

Chapter 16


Greedy Algorithms

- typically applied to optimization problems
- Intuitively appealing, but not optimal in most cases
- very efficient and simple to implement
- no backtracking: solution is built by a series of decisions and the decision are never reversed at a future point.
- choice of decision is based on local information (optimization applied at each step locally)
- problem is the cumulative effect of locally optimally decisions may not give a globally optimal solution.

Problem 1: Interval Scheduling

Interval scheduling.

- Job j starts at s_j and finishes at f_j .
- Two jobs **compatible** if they don't overlap.
- Goal: find maximum subset of mutually compatible jobs.



Interval Scheduling: Greedy Algorithms

Greedy template. Consider jobs in some order. Take each job provided it's compatible with the ones already taken.

- [Earliest start time] Consider jobs in ascending order of start time s_j .
- [Shortest interval] Consider jobs in ascending order of interval length $f_j - s_j$.
- [Fewest conflicts] For each job, count the number of conflicting jobs c_j . Schedule in ascending order of conflicts c_j .
- [Earliest finish time] Consider jobs in ascending order of finish time f_j .

Interval Scheduling: Greedy Algorithms


Greedy template. Consider jobs in some order. Take each job provided it's compatible with the ones already taken.



breaks earliest start time



breaks shortest interval



breaks fewest conflicts

Interval Scheduling: Greedy Algorithm

Greedy algorithm. Consider jobs in increasing order of finish time. Take each job provided it's compatible with the ones already taken.

```
Sort jobs by finish times so that  $f_1 \leq f_2 \leq \dots \leq f_n$ .  
  ↙ jobs selected  
A ← {1}  
for j = 2 to n {  
    if (job j compatible with A)  
        A ← A ∪ {j}  
}  
return A
```

Implementation. $O(n \log n)$.


- suppose job j^* that was added last to A then job j is compatible with A if $s_j \geq f_{j^*}$.

Interval Scheduling: Analysis

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

- Assume greedy is not optimal, and let's see what happens.
- Let i_1, i_2, \dots, i_k denote set of jobs selected by greedy.
- Let j_1, j_2, \dots, j_m denote set of jobs in the optimal solution with $i_1 = j_1, i_2 = j_2, \dots, i_r = j_r$ for the largest possible value of r .

