

Greedy Scheduling

<https://cs.pomona.edu/classes/cs140/>

Outline

Topics and Learning Objectives

- Introduce greedy algorithms
- ★ • Discuss the greedy scheduling algorithm
- ★ • Discuss exchange argument proofs

Exercise

- Greedy scheduling

Extra Resources

- Introduction to Algorithms, 3rd, chapter 16
- Algorithms Illuminated Part 3: Chapter 13
- Lots of examples: <https://www.geeksforgeeks.org/greedy-algorithms/>

Greedy Algorithms

- Iteratively make myopic (short-sighted) decisions and hope it works
- Never go back and recheck/reevaluate that you were correct

Contrasting with Divide and Conquer

- It is generally easier to create greedy algorithms (good and bad to this)
- It is typically easier to analyze greedy algorithms (e.g., no master theorem)
- It is often harder to prove/understand the correctness of greedy algorithms
- It is common for greedy algorithms to be **incorrect** (/ sub-optimal)


Greedy Algorithms

Proofs of correctness

- It can sometimes feel like more of an art than a science
1. Proof by induction on the greedy decision
 2. Proof by induction on an **exchange argument**
 1. Either by contraction
 2. Or by exchanging with the optimal solution
 3. Whatever works...

Example of a greedy algorithm

- We've seen one greedy algorithm before. What was it?



- What path length does Dijkstra's output for $S \rightarrow W$?
- What is the correct shortest path length for $S \rightarrow W$?

Scheduling (ignoring concurrency)

You have a shared resource

For example, a processor (or professor)

You have many jobs that need to use the resource (students at office hours)

Each job j has:


- A Priority P_j that stands for the job's importance
- A Duration D_j that stands for the length of time to run the job

In what sequence should we complete the jobs?

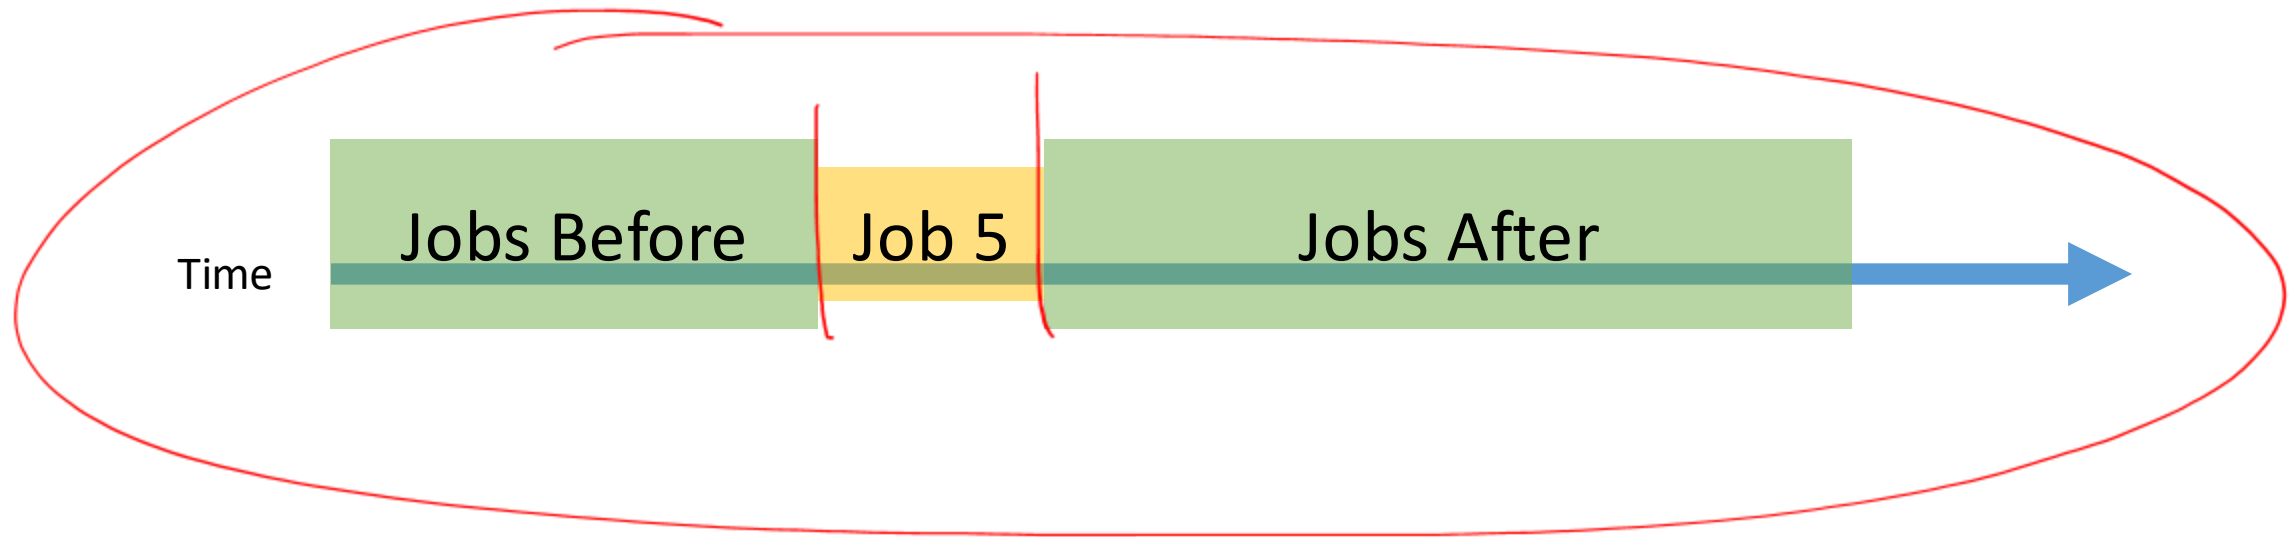
Scheduling (ignoring concurrency)

