Objectives

- Introduction to Greedy Algorithms
- Interval Scheduling

Greedy Algorithms

At each step, take as much as you can get
→ "local" optimizations

- Need a proof to show that the algorithm finds an optimal solution
- A counter example shows that a greedy algorithm does not provide an optimal solution

Example of Greedy Algorithm

- How do you make change to give out the *fewest* coins?
- Determine for 34¢

Example of Greedy Algorithm

- How do you make change to give out the *fewest* coins?

```
while change > 0:
        if change >= 25:
                print "Quarter"
                change -= 25
        elif change >= 10:
                print "Dime"
                change -= 10
        …
```

Let's generalize …

- Ex: 34¢. 

Coin Changing

- Goal. Given currency denominations: 1, 5, 10, 25, 100, devise a method to pay amount to customer using fewest number of coins.
- Ex: 34¢. 
- Cashier's algorithm. At each iteration, add coin of the largest value that does not take us past the amount to be paid.
- Ex: $2.89. 

Coin-Changing: Greedy Algorithm

- Cashier's algorithm. At each iteration, add coin of the largest value that does not take us past the amount to be paid.

Sort coins' denominations by value: $c_1 < c_2 < \ldots < c_n$.

coins selected

```
S ← ϕ
while x ≠ 0
   let k be largest integer such that c_k ≤ x
   if k = 0
      return "no solution found"   ← How could this happen?
   x ← x - c_k
   S ← S ∪ {k}
return S
```

Is cashier's algorithm ***optimal***?

Coin-Changing: Analysis of Greedy Algorithm

- Theorem. Greedy is optimal for U.S. coinage: 1, 5, 10, 25, 100
- Pf. (by induction on x)
  - Consider optimal way to change $c_k \le x < c_{k+1}$
    - Greedy takes coin *k*
  - Any optimal solution must also take coin *k*
    - If not, it needs enough coins of type $c_1, \ldots, c_{k-1}$ to add up to *x*
    - Table below indicates no optimal solution can do this
  - Problem reduces to coin-changing x - c_k cents, which, by induction, is optimally solved by greedy algorithm. ▪

Which coin: k	Coin value: c_k	All optimal solutions must satisfy	Max value of coins 1, 2, …, k-1 in any OPT (If don't take c_k)
1	1	P ≤ 4	-
2	5	N ≤ 1	4
3	10	N + D ≤ 2	4 + 5 = 9
4	25	Q ≤ 3	20 + 4 = 24
5	100	no limit	75 + 24 = 99

Coin-Changing: Analysis of Greedy Algorithm

- Observation. Greedy algorithm is sub-optimal for US postal denominations:
  - 500 300 200 100 86 85 79 78 66 65 46 44 33 32 20 4 3 2 1
- Counterexample. 158¢.
  - Greedy: 100, 44, 4, 4, 4, 2.
  - Optimal: 79, 79.

Proving Greedy Algorithms Work

- Specifically, produce an **optimal** solution
- Approaches:
  - Greedy algorithm stays ahead
    - Does better than any other algorithm at each step
  - Exchange argument
    - Transform any solution into a greedy solution
  - Structural argument
    - Figure out some structural bound that all solutions must meet

Greedy algorithm stays ahead

INTERVAL SCHEDULING

Interval Scheduling

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

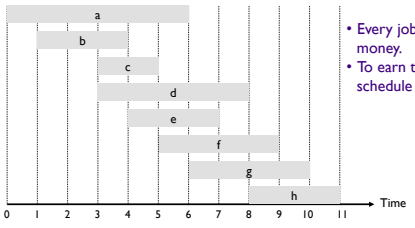

- Every job is worth equal money.
- To earn the most money → schedule the most jobs

Greedy Algorithm Template

- Consider jobs (or whatever) in some order
  - Decision: What order is best?
- Take each job provided it's compatible with the ones already taken

What are options for orders? (rhetorical for now)

What is our goal?
What are we trying to minimize/maximize?

What is the worst case?