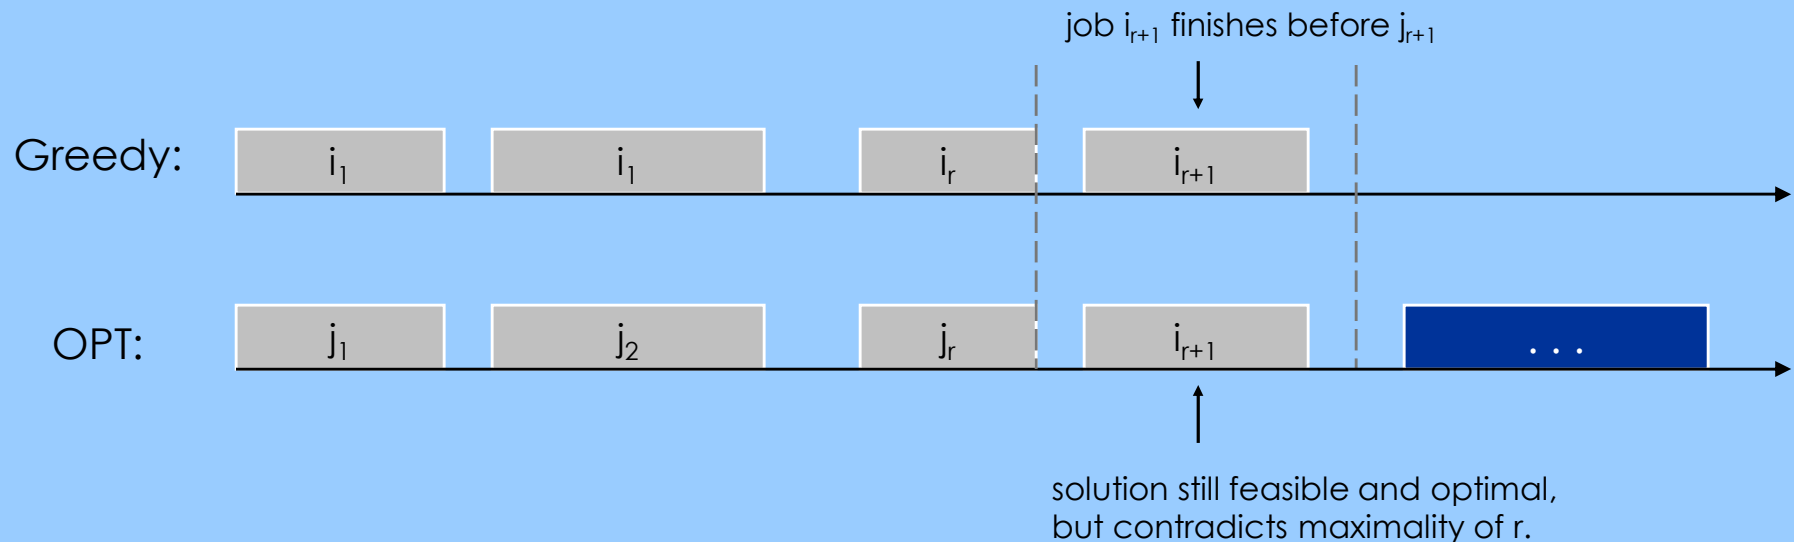


Interval Scheduling: Analysis

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

- Assume greedy is not optimal, and let's see what happens.
- Let i_1, i_2, \dots, i_k denote set of jobs selected by greedy.
- Let j_1, j_2, \dots, j_m denote set of jobs in the optimal solution with $i_1 = j_1, i_2 = j_2, \dots, i_r = j_r$ for the largest possible value of r .




Problem 2: Interval Partitioning

Interval partitioning.

- Lecture j starts at s_j and finishes at f_j .
- Goal: find minimum number of classrooms to schedule all lectures so that no two occur at the same time in the same room.

Ex: This schedule uses 4 classrooms to schedule 10 lectures.




Interval Partitioning

Interval partitioning.

- Lecture j starts at s_j and finishes at f_j .
- Goal: find minimum number of classrooms to schedule all lectures so that no two occur at the same time in the same room.

Ex: This schedule uses only 3.



Interval Partitioning: Lower Bound on Optimal Solution


Def. The **depth** of a set of open intervals is the maximum number that contain any given time.

Key observation. Number of classrooms needed \geq depth.

Ex: Depth of schedule below = 3 \Rightarrow schedule below is optimal.

↑
a, b, c all contain 9:30

Q. Does there always exist a schedule equal to depth of intervals?



Interval Partitioning: Greedy Algorithm

Greedy algorithm. Consider lectures in increasing order of start time: assign lecture to any compatible classroom.

```
Sort intervals by starting time so that  $s_1 \leq s_2 \leq \dots \leq s_n$ .  
d ← 0 ← number of allocated classrooms  
  
for j = 1 to n {  
  if (lecture j is compatible with some classroom k)  
    schedule lecture j in classroom k  
  else  
    allocate a new classroom d + 1  
    schedule lecture j in classroom d + 1  
    d ← d + 1  
}
```

Implementation. $O(n \log n)$.

- For each classroom k , maintain the finish time of the last job added.
- Keep the classrooms in a priority queue.

Implementation details

The implementation of the above algorithm is tricky, especially if we want to achieve an $O(n \log n)$ bound.

First, we note that the depth d can be computed in time $O(n \log n)$. We discussed this algorithm in an earlier lecture.

Input: n intervals $(s_1, t_1), (s_2, t_2), \dots, (s_n, t_n)$.

Output: The depth d = the maximum overlap among the intervals.

Step 1: Create two arrays A and B of size $2n$ where A contains the $2n$ inputs sorted. $B[j] = 0$ (1) if $A[j]$ is the opening (closing) of an interval.

```
Step 2: depth = 1; max = 1;
        for j from 2 to 2n do
          { if (B[j] == 0)
            { depth++; if (depth > max_depth) max_depth = depth;}
            else depth--;
          }
        return max_depth;
```

Implementation of algorithm for interval scheduling

Idea: We process the intervals in increasing order of start time. Suppose the next interval is $I = (s_j, t_j)$. We keep in a linked list L the available rooms for scheduling the interval I . We choose a room from this linked list and schedule I in that room. We want to make sure that this scheduling is valid and that the list L is never empty.

When we are considering I , if a room is in L , this means that this room is free for scheduling I . Before we choose a room from L and schedule the interval I , we will remove the rooms that are potentially available for scheduling I from a heap. The heap contains collection of pairs (key, room_number) where key is the finish time of the last interval scheduled at the room with number = room_number.

Thus, for each key $< s_j$, we delete the node from the heap, and insert the corresponding room into the list L .

Then, we pick a room r from L , delete it from L and schedule I in that room, and insert the pair (t_j, r) into the heap.

Interval Partitioning: Greedy Analysis

Observation. Greedy algorithm never schedules two incompatible lectures in the same classroom.


Theorem. Greedy algorithm is optimal.

Pf.


