New Algorithmic Results for Scheduling and Bin Packing

Klaus Jansen [1]

[1] University of Kiel

Joint Work with Lin Chen, Kim-Manuel Klein, Lars Rohwedder, José Verschae and Gouchuan Zhang

Overview:

Main Topics

- Scheduling on Identical Machines
- Integer Programming
- Bin Packing
- Open Problems

Scheduling on Identical Machines $P||C_{max}$:

- Given: n jobs with processing times p_j
- and m machines
- Objective: Minimize makespan (maximum machine load)

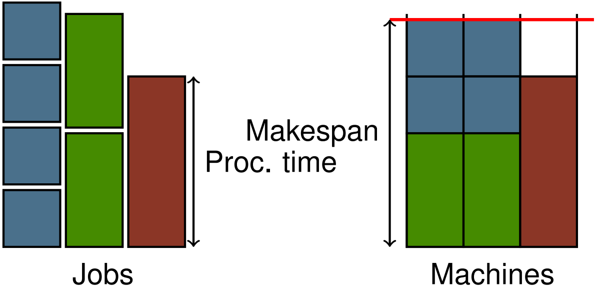

Literature:

Complexity

- Strongly NP-hard: If $P \neq NP$, then there is no FPTAS.

Known Algorithms

There is a PTAS with running time:

- $n^{O(\frac{1}{\varepsilon^2})}$ [Hochbaum & Shmoys '87]
- $n^{O(\frac{1}{\varepsilon}\log(\frac{1}{\varepsilon}))}$ [Leung 97]

There is an EPTAS with running time:

- $2^{2^{\widetilde{O}(\frac{1}{\varepsilon})}} + O(n \log n)$ [Alon et al. '98 & H. & S. '96]
- $2^{\widetilde{O}(\frac{1}{\varepsilon^2})} + O(n \log n)$ [Jansen '10]

Closing the Gap:

Lower Bound [Chen, Jansen, Zhang '13]

- If the *Exponential Time Hypothesis* holds, there is no EPTAS with running time $2^{(\frac{1}{\varepsilon})^{1-\delta}} + \text{poly}(n)$.

Our Main Result:

Theorem [Jansen, Klein, Verschae '16]

Minimum makespan scheduling admits an EPTAS with running time

$$2^{\widetilde{O}(\frac{1}{\varepsilon})} + O(n).$$

General Strategy:

General scheme for designing a PTAS:

1. Guess the makespan T of the optimal solution.
2. Round instance $\rightsquigarrow$ $(1+\varepsilon)$ multiplicative loss in objective.
3. Solve the rounded instance using an ILP formulation.

Rounding:

Lemma (Rounding and scaling)

$T = 1/\varepsilon^2$ and jobs sizes belong to $\Pi = \{\pi_1, \ldots, \pi_d\}$:

- $\Pi \subseteq \{\frac{1}{\varepsilon}, \frac{1}{\varepsilon}+1, \ldots, \frac{1}{\varepsilon^2}\}$ *and,* $\rightarrow$ *integer numbers*
- $|\Pi| = O(\frac{1}{\varepsilon}\log(\frac{1}{\varepsilon})) = \widetilde{O}(\frac{1}{\varepsilon})$. $\rightarrow$ *few sizes*

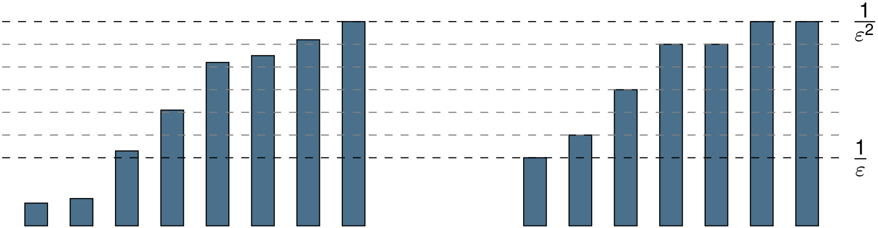

Configurations:

A *configuration* represents one possibility of assigning jobs from Π to a single machine.