Greedy Algorithm Pseudo-Code

In some specified order

```
Set Greedy (Set candidate){
   solution = new Set( );
   while candidate.isNotEmpty()
      next = candidate.select() //use selection criteria,
      //remove from candidate and return value
      if solution.isFeasible(next) //constraints satisfied
         solution.union(next)
      if solution.solves()
         return solution

   //No more candidates and no solution
   return null
}
```


Interval Scheduling

- Earliest start time. Consider jobs in ascending order of start time s_j
  - Utilize CPU as soon as possible
- Earliest finish time. Consider jobs in ascending order of finish time f_j
  - Resource becomes free ASAP
  - Maximize time left for other requests
- 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

Can we "break" any of these? i.e., prove they're not optimal?

Counterexamples to Optimality of Various Job Orders

Not optimal when …

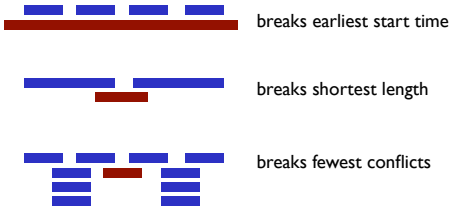

Interval Scheduling: Greedy Algorithm

- 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 f1 ≤ f2 ≤ ... ≤ fn
G = {}
for j = 1 to n
   if job j compatible with G
      G = G ∪ {j}
return G
```

jobs selected

Interval Scheduling

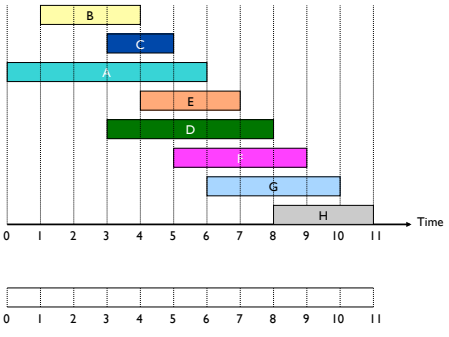

Interval Scheduling

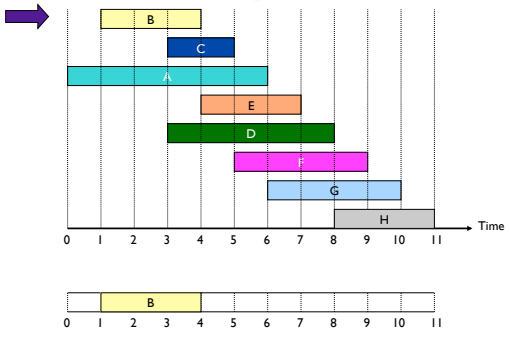

Interval Scheduling

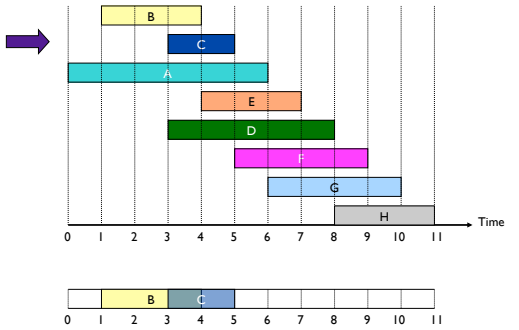

Interval Scheduling

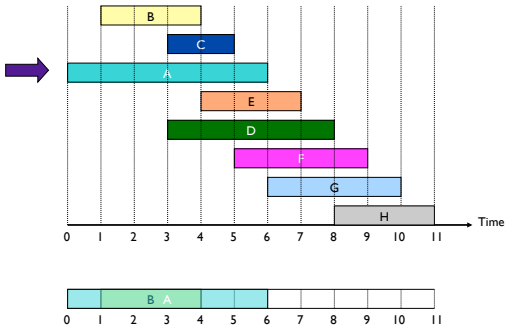

Interval Scheduling

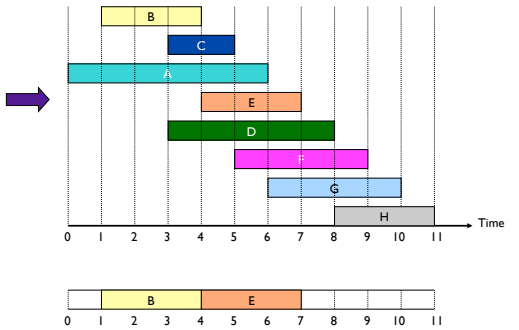

Interval Scheduling

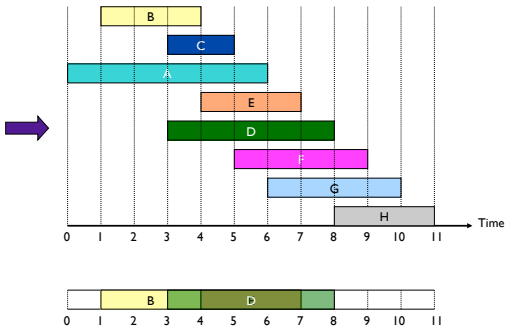

Interval Scheduling

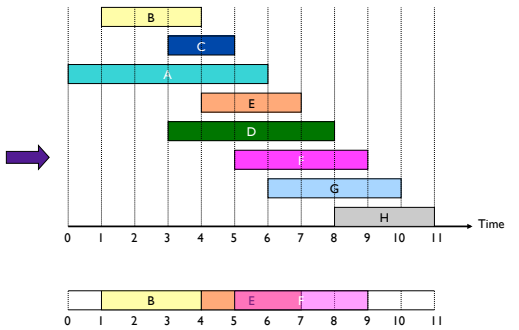

Interval Scheduling

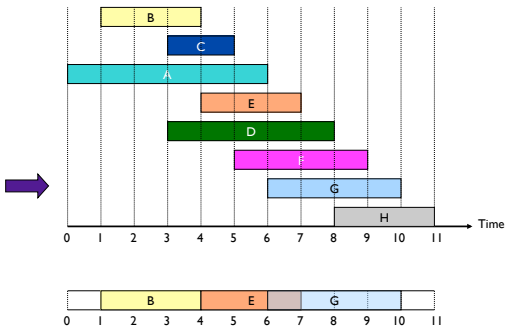

Interval Scheduling

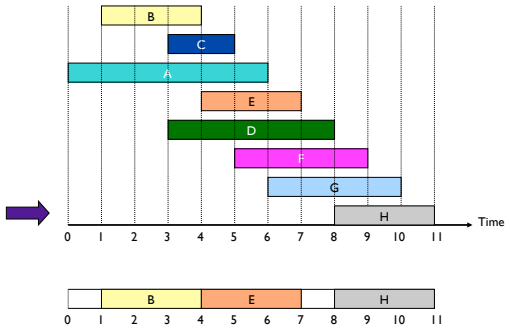

Interval Scheduling: Greedy Algorithm

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

jobs selected

