

Greedy Algorithms

T. M. Murali

February 11, 16, 2016

Algorithm Design


















- ▶ Start discussion of different ways of designing algorithms.
- ▶ Greedy algorithms, divide and conquer, dynamic programming.
- ▶ Discuss principles that can solve a variety of problem types.
- ▶ Design an algorithm, prove its correctness, analyse its complexity.






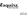











Algorithm Design

- ▶ Start discussion of different ways of designing algorithms.
- ▶ Greedy algorithms, divide and conquer, dynamic programming.
- ▶ Discuss principles that can solve a variety of problem types.
- ▶ Design an algorithm, prove its correctness, analyse its complexity.
- ▶ Greedy algorithms: make the current best choice.



Filters	9:00	9:30	10:00	10:30	11:00	11:30	12:00	12:30	1:00	1:30
XFL	XFLITY On Demand						XFLITY On Demand			
MeTV	Happy Days	Laverne & Shirley	Hogan's Heroes	Hogan's Heroes	Carol Burnett and Friends	Perry Mason	The Twilight Zone		The Rockford Files	
CBS	Monk	The Big Bang Theory	Elementary	Eyewitness News at 11	The Late Show With Stephen Colbert	The Late Late Show With James Corden	Comics Unleashed...			
USA	Gordon Chevrolet		Gordon Chevrolet	South Jersey News	J&D Furniture	UA Auto Sales	Single Puppy Training	We Say Yes	10 Ways to Help Arthritis Pain	
U-2	WWE SmackDown		Colony	Law & Order: Special Victims Unit	Law & Order: Special Victims Unit	Colony				
abc	Scandal	How to Get Away With Murder	Action News 11pm	Jimmy Kimmel Live	Nightline	Who Wants to Be a Millionaire	Action News 11pm			
ESPN	College Basketball: Iowa at Indiana			SportsCenter	SportsCenter	SportsCenter				
FCB	Fight Sports MMA	In Depth With Graham Bensell...	Orange Line	The 700Level Show	NHL Hockey: Sabres at Flyers	Paid Programming	Paid Programming			
FOX	American Idol	Fox 29 News at 10	TMC	Dish Nation	Chasing News	The Simpsons	Paid Programming	Paid Programming		
NBC	The Blacklist	Shades of Blue	NBC30 News at 11pm	The Tonight Show Starring Jimmy Fallon	Late Night With Seth Meyers	Last Call With Carson...				
Q	Shoe Shopping With Jane			France Santo - Postwar	Josie Maran Argan Oil Cosmetics	Inspired Style				
WPS	PBS NewsHour			Mercy Street	Charlie Rose	On the Psychiatrist's Couch With Dr. Daniel Amen, MD				
CW	The 100	Eyewitness News on The CW...	2 Broke Girls	Mike & Molly	The King of Queens	Mike & Molly	2 Broke Girls	Family Guy	Seinfeld	
ABC	The List with Colleen Lopez		The List with Colleen Lopez	Serena Williams Signature Statement Collection		Home Solutions Featuring Nellie's		Home Solutions Featuring Nellie's		
BET	Criminals at Work	Kevin Hart: Seriously Funny	Real Husbands of Hollywood	Real Husbands of Hollywood	The Wendy Williams Show		The Real			
ESPN2	College Basketball: Oregon at California			College Basketball Live	NBA Tonight	NFL Live	NBA Tonight	College Basketball Live		
ABC	The Mentalist	Action News at 10 on PBS17	Modern Family	Friends	Friends	Raising Hope	The Mentalist			

	9 am	9:30	10 am	10:30	11 am	11:30	12 am	12:30	1 am	1:30
 Netflix	Grown Up 2 (2015) 62:25	Baskets	Baskets	Baskets	Grown Up 2 (2015) 62:25	Baskets				
 Nick	Zookeeper (2011) 62:25	Full House	Full House	Friends	Friends	Friends	Friends	The Fresh Prince of Bel-Air	The Fresh Prince	
 Disney	Teen Beach 2 (2015) 62:25	... Liv and Maddie	Austin & Ally	Bunk'd	Get Ready World	Jessie	Jessie	Austin & Ally	Austin & Ally	
 Hulu	The Notebook (2004) 62:25			The 700 Club	Just Married (2005) 62:13					
 HBO	The 47th NAACP Image Awards	Trading Places (1983) 92						The 47th NAACP Image Awards		
 ABC	Love on the Suburbs (2016) 62:25	The Middle	The Middle	The Golden Girls	The Golden Girls	The Golden Girls	The Golden Girls	Frasier	Frasier	
 Fox	The Mechanic (2011) 62:25	The Mechanic (2011) 62:25			Team Ninja Warrior			Team Ninja Warrior		
 Syfy	Blade: Trinity (2004) 62:25			Snowyfl				Snowyfl	Fright Night (2011) 62:25	
 VH1	Hitch (2005) 62:25			Hitch (2005) 62:25						
 Eurosport	Initiation of Life (1994) 92	Gold Diggers of 1933 (1933) 92			Top Hat (1935) 92				The Moon and Sixpence (1942)	
 Pop	13 Going on 30 (2004) 62:25				Waiting to Exhale (1995) 62:25					
 Wild	Mission Critical	Behind Russia's Frozen Curtain		Monster Croc Hunt (2017) 92	Mission Critical			Behind Russia's Frozen Curtain		
 BBC	Black Hawk Down (2001) 62:25	London Spy		Black Hawk Down (2001) 62:25						
 Fox	The Mechanic (2011) 62:25	The Mechanic (2011) 62:25			Team Ninja Warrior			Team Ninja Warrior		
 LMA	Roseo Kiler: The Chris Purse Story (2015) 62:25	Beyond the Headlines		Beyond the Headlines	Roseo Kiler: The Chris Purse Story (2015) 62:25					
 Crypsis	Sex and the... Ghost (1996) 62:25				Keeping the Faith (2000) 62:25					
 Family	Who Framed Roger Rabbit (1988) 62:25			Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing	