- Let d = number of classrooms that the greedy algorithm allocates.
- Classroom d is opened because we needed to schedule a job, say j , that is incompatible with all $d-1$ other classrooms.
- Since we sorted by start time, all these incompatibilities are caused by lectures that start no later than s_j .
- Thus, we have d lectures overlapping at time $s_j + \epsilon$.
- Key observation \Rightarrow all schedules use $\geq d$ classrooms. ▪

Minimum Spanning Tree

Minimum spanning tree. Given a connected graph $G = (V, E)$ with real-valued edge weights c_e , an MST is a subset of the edges $T \subseteq E$ such that T is a spanning tree whose sum of edge weights is minimized.



$G = (V, E)$



$T, \sum_{e \in T} c_e = 50$

Applications

MST is fundamental problem with diverse applications.

- Network design.
 - telephone, electrical, hydraulic, TV cable, computer, road
- Approximation algorithms for NP-hard problems.
 - traveling salesperson problem, Steiner tree
- Indirect applications.
 - max bottleneck paths
 - LDPC codes for error correction
 - image registration with Renyi entropy
 - learning salient features for real-time face verification
 - reducing data storage in sequencing amino acids in a protein
 - model locality of particle interactions in turbulent fluid flows
 - autoconfig protocol for Ethernet bridging to avoid cycles in a network
- Cluster analysis.

Greedy Algorithms

Kruskal's algorithm. Start with $T = \phi$. Consider edges in ascending order of cost. Insert edge e in T unless doing so would create a cycle.

Reverse-Delete algorithm. Start with $T = E$. Consider edges in descending order of cost. Delete edge e from T unless doing so would disconnect T .

Prim's algorithm. Start with some root node s and greedily grow a tree T from s outward. At each step, add the cheapest edge e to T that has exactly one endpoint in T .

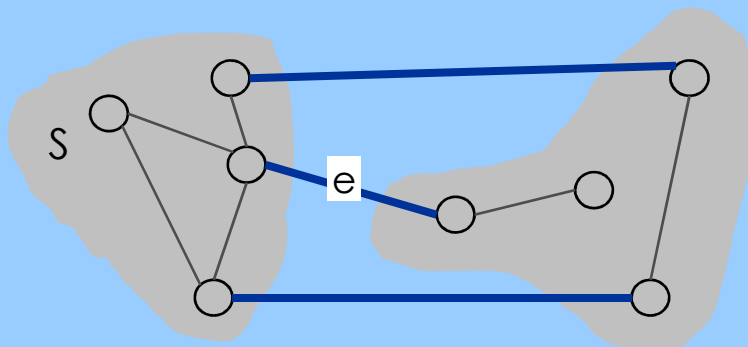
Remark. All three algorithms produce an MST.

Greedy Algorithms


Simplifying assumption. All edge costs c_e are distinct.

Cut property. Let S be any subset of nodes, and let e be the min cost edge with exactly one endpoint in S . Then the MST contains e .

Cycle property. Let C be any cycle, and let f be the max cost edge belonging to C . Then the MST does not contain f .




e is in the MST



f is not in the MST


Cycles and Cuts

Cycle. Set of edges the form a-b, b-c, c-d, ..., y-z, z-a.



Cycle C = 1-2, 2-3, 3-4, 4-5, 5-6, 6-1


Cutset. A cut is a subset of nodes S. The corresponding cutset D is the subset of edges with exactly one endpoint in S.



Cut S = { 4, 5, 8 }
Cutset D = 5-6, 5-7, 3-4, 3-5, 7-8

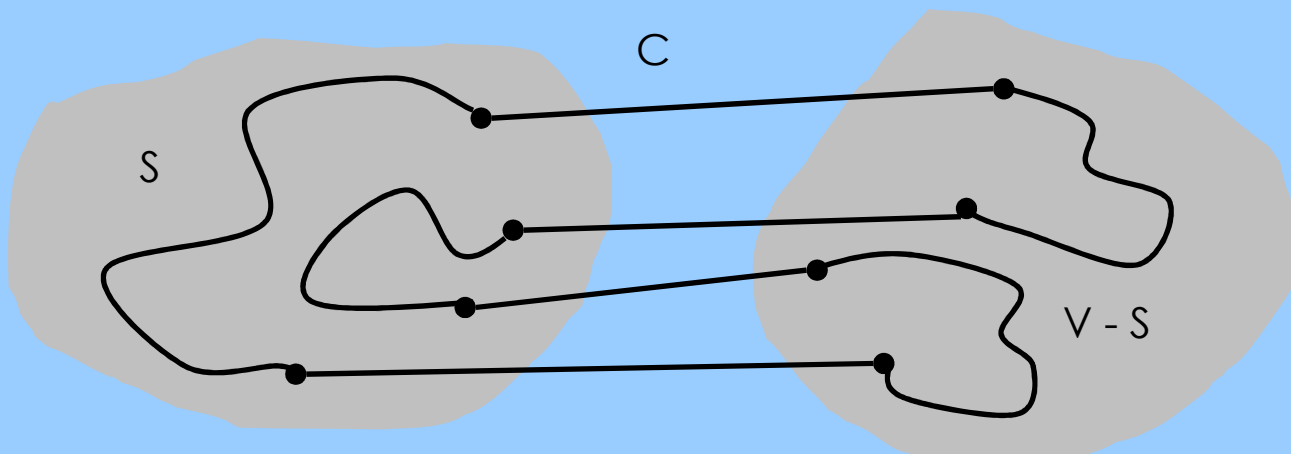
Cycle-Cut Intersection

Claim. A cycle and a cutset intersect in an even number of edges.



