

Online algorithms


Perform some sequence of operations op_1, op_2, \dots, op_m

When performing op_i , the future operations op_{i+1}, \dots, op_m are unknown.

- very broad topic
- most data structures can be seen this way

Example

Online paging problem



Sequence of requests

$$r_1, r_2, \dots, r_m \in \{1, \dots, N\}$$

for each request r_i , - if page r_i is in cache, "cache hit",
cost = 0, nothing to do

- if page r_i is not in cache, "cache miss / fault"
move r_i into cache

$$\text{cost} = 1$$


If cache memory is full, need to evict one page
from cache.

Question: how to decide which page to evict
from cache?

→ Caching / paging strategy.

(Ignore other bookkeeping
cost and computation)

e.g. $k=3$
 $N=100$



request (a)

a

 cost 1

request (b)

a
b

 cost 1

request (a)

a
b

 cost 0

request (c)

a
b
c

 cost 1

request (d)

a
d
c

 cost 1 (evict b)

request (b)

a
d
b

 cost 1

(In hindsight evicting b was not a good decision)

Possible strategies

1. LRU (evict least recently used)
2. LFU (evict least frequently used)
3. FIFO (first in first out → evict first inserted)
4. LIFO (last in first out → evict last inserted)
5. Random (evict a page chosen uniformly at random)




online algorithms

6. LFD (evict page with longest forward distance)

page which will be requested furthest in future
 (if some page never needed in future, evict that one)

offline algorithm


e.g. $k=3$
LRU



* each miss cost = 1

cost of req. with LRU is \neq

LFD



cost of seq. with LFD is 5.

How good are these strategies? Which strategy is best?
 How to analyze/compare different strategies?

request sequence $R = r_1, r_2, \dots, r_m$

$K < N$
 cache size # pages

Worst-case analysis:

for an arbitrary algorithm cost = m

(adversary always requests a page not in the cache)

e.g. $k=3$
 request: 1, 2, 3, 4, 5, 6, 7, ... m
 cost = m

even OPT cost is m

Competitive analysis of online algorithms

Let $Alg(R)$ denote cost of Alg on sequence R .

Let $OPT(R)$ denote the optimal cost of serving sequence R

knowing entire R in advance,
 infinite computation

Assume:

Initial state of algorithm: cache is full



"the price of not knowing the future"

Def. "competitive ratio" of Alg is \underline{c} , if $Alg(R) \leq c \cdot OPT(R) + O(1)$

for all sequences R .

(think of c as the maximum ratio between Alg. cost and optimum cost)

Theorem: $LFD(R) = OPT(R)$ for all R .

Proof


Requests r_1, \dots, r_m

$OPT \rightarrow$ optimal strategy

$LFD \rightarrow$ strategy of LFD algorithm

Proof strategy: we transform OPT step-by-step, until it matches LFD , the transformation will not increase cost.

Suppose that OPT and LFD do the same on r_1, \dots, r_i , do different on r_{i+1}



Case A. r_{i+1} is a cache hit (for both algorithms)

Case A1 OPT keeps p' in cache until it is requested

$OPT \rightarrow OPT'$
 s.t. $p \leftarrow p'$ postponed to time j
 at time $i+1$ do nothing

OPT and OPT' have same cost

but OPT' agrees with LFD on r_1, \dots, r_{i+1}


Case A2

OPT replaces p' with q before time j (at time $j' < j$)

$OPT \rightarrow OPT'$
 s.t. at time $i+1$ $p \leftarrow q$
 at time j' do nothing

OPT' is valid and cost smaller than cost of OPT
 contradiction.

Case B request r_{i+1} is a cache miss \rightarrow OPT: $P_1 \leftarrow P$
 \rightarrow LFD: $P_2 \leftarrow P$



Suppose r_{i+1} requests some page P

OPT first loads P_1 at time j
 ($X \leftarrow P_1$)

Case B1

OPT keeps P_2 until time j

OPT \rightarrow OPT'

at time $i+1$ do $P_2 \leftarrow P$

at time j do $X \leftarrow P_2$


OPT' valid, same cost as OPT,

OPT' agrees with LFD

on r_1, \dots, r_{i+1} requests

Case B2

OPT evicts P_2 before time j (at time j')
 ($P_2 \leftarrow Y$)



OPT \rightarrow OPT'

at time $i+1$
 $P_2 \leftarrow P$

at time j'
 $P_1 \leftarrow X$

OPT' valid, same cost as OPT

OPT' agrees with LFD on r_1, \dots, r_{i+1}

Summary: LFD is optimal.