Example (The set of configurations)

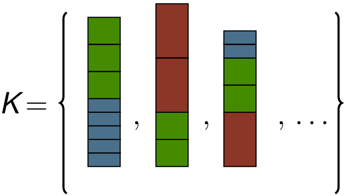

Configurations:

Knapsack polytope

$$P = \{k \in \mathbb{R}_{\geq 0}^{|\Pi|} : k^t \cdot \pi \leq T\}$$

Polyhedral view

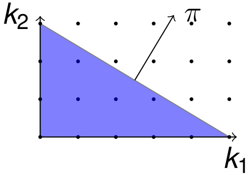

Configurations:

Set of configurations

$$K := P \cap \mathbb{Z}_{\geq 0}^{|\Pi|}$$

Observation 1

$$|K| \leq (T+1)^{|\Pi|} == 2^{O(\frac{1}{\varepsilon} \log^2(\frac{1}{\varepsilon}))} = 2^{\widetilde{O}(\frac{1}{\varepsilon})}.$$

Integer Programming Formulation

Observation 2:

The vector $(x_k)_{k\in K}$ belongs to the system

$$\left.\begin{aligned} \sum_{k\in K} x_k &= m \\ \sum_{k\in K} k_i x_k &= n_i \quad \text{for all } \pi_i \in \Pi \\ x &\in \mathbb{Z}^K_{\geq 0} \end{aligned}\right\}$$

# of constraints $= \widetilde{O}(\frac{1}{\varepsilon})$

variables $= 2^{\widetilde{O}(\frac{1}{\varepsilon})}$

Solving the ILP, first Approach:

Method [Alon et al. '98] and [Hochbaum & Shmoys '97] uses

Theorem [Kannan '87 / Lenstra '83]

An integer program with N variables can be solved in time $2^{\widetilde{O}(N)}\,s$ (where s is the length of the input).

In our case $N = |K| = 2^{\widetilde{O}(\frac{1}{\varepsilon})}$ and thus the running time is

$$2^{\widetilde{O}(N)} \log(n) = 2^{2^{\widetilde{O}(\frac{1}{\varepsilon})}} \log(n) \leftarrow \text{doubly exponential!}$$

Main Idea: Try to reduce the number of variables.

Solving the ILP, second Approach:

Guess the support [Jansen '10]

Theorem [Eisenbrand & Shmonin '06]

There is an optimum sol. x^* for $\{c^t x : Ax = b, x \geq 0, x \text{ integer}\}$ s.t. $|\text{support}(x^*)| \leq O(M(\log(M \cdot \Delta))$ where

- M = number of constraints,
- Δ = largest coefficient in A, c.

In our case:

- $M = |\Pi| = \widetilde{O}(\frac{1}{\varepsilon})$, and $\Delta = \frac{1}{\varepsilon}$
- $|\text{support}(x^*)| \leq \widetilde{O}(\frac{1}{\varepsilon})$

Solving the ILP, second Approach:

Guess the support [Jansen '10]

Algorithm:

1. Try each possible support: there are $\widetilde{O}(\frac{1}{\varepsilon}) \cdot \binom{|K|}{\widetilde{O}(\frac{1}{\varepsilon})} = 2^{\widetilde{O}(\frac{1}{\varepsilon^2})}$ many.
2. Solve ILP restricted to guessed variables with Kannan's algorithm (running time $2^{\widetilde{O}(\frac{1}{\varepsilon})} \log(n)$)
3. Total running time: $2^{\widetilde{O}(\frac{1}{\varepsilon^2})} \log(n)$.

Solving the ILP, third Approach:

Understanding the Optimum

Definition

A configuration k is *complex* if it contains more than $\log(T+1)$ different sizes; o.w. it is *simple*.

