Chapter 4

Greedy Algorithms

4.1 Interval Scheduling

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 natural order.
Take each job provided it's compatible with the ones already taken.

- [Earliest start time] Consider jobs in ascending order of s_j.

- [Earliest finish time] Consider jobs in ascending order of f_j.

- [Shortest interval] Consider jobs in ascending order of $f_j - s_j$.

- [Fewest conflicts] For each job j, count the number of conflicting jobs c_j. Schedule in ascending order of c_j.

Interval Scheduling: Greedy Algorithms

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

Greedy 1: **Earliest** Start Time

Does it work ?

1st Step of **Problem Solvers**: Find bad situations for the Algorithm.

Interval Scheduling: Greedy Algorithms

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

Greedy 1: **Earliest** Start Time

counterexample for earliest start time

Interval Scheduling: Greedy Algorithms

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

Greedy 2: **Shortest Interval**

Bad Situations ???

Interval Scheduling: Greedy Algorithms

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

Greedy 2: **Shortest Interval**

counterexample for shortest interval

Interval Scheduling: Greedy Algorithms

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

Greedy 3: **FEWEST** CONFLICTS

IDEA: Use GRAPH MODELLING;

Which is the problem in terms of GRAPHS?

Interval Scheduling: Greedy Algorithms

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

Greedy 3: **FEWEST** CONFLICTS

counterexample for fewest conflicts

Interval Scheduling: Greedy Algorithm

Greedy 4. 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.