In what sequence should we complete the jobs?

- What is our criteria? What do we want to optimize?
- Let's start by looking at a **derived property** job **j**'s **completion time C_j**
- Given three jobs: $D_1 = 1$, $D_2 = 2$, $D_3 = 3$
- What is the completion time for each if they are scheduled in order?



What is the completion time of Job 5?



$$C_S = D_S + D_B$$

Scheduling


→ 10
15 ★

Optimization objective: *minimize* the *weighted* sum of completion times

$$S_{\text{cost}} = \min \left[\sum_{j=1}^n P_j C_j \right]$$

What is the weighted sum of completion times if we schedule the following jobs *in order*?

Job	J ₁	J ₂	J ₃
Duration	D ₁ = 1	D ₂ = 2	D ₃ = 3
Priority	P ₁ = 3	P ₂ = 2	P ₃ = 1
Compl-T	1	3	6



Job₃



Job₁




Job₂

Time



Exercise Question 1, 2, and 3



Scheduling

Calculate the weighted sum of completion times for the following jobs if they are scheduled in the order: 1, 2, 3.

Job	J_1	J_2	J_3
Duration	$D_1 = 1$	$D_2 = 2$	$D_3 = 3$
Priority	$P_1 = 3$	$P_2 = 2$	$P_3 = 1$
Completion	1	3	6
Weight	3	6	6

Weighted sum of completion times: ?

15

Greedy Scheduling

Our process for creating a **greedy** scheduling algorithm

1. Look at some special cases for our problem
2. Describe some possible greedy criteria
3. Compare our greedy criteria
4. Select the “best” one
5. Prove correctness if possible

Greedy Scheduling

Goal: devise a **greedy algorithm** to **minimize** the **weighted** sum of completion times

Why are we approaching this problem with a greedy algorithm?

- It's a pretty easy way to start.
- Compare the approach we go through in these slides with a **Divide and Conquer** approach

1. What are some special cases to consider?

Consider two jobs with equal durations (D)


- These jobs have different priorities (P_H and P_L)
- Do we schedule the lower or higher priority job first?



1. What are some special cases to consider?

Consider two jobs with equal durations (D)

- These jobs have different priorities (P_H and P_L)
- Do we schedule the **lower** or **higher** priority job first?




$$S_L = P_L \cdot D + P_H(D + D)$$

1. What are some special cases to consider?

Consider two jobs with equal durations (D)

- These jobs have different priorities (P_H and P_L)
- **Do we schedule the lower or higher priority job first?**