Example ($\log(T+1)=1$)

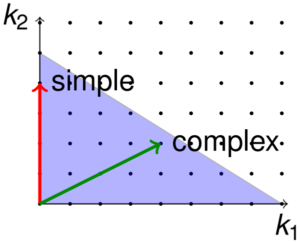

Example ($\log(T+1)=3$)

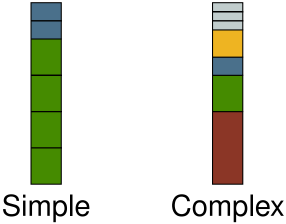

Solving the ILP, third Approach:

Understanding the Optimum

A “subconfiguration” $k' \leq k$ of configuration k is called *maximal* if it contains all possible jobs of each taken size.

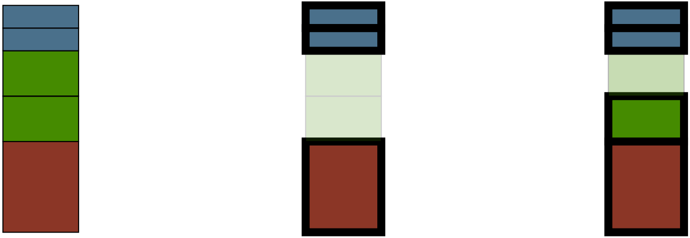

Original Configuration

Maximal Subconfiguration

Non-Maximal Subconfiguration

Lemma

Every complex conf. $k \in K$ contains two maximal disjoint subconfigurations k_1, k_2 s.t. the total size of k_1 and k_2 coincide.

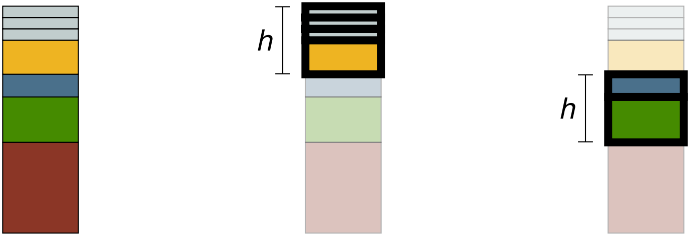

Complex Configuration k

Subconfiguration k_1

Subconfiguration k_2

Lemma

Every complex conf. $k \in K$ contains two maximal disjoint subconfigurations k_1, k_2 s.t. $\pi \cdot k_1 = \pi \cdot k_2$.

Proof.

- Let $C > \log(T + 1)$ be the number of sizes (colors) in k.
- Number of maximal subconfigurations $= 2^C > T + 1$.
- Total size of each configuration is in $\{0, 1, 2, \ldots, T\}$.
- Pigeonhole principle $\Rightarrow$ there are two maximal subconfigurations of same total size.

□

Solving the ILP, third Approach:

Lemma (Sparsification Lemma (informal))

If a complex configuration is taken twice in a solution, then we can replace it by two other “less complex” configurations.

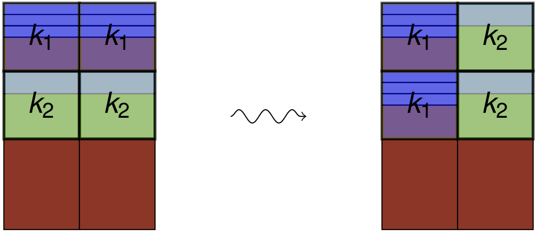

Solving the ILP, third Approach:

Theorem (Thin solutions)

If the ILP is feasible, then there is a solution x^* such that:

- *At most $\widetilde{O}(\frac{1}{\varepsilon})$ machines get complex configurations.*
- *Each complex configuration is used at most once.*
- $|support(x^*)| \leq O(|\Pi| \log(|\Pi| T)) = \widetilde{O}(\frac{1}{\varepsilon})$.

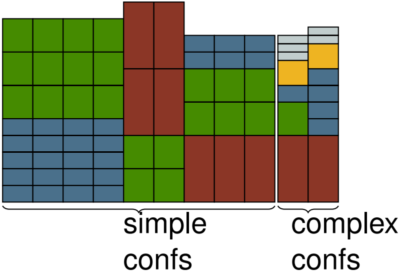

Lemma

The number of simple configurations in K is $2^{O(\log^2(\frac{1}{\varepsilon}))} = 2^{\widetilde{O}(1)}$.

