

CMSC 451 - Algorithm Design

Lecture 5 - Greedy Algorithms for Scheduling

Discrete Optimization:

Compute a discrete structure of a given class (e.g., subset, tree, path, partition) to maximize/minimize a given objective function (e.g., cost, distance, size) subject to a given set of constraints (e.g., disjointness, connectedness, completeness).

A feasible solution satisfies all constraints

An optimal solution is feasible and max/minimizes the objective function

Example:

Min. Spanning Tree: Given a connected, weighted graph compute:

structure - subset of edges

objective - min. total weight

constraints - connected, acyclic, cover all vertices

Common Strategies:

- Brute-force search - Try all possibilities - Slow!!
- Local search - Find an init. feasible solution + repeatedly make small improvements
- Dynamic programming - (Future lectures)
- Greedy - Build a solution by repeated additions (never revoked/reversed) each based on best choice subject to constraints

E.g., Kruskal's MST algorithm
(add min. weight edge that doesn't cause a cycle)

- ...

This lecture: Three scheduling problems

- Interval scheduling
- Interval partitioning
- Schedule to minimize lateness


Interval Scheduling:

Given a set $R = \{r_1, \dots, r_n\}$ of requests, each being an interval $r_i = [s_i, f_i]$ compute a subset of non-overlapping requests of max. cardinality

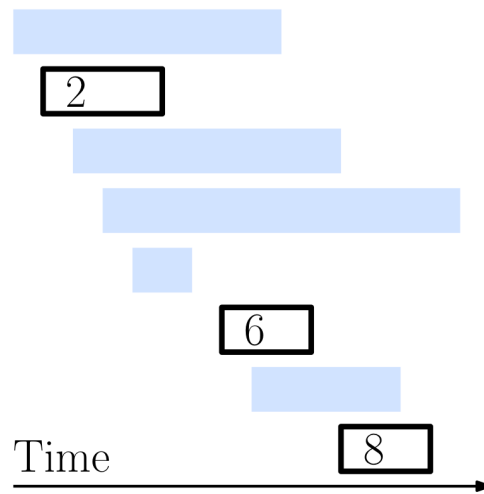
Application: Scheduling events at a facility

Example:

Requests:



Possible solution: {2, 6, 8} Also: {5, 6, 8}



3 requests satisfied

Greedy Approach - Select a request that does not conflict with prior selection + min. some measure:

- **Earliest Start** - request with smallest start, s_i
- **Earliest Finish** - " " " finish, f_i
- **Shortest Duration** - " " " duration, $f_i - s_i$
- **Min conflicts** - overlaps fewest among remaining requests

Optimal?

- **Earliest Start** - X
 - **Earliest Finish** - ✓
 - **Shortest Duration** - X
 - **Min conflicts** - X
- Exercise: Find counterexamples

Earliest Finish First:

greedy-interval-sched(s, f)

sort requests by f -times

$S \leftarrow \emptyset$ // init empty schedule

prevFinish $\leftarrow -\infty$

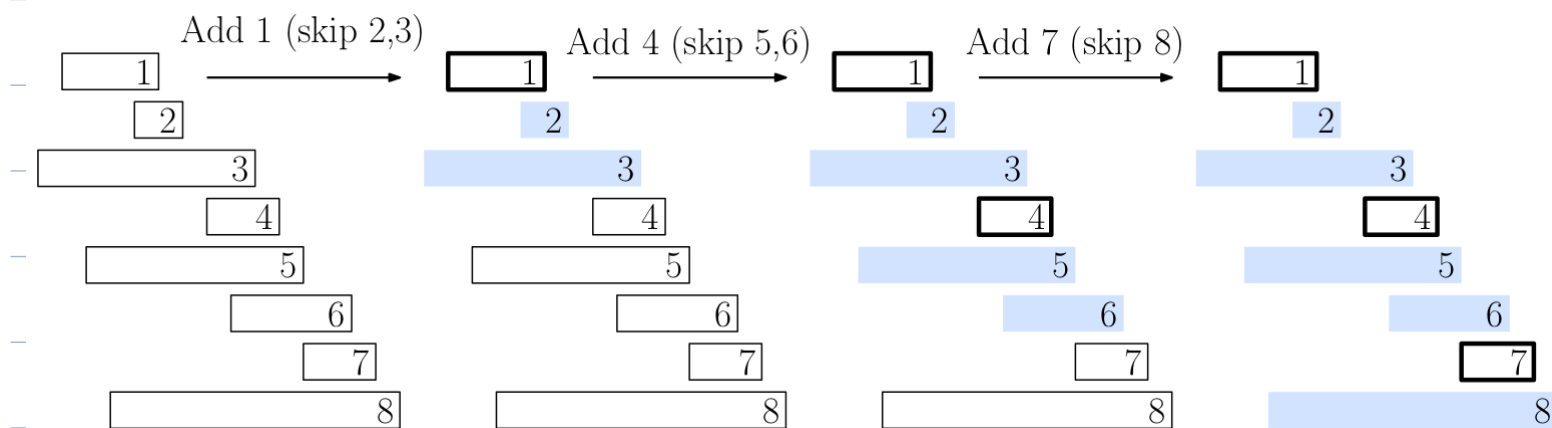
for ($i \leftarrow 1$ to n)

if ($s[i] > \text{prevFinish}$) // no conflict?

append i to S // ... schedule it

prevFinish $\leftarrow f[i]$

Example: (Presorted by finish times)



Running time:



$O(n \log n)$ - sort

$O(n)$

- remaining processing

} $O(n \log n)$
total

Correctness:

Must show: **Feasibility** - A valid schedule (no conflicts)
Optimality - Maximizes no. of requests

Feasibility: **Easy** - No request is scheduled until after prev request finishes

Optimality: Not so easy - Let's do this rigorously.