	9:00	9:30	10:00	10:30	11:00	11:30	12:00	12:30	1:00	1:30
 ETN	Grown Ups 2 (2015) 62:25		Baskets	Baskets	Baskets	Grown Ups 2 (2015) 62:25			Baskets	
 nick	Zookeeper (2011) 62:25		Full House	Full House	Friends	Friends	Friends	Friends	The Fresh Prince of Bel-Air	The Fresh Prince .
 Disney	Teen Beach 2 (2015) 62:25		Liv and Maddie	Austin & Ally	Bunk'd	Get Movies World	Jessie	Jessie	Austin & Ally	Austin & Ally
 netflix	The Notebook (2004) 62:25				The '70s Club		Just Married (2005) 62:13			
 none	The 47th NAACP Image Awards		Trading Places (1983) 92					The 47th NAACP Image Awards		
 ABC	Love on the Suburbs (2016) 62:25		The Middle	The Middle	The Golden Girls	The Golden Girls	The Golden Girls	The Golden Girls	Frasier	Frasier
 Netflix	The Mechanic (2011) 62:25		The Mechanic (2011) 62:25			Team Ninja Warrior			Team Ninja Warrior	
 Syfy	Blade: Trinity (2004) 62:25				Snowfall		Snowfall		Fight Night (2011) 62:25	
 VH1	Hush (2005) 62:25				Hush (2005) 62:25					
 ABC	Initiation of Life (1994) 92		Gold Diggers of 1933 (1935) 92			Top Hat (1935) 92			The Moon and Sincere (1942)	
 Pop	13 Going on 30 (2004) 62:25				Waiting to Exhale (1995) 62:25					
 Wild	Mission Critical		Behind Russia's Frozen Curtain		Monster Croc Hunt (2012) 92		Mission Critical		Behind Russia's Frozen Curtain	
 Netflix	Black Hawk Down (2001) 62:25		London Spy		Black Hawk Down (2001) 62:25					
 Netflix	The Mechanic (2011) 62:25		The Mechanic (2011) 62:25			Team Ninja Warrior			Team Ninja Warrior	
 U2	Roxoco Killer: The Chris Posco Story (2015) 62:25		Beyond the Headlines			Beyond the Headlines		Roxoco Killer: The Chris Posco Story (2015) 62:25		
 grypen	Sex and the... Ghost (1996) 62:25					Keeping the Faith (2000) 62:25				
 family	Who Framed Roger Rabbit (1988) 62:25				Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing	Extreme Couponing

- ▶ Input: Start and end time of each movie.
- ▶ Constraint: Only one TV \Rightarrow cannot watch two overlapping classes at the same time.
- ▶ Output: Compute the largest number of movies we can watch.

Interval Scheduling

INTERVAL SCHEDULING

INSTANCE: Nonempty set $\{(s(i), f(i)), 1 \leq i \leq n\}$ of start and finish times of n jobs.

SOLUTION: The largest subset of mutually compatible jobs.

- ▶ Two jobs are *compatible* if they do not overlap.
- ▶ This problem models the situation where you have a resource, a set of fixed jobs, and you want to schedule as many jobs as possible.
- ▶ For any input set of jobs, algorithm must provably compute the **largest** set of compatible jobs.

Template for Greedy Algorithm

- ▶ Process jobs in some order. Add next job to the result if it is compatible with the jobs already in the result.
- ▶ Key question: in what order should we process the jobs?

Template for Greedy Algorithm

- ▶ Process jobs in some order. Add next job to the result if it is compatible with the jobs already in the result.
- ▶ Key question: in what order should we process the jobs?
 - Earliest start time Increasing order of start time $s(i)$.
 - Earliest finish time Increasing order of finish time $f(i)$.
 - Shortest interval Increasing order of length $f(i) - s(i)$.
 - Fewest conflicts Increasing order of the number of conflicting jobs. How fast can you compute the number of conflicting jobs for each job?

Greedy Ideas that Do Not Work




Figure 4.1 Some instances of the Interval Scheduling Problem on which natural greedy algorithms fail to find the optimal solution. In (a), it does not work to select the interval that starts earliest; in (b), it does not work to select the shortest interval; and in (c), it does not work to select the interval with the fewest conflicts.

Interval Scheduling Algorithm: Earliest Finish Time

- Schedule jobs in order of earliest finish time (EFT).

Initially let R be the set of all requests, and let A be empty

While R is not yet empty

 Choose a request $i \in R$ that has the smallest finishing time

 Add request i to A

 Delete all requests from R that are not compatible with request i

EndWhile

Return the set A as the set of accepted requests

Interval Scheduling Algorithm: Earliest Finish Time

- Schedule jobs in order of earliest finish time (EFT).

Initially let R be the set of all requests, and let A be empty

While R is not yet empty

 Choose a request $i \in R$ that has the smallest finishing time

 Add request i to A

 Delete all requests from R that are not compatible with request i

EndWhile

Return the set A as the set of accepted requests

- Claim: A is a compatible set of jobs.