Proof.

Let $D = \log(T+1)$ and $T = 1/\varepsilon^2$.

$$\begin{aligned}
\#\text{ simple conf} &\leq \sum_{i=0}^{D} \binom{|\Pi|}{i} \times (T+1)^i \\
&\leq (D+1)|\Pi|^D \times (T+1)^D \\
&\leq (\frac{1}{\varepsilon}\log(\frac{1}{\varepsilon}))^{O(\log(\frac{1}{\varepsilon}))} \\
&\leq 2^{O(\log^2(\frac{1}{\varepsilon}))} \leq 2^{\widetilde{O}(1)}.
\end{aligned}$$

□

Solving the ILP, third Approach:

Algorithm

Part 1: Complex Configurations.

1. Guess jobs assigned to complex configurations and number of complex machines.
2. Solve that subinstance optimally with a dynamic program.

Solving the ILP: Third Approach

Algorithm

Part 2: Remaining Instance.

1. Guess the (simple!) configurations in support:

$$\#\text{ possibilities} \leq \binom{2^{\widetilde{O}(1)}}{\widetilde{O}(\frac{1}{\varepsilon})} = 2^{\widetilde{O}(\frac{1}{\varepsilon})}$$

2. For each possibility solve the ILP restricted to those variables with Kannan's algorithm.

Total running time: $2^{\widetilde{O}(\frac{1}{\varepsilon})} \log(n)$

Main Result:

Algorithm

Theorem [Jansen, Klein, Verschae '16]

The minimum makespan problem on identical machines admits an EPTAS with running time

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

Integer Linear Programming

$$\max c^t x$$
$$Ax = b$$
$$x \in \mathbb{Z}^n_{\geq 0}$$

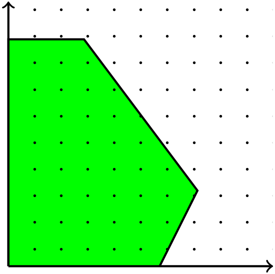

where $A \in \mathbb{Z}^{M \times N}$, $b \in \mathbb{Z}^M$, $c \in \mathbb{Z}^N$.

Considered case

M (#constraints) is a constant, entries of A are small ($\leq \Delta$).

Pseudo-polynomial Algorithms

Known Algorithms

There is an algorithm with running time:

- $(M(\Delta + \|b\|_\infty))^{O(M^2)}$ [Papadimitrou '81]
- $N \cdot O(M\Delta)^{2M} \cdot \|b\|_\infty^2$. [Eisenbrand & Weismantel '18]

Theorem [Jansen & Rohwedder '19]

IP can be solved in time $O(M\Delta)^{2M} \cdot \log(\|b\|_\infty) + O(NM)$.
Moreover, improving the exponent to $2M - \delta$ is equivalent to finding a truly subquadratic algorithm for $(\min, +)$-convolution.

Feasibility problem

Theorem [Jansen & Rohwedder '19]

Algorithm with running time:
$O(M\Delta)^{M} \cdot \log(\Delta) \cdot \log(\Delta + \|b\|_\infty) + O(NM)$. Improving exponent to $M - \delta$ would contradict the Strong Exponential Time Hypothesis (SETH).

Previous best result

$N \cdot O(M\Delta)^M \cdot \|b\|_\infty$. [Eisenbrand & Weismantel '18]

Application $P||C_{max}$

Configuration IP

$$\begin{aligned} \textstyle\sum_{k \in K} x_k &= m & \\ \textstyle\sum_{k \in K} k_i x_k &= n_i & \forall \pi_i \in \Pi \\ x_k &\in \mathbb{Z}_{\geq 0} & \forall k \in K \end{aligned}$$

has $M + 1 = O(\frac{1}{\epsilon}\log(\frac{1}{\epsilon}))$ constraints and $N = |K| = 2^{O(\frac{1}{\epsilon})}$ many variables. The value $\Delta = \max_{k,i} k_i \leq \frac{1}{\epsilon}$ and $\|b\|_\infty \leq n$.