Thm A For $R = r_1, \dots, r_m,$

$N = k + 1,$

$OPT(R) \leq \frac{m}{k}$

"impossibility result"

Thm A. OPT has low cost on R
Thm B. Arbitrary alg. has high cost on R } $\frac{Alg(R)}{OPT(R)}$ is large

Proof



- Suppose r_i is a cache miss

- load r_i

- decide which page to evict next $k+1$ requests

future requests are $r_{i+1}, r_{i+2}, \dots, r_{i+k+1}$

a) if all these $k+1$ requests are different, then evict r_{i+k+1} .

\Rightarrow Next k requests will have 0 miss.

b) if next $k+1$ requests have some duplicates \rightarrow at most k different requests away then \rightarrow keep all these k pages in cache \Rightarrow next $k+1$ requests will have 0 miss

For every cost 1 request, we have $\geq k$ consecutive cost 0 requests.

\Rightarrow Total cost $\leq \frac{m}{k}$



Thm B $N = k + 1$

Let Alg be an arbitrary deterministic online paging algorithm.

There is a sequence $R = r_1, \dots, r_m$ s.t. $Alg(R) = m$.

Proof

Cache size is k , Adversary requests exactly the missing page.

Here we use that Alg is deterministic so adversary knows which page is missing

□

Thm A } \Rightarrow For every deterministic online paging algorithm
Thm B } competitive ratio $\geq k$.


[Lower bound on competitive ratio]

Thm. LRU is k -competitive. \rightarrow at a cache miss, evict the page that was not requested for the longest time.

Arbitrary $R = r_1, \dots, r_m$ \rightarrow cache size.

Claim. $LRU(R) \leq k \cdot OPT(R)$

Proof Split the sequence R into phases P_0, P_1, \dots st. in each phase LRU has exactly k misses (except for last phase)



In each phase LRU cost = k

Need to show that OPT also makes at least one miss in each phase.

True in P_0 (LRU, OPT start from same initial state)

Look at phase P_i : $r_{t-1} \quad r_t \quad r_{t+1} \quad \dots \quad r_{t+e}$
 \downarrow
 Page P
 P_i

Lemma: Phase P_i contains requests to k different pages (other than P)


Observe: After time $t-1$, OPT has P in the cache

Lemma \Rightarrow at least one of the k different pages requested in P_i (other than P) is currently not in the cache \Rightarrow OPT will make a miss.

Proof of Lemma

Case A. If LRU makes k misses on k different pages (other than P) \Rightarrow Done

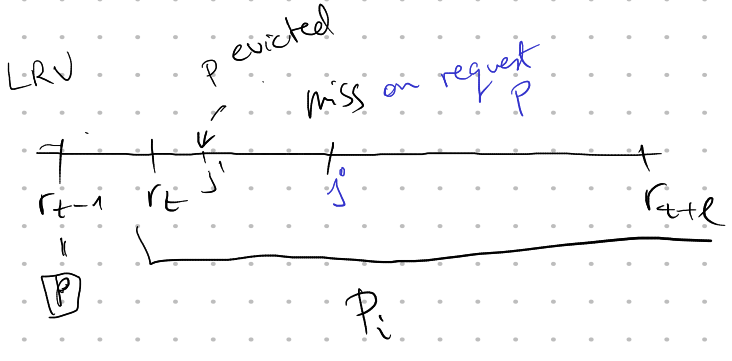
Case B LRU makes two misses on same page $Q \neq P$



- r_{i_1}, r_{i_2} two requests for q , where LRU misses
- q must be evicted between r_{i_1}, r_{i_2}
- Why did LRU evict q at time j ? ($j \in (i_1, i_2)$)
 - \Rightarrow because k other pages were requested in time interval $(i_1, j]$ (other than q) (these are "more recent" than q)
 - $\Rightarrow k+1$ diff. requests in the $[i_1, j]$ (including q)
 - $\Rightarrow k$ diff. requests (other than p) in this interval.