Interval Scheduling Algorithm: Earliest Finish Time

- Schedule jobs in order of earliest finish time (EFT).

Initially let R be the set of all requests, and let A be empty

While R is not yet empty

 Choose a request $i \in R$ that has the smallest finishing time

 Add request i to A

 Delete all requests from R that are not compatible with request i

EndWhile

Return the set A as the set of accepted requests

- Claim: A is a compatible set of jobs. Proof follows by construction, i.e., the algorithm computes a compatible set of jobs.

Ideas for Analysing the EFT Algorithm

- ▶ We need to prove that $|A|$ (the number of jobs in A) is the largest possible in *any* set of mutually compatible jobs.

Ideas for Analysing the EFT Algorithm

- ▶ We need to prove that $|A|$ (the number of jobs in A) is the largest possible in *any* set of mutually compatible jobs.
- ▶ Proof idea 1: algorithm makes the best choice at each step, so it must choose the largest number of mutually compatible jobs.

Ideas for Analysing the EFT Algorithm

- ▶ We need to prove that $|A|$ (the number of jobs in A) is the largest possible in *any* set of mutually compatible jobs.
- ▶ Proof idea 1: algorithm makes the best choice at each step, so it must choose the largest number of mutually compatible jobs.
 - ▶ What does “best” mean?
 - ▶ This idea is too generic. It can be applied even to algorithms that we know do not work correctly.

Ideas for Analysing the EFT Algorithm

- ▶ We need to prove that $|A|$ (the number of jobs in A) is the largest possible in *any* set of mutually compatible jobs.
- ▶ Proof idea 1: algorithm makes the best choice at each step, so it must choose the largest number of mutually compatible jobs.
 - ▶ What does “best” mean?
 - ▶ This idea is too generic. It can be applied even to algorithms that we know do not work correctly.
- ▶ Proof idea 2: at each step, can we show algorithm has the “better” solution than any other answer?
 - ▶ What does “better” mean?
 - ▶ How do we measure progress of the algorithm?

Ideas for Analysing the EFT Algorithm

- ▶ We need to prove that $|A|$ (the number of jobs in A) is the largest possible in *any* set of mutually compatible jobs.
- ▶ Proof idea 1: algorithm makes the best choice at each step, so it must choose the largest number of mutually compatible jobs.
 - ▶ What does “best” mean?
 - ▶ This idea is too generic. It can be applied even to algorithms that we know do not work correctly.
- ▶ Proof idea 2: at each step, can we show algorithm has the “better” solution than any other answer?
 - ▶ What does “better” mean?
 - ▶ How do we measure progress of the algorithm?
- ▶ Basic idea of proof:
 - ▶ We can sort jobs in any solution in increasing order of their finishing time.
 - ▶ Finishing time of job number r selected by $A \leq$ finishing time of job number r selected by any other algorithm.

Analysing the EFT Algorithm

- ▶ Let O be an optimal set of jobs. We will show that $|A| = |O|$.
- ▶ Let i_1, i_2, \dots, i_k be the set of jobs in A in order.
- ▶ Let j_1, j_2, \dots, j_m be the set of jobs in O in order, $m \geq k$.
- ▶ Claim: For all indices $r \leq k$, $f(i_r) \leq f(j_r)$.

Analysing the EFT Algorithm

- ▶ Let O be an optimal set of jobs. We will show that $|A| = |O|$.
- ▶ Let i_1, i_2, \dots, i_k be the set of jobs in A in order.
- ▶ Let j_1, j_2, \dots, j_m be the set of jobs in O in order, $m \geq k$.
- ▶ Claim: For all indices $r \leq k$, $f(i_r) \leq f(j_r)$. Prove by induction on r .

Analysing the EFT Algorithm

- ▶ Let O be an optimal set of jobs. We will show that $|A| = |O|$.
- ▶ Let i_1, i_2, \dots, i_k be the set of jobs in A in order.
- ▶ Let j_1, j_2, \dots, j_m be the set of jobs in O in order, $m \geq k$.
- ▶ Claim: For all indices $r \leq k$, $f(i_r) \leq f(j_r)$. Prove by induction on r .




Figure 4.3 The inductive step in the proof that the greedy algorithm stays ahead.

Analysing the EFT Algorithm

- ▶ Let O be an optimal set of jobs. We will show that $|A| = |O|$.
- ▶ Let i_1, i_2, \dots, i_k be the set of jobs in A in order.
- ▶ Let j_1, j_2, \dots, j_m be the set of jobs in O in order, $m \geq k$.
- ▶ Claim: For all indices $r \leq k$, $f(i_r) \leq f(j_r)$. Prove by induction on r .




Figure 4.3 The inductive step in the proof that the greedy algorithm stays ahead.

- ▶ Claim: $m = k$.

Analysing the EFT Algorithm

- ▶ Let O be an optimal set of jobs. We will show that $|A| = |O|$.
- ▶ Let i_1, i_2, \dots, i_k be the set of jobs in A in order.
- ▶ Let j_1, j_2, \dots, j_m be the set of jobs in O in order, $m \geq k$.
- ▶ Claim: For all indices $r \leq k$, $f(i_r) \leq f(j_r)$. Prove by induction on r .




Figure 4.3 The inductive step in the proof that the greedy algorithm stays ahead.

- ▶ Claim: $m = k$.
- ▶ Claim: The greedy algorithm returns an optimal set A .