New result: Including preprocessing $O(n + \frac{1}{\epsilon}\log(\frac{1}{\epsilon}))$, we get:

$$\begin{aligned} &O(M\Delta)^M \cdot \log(\Delta) \cdot \log(\Delta + \|b\|_\infty) + O(NM) + O(n + \tfrac{1}{\epsilon}\log(\tfrac{1}{\epsilon})) \\ &\leq 2^{O(\frac{1}{\epsilon}\log^2(\frac{1}{\epsilon}))} \log(n) + O(n) \leq 2^{O(\frac{1}{\epsilon}\log^2(\frac{1}{\epsilon}))} + O(n). \end{aligned}$$

Steinitz Lemma

Let $\|\cdot\|$ be a norm in $\mathbb{R}^M$ and $v^{(1)}, \dots, v^{(t)} \in \mathbb{R}^M$ with $\|v^{(i)}\| \leq 1$ $\forall i$ and $v^{(1)} + \dots + v^{(t)} = 0$. Then there is a permutation $\pi \in S_t$ with $\|\sum_{i=1}^{j} v^{(\pi(i))}\| \leq M$ for all $j = 1, \dots, t$.

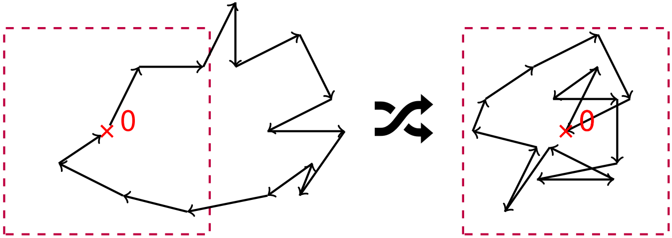

Consider an optimal solution x^* of (IP) and the sequence of column vectors

$$\underbrace{A_1, \ldots, A_1}_{x_1^* \text{ times}}, \underbrace{A_2, \ldots, A_2}_{x_2^* \text{ times}}, \ldots$$

$$\max c^t x \qquad Ax = b \qquad x \in \mathbb{Z}^N_{\geq 0} \qquad \text{(IP)}$$

Recall that $\|A_i\|_\infty \leq \Delta$.

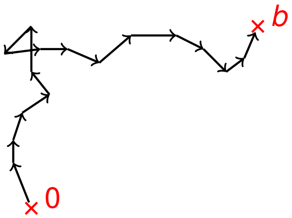

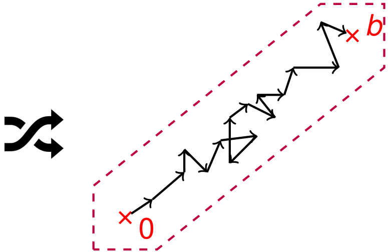

Steinitz for IP (formally)

Corollary

Let $v^{(1)}, \ldots, v^{(t)}$ denote columns of A with $\sum_{i=1}^{t} v^{(i)} = b$. Then there exists a permutation $\pi \in S_t$ such that for all $j \in \{1, \ldots, t\}$

$$\left\| \sum_{i=1}^{j} v^{(\pi(i))} - j \cdot b/t \right\|_\infty \leq 2M\Delta.$$

This follows easily from the Steinitz Lemma: Insert vectors $\frac{v^{(i)} - b/t}{\Delta}$ in the Steinitz Lemma. Notice $\|\frac{v^{(i)} - b/t}{\Delta}\|_\infty \leq 2$.