       ↙  set of jobs selected

A ← φ
for j = 1 to n {
    if (job j compatible with A)
        A ← A ∪ {j}
}
return A
```

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

Interval Scheduling: Analysis

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

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

job i_{r+1} finishes before j_{r+1}

Greedy: i_1 i_2 i_r i_{r+1}

OPT: j_1 j_2 j_r j_{r+1} . . .

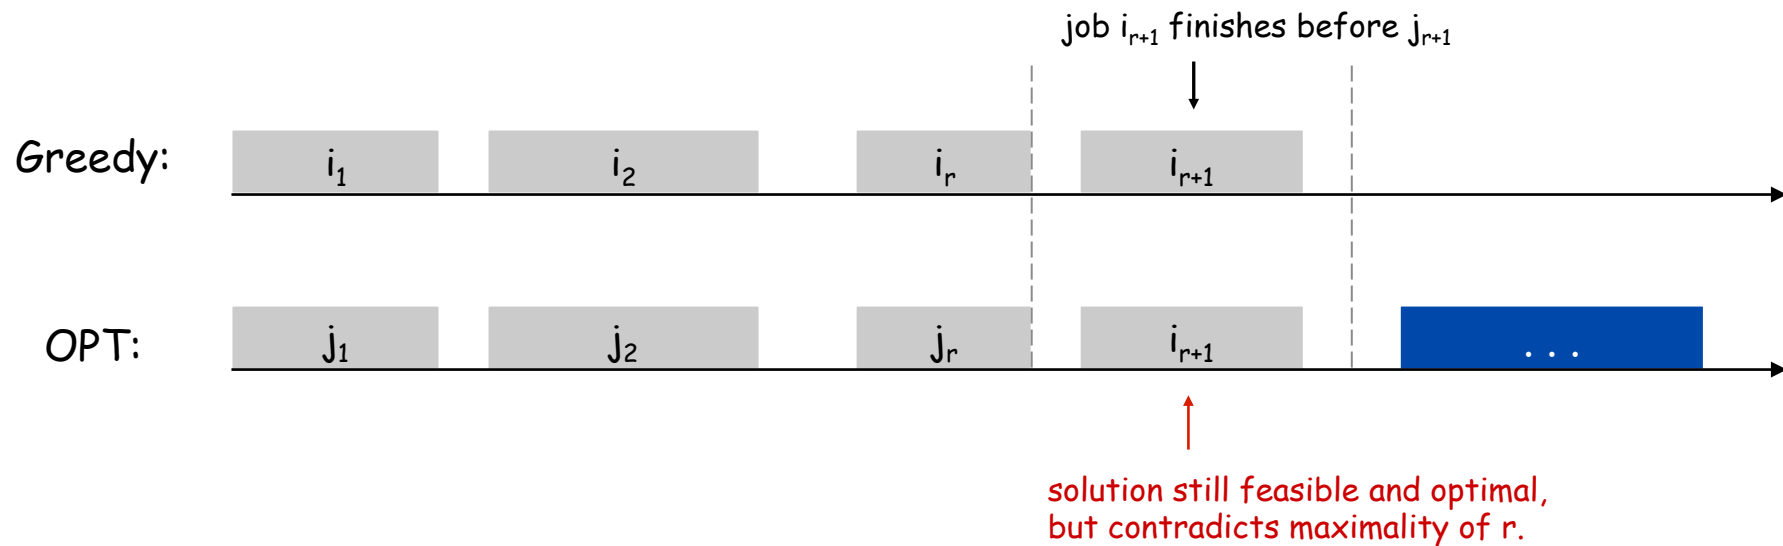why not replace job j_{r+1} with job i_{r+1}?

Interval Scheduling: Analysis

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

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

job i_{r+1} finishes before j_{r+1}

Greedy:
| i_1 | i_2 | i_r | i_{r+1} |

OPT:
| j_1 | j_2 | j_r | i_{r+1} | . . . |

solution still feasible and optimal,
but contradicts maximality of r.

4.1 Interval Partitioning

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** ≥ depth.

Ex: depth of schedule below = 3 ⇒ 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 s₁ ≤ s₂ ≤ ... ≤ sₙ.
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.

Interval Partitioning: Greedy Analysis

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

Theorem. Greedy algorithm is optimal.
Pf(Idea).

- 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 (i.e. $d-1$ jobs)
- These d jobs, **each** ends after s_j.
- Since we sorted by start time, all these incompatibilities are caused by **jobs** *that start no later than* s_j.
- Thus, we have d jobs overlapping at time $s_j + \varepsilon$.
- **Key observation (Lower Bound)** $\Rightarrow$ all schedules use $\geq d$ classrooms.
  - ▪

Partitioning

Formal PROOF.

Do as Excercise.

EXCERCISE

-The Car Traveling problem.

4.2 Scheduling to Minimize Lateness

Scheduling to Minimizing Lateness

Minimizing lateness problem.

- **Single** resource processes **one** job at a time.
- Job **j** requires t_j units of processing time and is due at time d_j.
- If **j** starts at time s_j, it finishes at time $f_j = s_j + t_j$.
- Lateness: $l_j = \max\{0, f_j - d_j\}$.
- *Goal: schedule all jobs to minimize maximum lateness* $L = \max l_j$.

Ex:

	1	2	3	4	5	6
t_j	3	2	1	4	3	2
d_j	6	8	9	9	14	15

lateness = 2   lateness = 0   max lateness = 6

| $d_3 = 9$ | $d_2 = 8$ | $d_6 = 15$ | $d_1 = 6$ | $d_5 = 14$ | $d_4 = 9$ |

0   1   2   3   4   5   6   7   8   9   10   11   12   13   14   15

Minimizing Lateness: Greedy Algorithms

Greedy template. Consider jobs in some order.

- [Shortest processing time first] Consider jobs in ascending order of processing time t_j.

- [Earliest deadline first] Consider jobs in ascending order of deadline d_j.

- [Smallest slack] Consider jobs in ascending order of slack $d_j - t_j$.

Minimizing Lateness: Greedy Algorithms

Greedy template. Consider jobs in some order.

- [G1: Shortest processing time first] Consider jobs in **ascending** order of processing time t_j.

	1	2
t_j	1	10
d_j	100	10

counterexample

G1 solution: Job 1; Job 2 --> Latency = 1
Optimal Solution: Job 2; Job 1 --> Latency = 0

Minimizing Lateness: Greedy Algorithms

Greedy template. Consider jobs in some order.

- [G2 Smallest slack] Consider jobs in **ascending** order of slack $d_j - t_j$.

G2 Solution: Job 2; Job 1. Latency = 10

Optimal: Job 1; Job 2. Latency = 1

	1	2
t_j	1	10
d_j	2	10

counterexample

Minimizing Lateness: Greedy Algorithm

Greedy algorithm. Earliest deadline **d** first.

