CSCI 3100: Algorithms

Greedy Algs (pt 2)

Announcements

Overall greedy strategy:

- Assume optimal is different than greedy
- Find the "first" place they differ.
- Argue that we can exchange the two without making optimal worse.

⟹ there is no "first place" where they must differ, so greedy in fact is an optimal solution.

Another example in notes: storing the most files on a tape

Intuition: (check notes)

A different scheduling problem:

Setting: a single resource (ie a processor / CPU)

Input: n requests, each with:

$D[1..n]$: deadline $D[i]$ is the time when request i should be completed

$T[1..n]$: $T[i]$ is amount of time that request i needs on the resource

Goal: Run everything.

If not everything can be finished by its deadline, then minimize the worst "lateness"

Lateness: Let's formalize

If i gets finish time $F[i]$,

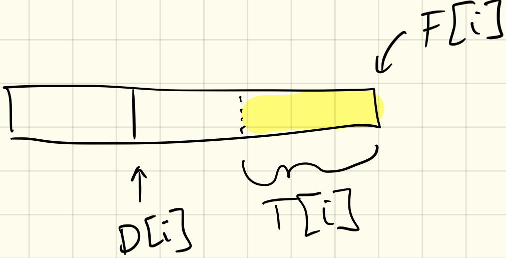

lateness $L[i] = F[i] - D[i]$

Goal:

Example: Job 1:

Job 2:

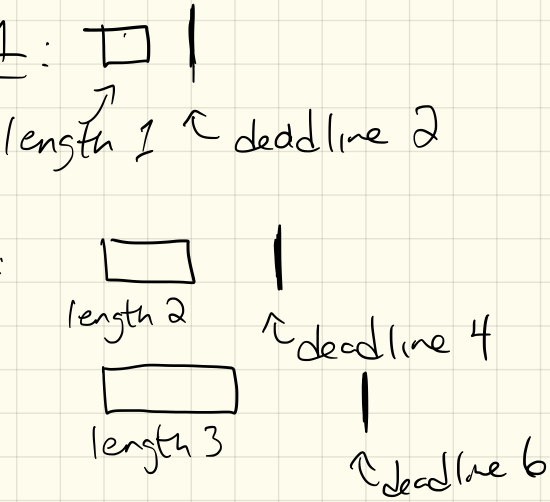

Input: D

T:

Schedule:

lateness?

Example: Job 1: length 1 deadline 2

Job 2: length 2 deadline 3

length 3 deadline 4

Schedule?

Ideas for how to be greedy:

Earliest deadline first (EDF)

Sort by $D[i]$ & run in this order!

Sort of hard to believe this works!

That's why the proof is key...

But first: run time?

Correctness:

First: We may assume the optimal will have no idle time.

Why?

schedule:

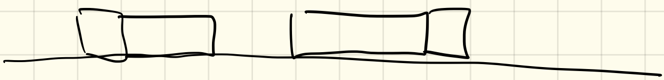

Definition: I'll say 2 jobs are inverted if job i goes before job j in the schedule, but $D[i] > D[j]$.

Note: our greedy schedule has no inversions.

Lemma: All schedules with no inversions and no idle time have the same max lateness.

pf:

Thm: There is an optimal schedule with no inversions.

pf: Suppose opt must have inversions.

Then $D[a] > D[b]$
but

opt: 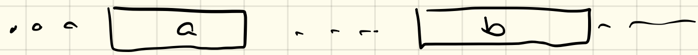

pf cont: So consider 1st adjacent inversions, i + j.

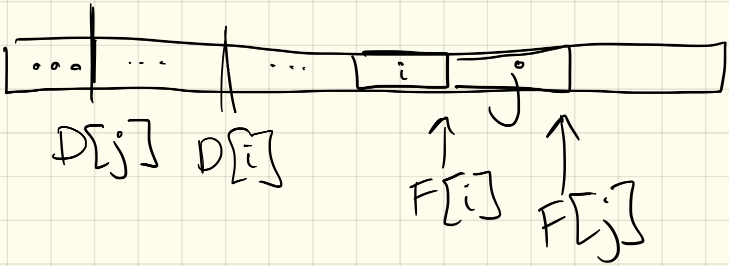

idea: Swap them!

Know j gets better.

What about i?

<u>formalize</u>: Worried about i:

After swap, i finishes at $F[j]$ from 1st schedule.

New lateness for i:

$$F[j] - D[i]$$

What was j's before the swap?

Finally:

Since things only get better if we fix inversions, can just keep swapping.

Will we reach a schedule with none?

ie - How many inversions can there be in the worst case?