Cycle $C = 1-2, 2-3, 3-4, 4-5, 5-6, 6-1$
Cutset $D = 3-4, 3-5, 5-6, 5-7, 7-8$
Intersection = $3-4, 5-6$

Pf. (by picture)




Greedy Algorithms

Simplifying assumption. All edge costs c_e are distinct.

Cut property. Let S be any subset of nodes, and let e be the min cost edge with exactly one endpoint in S . Then the MST T^* contains e .

Pf. (exchange argument)

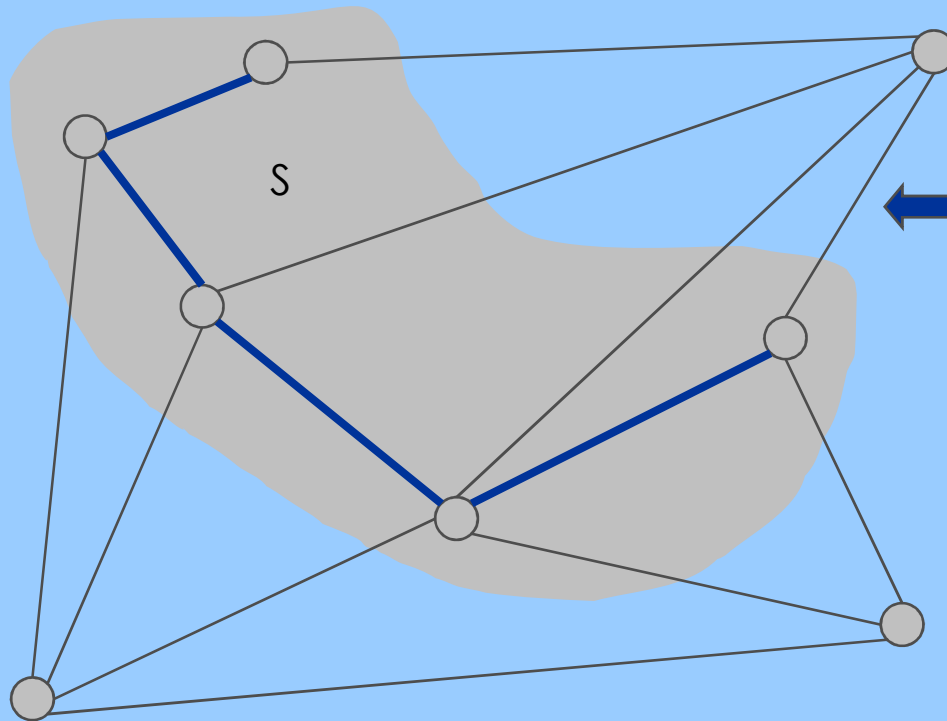
- Suppose e does not belong to T^* , and let's see what happens.
- Adding e to T^* creates a cycle C in T^* .
- Edge e is both in the cycle C and in the cutset D corresponding to $S \Rightarrow$ there exists another edge, say f , that is in both C and D .
- $T' = T^* \cup \{e\} - \{f\}$ is also a spanning tree.
- Since $c_e < c_f$, $\text{cost}(T') < \text{cost}(T^*)$.
- This is a contradiction. ▪



Prim's Algorithm: Proof of Correctness

Prim's algorithm. [Jarník 1930, Dijkstra 1957, Prim 1959]

- Initialize $S =$ any node.
- Apply cut property to S .
- Add min cost edge in cutset corresponding to S to T , and add one new explored node u to S .



Implementation: Prim's Algorithm

Implementation. Use a priority queue just as in shortest path algorithm

- Maintain set of explored nodes S .
- For each unexplored node v , maintain attachment cost $a[v] =$ cost of cheapest edge v to a node in S .
- $O(n^2)$ with an array; $O(m \log n)$ with a binary heap.

```
Prim(G, c) {  
    foreach (v ∈ V) a[v] ← ∞  
    Initialize an empty priority queue Q  
    foreach (v ∈ V) insert v onto Q  
    Initialize set of explored nodes S ← ∅  
  
    while (Q is not empty) {  
        u ← delete min element from Q  
        S ← S ∪ { u }  
        foreach (edge e = (u, v) incident to u)  
            if ((v ∉ S) and (ce < a[v]))  
                decrease priority a[v] to ce  
    }  
}
```