```
Sort n jobs by deadline so that d₁ ≤ d₂ ≤ … ≤ dₙ

t ← 0
for j = 1 to n
    Assign job j to interval [t, t + tⱼ]
    sⱼ ← t, fⱼ ← t + tⱼ
    t ← t + tⱼ
output intervals [sⱼ, fⱼ]
```

max lateness = 1

$d_1 = 6$			$d_2 = 8$	$d_3 = 9$	$d_4 = 9$			$d_5 = 14$		$d_6 = 15$

0   1   2   3   4   5   6   7   8   9   10   11   12   13   14   15

Minimizing Lateness: No Idle Time

Observation. There exists an optimal schedule with no idle time.

d = 4			d = 6					d = 12		
0	1	2	3	4	5	6	7	8	9	10	11

d = 4		d = 6		d = 12						
0	1	2	3	4	5	6	7	8	9	10	11

Observation. *The greedy schedule has no idle time.*

Minimizing Lateness: Inversions

Def. Given a **schedule** S, an inversion is a pair of jobs **i** and **j** such that:
i < **j** but **j** scheduled before i.

[as before, we assume jobs are numbered so that $d_1 \leq d_2 \leq \ldots \leq d_n$]

Observation. *Greedy schedule has no inversions.*

Observation. *If a schedule (with no idle time) has an inversion, it has one with a pair of inverted jobs scheduled* **consecutively**.

Minimizing Lateness: Inversions

Def. Given a schedule S, an inversion is a pair of jobs i and j such that: i < j but j scheduled before i.

Claim. Swapping two consecutive, inverted jobs **reduces** the number of inversions by **one** and does not increase the max lateness (the sum is commutative!).

- **Pf.** Let L be the lateness before the swap, and let L' be it afterwards.
- $\ell'_k = \ell_k$ for all $k \neq i, j$
- $\ell'_i \leq \ell_i$
- If job j is late:

$$
\begin{aligned}
\ell'_j &= f'_j - d_j && \text{(definition)} \\
&= f_i - d_j && (j \text{ finishes at time } f_i) \\
&\leq f_i - d_i && (i < j) \\
&\leq \ell_i && \text{(definition)}
\end{aligned}
$$

Theorem. Greedy schedule **S** is optimal.

Pf. Define **S*** to be an optimal schedule that has the **fewest°** number of **inversions**, and let's see what happens.

- Can assume **S*** has no idle time.
- If S* has **no** inversions, then **S** = **S***.
- If S* has an inversion, let i-j be an **adjacent** inversion.
  - swapping i and j does not increase the maximum lateness and strictly decreases the number of inversions
  - this contradicts definition° of S* ▪

Greedy Analysis Strategies

Greedy algorithm stays ahead. Show that after each step of the greedy algorithm, its solution is at least as good as any other algorithm's.

Structural. Discover a simple "structural" bound asserting that every possible solution must have a certain value. Then show that your algorithm always achieves this bound.

Exchange argument. Gradually transform any solution to the one found by the greedy algorithm without hurting its quality.

