

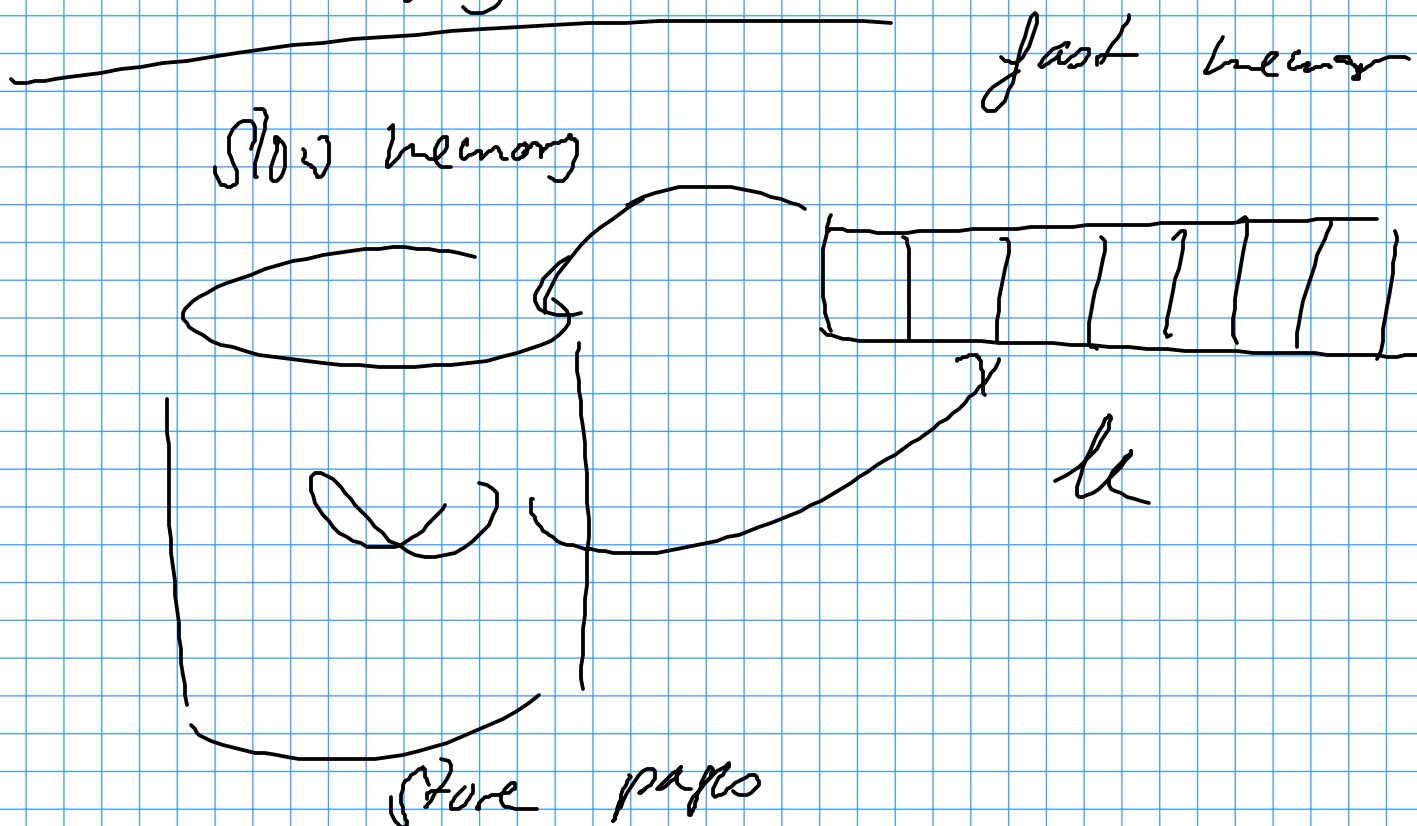
ONLINE ALGORITHMS

request sequence

$$\sigma = \sigma(1), \sigma(2), \dots, \sigma(n)$$

comes in an online fashion, nothing is known about future requests, have to find strategies which perform good compared to an optimal strategy with 'a look in the future'

The Paging Problem



if a page is in fast memory at the time of request, the request is served immediately, otherwise:

- select page in fast memory
- evict it (put it in slow memory)
- load the desired page in fast memory

LRU

- Last Recently Used - with the page that was accessed last of all pages in actual fast memory

• Least Frequently Used - evict the page in fast memory which has the lowest access count since it was loaded there

- Random

- FIFO - first in first out

Theorem: LRU and FIFO are h -competitive.