Schedule with Lower Priority First

$$S_L = P_L D + P_H \underbrace{(D+D)}_{C_H}$$

C_L C_H

S_L

$?$
 0

S_H

$$\cancel{P_L D} + 2 \cancel{P_H D}$$

$-P_L$ $-P_H$

P_H

$?$
 0

$$\cancel{P_H D} + 2 \cancel{P_L D}$$

$-P_H$ $-P_L$

$>$

P_L

Schedule with Higher Priority First

$$S_H = P_H D + P_L \underbrace{(D+D)}_{C_L}$$

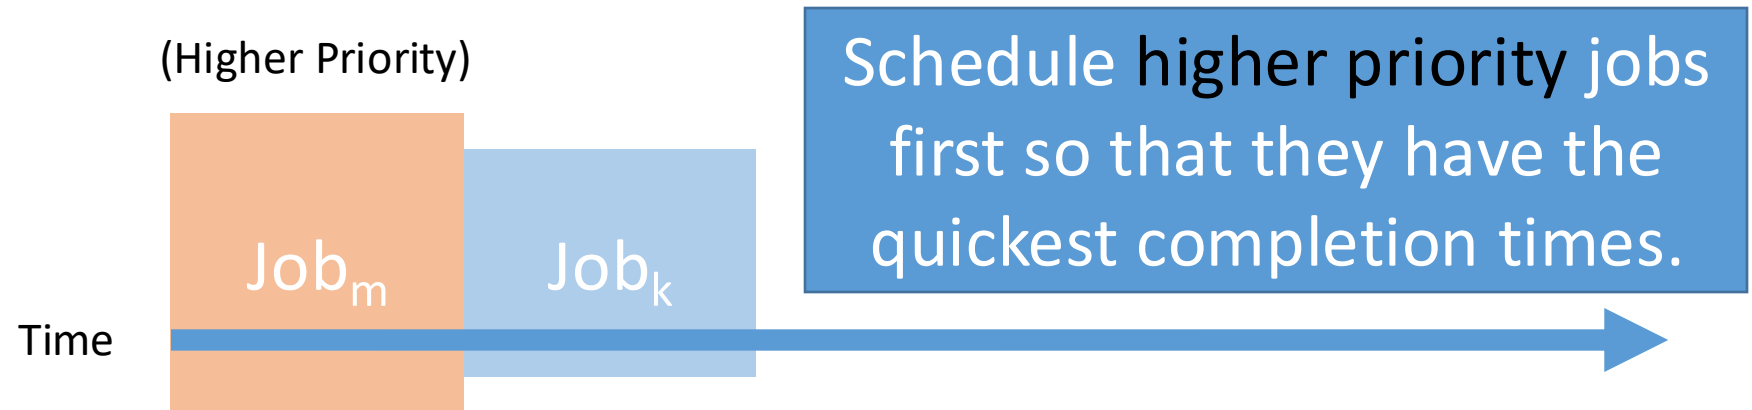
C_H C_L



1. What are some special cases to consider?

Consider two jobs with equal durations (D)

- These jobs have different priorities (P_H and P_L)
- **Do we schedule the lower or higher priority job first?**



1. What are some special cases to consider?

Consider two jobs with equal priorities (P)


- These jobs have different durations (D_E and D_S)
- Do we schedule the shorter or longer (Extended) job first?



1. What are some special cases to consider?

Consider two jobs with equal priorities (P)

- These jobs have different durations (D_E and D_S)
- Do we schedule the **shorter** or **longer** (**E**xtended) job first?



1. What are some special cases to consider?

Consider two jobs with equal priorities (P)

- These jobs have different durations (D_E and D_S)
- Do we schedule the **shorter** or **longer** (**E**xtended) job first?



Schedule with Shorter Job First

Schedule with Longer Job First

S_S

?


S_E

$D_S < D_E$

1. What are some special cases to consider?

Consider two jobs with equal priorities (P)

- These jobs have different durations (D_E and D_S)
- Do we schedule the **shorter** or **longer** (**E**xtended) job first?



$$f(P, D) = \langle \text{score} \rangle = P - D \quad \left\} \quad \frac{P}{D}$$

2. Describe some possible greedy criteria

What do we do when in the more general case:

1. Schedule highest priority first
2. Schedule shortest duration first

$P_i > P_j$ and $D_i > D_j$ (job i has higher priority and longer duration)

What are some simple **scoring functions** that *aggregate* our criteria?

We want a function for which jobs with a bigger score are scheduled first:

- Score increases for higher priorities
- Score increases for shorter times

1. Greedy Criterion 1: $P_i - D_i$ (take the difference)

2. Greedy Criterion 2: P_i / D_i (take the ratio)

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest value

Job with same duration	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=2, D=1		
Job 2: P=5, D=1		

Which job should be scheduled first?

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest value

Highest priority	Job with same duration	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
	Job 1: P=2, D=1	1	2
	Job 2: P=5, D=1	4	5
	Total weighted sum		

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest value

Highest priority	Job with same duration	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)	Same Result
	Job 1: P=2, D=1	1	2	
	Job 2: P=5, D=1	4	5	
	Total weighted sum	$5*1 + 2*2 = 9$	$5*1 + 2*2 = 9$	

Job with same priority	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=1, D=3		
Job 2: P=1, D=4		
Total weighted sum		

Which job should be scheduled first?