Let $O = \langle x_1, \dots, x_k \rangle$ be any optimal schedule
(may be many)

Let $G = \langle g_1, \dots, g_{k'} \rangle$ be EFF greedy schedule
($k' \leq k$)


If $G = O$, we're done!

otherwise, let j be smallest index s.t. $x_j \neq g_j$

- By definition of EFF, $f[g_j] \leq f[x_j]$

- Form a new schedule O' by replacing

x_j by g_j - Still feasible + same size



Repeat (hidden induction) until $O'''' = G \Rightarrow G$ is optimal

□

Next problem:



Interval Partitioning: Suppose we have ∞ -many resources, but they cost \$\$\$

Want to satisfy all requests with fewest resources

Problem:

Given a set $R = \{r_1, \dots, r_n\}$ of requests, each with start/finish times s_i / f_i , partition R into the smallest num. of sets R_1, \dots, R_d such that all requests in R_i are pairwise non-conflicting

Example:

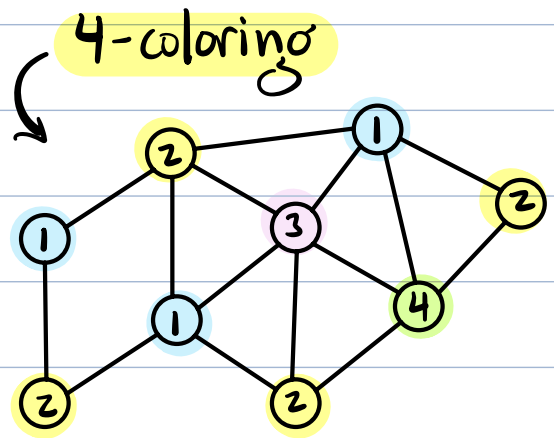



This is an example coloring -

Given a set of objects + a conflict relation, partition into smallest number of conflict-free subsets

Number of subsets = coloring number

Graph coloring:





Graph coloring is NP-hard, but interval coloring is much easier.

Greedy Strategy for Interval Partitioning Lowest Available Color

```
greedy-interval-partition(s, f) // s=start
                                // f=finish
  sort by start times
  for (i ← 1 to n)
    X ← ∅ // X = excluded colors
    for (j ← 1 to i-1)
      if ([sj, fj] overlaps [si, fi])
        add color [j] to X
    color[i] ← smallest color not in X
  return color array
```

Example:



$X = \{2, 3\}$ ↑
↑ smallest available = 1

Running Time:



- Sorting - $O(n \log n)$

- 2 nested loops $1 \dots n + 1 \dots i-1$

$$= \sum_{i=1}^n (i-1) = O(n^2)$$

$O(n^2)$
total

There is smarter approach
that runs in $O(n)$ time after
sorting - see lecture notes.

Correctness: Need to show

- Feasibility - Obvious (Avoid conflicting colors)
- Optimality - ??


Optimality - Approach

- Define a statistic - depth
- Lower bound - Any solution must use \geq depth colors
- Upper bound - Greedy algorithm uses \leq depth colors

\Rightarrow Greedy is optimal

Depth - Given request set $R = \{r_1, \dots, r_n\}$, $r_i = [s_i, f_i]$
+ time t define

$$\text{depth}_R(t) = \text{num. intervals overlap } t \\ = |\{i \mid t \in [s_i, f_i]\}|$$



and: $\text{depth}(R) = \max_{t \geq 0} \text{depth}_R(t)$

Clearly, all requests contributing to depth conflict...

Lemma: For any d -coloring of R , $d \geq \text{depth}(R)$

The following implies that greedy is optimal

Lemma: The greedy algorithm generates a d -coloring, where $d \leq \text{depth}(R)$

Proof: Suppose towards a contradiction that there is a first time s_i where greedy uses more than $\text{depth}_R(s_i)$ colors.

Consider time t just prior to s_i

$$\text{depth}_R(s_i) = \text{depth}_R(t) + 1 \quad \textcircled{a}$$


Let $d = \text{num. greedy colors at time } t$. Since s_i is the first error:

$$d \leq \text{depth}_R(t) \quad \textcircled{b}$$

Thus, there are d excluded colors in X when s_i is processed. Implies that greedy uses

$$d + 1 \leq \text{depth}_R(t) + 1 = \text{depth}_R(s_i) \quad \textcircled{a}$$

Contradicting hyp. that error at s_i \square





Q: If requests were sorted by a different criterion, would the algorithm still be optimal?