Case C.

LRU makes a miss on p



- After time $t-1$, p is in cache
- At time j , p is a miss
- \Rightarrow at some time $j' \in (t-1, j)$ p was evicted.
- Why was p evicted?
 - \Rightarrow there were k diff. pages (other than p) that were "more recent" than p .
 - \rightarrow these were requested in interval $[t, j']$

Thm. FIFO is k -competitive (proof similar)

This suggests FIFO and LRU similar.

But in practice LRU better than FIFO.

Thm. LFU / LIFO are not k -competitive (or even $f(k)$ -competitive for any function $f(\cdot)$)


Proof (LIFO) evicts page that was most recently added

$R = 1, 2, \dots, k, k+1, (k, k+1)^m$ \xrightarrow{m} repeated m times

(length of R is $2m + k + 1$)


\downarrow
 m can be arbitrarily large

OPT



OPT cost = $k+1$

LIFO



LIFO cost = $2m + k + 1$

We cannot bound $\frac{\text{LIFO}(R)}{\text{OPT}(R)} = \frac{2m + k + 1}{k + 1}$

\Rightarrow LIFO is not competitive.

Randomized algorithms

\rightarrow in expectation w.r. random choices
 \rightarrow adversary knows algorithm, but not random choices.
"oblivious adversary"

Thm. RANDOM is k -competitive

Proof (one side)

RANDOM is not better than k -competitive

$R = (k+1), 1, 2, \dots, k, (1, 2, \dots, k)^*$

OPT(R) = 1

OPT makes one miss evict $k+1$



How many misses by RANDOM for R?

Obs. After $(k+1)$ is evicted, algorithm makes no more miss.

misses until $k+1$ evicted:

when make a miss, $\Pr(k+1 \text{ is evicted}) = \frac{1}{k}$ (each page in cache equally likely to be evicted)

$E[\# \text{ misses}] = \text{Coinflips} \begin{cases} \frac{1}{k} \text{ pr. tail} \\ 1 - \frac{1}{k} \text{ pr. head} \end{cases} \Bigg| \text{expected \# coinflips until first tail.}$
 $= k$
 Expected cost $\text{RANDOM}(R) = k$

$\Rightarrow \frac{\text{RANDOM}}{\text{OPT}} \geq k$

Can we go beyond k -competitiveness? Answer: YES

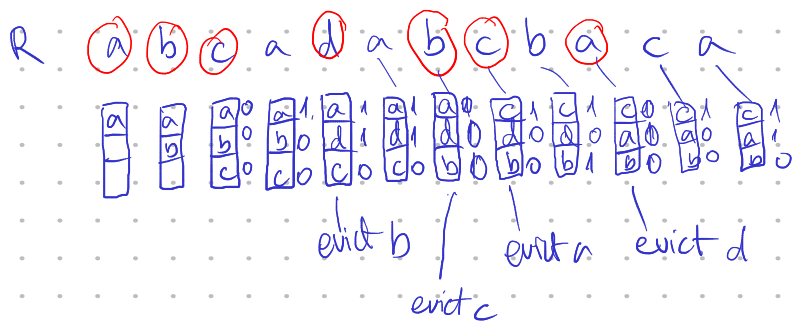
$O(\log k)$ - competitive paging algorithm. (Randomized)

Marking algorithm $\rightarrow O(\log k)$ competitive

hybrid between RANDOM / LRU

- Each k locations in cache are marked 0/1
- Initially all are marked 0
- cache hit \Rightarrow mark location 1
- cache miss \Rightarrow load requested page, mark 1
- evict randomly chosen page away those that are marked 0
- once all k locations are marked 1, we reset all to 0

e.g.



$k=3$

Cost = 7

Thm. No randomized Alg can have competitive ratio $o(\log K)$.