Eisenbrand & Weismantel

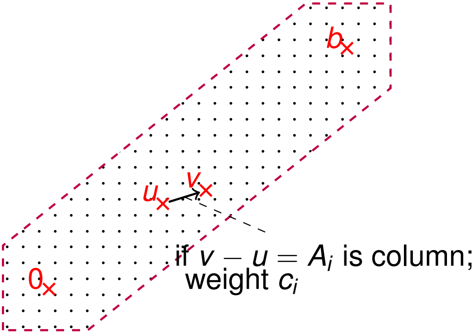

- Every $0 - b$ path gives a feasible solution
- Longest path is optimal
- $O(M\Delta)^M \cdot \|b\|_\infty$ vertices
- $N \cdot O(M\Delta)^M \cdot \|b\|_\infty$ edges
- Running time: $N \cdot O(M\Delta)^{2M} \cdot \|b\|_\infty^2$

Observation: There is an optimal solution of bounded norm, i.e., $\|x\|_1 \leq O(M\Delta)^M \cdot \|b\|_\infty$.

Our Approach

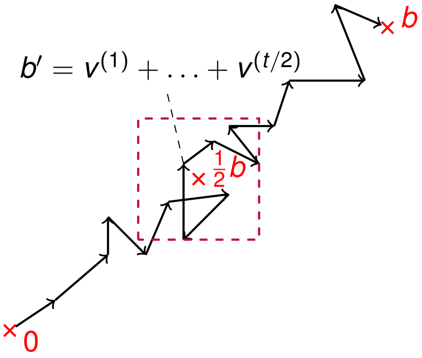

Let $v^{(1)} + \ldots + v^{(t)} = b$ be columns corresponding to an optimal solution of (IP).

Equivalent:

$v^{(1)} + \ldots + v^{(t/2)}$ is optimal for

$$\{\max c^t x, Ax = b', x \in \mathbb{Z}^N_{\geq 0}\}$$

and $v^{(t/2+1)} + \ldots + v^{(t)}$ is for

$$\{\max c^t x, Ax = b - b', x \in \mathbb{Z}^N_{\geq 0}\}.$$

If ordered via Steinitz Lemma, b' and $b - b'$ are not far from $\frac{1}{2}b$.

Assume w.l.o.g. there is an optimal solution x with $\|x\|_1 = 2^K$, where $K \in \log(O(M\Delta)^M \cdot \|b\|_\infty) = O(M \log(M\Delta) + \log(\|b\|_\infty))$

Solve for every $i = K, K-1, \dots, 0$ and every b' with

$$\left\| b' - \frac{1}{2^i} b \right\|_\infty \leq 4M\Delta$$

the problem

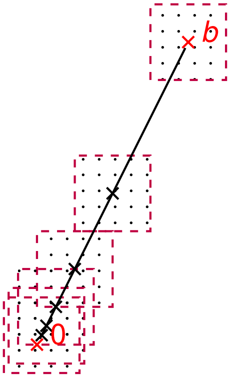

$$\begin{aligned} \max\ & c^t x \\ Ax &= b' \\ \|x\|_1 &= 2^{K-i} \\ x &\in \mathbb{Z}^N_{\geq 0}. \end{aligned}$$

Original problem for $i = 0$ and $b' = b$.

Consider iteration $i < K$ and b' with $\|b' - 1/2^i \cdot b\|_\infty \leq 4M\Delta$.

Let $v^{(1)}, \ldots, v^{(2^{K-i})}$ be a solution of $\max\{c^t x, Ax = b', \|x\|_1 = 2^{K-i}, x \in \mathbb{Z}^N_{\geq 0}\}$ ordered via Steinitz. Set $b'' := v^{(1)} + \ldots + v^{(2^{K-i-1})}$. Then we obtain

$$\left\| b'' - \frac{1}{2^{i+1}} b \right\|_\infty \leq \underbrace{\left\| b'' - \frac{1}{2} b' \right\|_\infty}_{\leq 2m\Delta} + \underbrace{\left\| \frac{1}{2} b' - \frac{1}{2^{i+1}} b \right\|_\infty}_{\leq 1/2 \cdot 4M\Delta} \leq 4M\Delta$$

Similarly, $\left\|(b' - b'') - \frac{1}{2^{i+1}} b\right\|_\infty \leq 4M\Delta$.
Our algorithm: Guess b'' ($O(M\Delta)^M$ candidates), look up solutions for $(i+1, b'')$ and $(i+1, b' - b'')$, and take the best.