Implementing the EFT Algorithm

1. Reorder jobs so that they are in increasing order of finish time.
2. Store starting time of jobs in an array S .
3. $k = 1$.
4. While $k \leq |S|$,
 - 4.1 Output job k .
 - 4.2 Let finish time of job k be f .
 - 4.3 Iterate over S from index k onwards to find the first index i such that $S[i] \geq f$.
 - 4.4 $k = i$

Implementing the EFT Algorithm

1. Reorder jobs so that they are in increasing order of finish time.
 2. Store starting time of jobs in an array S .
 3. $k = 1$.
 4. While $k \leq |S|$,
 - 4.1 Output job k .
 - 4.2 Let finish time of job k be f .
 - 4.3 Iterate over S from index k onwards to find the first index i such that $S[i] \geq f$.
 - 4.4 $k = i$
- ▶ Must be careful to iterate over S such that we never scan same index more than once.
 - ▶ Running time is $O(n \log n)$, dominated by sorting. advance from the

12528	CS-3604	Professionalism in Computing	L	3	Full -5	54	DR Dunlap	T R	11:00AM	12:15PM	PAM 32	11T
Comments for CRN 12528: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
20110	CS-3634	Comp Sci Foundations for CMDA	L	3	Full -8	30	T Warburton	T R	5:00PM	6:15PM	SAUND 408	12T
Comments for CRN 20110: Not for CS major credit. To submit a force/add request, send e-mail to Ms. Nora Sullivan: nora84@vt.edu												
12529	CS-3654	Intro Data Analytics & Visual	L	3	2	60	CL North	M W	4:00PM	5:15PM	WLH 350	16M
Comments for CRN 12529: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
19527	CS-3704	Internet Software Des	L	3	33	75	FJ Servant Cortes	M W	2:30PM	3:45PM	WLH 350	14M
Comments for CRN 19527: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12530	CS-3714	Mobile Software Development	L	3	Full -12	65	DS McCrickard	T R	3:30PM	4:45PM	MCB 113	15T
Comments for CRN 12530: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12531	CS-3724	Human-Computr Interac	L	3	Full -24	50	SR Harrison	T R	12:30PM	1:45PM	SURGE 109	12T
Comments for CRN 12531: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12533	CS-3744	Intro GUI Programming/Graphics	L	3	Full -9	75	K Luther	T R	5:00PM	6:15PM	MCB 129	12T
Comments for CRN 12533: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12534	CS-4104	Data and Algorithm Analysis	L	3	Full -1	65	TM Murali	T R	2:00PM	3:15PM	GYM 124	14T
Comments for CRN 12534: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12535	CS-4114	Formal Languages	L	3	Full -16	50	L Zhang	T R	3:30PM	4:45PM	TORG 1060	15T
Comments for CRN 12535: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12537	CS-4284	Systems & Networking Capstone	L	3	5	29	AR Butt	M W	2:30PM	3:45PM	RAND 222	14M
Comments for CRN 12537: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
19512	CS-4504	Computer Organization	L	3	2	36	W Feng	T R	12:30PM	1:45PM	WHIT 257	12T
Comments for CRN 19512: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12540	CS-4604	Int Data Base Mgt Sys	L	3	7	60	BA Prakash	M W	4:00PM	5:15PM	WLH 340	16M
Comments for CRN 12540: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
12541	CS-4624	Multimedia/Hypertext	L	3	Full -11	45	EA Fox	T R	9:30AM	10:45AM	WLH 340	09T
Comments for CRN 12541: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey .												
20098	CS-4644	Creative Computing Studio	L	3	Full -10	20	A Kelliber	T R	11:00AM	12:15PM	MCB 238	11T

12528	CS-3604	Professionalism in Computing	L	3	Full -5	54	DR Dunlap	T R	11:00AM	12:15PM	PAM 32	11T
Comments for CRN 12528: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
20110	CS-3634	Comp Sci Foundations for CMDA	L	3	Full -8	30	T Warburton	T R	5:00PM	6:15PM	SAUND 408	12T
Comments for CRN 20110: Not for CS major credit. To submit a force/add request, send e-mail to Ms. Nora Sullivan: nora84@vt.edu												
12529	CS-3654	Intro Data Analytics & Visual	L	3	2	60	CL North	M W	4:00PM	5:15PM	WLH 350	16M
Comments for CRN 12529: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
19527	CS-3704	Internet Software Des	L	3	33	75	FJ Servant Cortes	M W	2:30PM	3:45PM	WLH 350	14M
Comments for CRN 19527: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12530	CS-3714	Mobile Software Development	L	3	Full -42	65	DS McCrickard	T R	3:30PM	4:45PM	MCB 113	15T
Comments for CRN 12530: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12531	CS-3724	Human-Computer Interact	L	3	Full -24	50	SR Harrison	T R	12:30PM	1:45PM	SURGE 109	12T
Comments for CRN 12531: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12533	CS-3744	Intro GUI Programming/Graphics	L	3	Full -9	75	K Luther	T R	5:00PM	6:15PM	MCB 129	12T
Comments for CRN 12533: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12534	CS-4104	Data and Algorithm Analysis	L	3	Full -1	65	TM Murali	T R	2:00PM	3:15PM	GYM 124	14T
Comments for CRN 12534: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12535	CS-4114	Formal Languages	L	3	Full -16	50	L Zhang	T R	3:30PM	4:45PM	TORG 1060	15T
Comments for CRN 12535: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12537	CS-4284	Systems & Networking Capstone	L	3	5	29	AR Butt	M W	2:30PM	3:45PM	RAND 222	14M
Comments for CRN 12537: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
19512	CS-4504	Computer Organization	L	3	2	36	W Feng	T R	12:30PM	1:45PM	WHIT 257	12T
Comments for CRN 19512: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12540	CS-4604	Int Data Base Mgr Sys	L	3	7	60	BA Prakash	M W	4:00PM	5:15PM	WLH 340	16M
Comments for CRN 12540: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
12541	CS-4624	Multimedia/Hypertext	L	3	Full -11	45	EA Fox	T R	9:30AM	10:45AM	WLH 340	09T
Comments for CRN 12541: Force/add subject to availability; for information, see www.cs.vt.edu/undergraduate/survey.												
20098	CS-4644	Creative Computing Studio	L	3	Full -10	20	A Kelliber	T R	11:00AM	12:15PM	MCB 238	11T

