

Makespan minimization

Largest Processing time rule (LPTR) algorithm

Input: m , n jobs a_1, \dots, a_n .

Algorithm: (1) Consider the jobs in non-increasing order of processing times. Assume that the order is $a^{(1)}, a^{(2)}, \dots, a^{(n)}$.

(2) In the i -th iteration, assign $a^{(i)}$ to the least loaded machine.

Theorem 1: LPTR is a $\frac{4}{3}$ -approximation algorithm.

Proof: Let j be the machine of maximum makespan and let i be the last job that is assigned to it.

Just before i is assigned to it, its makespan is at most $\sum_{k=1}^n a_k / m \leq \text{OPT}$.

Case 1: $a_i \leq \frac{\text{OPT}}{3}$. Then $M(j) \leq \text{OPT} + \frac{\text{OPT}}{3} = \frac{4}{3} \cdot \text{OPT}$

Case 2: $a_i > \frac{\text{OPT}}{3}$.

Claim: In this case LPTR computes an optimal schedule.

Proof: $a^{(1)} \geq a^{(2)} \geq \dots \geq a^{(n)}$. Towards a contradiction, assume that LPTR does not compute an optimal schedule. Let $a^{(k)}$ be the first job that finishes after OPT. $k \leq i$ (since i finishes after OPT). Therefore, $a^{(k)} \geq a^{(i)} > \frac{\text{OPT}}{3}$.

Just before a_k is processed, each machine contains one or two jobs. (No machine has 0 job because then the k -th job can be assigned to that machine and made to finish before OPT. No machine can have three or more jobs, as then one of those jobs will finish after OPT.) Let machines 1 to l contain long jobs and machines l to m contain two jobs after iteration $k-1$.
short jobs (including k).

$$k-1 = l + 2(m-l).$$

Consider an optimal schedule. I claim that each long job is the only job in its machine. Let t be the processing time of a long job. $t + a_k > \text{OPT}$. (otherwise k would finish before OPT in the schedule computed by LPT). This implies that $t + a_j > \text{OPT}$ for all $j=1, \dots, k$. Thus t cannot be paired with any other job in an optimal schedule. Also no machine can have more than two jobs from $\{1, \dots, k\}$. The total number of jobs must be at most $l + (m-l) \cdot 2 = k-1 < k$. This is a contradiction.

Theorem 1 is tight; there is an instance on which LPT produces a solution with makespan $\approx \frac{4}{3} \cdot \text{OPT}$.

A PTAS for makespan minimization: Makespan minimization is strongly NP-hard. So, an FPTAS does not exist (unless $P=NP$). We will obtain a PTAS for it.

Let $\bar{a} = a_1, \dots, a_n$, m be an input to makespan minimization. Let $\text{BIN}(\bar{a}, t)$ be the minimum number of bins of capacity t that are needed to pack a_1, \dots, a_n . Then,

$$\text{OPT}(\bar{a}, m) = \min_{t \in \mathbb{R}^{\geq 0}} \left\{ t : \text{BIN}(\bar{a}, t) \leq m \right\} \quad \textcircled{1}$$

Let s be the makespan in some schedule.

Step 1: An efficient Bin-Packing algorithm for instances having at most k distinct sizes.

Input: A Bin-Packing instance where there are n_1 items of size s_1 , n_2 items of size s_2 , ..., n_k items of size s_k . $n_1 + n_2 + \dots + n_k = n$. Capacity of each bin is t .

$$\text{Let } \mathcal{Q} = \left\{ (q_1, \dots, q_k) : 0 \leq q_i \leq n_i, 1 \leq i \leq k, \text{BIN}((q_1, \dots, q_k), t) = 1 \right\}$$

$$|\mathcal{Q}| \leq (n_1 + 1) \cdot (n_2 + 1) \cdot \dots \cdot (n_k + 1) = O(n^k).$$

$$T(a_1, \dots, a_k) := 1 \text{ if } (a_1, \dots, a_k) \in \mathcal{Q}.$$

$$\text{BIN}((i_1, \dots, i_k), t) = 1 + \min_{\substack{(a_1, \dots, a_k) \\ \in \mathcal{Q}}} \text{BIN}(i_1 - a_1, \dots, i_k - a_k, t).$$

$O(n^{2k})$ time.

Step 2:

Input: $a_1, \dots, a_n, m =: \bar{a}$

$$\text{OPT} \geq \begin{cases} \max_{i=1}^n a_i \\ \frac{1}{m} \sum_{i=1}^n a_i \end{cases}$$


$$\text{LB} := \max \left\{ \max_{i=1}^n a_i, \frac{1}{m} \sum_{i=1}^n a_i \right\}$$

The greedy algo returns a solution with cost $\leq 2 \cdot \text{LB}$.

$$\Rightarrow \boxed{\text{LB} \leq \text{OPT} \leq 2 \cdot \text{LB}}$$

Let $t \in [\text{LB}, 2 \cdot \text{LB}]$.

