th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1571–1581 (2017). arXiv:1604.07286
16. Lenstra, H.W.: Integer programming with a fixed number of variables. Math. Oper. Res. **8**, 538–548 (1983)
17. Leung, J.Y.-T.: Bin packing with restricted piece sizes. Inf. Process. Lett. **31**, 145–149 (1989)
18. Kannan, R.: Minkowski's convex body theorem and integer programming. Math. Oper. Res. **12**, 415–440 (1987)

Springer

http://www.springer.com/978-3-319-57585-8

Algorithms and Complexity
10th International Conference, CIAC 2017, Athens, Greece, May 24-26, 2017, Proceedings
Fotakis, D.; Pagourtzis, A.; Paschos, V.T. (Eds.)
2017, XX, 486 p. 59 illus., Softcover
ISBN: 978-3-319-57585-8