OPTIMAL STRATEGY 'FUR'

⇒ On a fault, evict the page whose next request occurs furthest in the future.

(Randomized alg. can achieve $2/k$ competitiveness, while det. LB says no better than k -comp possible)

Proof of Theorem for LRU:

Consider an arbitrary sequence of requests $\sigma = \sigma(1), \sigma(2), \dots, \sigma(n)$.




We will prove that $C_{LRU}(\sigma) \leq k \cdot C_{OPT}(\sigma)$.

W.l.o.g. we assume the fast memory to contain the same pages at the start of the algorithms.

We partition σ into phases P such that LRU has at most k faults in P and exactly k faults in

every subsequent P_i . In the remainder we will show that the optimal strategy will have at least one fault in every phase.


P_0 : The first fault for LRU implies that a page is requested which is not in fast memory. But as LRU and MIN start with the same pages there, this also implies a fault for MIN.



P_i : Let $\sigma(t_i)$ be the first request in P_i and $\sigma(t_{i+n-1})$ the last. Furthermore let p be the last page requested in P_{i+n} .

Want to show: P_i contains requests to k distinct pages that are different from p .

if this is true, the MIN has to have a fault in P_i . (Because at the end of P_{i-1} the fast memory of MIN contains p for sure and any $k-1$ other pages, so k other requests force a fault).



Suppose for contradiction that LRU faults twice on the same page q . Assume that LRU has a fault on $\sigma(t) = q$ and $\sigma(t') = q$ with $t_i \leq t < t' \leq t_{i+n-1}$

Page q is in fast memory right after $\sigma(t)$ is served.

To create a fault at t' it needs to be evicted from fast memory at some time

t^* with $t < t^* < t'$.

When q is evicted, it is the least recently used requested page in fast memory.

Therefore the subsequence $\sigma(t) \dots \sigma(t^*)$ has to contain t distinct requests to pages of which q must be different from q .

Finally suppose that within p_i LRU does not fault twice on a page but on one of the faults page p was requested. Let $t \geq t_i$ be the first time that p is evicted. Using the same argument as above, we obtain a request sequence $\sigma(t_i - 1), \sigma(t_i), \dots, \sigma(t)$

which must contain $k-1$
distinct pages. //

\Rightarrow LRU is k -competitive
(and FIFO as well with very
similar arguments)

Theorem: Let A be some deterministic
online paging algorithm.

If A is C -competitive,
then $C \geq k$.

Proof: Let $S = \{p_1, p_2, \dots, p_{k+1}\}$ be
a set of $k+1$ arbitrary pages.

We assume w.l.o.g. that
 A and OPT initially have
 p_1, \dots, p_k in fast memory.

Consider the following request
sequence: Request always the
page that is not in A 's memory
right now.

The algorithm A has a page fault on any request. Suppose that MIN has a fault on some request $\sigma(\epsilon)$, then it lacks a page that is not requested during the next $k-1$ steps. Thus, on any k consecutive requests, MIN has at most one fault.

From a practical point of view, LRU and FIFO get typically better if k increases. This is exactly the opposite of what's shown in the theoretical analysis.

Randomized online algorithm

MARKING: The algorithm processes requests in phases, and marks some of them to remember accesses.


→ Whenever a page is requested, it gets marked (at the beginning of a phase everything is unmarked).

→ On a fault 1 a page is chosen u.a.v. from the unmarked elements to be evicted.