- ▶ Input: Start and end time of each class.
- ▶ Constraint: Cannot schedule two overlapping classes to the same room.
- ▶ Output: Assign each class to a room and use smallest number of rooms possible.

Interval Partitioning

INTERVAL PARTITIONING

INSTANCE: Set $\{(s(i), f(i)), 1 \leq i \leq n\}$ of start and finish times of n jobs.

SOLUTION: A partition of the jobs into k sets, where each set of jobs is mutually compatible, and k is minimised.

- ▶ This problem models the situation where you have a set of fixed jobs, and you want to schedule all jobs using as few resources as possible.

Depth of Intervals




Figure 4.4 (a) An instance of the Interval Partitioning Problem with ten intervals (*a* through *j*). (b) A solution in which all intervals are scheduled using three resources: each row represents a set of intervals that can all be scheduled on a single resource.

- ▶ The *depth* of a set of intervals is the maximum number of intervals that contain any time point.

Depth of Intervals




Figure 4.4 (a) An instance of the Interval Partitioning Problem with ten intervals (*a* through *j*). (b) A solution in which all intervals are scheduled using three resources: each row represents a set of intervals that can all be scheduled on a single resource.

- ▶ The *depth* of a set of intervals is the maximum number of intervals that contain any time point.
- ▶ Claim: In any instance of INTERVAL PARTITIONING, $k \geq \text{depth}$.

Depth of Intervals




Figure 4.4 (a) An instance of the Interval Partitioning Problem with ten intervals (*a* through *j*). (b) A solution in which all intervals are scheduled using three resources: each row represents a set of intervals that can all be scheduled on a single resource.

- ▶ The *depth* of a set of intervals is the maximum number of intervals that contain any time point.
- ▶ Claim: In any instance of INTERVAL PARTITIONING, $k \geq \text{depth}$.
- ▶ Is it possible to compute the depth efficiently? Is $k = \text{depth}$?

Computing the Depth of the Intervals

- ▶ How efficiently can we compute the depth of a set of intervals?

Computing the Depth of the Intervals

- ▶ How efficiently can we compute the depth of a set of intervals?
- 1. Sort the start times and finish times of the jobs into a single list L .
- 2. $d \leftarrow 0$.
- 3. For i ranging from 1 to $2n$
 - 3.1 If L_i is a start time, increment d by 1.
 - 3.2 If L_i is a finish time, decrement d by 1.
- 4. Return the largest value of d computed in the loop.

Computing the Depth of the Intervals

- ▶ How efficiently can we compute the depth of a set of intervals?
- 1. Sort the start times and finish times of the jobs into a single list L .
- 2. $d \leftarrow 0$.
- 3. For i ranging from 1 to $2n$
 - 3.1 If L_i is a start time, increment d by 1.
 - 3.2 If L_i is a finish time, decrement d by 1.
- 4. Return the largest value of d computed in the loop.
- ▶ Algorithm runs in $O(n \log n)$ time.

Interval Partitioning Algorithm

- ▶ Compute the depth d of the intervals.

Interval Partitioning Algorithm

- Compute the depth d of the intervals.

Sort the intervals by their start times, breaking ties arbitrarily

Let I_1, I_2, \dots, I_n denote the intervals in this order

For $j = 1, 2, 3, \dots, n$

 For each interval I_i that precedes I_j in sorted order and overlaps it

 Exclude the label of I_i from consideration for I_j

 Endfor

 If there is any label from $\{1, 2, \dots, d\}$ that has not been excluded then

 Assign a nonexcluded label to I_j

 Else

 Leave I_j unlabeled

 Endif

Endfor

- Claim: Every interval gets a label and no pair of overlapping intervals get the same label.

Interval Partitioning Algorithm

- Compute the depth d of the intervals.

Sort the intervals by their start times, breaking ties arbitrarily

Let I_1, I_2, \dots, I_n denote the intervals in this order

For $j = 1, 2, 3, \dots, n$

 For each interval I_i that precedes I_j in sorted order and overlaps it

 Exclude the label of I_i from consideration for I_j

 Endfor

 If there is any label from $\{1, 2, \dots, d\}$ that has not been excluded then

 Assign a nonexcluded label to I_j

 Else

 Leave I_j unlabeled

 Endif

Endfor

- Claim: Every interval gets a label and no pair of overlapping intervals get the same label.
- Claim: The greedy algorithm is optimal.

Interval Partitioning Algorithm

- Compute the depth d of the intervals.