```
Sort jobs by finish times so that f1 ≤ f2 ≤ ... ≤ fn
G = {}
for j = 1 to n
   if job j compatible with G
      G = G ∪ {j}
return G
```

Runtime of algorithm?
- Where/what are the costs?

Interval Scheduling: Greedy Algorithm

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

O(n logn)

jobs selected

```
Sort jobs by finish times so that f1 ≤ f2 ≤ ... ≤ fn
G = {}
for j = 1 to n
   if job j compatible with G   O(1)   } O(n)
      G = G ∪ {j}
return G
```

- Implementation. O(n log n)
  - Remember job j* that was added last to A
  - Job j is compatible with A if $s_j \geq f_j^*$

Analyzing Interval Scheduling

- Know that the intervals are compatible
  - Handled by the if statement
- But is it optimal?
  - What does it mean to be optimal?
  - Recall our goal for maximization

Greedy Stays Ahead Proofs

1. Define your solutions
   - Describe the form of your greedy solution (**A**) and of some other solution (possibly the optimal solution, **O**)
2. Find a measure
   - Find a measure by which greedy ***stays ahead*** of the optimal solution
     - Ex: Let $a_1, \ldots, a_k$ be the first k measures of greedy algorithm and $o_1, \ldots, o_m$ be the first m measures of other solution (sometimes $m = k$)
3. Prove greedy stays ahead
   - Show that greedy's partial solutions constructed are always just as good as the optimal solution's initial segments based on the measure
     - Ex: for all indices $r \leq \min(k,m)$, prove by induction that $a_r \geq o_r$ or $a_r \leq o_r$
   - Use the greedy algorithm to help you argue the inductive step
4. Prove optimality
   - Prove that since greedy stays ahead of the other solution with respect to the measure, then the greedy solution is optimal

Interval Scheduling: Analysis

- Theorem. Greedy algorithm is optimal.
- Pf. (by contradiction)
  - Assume greedy is not optimal
  - Let $a_1, a_2, \ldots, a_k$ denote set of jobs selected by *greedy* (k jobs)
  - Let $o_1, o_2, \ldots, o_m$ denote set of jobs in *optimal* solution (m jobs)
  - Both sets ordered by finish time for comparison ordering
  - ➡ Want to show that $k = m$

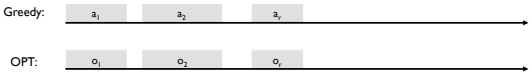

What can we say about a_1 and o_1? $f(a_1) \leq f(o_1)$

Interval Scheduling: Analysis

- Theorem. Greedy algorithm is optimal.
- Pf. (by contradiction)
  - Since we picked the first job to have the first finishing time, we know that $f(a_1) <= f(o_1)$
  - Want to show that Greedy "stays ahead"
  - Each interval finishes at least as soon as Optimal's
  - **Induction hypothesis**: for all indices r <= k, $f(a_r) <= f(o_r)$

Prove for r+1

Interval Scheduling: Analysis

- Theorem. Greedy algorithm is optimal.
- Pf. (by contradiction)
  - Since we picked the first job to have the first finishing time, we know that $f(a_1) <= f(o_1)$
  - Want to show that Greedy "stays ahead"
  - Each interval finishes at least as soon as Optimal's
  - **Induction hypothesis**: for all indices r <= k, $f(a_r) <= f(o_r)$

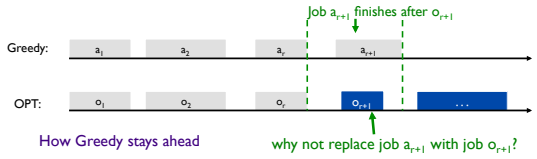

Interval Scheduling: Analysis

- Theorem. Greedy algorithm is optimal.
- Pf. (by contradiction)
  - Assume Greedy is not optimal (i.e., m > k)
    - Optimal solution has more jobs than Greedy
  - We already showed that for all indices $r \le k$, $f(a_r) \le f(o_r)$
  - Since m > k, there is a request o_{k+1} in Optimal

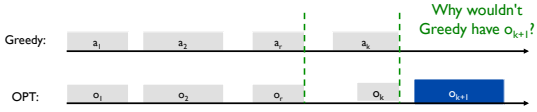

Interval Scheduling: Analysis

- Theorem. Greedy algorithm is optimal.
- Pf. (by contradiction)
  - Assume Greedy is not optimal (i.e., m > k)
  - We already showed that for all indices $r \le k$, $f(i_r) \le f(j_r)$
  - Since m > k, there is a request o_{k+1} in Optimal
    - Starts after o_k ends → after a_k ends
  - So, Greedy could also add o_k
    - Contradiction because now Greedy has another job

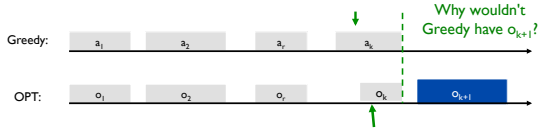

Greedy Algorithm Pseudo-Code

In some specified order

```
Set Greedy (Set candidate){
   solution = new Set( );
   while candidate.isNotEmpty()
      next = candidate.select() //use selection criteria,
      //remove from candidate and return value
      if solution.isFeasible(next) //constraints satisfied
            solution.union(next)
      if solution.solves()
            return solution

   //No more candidates and no solution
   return null
}
```


Problem Assumptions

- All requests were known to scheduling algorithm
  - Online algorithms: make decisions without knowledge of future input
- Each job was worth the same amount
  - What if jobs had *different* values?
    - E.g., scaled with size
- Single resource requested
  - Rejected requests that didn't fit

Assignments

- Exam 1
  - Can use book, lecture notes, your notes
  - Covers chapters 1-3
  - No "outside" resources
  - Limited access to me
  - Consider typing up answers
  - Due Friday at 5 p.m.
- No journal for Tuesday