Other greedy algorithms. Kruskal, Prim, Dijkstra, Huffman, …

EXCERCISE

-BUYING ITEMS OF INCREASING COSTS
(excercise 2 of the Book)

4.3 Optimal Caching

Optimal Offline Caching

Caching.

- Cache with capacity to store k items.
- Sequence of m item requests $d_1, d_2, ..., d_m$.
- Cache hit: item already in cache when requested.
- Cache miss: item not already in cache when requested: must bring requested item into cache, and evict some existing item, if full.

Goal. Eviction schedule that minimizes number of cache misses.

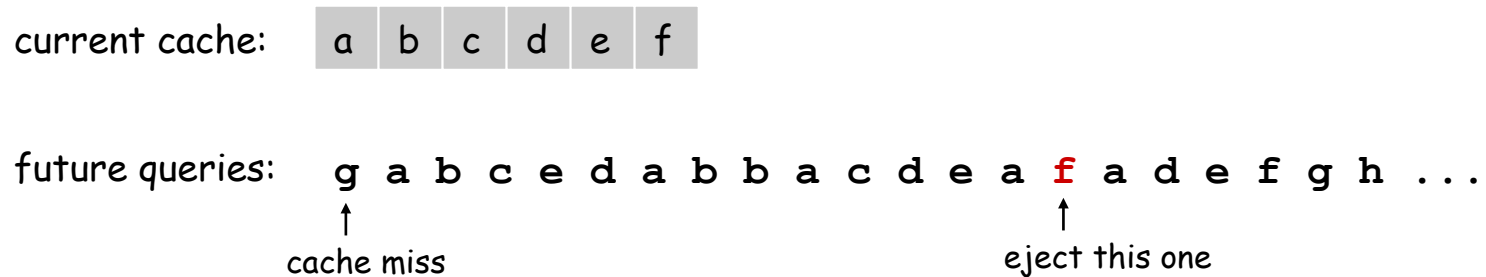Ex: k = 2, initial cache = ab,
    requests: a, b, c, b, c, a, a, b.

Optimal eviction schedule: 2 cache misses.

red = cache miss

requests	cache	
a	a	b
b	a	b
c	c	b
b	c	b
c	c	b
a	a	b
a	a	b
b	a	b

Optimal Offline Caching: Farthest-In-Future

Farthest-in-future. Evict item in the cache that is not requested until farthest in the future.

current cache: | a | b | c | d | e | f |

future queries: g a b c e d a b b a c d e a **f** a d e f g h ...

        ↑   ↑

    cache miss   eject this one

Theorem. [Bellady, 1960s] FF is optimal eviction schedule.
Pf. Algorithm and theorem are intuitive; proof is subtle.

Reduced Eviction Schedules

Def. A reduced schedule is a schedule that only inserts an item into the cache in a step in which that item is requested.

Intuition. Can transform an unreduced schedule into a reduced one with no more cache misses.


an unreduced schedule

a reduced schedule

Reduced Eviction Schedules

Claim. Given any unreduced schedule S, can transform it into a reduced schedule S' with no more cache misses.

Pf. (by induction on number of unreduced items)

doesn't enter cache at requested time

- Suppose S brings d into the cache at time t, without a request.
- Let c be the item S evicts when it brings d into the cache.
- Case 1: d evicted at time t', before next request for d.
- Case 2: d requested at time t' before d is evicted. ▪

Case 1

Case 2

Farthest-In-Future: Analysis

Theorem. FF is optimal eviction algorithm.
Pf. (by induction on number or requests j)

> Invariant: There exists an optimal reduced schedule S that makes the same eviction schedule as S_{FF} through the first j+1 requests.

Let S be reduced schedule that satisfies invariant through j requests.
We produce S' that satisfies invariant after j+1 requests.

- Consider $(j+1)^{st}$ request $d = d_{j+1}$.
- Since S and S_{FF} have agreed up until now, they have the same cache contents before request j+1.
- Case 1: (d is already in the cache). S' = S satisfies invariant.
- Case 2: (d is not in the cache and S and S_{FF} evict the same element). S' = S satisfies invariant.