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest value

Highest priority	Job with same duration	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)	Same Result
	Job 1: P=2, D=1	1	2	
	Job 2: P=5, D=1	4	5	
	Total weighted sum	$5*1 + 2*2 = 9$	$5*1 + 2*2 = 9$	

Shortest time	Job with same priority	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
	Job 1: P=1, D=3	-2	1/3
	Job 2: P=1, D=4	-3	1/4
	Total weighted sum		

Which job should be scheduled first?

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest value

Highest priority	Job with same duration	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)	Same Result
	Job 1: P=2, D=1	1	2	
	Job 2: P=5, D=1	4	5	
	Total weighted sum	$5*1 + 2*2 = 9$	$5*1 + 2*2 = 9$	

Shortest time	Job with same priority	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)	Same Result
	Job 1: P=1, D=3	-2	1/3	
	Job 2: P=1, D=4	-3	1/4	
	Total weighted sum	$1*3 + 1*7 = 10$	$1*3 + 1*7 = 10$	

Which job should be scheduled first?

3. Compare our greedy criteria

- Let's try to get them to disagree.
 - Why does it matter if they don't produce the same result?
 - One scoring metric must be better than the other
-
- Apply the two greedy algorithms and calculate their weighted sum of completion times

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest metric value

Job	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=3, D=5		
Job 2: P=1, D=2		
Total weighted sum		

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest metric value

Job	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=3, D=5	-2	3/5
Job 2: P=1, D=2	-1	1/2
Total weighted sum		

Which job goes first?

3. Compare our greedy criteria

- Jobs will be ordered from biggest to smallest metric value

Job	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=3, D=5	-2	3/5
Job 2: P=1, D=2	-1	1/2
Total weighted sum		

Which job goes first?

What is the priority sum?

4. Select the “best” one

- Jobs will be ordered from biggest to smallest metric value

Job	Difference Metric ($P_i - D_i$)	Ratio Metric (P_i/D_i)
Job 1: P=3, D=5	-2	3/5
Job 2: P=1, D=2	-1	1/2
Total weighted sum	$1*2 + 3*7 = 23$	$3*5 + 1*7 = 22$

Which job goes first?

What is the priority sum?

Which criteria is better?

5. Prove correctness if possible

ratio
Is criteria 2 optimal?

- We don't know yet.

Criteria 2
Proved to be
not optimal

Claim: Criteria 2 is optimal for minimizing the weighted sum of completion times.

- We're going to prove this using an exchange argument!

Exchange Arguments

- Consider your greedy solution, G
- Consider an alternative solution, A
 - A can be anything that is not G
 - Create A by changing G in some way
- Compare these solutions
 - Show that turning A into G makes A get better

Proof

- Assume that we have no ties (all P_i/D_i are distinct numbers)
- Fix an arbitrary input with n jobs
- Let's perform a proof using an exchange argument **contradiction**

Let G = the greedy schedule and A = the (alternative) optimal schedule

- Let's **assume** that A must be better than G (**assume greedy is not optimal**)
- To perform the contradiction, we must show that G is better than A , thus contradicting the purported optimality of A

$$R = \frac{P}{D}$$

Proof

Let **G** = the greedy schedule and **A** = the optimal schedule

- Assume that: $P_1/D_1 > P_2/D_2 > \dots > P_n/D_n$
- We can just rename all jobs after we calculate their scores...
- Thus, **G** is just job 1 followed by job 2 etc. (1, 2, ..., n)

Job ID	Priority	Duration	Ratio	Reorder
1	1	4	0.3	4
2	8	6	1.3	3
3	6	1	6.0	1
4	1	5	0.2	5
5	1	9	0.1	6
6	7	3	2.3	2

Proof

Let **G** = the greedy schedule and **A** = the optimal schedule


- Assume that: $P_1/D_1 > P_2/D_2 > \dots > P_n/D_n$
- We can just rename all jobs after we calculate their scores...
- Thus, **G** is just job 1 followed by job 2 etc. (1, 2, ..., n)
- For **A** there must be at least two jobs that are “out of order”
 - Specifically, jobs i and j where i is scheduled after j , but $S_i > S_j$ (for example, Job₅ after Job₆)
- The greedy schedule is the **only schedule** where the jobs are in order

G VS A

(jobs i and j where i is scheduled after j, but $P_i/D_i > P_j/D_j$)

