

A Greedy Strategy

List Scheduling:

Order all processes in a list. When a machine runs empty assign the next yet unprocessed job to it.

Alternatively:

Consider processes in some order. Assign the i -th process to the least loaded machine.

It is easy to see that the result of these greedy strategies fulfill the local optimality condition of our local search algorithm. Hence, these also give 2-approximations.

A Greedy Strategy

List Scheduling:

Order all processes in a list. When a machine runs empty assign the next yet unprocessed job to it.

Alternatively:

Consider processes in some order. Assign the i -th process to the least loaded machine.

It is easy to see that the result of these greedy strategies fulfill the local optimality condition of our local search algorithm. Hence, these also give 2-approximations.

A Greedy Strategy

List Scheduling:

Order all processes in a list. When a machine runs empty assign the next yet unprocessed job to it.

Alternatively:

Consider processes in some order. Assign the i -th process to the least loaded machine.

It is easy to see that the result of these greedy strategies fulfill the local optimality condition of our local search algorithm. Hence, these also give 2-approximations.

A Greedy Strategy

List Scheduling:

Order all processes in a list. When a machine runs empty assign the next yet unprocessed job to it.

Alternatively:

Consider processes in some order. Assign the i -th process to the least loaded machine.

It is easy to see that the result of these greedy strategies fulfill the local optimally condition of our local search algorithm. Hence, these also give 2-approximations.

A Greedy Strategy

Lemma 2

If we order the list according to non-increasing processing times the approximation guarantee of the list scheduling strategy improves to $4/3$.

Proof:

- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3} C_{\max}^* .$$

Therefore $p_n > C_{\max}^*/3$.

This means that all jobs must have a processing time

at least

at least $C_{\max}^*/3$ and any machine in the optimum schedule can handle

at most

two jobs. Therefore, the optimum schedule has length at least

Proof:

- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3} C_{\max}^* .$$

Proof:

- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3}C_{\max}^* .$$

Hence, $p_n > C_{\max}^*/3$.

Proof:

- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3}C_{\max}^* .$$

Hence, $p_n > C_{\max}^*/3$.

- ▶ This means that all jobs must have a processing time $> C_{\max}^*/3$.
- ▶ But then any machine in the optimum schedule can handle at most two jobs.
- ▶ For such instances Longest-Processing-Time-First is optimal.

Proof:

- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3}C_{\max}^* .$$

Hence, $p_n > C_{\max}^*/3$.

- ▶ This means that all jobs must have a processing time $> C_{\max}^*/3$.
- ▶ But then any machine in the optimum schedule can handle at most two jobs.
- ▶ For such instances Longest-Processing-Time-First is optimal.

Proof:


- ▶ Let $p_1 \geq \dots \geq p_n$ denote the processing times of a set of jobs that form a counter-example.
- ▶ Wlog. the last job to finish is n (otw. deleting this job gives another counter-example with fewer jobs).
- ▶ If $p_n \leq C_{\max}^*/3$ the previous analysis gives us a schedule length of at most

$$C_{\max}^* + p_n \leq \frac{4}{3}C_{\max}^* .$$

Hence, $p_n > C_{\max}^*/3$.

- ▶ This means that all jobs must have a processing time $> C_{\max}^*/3$.
- ▶ But then any machine in the optimum schedule can handle at most two jobs.
- ▶ For such instances Longest-Processing-Time-First is optimal.

When in an optimal solution a machine can have at most 2 jobs the optimal solution looks as follows.



- ▶ We can assume that one machine schedules p_1 and p_n (the largest and smallest job).
- ▶ If not assume wlog. that p_1 is scheduled on machine A and p_n on machine B .
- ▶ Let p_A and p_B be the other job scheduled on A and B , respectively.
- ▶ $p_1 + p_n \leq p_1 + p_A$ and $p_A + p_B \leq p_1 + p_A$, hence scheduling p_1 and p_n on one machine and p_A and p_B on the other, cannot increase the Makespan.
- ▶ Repeat the above argument for the remaining machines.

- ▶ We can assume that one machine schedules p_1 and p_n (the largest and smallest job).
- ▶ If not assume wlog. that p_1 is scheduled on machine A and p_n on machine B .
- ▶ Let p_A and p_B be the other job scheduled on A and B , respectively.
- ▶ $p_1 + p_n \leq p_1 + p_A$ and $p_A + p_B \leq p_1 + p_A$, hence scheduling p_1 and p_n on one machine and p_A and p_B on the other, cannot increase the Makespan.
- ▶ Repeat the above argument for the remaining machines.

- ▶ We can assume that one machine schedules p_1 and p_n (the largest and smallest job).
- ▶ If not assume wlog. that p_1 is scheduled on machine A and p_n on machine B .
- ▶ Let p_A and p_B be the other job scheduled on A and B , respectively.
- ▶ $p_1 + p_n \leq p_1 + p_A$ and $p_A + p_B \leq p_1 + p_A$, hence scheduling p_1 and p_n on one machine and p_A and p_B on the other, cannot increase the Makespan.
- ▶ Repeat the above argument for the remaining machines.

- ▶ We can assume that one machine schedules p_1 and p_n (the largest and smallest job).
- ▶ If not assume wlog. that p_1 is scheduled on machine A and p_n on machine B .
- ▶ Let p_A and p_B be the other job scheduled on A and B , respectively.
- ▶ $p_1 + p_n \leq p_1 + p_A$ and $p_A + p_B \leq p_1 + p_A$, hence scheduling p_1 and p_n on one machine and p_A and p_B on the other, cannot increase the Makespan.
- ▶ Repeat the above argument for the remaining machines.

- ▶ We can assume that one machine schedules p_1 and p_n (the largest and smallest job).
- ▶ If not assume wlog. that p_1 is scheduled on machine A and p_n on machine B .
- ▶ Let p_A and p_B be the other job scheduled on A and B , respectively.
- ▶ $p_1 + p_n \leq p_1 + p_A$ and $p_A + p_B \leq p_1 + p_A$, hence scheduling p_1 and p_n on one machine and p_A and p_B on the other, cannot increase the Makespan.
- ▶ Repeat the above argument for the remaining machines.