Farthest-In-Future: Analysis

Pf. (continued)

- Case 3: (d is not in the cache; S_{FF} evicts e; S evicts f ≠ e).

  - begin construction of S' from S by evicting e instead of f

j		same	e	f		same	e	f

  S   S'

j+1		same	e	d		same	d	f

  S   S'

  - now S' agrees with S_{FF} on first j+1 requests; we show that having element f in cache is no worse than having element e

Farthest-In-Future: Analysis

Let j' be the **first** time after j+1 that S and S' take a different action, and let g be item requested at time j'.

↑
must involve e or f (or both)

j' | same | e | | same | f |
       S   S'

- Case 3a: g = e. Can't happen with Farthest-In-Future since there must be a request for f before e.

- Case 3b: g = f. Element f can't be in cache of S, so let e' be the element that S evicts.
  - if e' = e, S' accesses f from cache; now S and S' have same cache
  - if e' ≠ e, S' evicts e' and brings e into the cache; now S and S' have the same cache
    ↑

    Note: S' is no longer reduced, but can be transformed into a reduced schedule that agrees with S_{FF} through step j+1

Farthest-In-Future: Analysis

Let j' be the first time after j+1 that S and S' take a different action, and let g be item requested at time j'.

↑
must involve e or f (or both)

j' | same | e | | same | f |
         S   S'

otherwise S' would take the same action
↓

- Case 3c: g ≠ e, f. S must evict e.
  Make S' evict f; now S and S' have the same cache. ▪

j' | same | g | | same | g |
         S   S'

Caching Perspective

Online vs. offline algorithms.

- Offline: full sequence of requests is known a priori.
- Online (reality): requests are not known in advance.
- Caching is among most fundamental online problems in CS.

LIFO. Evict page brought in most recently.

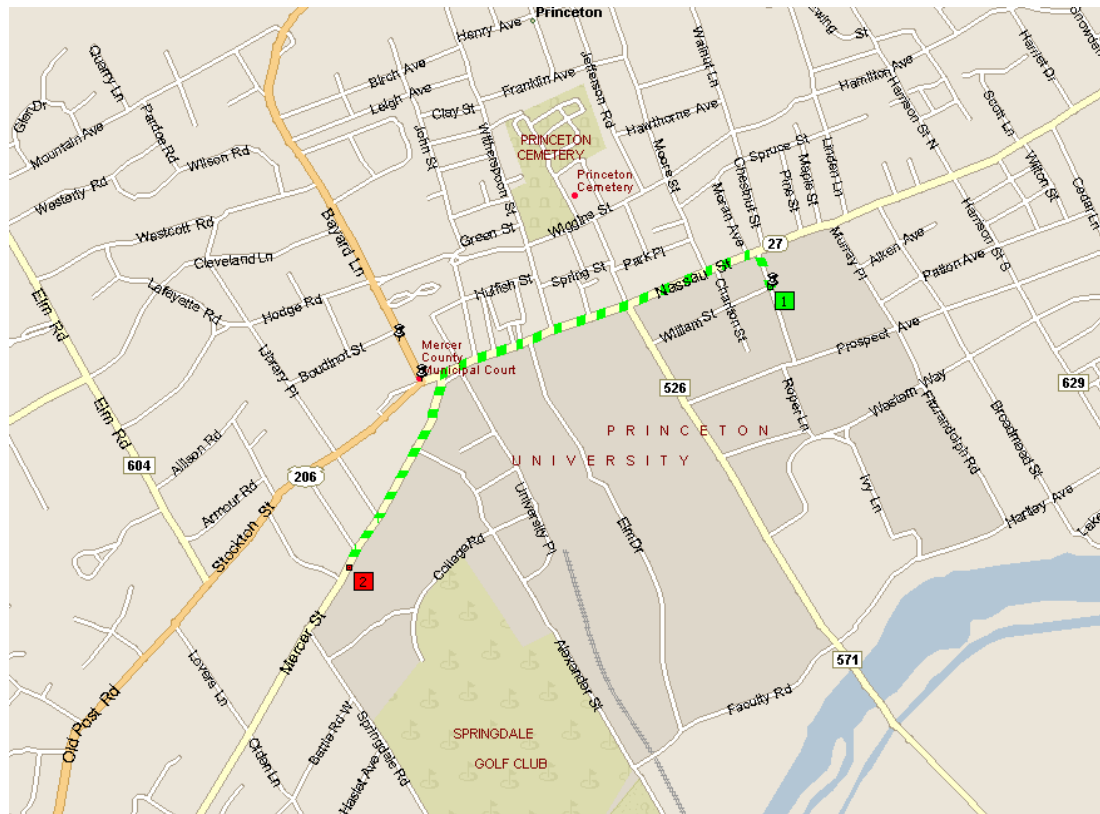LRU. Evict page whose most recent access was earliest.

$\uparrow$

FF with direction of time reversed!

Theorem. FF is optimal offline eviction algorithm.

- Provides basis for understanding and analyzing online algorithms.
- LRU is k-competitive. [Section 13.8]
- LIFO is arbitrarily bad.

4.4 Shortest Paths in a Graph

shortest path from Princeton CS department to Einstein's house

Shortest Path Problem

Shortest path network.
- Directed graph G = (V, E).
- Source s, destination t.
- Length ℓ_e = length of edge e.

Shortest path problem: find shortest directed path from s to t.

cost of path = sum of edge costs in path

Cost of path s-2-3-5-t
  = 9 + 23 + 2 + 16
  = 50.

Dijkstra's Algorithm

Dijkstra's algorithm.

- Maintain a set of explored nodes S for which we have determined the shortest path distance d(u) from s to u.