Merging solutions

(MAX, +)-CONVOLUTION

Input: $r_1, \ldots, r_n \in \mathbb{R}$,
$s_1, \ldots, s_n \in \mathbb{R}$

Output: $t_1, \ldots, t_n \in \mathbb{R}$ with
$t_i = \max_j [r_j + s_{i-j}]$

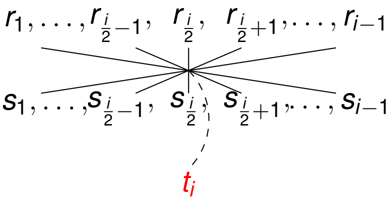

$T(n)$ time algorithm for (min, +)-convolution $\Rightarrow$
$T(O(M\Delta)^M) \cdot O(M \log(M\Delta) + \log(\|b\|_\infty)) + O(NM)$ for IP.

With $T(n) = O(n^2/\log(n))$: $O(M\Delta)^{2M} \cdot \log(\|b\|_\infty) + O(NM)$.

Feasibility of IP

BOOLEAN-CONVOLUTION

Input: $r_1, \ldots, r_n \in \{0,1\}$,
$s_1, \ldots, s_n \in \{0,1\}$

Output: $t_1, \ldots, t_n \in \{0,1\}$ s.t.
$t_i = \bigvee_j [r_j \wedge s_{i-j}]$

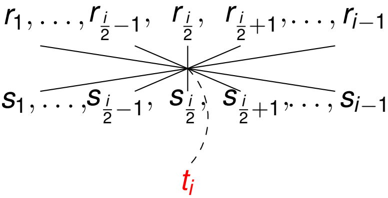

Boolean Convolution can be computed in $T(n) = O(n \log n)$.

$\Rightarrow$ Feasibility of IP in time

$$T(O(M\Delta)^M) \cdot (M \log(M\Delta) + \log(\|b\|_\infty)) + O(NM)$$
$$= O(M\Delta)^M \cdot \log(\Delta) \cdot \log(\Delta + \|b\|_\infty) + O(NM).$$

Bin Packing:

Problem Definition

- d item sizes
- s_i: size of item
- b_i: multiplicity of item size s_i
- Objective: Find a packing into a minimum number of unit bins.

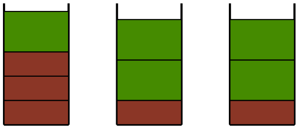

Cone:

Given a set of points $P \subset \mathbb{Z}^d$ then

$$Cone(P) = \{\sum_{p \in P} \lambda_p p \mid \lambda \in \mathbb{R}^P_{\geq 0}\}$$

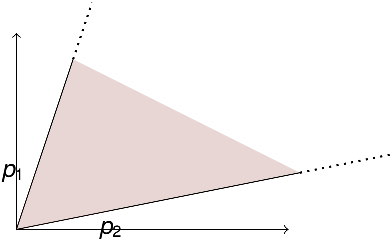

Integer Cone:

Given a set of points $P \subset \mathbb{Z}^d$ then

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

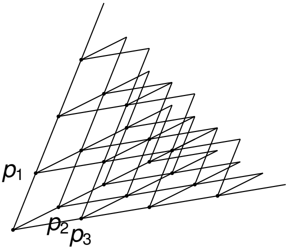

Integer Cones of Polytopes:

Given 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 $c \in \mathbb{Z}^d$.

Knapsack polytope $\mathcal{P} = \{x \in \mathbb{R}^d \mid s_1 x_1 + \ldots s_d x_d \leq 1, x \geq 0\}$ for sizes $s_1, \ldots, s_d$.

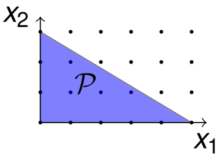

We consider $int.cone(\mathcal{P} \cap \mathbb{Z}^d)$.

The Bin Packing Problem:

Given: a set of item sizes $s_1, \ldots, s_d \in (0, 1]$ and multiplicities $b_1, \ldots, b_d$ of the corresponding item sizes.
Objective: Find a packing into a minimum number of unit bins.