Sort the intervals by their start times, breaking ties arbitrarily

Let I_1, I_2, \dots, I_n denote the intervals in this order

For $j = 1, 2, 3, \dots, n$

 For each interval I_i that precedes I_j in sorted order and overlaps it

 Exclude the label of I_i from consideration for I_j

 Endfor

 If there is any label from $\{1, 2, \dots, d\}$ that has not been excluded then

 Assign a nonexcluded label to I_j

 Else

 Leave I_j unlabeled

 Endif

Endfor

- Claim: Every interval gets a label and no pair of overlapping intervals get the same label.
- Claim: The greedy algorithm is optimal.
- The running time of the algorithm is $O(n \log n)$.

Interval Partitioning Algorithm

- Compute the depth d of the intervals.

Sort the intervals by their start times, breaking ties arbitrarily

Let I_1, I_2, \dots, I_n denote the intervals in this order

For $j = 1, 2, 3, \dots, n$

 For each interval I_i that precedes I_j in sorted order and overlaps it

 Exclude the label of I_i from consideration for I_j

 Endfor

 If there is any label from $\{1, 2, \dots, d\}$ that has not been excluded then

 Assign a nonexcluded label to I_j

 Else

 Leave I_j unlabeled

 Endif

Endfor

- Claim: Every interval gets a label and no pair of overlapping intervals get the same label.
- Claim: The greedy algorithm is optimal.
- The running time of the algorithm is $O(n \log n)$. Can modify algorithm for computing depth to maintain set of available labels and to assign them efficiently.

Scheduling to Minimise Lateness

- ▶ Study different model: job i has a length $t(i)$ and a deadline $d(i)$.
- ▶ We want to schedule all n jobs on one resource.
- ▶ Our goal is to assign a starting time $s(i)$ to each job such that each job is delayed as little as possible.

Scheduling to Minimise Lateness

- ▶ Study different model: job i has a length $t(i)$ and a deadline $d(i)$.
- ▶ We want to schedule all n jobs on one resource.
- ▶ Our goal is to assign a starting time $s(i)$ to each job such that each job is delayed as little as possible.
- ▶ A job i is *delayed* if $f(i) > d(i)$; the *lateness of the job* is

$$\max(0, f(i) - d(i)).$$

- ▶ The *lateness of a schedule* is

$$\max_{1 \leq i \leq n} (\max(0, f(i) - d(i))).$$

Scheduling to Minimise Lateness

- ▶ Study different model: job i has a length $t(i)$ and a deadline $d(i)$.
- ▶ We want to schedule all n jobs on one resource.
- ▶ Our goal is to assign a starting time $s(i)$ to each job such that each job is delayed as little as possible.
- ▶ A job i is *delayed* if $f(i) > d(i)$; the *lateness of the job* is

$$\max(0, f(i) - d(i)).$$

- ▶ The *lateness of a schedule* is

$$\max_{1 \leq i \leq n} (\max(0, f(i) - d(i))).$$

MINIMISE LATENESS

INSTANCE: Set $\{(t(i), d(i)), 1 \leq i \leq n\}$ of lengths and deadlines of n jobs.

SOLUTION: Set $\{s(i), 1 \leq i \leq n\}$ of start times such that $\max_{1 \leq i \leq n} (\max(0, s(i) + t(i) - d(i)))$ is as small as possible.

Scheduling to Minimise Lateness

MINIMISE LATENESS

INSTANCE: Set $\{(t(i), d(i)), 1 \leq i \leq n\}$ of lengths and deadlines of n jobs.

SOLUTION: Set $\{s(i), 1 \leq i \leq n\}$ of start times such that $\max_{1 \leq i \leq n} (\max(0, s(i) + t(i) - d(i)))$ is as small as possible.


Scheduling to Minimise Lateness

MINIMISE LATENESS

INSTANCE: Set $\{(t(i), d(i)), 1 \leq i \leq n\}$ of lengths and deadlines of n jobs.

SOLUTION: Set $\{s(i), 1 \leq i \leq n\}$ of start times such that $\max_{1 \leq i \leq n} (\max(0, s(i) + t(i) - d(i)))$ is as small as possible.

i	1	2
$t(i)$	3	2
$d(i)$	1	3



Template for Greedy Algorithm

- ▶ Key question: In what order should we schedule the jobs?

Template for Greedy Algorithm

- ▶ Key question: In what order should we schedule the jobs?

Shortest length Increasing order of length $t(i)$.

Shortest slack time Increasing order of $d(i) - t(i)$.

Earliest deadline Increasing order of deadline $d(i)$.

Template for Greedy Algorithm

- Key question: In what order should we schedule the jobs?

Shortest length Increasing order of length $t(i)$. Ignores deadlines completely!

Shortest job may have a very late deadline.

i	1	2
$t(i)$	1	10
$d(i)$	100	10

Shortest slack time Increasing order of $d(i) - t(i)$.

Earliest deadline Increasing order of deadline $d(i)$.

Template for Greedy Algorithm

- Key question: In what order should we schedule the jobs?

Shortest length Increasing order of length $t(i)$. Ignores deadlines completely!

Shortest job may have a very late deadline.

i	1	2
$t(i)$	1	10
$d(i)$	100	10

Shortest slack time Increasing order of $d(i) - t(i)$. Job with smallest slack may take a long time.

i	1	2
$t(i)$	1	10
$d(i)$	2	10

Earliest deadline Increasing order of deadline $d(i)$.

Template for Greedy Algorithm

- Key question: In what order should we schedule the jobs?

Shortest length Increasing order of length $t(i)$. Ignores deadlines completely!