→ if all elements in fast memory are marked, start a new phase
phase i phase $i+1$



→



11, 7, 19, 5, 19, 100

and unmark all elements.

Theorem: MARKING IS $2H_n$
 $\approx 2 \ln n$
competitive.

$$H_n = \sum_{i=1}^n \frac{1}{i} \quad \ln n \leq H_n \leq \ln n + 1$$

Proof: Given a request sequence

$\sigma = \sigma(1), \dots, \sigma(n)$. w.l.o.g.

we assume that MARKING al-



red dy facts on $O(N)$.

MARKING divides requests in phases. A phase starting with $\sigma(i)$ ends with $\sigma(j)$ where $j \geq i$ and j is the smallest index such that the set


$$\{\sigma(i), \sigma(i+1), \dots, \sigma(j)\}$$

contains $n+1$ distinct pages.

(As at the end of the phase all elements in fast memory need to be marked)



Consider an arbitrary phase P_i



Call a page **State** if it is unmarked but was marked in P_{i-1} . Call a page **Clean** if it is neither marked nor State.

\Rightarrow going to show:

Let c be the # of clean pages requested in the phase, then:

- the amortized # faults made by MIN is at least $\frac{C}{2}$.

- the expected number of faults made by MARKING is at most CH_e .

The two statements imply the theorem as $\frac{CH_e}{\frac{C}{2}} = 2H_e$.

MIN

7	10	100	9	8
---	----	-----	---	---

MARKING

5	100	7	27	1000
---	-----	---	----	------

We first analyze MIN. Let S_{opt} be the set of pages contained in the fast memory, S_M be the ones in the MARKING phase.

$$d_i := |S_{opt} \setminus S_M|$$

$$d_i = 5 - 2 = 3$$

$$d_F := |S_{opt} \setminus S_M|$$

number of elements in S_{opt} but not in S_M at the beginning of a phase

↘ at the end

MIN has at least $C \cdot d_F$ faults during the phase.

(Because at least $c - d_I$ of c clean rays are not in the fast memory of MIN).

Also, MIN has at least d_F faults during the phase, because d_F rays requested during the phase are not in MIN's fast memory (at the end of the phase)

We conclude that MIN incurs at least

$$\begin{aligned} & \max \{c - d_I, d_F\} \\ & \geq \frac{1}{2} (c - d_I + d_F) \\ & = \frac{c}{2} - \frac{d_I}{2} + \frac{d_F}{2} \end{aligned}$$

Summing over all phases, the terms $\frac{d_F}{2}$ and $\frac{d_I}{2}$ telescope,

except for the first and last term.

$\frac{d_I}{2}$ is 0 (because $\text{sort} = \text{sort}$)

$$\Rightarrow \frac{c}{2} + \frac{d_F}{2} \geq \frac{c}{2}$$

Next we analyze MARKING.

Serving c requests to clean pages obviously incurs c faults.

There are $s = k - c \leq k - 1$ requests to stall pages. For $i = n - 1, \dots$

we compute the expected cost of the i -th request to a stall page. Let $c(i)$ be the # of

clean pages that were requested in the phase before request # i .


Let $s(i)$ denote the # of stall pages that remain before the request i .

$$s(i) = k - (i - 1)$$


When MARKING serves the i -th request to a stall page,

exactly $s(i) - c(i)$ of the

$s(i)$ stall pages are in fast memory, each of them with



~~Let $s(i)$ denote the # of stall pages that remain before the request i .~~
 state class
 7, 5, 10, 3



20, 5, 10, 3

equal probability.

So the expected cost of the request is

$$\frac{p(i) \cdot c(i)}{s(i)} \cdot 0 + \frac{c(i)}{s(i)} \cdot 1$$

$$\leq \frac{c}{h-i+1}$$

The total expected cost for routing requests to state pages

is $\sum_{i=1}^h \frac{c}{h-i+1} \leq \sum_{i=2}^h \frac{c}{i} = c(H_h - 1)$

In total: cost c for clear pages
cost $(H_h - 1)c$ for state pages

cost $c H_h$

$c H_h$