

Farthest-In-Future: Analysis

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

Feb 11, 2009 CS211 7

Farthest-In-Future: Analysis

Pf. (continued)

- Case 3: d is not in the cache; S_{FF} evicts e ; S evicts $f \neq e$
 - 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 requests; we show that having element f in cache is *no worse* than having element e
 - Need to get schedules' caches back in sync again
 - All decisions will be the same until decision involves e or f

Feb 11, 2009 CS211 8

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' .

j'	same	e	→	same	f
	S			S'	

↑ must involve e or f (or both)

- What are the possibilities for g ?
 - Is g in the cache for S ? For S' ?
 - What does their caches look like afterwards?

Feb 11, 2009 CS211 9

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' .

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

Feb 11, 2009 CS211 10

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'

j'	same	e	→	same	f
	S			S'	

- Case 3b: $g \neq e, f$
 - g is not in either cache
 - S must evict e
 - otherwise S' would take the same action
 - Make S' evict f ; now S and S' have the same cache:

j'	same	g	→	same	g
	S			S'	

Feb 11, 2009 CS211 11

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' .

j'	same	e	→	same	f
	S			S'	

- Case 3c: $g = f$
 - Element f can't be in cache of S , so let e' be the element that S evicts
 - If $e' = e$, now S and S' have same cache
 - If $e' \neq e$, 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 $j+1$

Feb 11, 2009 CS211 12

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' .



For both cases (3b, 3c), have reduced schedule S' that agrees with S_{FF} for first $j+1$ items

Farthest-in-Future: Analysis

Theorem. FF is optimal eviction algorithm

Pf. (by induction on number of requests j)

Let S^* be an optimal schedule

Construct an optimal schedule S_1 that agrees with S_{FF} through the first step

Apply previous proof inductively for $j = 1, 2, 3, \dots, m$, producing schedules S_j that agree with S_{FF} through first j steps

Each schedule S_j incurs no more misses than the corresponding S_{FF} one

$S_m = S_{FF}$ because agrees through whole sequence

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

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

FF with direction of time reversed!


SHORTEST PATHS IN A GRAPH

Shortest Path Problem

Given

- Directed graph $G = (V, E)$
- Source s , destination t
- Length ℓ_e = length of edge e (non-negative)

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




Shortest Path Problem

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

Towards algorithm ideas:

- What is shortest path from s to 2? To 6?
- What is the shortest path to 3? 5? 7?



Dijkstra's Algorithm

Maintain a set of explored nodes S

- Know the shortest path distance $d(u)$ from s to u

Initialize $S=\{s\}$, $d(s)=0$

Repeatedly choose unexplored node v which minimizes $\pi(v) = \min_{e=(u,v): u \in S} d(u) + \ell_e$, ← shortest path to some u in explored part, followed by a single edge (u, v)

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

Dijkstra's Algorithm

Before

After

Feb 11, 2009 CS211 20

Dijkstra's Shortest Path Algorithm

Find shortest path from s to t .

21

Dijkstra's Shortest Path Algorithm

$S = \{ \}$
 $PQ = \{ s, 2, 3, 4, 5, 6, 7, t \}$

Initialize distances to all nodes to infinity

distance label → ∞

22

Dijkstra's Shortest Path Algorithm

$S = \{ \}$
 $PQ = \{ s, 2, 3, 4, 5, 6, 7, t \}$

delmin

distance label → ∞

23

Dijkstra's Shortest Path Algorithm







$S = \{ s \}$
 $PQ = \{ 2, 3, 4, 5, 6, 7, t \}$







decrease key





Add node s to explored set
Update distances to nodes it points to

distance label → 15

24







Dijkstra's Algorithm

Maintain a set of explored nodes S

- Know the shortest path distance $d(u)$ from s to u

Initialize $S = \{s\}$, $d(s) = 0$

Repeatedly choose unexplored node v which minimizes $\pi(v) = \min_{e = (u,v): u \in S} d(u) + \ell_e$ ← shortest path to some u in explored part, followed by a single edge (u, v)

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

Running time?
 Implementation?
 Data structures?

Dijkstra's Algorithm

Maintain a set of explored nodes S

- Know the shortest path distance $d(u)$ from s to u

Initialize $S = \{s\}$, $d(s) = 0$

Repeatedly choose unexplored node v which minimizes $\pi(v) = \min_{e = (u,v): u \in S} d(u) + \ell_e$ ← shortest path to some u in explored part, followed by a single edge (u, v)

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

Using a priority queue, how many

- Inserts?
- Finding minimum?
- Deletions?
- Updating the key?
- Determining if empty?

How long does each operation take?

Dijkstra's Algorithm: Implementation

For each unexplored node, explicitly maintain

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

- 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	Binary heap
Insert	n	$\log n$
ExtractMin	n	$\log n$
ChangeKey	m	$\log n$
IsEmpty	n	1
Total		$m \log n$

Feb 11, 2009

CS211

43

How Greedy?

How Greedy?

We always form shortest new $s-v$ path from a path in S followed by a *single* edge

Proof of optimality: *Stays ahead* of all other solutions

- Each time selects a path to a node v , that path is shorter than every other possible path to 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: $|S|=1$...

Inductive hypothesis?

Next step?

Feb 11, 2009

CS211

46

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: For $|S| = 1$, $S=\{s\}$; $d(s) = 0$

Inductive hypothesis: Assume true for $|S| = k$, $k \geq 1$

- Grow $|S|$ to $k+1$
- Adding next node v by $u \rightarrow v$
- What do we know about $s \rightarrow u$?
- What can we say about other $s \rightarrow v$ paths?
- Why didn't we pick y as the next node?



Feb 11, 2009

CS211

47