Shortest job may have a very late deadline.

i	1	2
$t(i)$	1	10
$d(i)$	100	10

Shortest slack time Increasing order of $d(i) - t(i)$. Job with smallest slack may take a long time.

i	1	2
$t(i)$	1	10
$d(i)$	2	10

Earliest deadline Increasing order of deadline $d(i)$. Correct? Does it make sense to tackle jobs with earliest deadlines first?

Minimising Lateness: Earliest Deadline First

Order the jobs in order of their deadlines

Assume for simplicity of notation that $d_1 \leq \dots \leq d_n$

Initially, $f = s$

Consider the jobs $i = 1, \dots, n$ in this order

Assign job i to the time interval from $s(i) = f$ to $f(i) = f + t_i$


Let $f = f + t_i$

End

Return the set of scheduled intervals $[s(i), f(i)]$ for $i = 1, \dots, n$


- ▶ Proof of correctness is more complex.
- ▶ We will use an exchange argument: gradually modify the optimal schedule O till it is the same as the schedule A computed by the algorithm.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.
- ▶ Claim 3: There is an optimal schedule with no idle time.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.
- ▶ Claim 3: There is an optimal schedule with no idle time.
- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.

Properties of Schedules




- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.
- ▶ Claim 3: There is an optimal schedule with no idle time.
- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.?!

Properties of Schedules



- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.
- ▶ Claim 3: There is an optimal schedule with no idle time.
- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.?!
- ▶ Claim 5: The greedy algorithm produces an optimal schedule.

Properties of Schedules



- ▶ A schedule has an *inversion* if a job i with deadline $d(i)$ is scheduled before a job j with an earlier deadline $d(j)$, i.e., $d(j) < d(i)$ and $s(i) < s(j)$.
- ▶ Claim 1: The algorithm produces a schedule with no inversions and no idle time.
- ▶ Claim 2: All schedules with no inversions and no idle time have the same lateness.
 - ▶ Case 1: All jobs have distinct deadlines. There is a unique schedule with no inversions and no idle time.
 - ▶ Case 2: Some jobs have the same deadline. Ordering of the jobs does not change the maximum lateness of these jobs.
- ▶ Claim 3: There is an optimal schedule with no idle time.
- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time. ?!
- ▶ Claim 5: The greedy algorithm produces an optimal schedule. Follows from Claims 1, 2 and 4.

Proving Claim 4

- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.

Proving Claim 4

- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.
- ▶ Approach: Start with an optimal schedule O and use an *exchange argument* to convert O into a schedule that satisfies Claim 4 and has lateness not larger than O .

Proving Claim 4

- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.
- ▶ Approach: Start with an optimal schedule O and use an *exchange argument* to convert O into a schedule that satisfies Claim 4 and has lateness not larger than O .
 1. If O has an inversion, then there is a pair of jobs i and j such that j is scheduled just after i and $d(j) < d(i)$.

Proving Claim 4

- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.
- ▶ Approach: Start with an optimal schedule O and use an *exchange argument* to convert O into a schedule that satisfies Claim 4 and has lateness not larger than O .
 1. If O has an inversion, then there is a pair of jobs i and j such that j is scheduled just after i and $d(j) < d(i)$.
 2. Let i and j be consecutive inverted jobs in O . After swapping i and j , we get a schedule O' with one less inversion.


Proving Claim 4

- ▶ Claim 4: There is an optimal schedule with no inversions and no idle time.
- ▶ Approach: Start with an optimal schedule O and use an *exchange argument* to convert O into a schedule that satisfies Claim 4 and has lateness not larger than O .
 1. If O has an inversion, then there is a pair of jobs i and j such that j is scheduled just after i and $d(j) < d(i)$.
 2. Let i and j be consecutive inverted jobs in O . After swapping i and j , we get a schedule O' with one less inversion.
 3. The lateness of O' is no larger than the lateness of O .
- ▶ It is enough to prove the last item, since after $\binom{n}{2}$ swaps, we obtain a schedule with no inversions whose lateness is no larger than that of O .

Swapping Inverted Jobs

Only the finishing times of i and j are affected by the swap.

Before swapping:



(a)

After swapping:



(b)

Figure 4.6 The effect of swapping two consecutive, inverted jobs.

- In O , assume each job r is scheduled for the interval $[s(r), f(r)]$ and has lateness $l(r)$. For O' , let the lateness of job r be $l'(r)$.

Swapping Inverted Jobs

Only the finishing times of i and j are affected by the swap.




Figure 4.6 The effect of swapping two consecutive, inverted jobs.

- ▶ In O , assume each job r is scheduled for the interval $[s(r), f(r)]$ and has lateness $l(r)$. For O' , let the lateness of job r be $l'(r)$.
- ▶ Claim: $l'(k) = l(k)$, for all $k \neq i, j$.

Swapping Inverted Jobs

Only the finishing times of i and j are affected by the swap.




Figure 4.6 The effect of swapping two consecutive, inverted jobs.

- ▶ In O , assume each job r is scheduled for the interval $[s(r), f(r)]$ and has lateness $l(r)$. For O' , let the lateness of job r be $l'(r)$.
- ▶ Claim: $l'(k) = l(k)$, for all $k \neq i, j$.
- ▶ Claim: $l'(j) \leq l(j)$.

Swapping Inverted Jobs

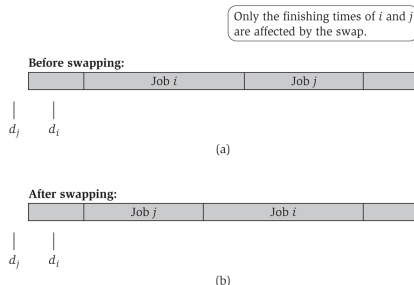


Figure 4.6 The effect of swapping two consecutive, inverted jobs.