Example:

Item sizes: $s_1 = \frac{1}{5}, s_2 = \frac{1}{3}$ with multiplicities: $b_1 = 5, b_2 = 5$

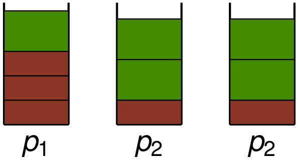

The Bin Packing Problem:

Each vector $\lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d}$ with $\sum_{p \in \mathcal{P} \cap \mathbb{Z}^d} \lambda_p p = b \in int.cone(\mathcal{P} \cap \mathbb{Z}^d)$ represents a possible solution of the bin packing problem.

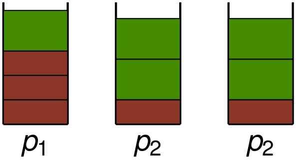

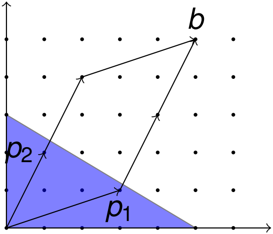

Structural Properties:

Arguments about the set of possible solutions $\lambda \in \mathbb{Z}_{\geq 0}^{\mathcal{P} \cap \mathbb{Z}^d}$.

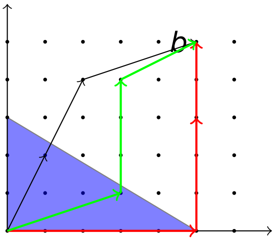

The Structure of the Integer Cone:

Theorem [Eisenbrand, Shmonin '06]

For any integral point $b \in int.cone(\mathcal{P} \cap \mathbb{Z}^d)$, 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 $|supp(\lambda)| \leq 2^d$.

The Structure of the Integer Cone:

Theorem [Goemans & Rothvoß'14]

There exists a set $X \subseteq \mathcal{P} \cap \mathbb{Z}^d$ with $|X| \leq m^d d^{O(d)} (\log \Delta)^d$ such that for any point $b \in int.cone(\mathcal{P} \cap \mathbb{Z}^d)$, 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. $\lambda_p \leq 1 \qquad \forall p \in (\mathcal{P} \cap \mathbb{Z}^d) \setminus X$
2. $|supp(\lambda) \cap X| \leq 2^{2d}$
3. $|supp(\lambda) \setminus X| \leq 2^{2d}$

Theorem [Goemans & Rothvoß'14]

Bin packing with d different item sizes can be solved in time $(\log \Delta)^{2^{O(d)}}$, where Δ is the maximum over all multiplicities b and denominators in s.

Vertices of the Integer Polytope:

Integer Polytope $\mathcal{P}_I = Conv(\mathcal{P} \cap \mathbb{Z}^d)$

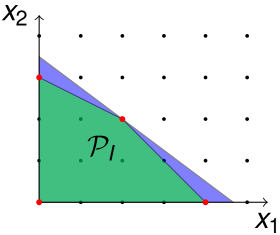

$s_1 = \frac{3}{14}$
$s_2 = \frac{2}{7}$

Theorem [Cook et al. '92]

For a polytope $\mathcal{P} = \{x \in \mathbb{R}^d \mid Ax \leq c\}$ the integer polytope $\mathcal{P}_I$ has at most $m^d \cdot O((\log \Delta)^d)$ vertices.

Our Structure Theorem:

Theorem [Jansen & Klein '17]

Let $V_I \subseteq \mathcal{P} \cap \mathbb{Z}^d$ be the set of vertices of the integer polytope $\mathcal{P}_I$. 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. $|supp(\lambda) \setminus V_I| \leq 2^{2d}$

Theorem [Jansen & Klein '17]

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.

Main Open Questions:

- Is there an EPTAS for scheduling on identical machines with running time $2^{O(1/\epsilon)} + O(n)$?
- Is there an FPT-algorithm for bin packing parameterized by the number d of different sizes?

Thanks for your attention!