- Initialize S = { s }, d(s) = 0.
- Repeatedly choose unexplored node v which minimizes

$$\pi(v) = \min_{e = (u,v) \,:\, u \in S} d(u) + \ell_e,$$

add v to S, and set d(v) = $\pi(v)$.

shortest path to some u in explored part, followed by a single edge (u, v)

Dijkstra's Algorithm

Dijkstra's algorithm.

- Maintain a set of explored nodes S for which we have determined the shortest path distance d(u) from s to u.
- Initialize S = { s }, d(s) = 0.
- Repeatedly choose unexplored node v which minimizes

$$\pi(v) = \min_{e = (u,v)\,:\,u \in S} d(u) + \ell_e,$$

add v to S, and set d(v) = π(v).

shortest path to some u in explored part, followed by a single edge (u, v)

Dijkstra's Algorithm: Proof of Correctness

Invariant. For each node u $\in$ S, d(u) is the length of the shortest s-u path.
Pf. (by induction on |S|)
Base case: |S| = 1 is trivial.
Inductive hypothesis: Assume true for |S| = k $\geq$ 1.

- Let v be next node added to S, and let u-v be the chosen edge.
- The shortest s-u path plus (u, v) is an s-v path of length $\pi(v)$.
- Consider any s-v path P. We'll see that it's no shorter than $\pi(v)$.
- Let x-y be the first edge in P that leaves S, and let P' be the subpath to x.
- P is already too long as soon as it leaves S.

$$\ell(P) \geq \ell(P') + \ell(x,y) \geq d(x) + \ell(x, y) \geq \pi(y) \geq \pi(v)$$

$\uparrow$ nonnegative weights

$\uparrow$ inductive hypothesis

$\uparrow$ defn of $\pi(y)$

$\uparrow$ Dijkstra chose v instead of y

Dijkstra's Algorithm: Implementation

For each unexplored node, explicitly maintain $\pi(v) = \min\limits_{e=(u,v)\,:\,u\in S} d(u) + \ell_e$.

- Next node to explore = node with minimum $\pi(v)$.
- When exploring v, for each incident edge e = (v, w), update

$$\pi(w) = \min\{\ \pi(w),\ \pi(v)+\ell_e\ \}.$$

Efficient implementation. Maintain a priority queue of unexplored nodes, prioritized by $\pi(v)$.

PQ Operation	Dijkstra	Array	Binary heap	d-way Heap	Fib heap †
Insert	n	n	log n	$d \log_d n$	1
ExtractMin	n	n	log n	$d \log_d n$	log n
ChangeKey	m	1	log n	$\log_d n$	1
IsEmpty	n	1	1	1	1
Total		n^2	m log n	$m \log_{m/n} n$	m + n log n

† Individual ops are amortized bounds

Extra Slides

Coin Changing

Greed is good. Greed is right. Greed works. Greed clarifies, cuts through, and captures the essence of the evolutionary spirit.

- Gordon Gecko (Michael Douglas)

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: c₁ < c₂ < … < cₙ.

        coins selected
       ↙
S ← ∅
while (x ≠ 0) {
    let k be largest integer such that cₖ ≤ x
    if (k = 0)
        return "no solution found"
    x ← x - cₖ
    S ← S ∪ {k}
}
return S
```