- ▶ In O , assume each job r is scheduled for the interval $[s(r), f(r)]$ and has lateness $l(r)$. For O' , let the lateness of job r be $l'(r)$.
- ▶ Claim: $l'(k) = l(k)$, for all $k \neq i, j$.
- ▶ Claim: $l'(j) \leq l(j)$.
- ▶ Claim: $l'(i) \leq l(j)$

Swapping Inverted Jobs




Figure 4.6 The effect of swapping two consecutive, inverted jobs.

- ▶ In O , assume each job r is scheduled for the interval $[s(r), f(r)]$ and has lateness $l(r)$. For O' , let the lateness of job r be $l'(r)$.
- ▶ Claim: $l'(k) = l(k)$, for all $k \neq i, j$.
- ▶ Claim: $l'(j) \leq l(j)$.
- ▶ Claim: $l'(i) \leq l(j)$ because $l'(i) = f(j) - d_i \leq f(j) - d_j = l(j)$.

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie?

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie?

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie?

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$
5. Find an inversion in X and then isolate the inversion to be between consecutive jobs in X .
6. Swap the jobs to get a new schedule X_{i-1} . Where does X_{i-1} lie?

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$
5. Find an inversion in X and then isolate the inversion to be between consecutive jobs in X .
6. Swap the jobs to get a new schedule X_{i-1} . Where does X_{i-1} lie? X_{i-1} has one fewer inversion! Lateness cannot increase! So X_{i-1} is at $(i-1, l_X)$ or "below."
7. Repeat until we have X_1 with one inversion at

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$
5. Find an inversion in X and then isolate the inversion to be between consecutive jobs in X .
6. Swap the jobs to get a new schedule X_{i-1} . Where does X_{i-1} lie? X_{i-1} has one fewer inversion! Lateness cannot increase! So X_{i-1} is at $(i-1, l_X)$ or "below."
7. Repeat until we have X_1 with one inversion at $(1, l_X)$ or "below", where $l_X < l_A$.
8. Repeat one more step: X_0 has no inversions. What is X_0 's location?

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$
5. Find an inversion in X and then isolate the inversion to be between consecutive jobs in X .
6. Swap the jobs to get a new schedule X_{i-1} . Where does X_{i-1} lie? X_{i-1} has one fewer inversion! Lateness cannot increase! So X_{i-1} is at $(i-1, l_X)$ or "below."
7. Repeat until we have X_1 with one inversion at $(1, l_X)$ or "below", where $l_X < l_A$.
8. Repeat one more step: X_0 has no inversions. What is X_0 's location? $(0, l_X)$ or "below" because of #7 and $(0, l_A)$ because of #3.

Summary of Proof

1. Think of a schedule as a 2D point: x -coordinate is the number of inversions in the schedule and y -coordinate is the lateness of the schedule.
2. Where does the schedule A produced by the algorithm lie? Somewhere on the y -axis since it has no inversions, say $(0, l_A)$.
3. Where does some other schedule B with no inversions lie? Also at $(0, l_A)$ since all schedules with no inversions have the same lateness.
4. Let X be any schedule that is supposed to be optimal (and better than A). Where does X lie? At some point (i, l_X) , where $i > 0$ and l_X are the number of inversions in and lateness of X , respectively. $l_X < l_A$
5. Find an inversion in X and then isolate the inversion to be between consecutive jobs in X .
6. Swap the jobs to get a new schedule X_{i-1} . Where does X_{i-1} lie? X_{i-1} has one fewer inversion! Lateness cannot increase! So X_{i-1} is at $(i-1, l_X)$ or "below."
7. Repeat until we have X_1 with one inversion at $(1, l_X)$ or "below", where $l_X < l_A$.
8. Repeat one more step: X_0 has no inversions. What is X_0 's location? $(0, l_X)$ or "below" because of #7 and $(0, l_A)$ because of #3.
9. We have a contradiction!
10. Lateness of A cannot be larger than that of O !

Common Mistakes with Exchange Arguments

- ▶ Wrong: start with algorithm's schedule A and argue that A cannot be improved by swapping two jobs.
- ▶ Correct: Start with an arbitrary schedule O (which can be the optimal one) and argue that O can be converted into the schedule that is essentially the same as the one the algorithm produces without increasing the lateness.

Common Mistakes with Exchange Arguments

- ▶ Wrong: start with algorithm's schedule A and argue that A cannot be improved by swapping two jobs.
- ▶ Correct: Start with an arbitrary schedule O (which can be the optimal one) and argue that O can be converted into the schedule that is essentially the same as the one the algorithm produces without increasing the lateness.
- ▶ Wrong: Swap two jobs that are not neighbouring in O . Pitfall is that the completion times of all intervening jobs changes.
- ▶ Correct: Show that an inversion exists between two neighbouring jobs and swap them.

Summary

- ▶ Greedy algorithms make local decisions.
- ▶ Three analysis strategies:

Greedy algorithm stays ahead Show that after each step in the greedy algorithm, its solution is at least as good as that produced by any other algorithm.

Structural bound First, discover a property that must be satisfied by every possible solution. Then show that the (greedy) algorithm produces a solution with this property.

Exchange argument Transform the optimal solution in steps into the solution by the greedy algorithm without worsening the quality of the optimal solution.