Job i has a larger greedy score

A Schedule



exchange

G Schedule




Ordered based on greedy scores

G VS A


(jobs i and j where i is scheduled after j, but $P_i/D_i > P_j/D_j$)

Job i has a larger greedy score

A Schedule



G Schedule



How does the exchange affect the completion time for:

1. Jobs other than i and j? *Ungr*

2. Job i *down*


3. Job j *up*

G VS A


(jobs i and j where i is scheduled after j, but $P_i/D_i > P_j/D_j$)

Job i has a larger greedy score


A Schedule



G Schedule



What is the weighted sum of completion times for each schedule?




$$\text{Cost}(\text{A}) = \text{Cost}(\text{Before}) + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost}(\text{After})$$

$$\text{Cost}(\text{G}) = \text{Cost}(\text{Before}) + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost}(\text{After})$$

$$\text{Cost}(\text{A}) < \text{Cost}(\text{G}) \quad \boxed{\text{Implied by optimality of A}}$$

$$\begin{aligned} & \text{Cost}(\text{Before}) + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost}(\text{After}) \\ & < \text{Cost}(\text{Before}) + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost}(\text{After}) \end{aligned}$$



$$\text{Cost}(\text{A}) = \text{Cost}(\text{Before}) + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost}(\text{After})$$

$$\text{Cost}(\text{G}) = \text{Cost}(\text{Before}) + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost}(\text{After})$$


$$\text{Cost}(\text{A}) < \text{Cost}(\text{G}) \quad \boxed{\text{Implied by optimality of A}}$$

$$\begin{aligned} &\text{Cost}(\text{Before}) + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost}(\text{After}) \\ &< \text{Cost}(\text{Before}) + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost}(\text{After}) \end{aligned}$$

$$\begin{aligned} &P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) \\ &< P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) \end{aligned}$$

$$\begin{aligned} &P_j * T_b + P_j * D_j + P_i * T_b + P_i * D_j + P_i * D_i \\ &< P_i * T_b + P_i * D_i + P_j * T_b + P_j * D_i + P_j * D_j \end{aligned}$$

A Schedule



G Schedule



$$\text{Cost(A)} = \text{Cost(Before)} + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost(After)}$$

$$\text{Cost(G)} = \text{Cost(Before)} + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost(After)}$$

$$\text{Cost(A)} < \text{Cost(G)}$$

Implied by optimality of A

$$\begin{aligned} &\text{Cost(Before)} + P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) + \text{Cost(After)} \\ &< \text{Cost(Before)} + P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) + \text{Cost(After)} \end{aligned}$$

$$\begin{aligned} &P_j * (T_b + D_j) + P_i * (T_b + D_j + D_i) \\ &< P_i * (T_b + D_i) + P_j * (T_b + D_i + D_j) \end{aligned}$$

$$\begin{aligned} &P_j * T_b + P_j * D_j + P_i * T_b + P_i * D_j + P_i * D_i \\ &< P_i * T_b + P_i * D_i + P_j * T_b + P_j * D_i + P_j * D_j \end{aligned}$$

$$P_i * D_j < P_j * D_i$$

$$P_i / D_i < P_j / D_j$$

Contradiction to how they were ordered by our greedy criteria

G is always Better

$$i < j$$

$$\frac{P_i}{D_i} > \frac{P_j}{D_j}$$


$$\frac{P_i D_j}{D_i D_j} < \frac{P_j D_i}{D_i D_j}$$

contradiction

$$\frac{P_i}{D_i} < \frac{P_j}{D_j}$$


Multiple Re-orderings

A



Our proof doesn't account for this

Multiple Re-orderings



Example with Randomly Generated Jobs

					Greedy		Unoptimized	
Job ID	Priority	Duration	Ratio	Reorder	Time	Weighted	Time	Weighted
1	1	4	0.3	4	14	14	4	4
2	8	6	1.3	3	10	80	10	80
3	6	1	6.0	1	1	6	11	66
4	1	5	0.2	5	19	19	16	16
5	1	9	0.1	6	28	28	25	25
6	7	3	2.3	2	4	28	28	196
						175		387

low scoring
job

Summary of Greedy Scheduling

- Given n jobs, each with a **priority** and a **duration**
- Give each job a **score** based on their ratio of **priority** to **duration**
- Schedule jobs in decreasing order of their **score**
- This gives us an optimal schedule

W.r.t. minimizing the weighted sum
of completion times

- What do we do if we're given more jobs while these are running?
- Any issues with this scheme?
 - Some jobs might always be postponed.

Starvation