- Set aside all a_i 's that are less than $t \cdot \epsilon$.



Let there be u circled points.

$$\Rightarrow t\varepsilon(1+\varepsilon)^u \leq t \quad \& \quad t\varepsilon(1+\varepsilon)^{u+1} \geq t.$$

$$(1+\varepsilon)^u \leq \frac{1}{\varepsilon} \Rightarrow u \leq \log_{1+\varepsilon} \frac{1}{\varepsilon}.$$

$$\text{Let } t\varepsilon(1+\varepsilon)^l \leq a_i \leq t\varepsilon(1+\varepsilon)^{l+1}$$

$$a'_i := t\varepsilon(1+\varepsilon)^l.$$

$a' = (a'_1, \dots, a'_{n'})$. Number of distinct item sizes is at most $\log_{1+\varepsilon} \frac{1}{\varepsilon}$. Pack a' optimally in bins of size t .

[*] Now, pack the smaller items greedily, open new bins if necessary. Let the number of bins used be $\alpha(a, t, \varepsilon)$.

Imagine replacing each a'_i by a_i and the capacity of each bin is $t(1+\varepsilon)$. Now, it's a valid packing of a_1, \dots, a_n in bins of size $t(1+\varepsilon)$.

Claim 1: $\alpha(a, t, \varepsilon) \leq \text{BIN}(a, t)$

Proof: Case 1: No new bin opened to pack the smaller items.

$\alpha(a, t, \epsilon) =$ optimal no. of bins needed to pack items a'_1, \dots, a'_n in bins of size t .

$$\leq \text{BIN}(a, t)$$

Case 2: All but possibly one bin are full upto the capacity of t , and all bins are non-empty.

$$\therefore \sum_{i=1}^n a_i > (\alpha(a, t, \epsilon) - 1) \cdot t$$

\Rightarrow To pack (a_1, \dots, a_n) using bins of capacity t , at least $\alpha(a, t, \epsilon)$ bins are needed

$$\Rightarrow \text{BIN}(a, t) \geq \alpha(a, t, \epsilon). \quad \square$$

Above, we designed an algorithm that has the following input-output behavior.

Input: ① a_1, \dots, a_n . This defines LB .

② $t \in [LB, 2LB]$

③ ε

Output: A packing of a_1, \dots, a_n using at most $BIN(a, t)$ many bins of capacity $t(1 + \varepsilon)$.

Step 3: Recall

$$OPT(\bar{a}, m) = \min_{\substack{t \in \mathbb{R}^+ \\ t \in [LB, 2 \cdot LB]}} \left\{ t : BIN(a, t) \leq m \right\}.$$

Idea: Do binary search in the interval $[LB, 2 \cdot LB]$.

Problems: ① $BIN(a, t)$ cannot be computed efficiently & exactly.

② (if a_1, \dots, a_n are not integers) convergence of the binary search.

Final algorithm A :

- Do binary search on $[LB, 2 \cdot LB]$ to

find the minimum t st. $\alpha(a, t, \varepsilon) \leq m$.

- Stop when the length of the current interval becomes at most $\varepsilon \cdot LB$. Set the current interval by $[c, d]$.

- Pass (a, d, ε) to the algorithm in step 2 and output the computed schedule.

$$d \geq \min_{t \in [LB, 2 \cdot LB]} \{t : \alpha(a, t, \varepsilon) \leq m\} =: t'$$

$$\alpha(a, d, \varepsilon) \leq \alpha(a, t', \varepsilon) \quad (\text{Claim 1})$$

(check).

$$\leq m.$$

\Rightarrow Algorithm A computes a valid schedule.

Makespan of the computed solution

$$\leq d(1 + \varepsilon).$$

$$t' \in [c, d]. \quad d - c \leq \varepsilon \cdot LB$$

$$\Rightarrow d - t' \leq \varepsilon \cdot LB$$

$$\Rightarrow d \leq t' + \varepsilon \cdot LB \leq \text{OPT} \cdot (1 + \varepsilon).$$

A horizontal line segment representing an interval from c to d . A point t' is marked on the segment between c and d .

$[t' \leq \text{OPT} \text{ (SEE LATER)}]$ *

\therefore Makespan of the computed solution is at most

$$\begin{aligned}d(1+\varepsilon) &\leq \text{OPT} (1+\varepsilon)^2 \\ &= \text{OPT} (1+2\varepsilon+\varepsilon^2) \\ &\leq \text{OPT} (1+3\varepsilon) \quad [\because \varepsilon \leq 1].\end{aligned}$$

$$\begin{aligned}(\ast) \left[t' = \min_t \{t : \alpha(a, t, \varepsilon) \leq m\} \right. \\ \leq \min_t \{t : \text{BIN}(a, t) \leq m\} \\ \left. = \text{OPT}. \right]\end{aligned}$$

To get a solⁿ of cost $\leq \text{OPT} \cdot (1+\varepsilon)$, replace ε by $\frac{\varepsilon}{3}$ in the algorithm.