Q. Is cashier's algorithm optimal?

Coin-Changing: Analysis of Greedy Algorithm

Theorem. Greed is optimal for U.S. coinage: 1, 5, 10, 25, 100.

Pf. (by induction on x)

- Consider optimal way to change $c_k \leq x < c_{k+1}$: greedy takes coin k.
- We claim that any optimal solution must also take coin k.
    - if not, it needs enough coins of type $c_1, ..., c_{k-1}$ 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. ▪

k	c_k	All optimal solutions must satisfy	Max value of coins 1, 2, ..., k-1 in any OPT
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: 1, 10, 21, 34, 70, 100, 350, 1225, 1500.

Counterexample. 140¢.
- Greedy: 100, 34, 1, 1, 1, 1, 1, 1.
- Optimal: 70, 70.

Selecting Breakpoints

Selecting Breakpoints

Selecting breakpoints.

- Road trip from Princeton to Palo Alto along fixed route.
- Refueling stations at certain points along the way.
- Fuel capacity = C.
- Goal: makes as few refueling stops as possible.

Greedy algorithm. Go as far as you can before refueling.

Selecting Breakpoints: Greedy Algorithm

Truck driver's algorithm.

```
Sort breakpoints so that: 0 = b₀ < b₁ < b₂ < ... < bₙ = L

S ← {0}   ←──── breakpoints selected
x ← 0     ←──── current location


while (x ≠ bₙ)
    let p be largest integer such that bₚ ≤ x + C
    if (bₚ = x)
        return "no solution"
    x ← bₚ
    S ← S ∪ {p}
return S
```

Implementation. O(n log n)

- Use binary search to select each breakpoint p.

Selecting Breakpoints: Correctness

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

- Assume greedy is not optimal, and let's see what happens.
- Let $0 = g_0 < g_1 < \ldots < g_p = L$ denote set of breakpoints chosen by greedy.
- Let $0 = f_0 < f_1 < \ldots < f_q = L$ denote set of breakpoints in an optimal solution with $f_0 = g_0$, $f_1 = g_1$, $\ldots$, $f_r = g_r$ for largest possible value of r.
- Note: $g_{r+1} > f_{r+1}$ by greedy choice of algorithm.


why doesn't optimal solution drive a little further?

Selecting Breakpoints: Correctness

Theorem. Greedy algorithm is optimal.

Pf. (by contradiction)

- Assume greedy is not optimal, and let's see what happens.
- Let $0 = g_0 < g_1 < \ldots < g_p = L$ denote set of breakpoints chosen by greedy.
- Let $0 = f_0 < f_1 < \ldots < f_q = L$ denote set of breakpoints in an optimal solution with $f_0 = g_0,\ f_1 = g_1,\ \ldots,\ f_r = g_r$ for largest possible value of r.
- Note: $g_{r+1} > f_{r+1}$ by greedy choice of algorithm.

Greedy: g_0 g_1 g_2 g_r g_{r+1}

OPT: f_0 f_1 f_2 f_r f_q

another optimal solution has
one more breakpoint in common
$\Rightarrow$ contradiction

Edsger W. Dijkstra

The question of whether computers can think is like the question of whether submarines can swim.

Do only what only you can do.

In their capacity as a tool, computers will be but a ripple on the surface of our culture. In their capacity as intellectual challenge, they are without precedent in the cultural history of mankind.

The use of COBOL cripples the mind; its teaching should, therefore, be regarded as a criminal offence.

APL is a mistake, carried through to perfection. It is the language of the future for the programming techniques of the past: it creates a new generation of coding bums.