Q: Can we modify this to color graphs? (NP-hard!)

Scheduling to Minimize Lateness

- Tasks rather than requests - $X = \{x_1, \dots, x_n\}$
- Execution time - t_i
- Deadline - d_i
- Application - Homework assignments

Objective:

- Compute start times $S = \{s_1, \dots, s_n\}$ s.t.
- Tasks do not overlap $[s_i, f_i] \cap [s_j, f_j] = \emptyset$
where $f_i = s_i + t_i$
- Minimize lateness - max. deadline excess

$$l_i = \max(0, f_i - d_i)$$

How far beyond the deadline did you finish?


max. lateness:

$$L(S) = \max_{1 \leq i \leq n} l_i$$

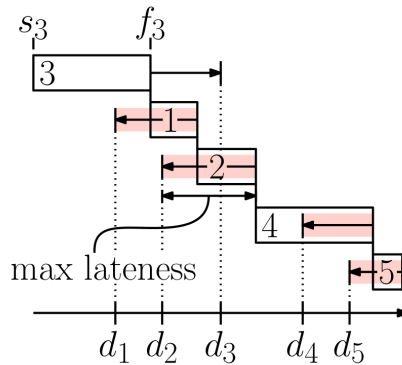
your worst deadline excess (not sum)

Example:

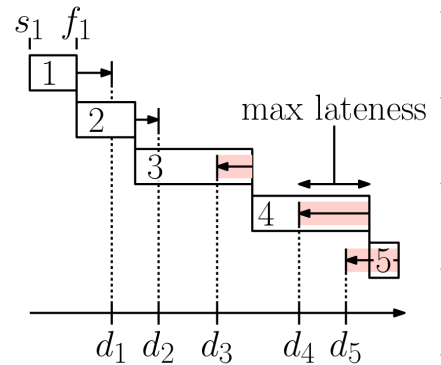
Input:



Possible solution:



Optimal solution:



Greedy approach

- Schedule based on some **criteria**

- Shortest duration first - Sort by t_i \times
- Earliest deadline first - Sort by d_i \checkmark
- Smallest slack first - Sort by $d_i - t_i$ \times

greedy-lateness-sched(t, d)

sort by **deadlines** ($d_1 \leq \dots \leq d_n$)

prevFinish $\leftarrow 0$ // when prev task finished

for ($i \leftarrow 1$ to n)

$s_i \leftarrow$ prevFinish // start after prev finish

prevFinish $\leftarrow f_i \leftarrow s_i + t_i$ // update finish

$l_i \leftarrow \max(0, f_i - d_i)$ // lateness

return $[s_1, \dots, s_n]$ $L = \max l_i$ // return start times

Running Time:

- Sorting - $O(n \log n)$
 - Processing - $O(n)$
- } Total: $O(n \log n)$

Correctness: Need to show


- Feasibility - Easy - No conflicting tasks
- Optimality - ??

Lemma: The greedy algorithm minimizes max. lateness.

Proof: We may limit consideration to schedules that are "slack-free"

- no gaps between tasks. (Greedy is slack-free)

- Let σ be any lateness-optimal, slack-free schedule.
- If σ is deadline sorted - we're done! $\sigma = G$
- o.w. let $x_i + x_j$ be first consecutive pair not in deadline order $d_i > d_j$
- We'll swap them + show that max lateness can only decrease.



- Let $l_i^\sigma + l_j^\sigma$ be latenesses before swap +
 $l_i^G + l_j^G$ " " after swap

- These are the only latenesses affected

- Want to show:

$$\max(l_i^G, l_j^G) \leq \max(l_i^\sigma, l_j^\sigma)$$

↪ Swap improves max lateness

- Let t be current time

- For simplicity, assume $d_i + d_j$ passed: $d_j < d_i \leq t$
 (Exercise - Fix this)

↪ We can ignore max

- Observe:

$$l_i^\sigma = (t + t_i) - d_i \quad l_j^\sigma = (t + t_i + t_j) - d_j$$

$$\text{Since } d_i > d_j + t_j \geq 0$$

$$l_j^\sigma = (t + t_i) + t_j - d_j > (t + t_i) - d_i = l_i^\sigma$$

$$\Rightarrow \max(l_i^\sigma, l_j^\sigma) = l_j^\sigma$$

- Also:

$$l_i^G = (t + t_i + t_j) - d_i < (t + t_i + t_j) - d_j = l_j^\sigma$$

$$l_j^G = (t + t_j) - d_j \leq (t + t_i + t_j) - d_j = l_j^\sigma$$

- Thus:

$$\max(l_i^G, l_j^G) \leq \max(l_j^\sigma, l_j^\sigma) = \max(l_i^\sigma, l_j^\sigma) \quad \square$$

Summary: 3 examples of greedy solutions to simple scheduling problems.

Interval scheduling - Earliest finish first

Interval partitioning - Smallest available color

Minimizing Max Lateness - Earliest deadline first

↳ All $O(n \log n)$