

Recitation 3: Competitive Analysis

SETUP

input : sequence of requests $R = r_1, r_2, \dots, r_k$

online algorithm \rightarrow processes requests one at a time w/o knowledge of subsequent requests

"myopic"

optimal offline (OPT) algorithm \rightarrow knows entire sequence from the get-go and processes each request optimally w/ this knowledge

Competitive analysis

algorithm A is α -competitive if for any input R

$$C_A(R) \leq \alpha C_{OPT}(R)$$

cost of A on R cost of OPT on R

competitive ratio
(we want α to be small!)

illustrating example: rent or buy?

context : renting ski gear $r = \$50$ per session ($b = 10r$)
 buying ski gear $b = \$500$

ski ≤ 10 times \rightarrow rent gear
 ski ≥ 10 times \rightarrow buy gear

|| but: not sure how many times I'll ski!
 do I rent or buy my gear?

strategy 1: buy at start \rightarrow terrible if I ski only once
 (renting is waaaaay better)

strategy 2: always rent! \rightarrow bad if I end up skiing a TON ($k \gg 10$)

$\alpha = 10$

strategy 3: rent first $\lceil b/r \rceil - 1 = g$ times and buy on the next visit after that ("~~better-late-than-never~~" strategy)

claim: this strategy is 2-competitive!

- if I end up skiing $k \leq \lceil b/r \rceil - 1 = g$ times, then I am optimal — wahoo!
- if $k > \lceil b/r \rceil - 1 = g$, then I should have bought gear right at the start; $OPT = b$.

worst case: I buy on my $\lceil b/r \rceil$ th visit and I never ski again

$$\alpha = \frac{r(\lceil b/r \rceil - 1) + b}{r} \quad \begin{array}{l} \leftarrow \text{my cost} \\ \leftarrow \text{OPT cost} \end{array}$$

< 2 (given b is multiple of r)

claim: this strategy gives the best possible competitive ratio (of all deterministic algorithms)

Worked example: competitive scheduling

setup: n identical machines M_1, \dots, M_n that can process jobs

input: sequence of jobs $\sigma = J_1, \dots, J_k$ arriving all at once ($t = 0$)

J_i has processing time p_i

goal: schedule jobs on machine s.t. the time at which the last job finishes (i.e., the "makespan") is minimized

Our strategy: always assign incoming job to the least loaded machine

claim: this strategy is 2-competitive

notation:

- $T_G(\sigma)$ makespan of our greedy approach
- $T_{OPT}(\sigma)$ makespan of the optimal schedule
- p_{max} processing time of the longest job
- t_i time when some machine finishes (i.e., all other machines run for $\geq t_i$ time)

Observations:

(B1) $T_{OPT}(W) \geq p_{max}$ "total time \geq time required for the longest job"

(B2) $T_{OPT}(W) \geq \frac{1}{n} \sum_{i=1}^k p_i$ "best possible schedule = distribute work completely evenly across machines"


total work up to t_1

total "work" over interval $[0, t_1]$:

$$n \cdot t_1 \leq \sum_{i=1}^k p_i$$

total work over all time

and so $t_1 \leq \frac{1}{n} \sum_{i=1}^k p_i$



Let

- $t_2 = T_G(W) - t_1$ "the amount of extra work done by the busiest machine compared to the least busy one"
- t_L : start time of LAST job

claim: $t_L \leq t_1$

by contradiction: suppose $t_L > t_1$.

\Rightarrow greedy approach assigns LAST to least busy machine
all machines busy until t_1

violates definition of t_1 \Downarrow

putting it all together:

$$P_{LAST} = T_G(W) - t_L \geq T_G(W) - t_1 = t_2$$

Since $P_{LAST} \leq p_{max}$, $t_2 \leq p_{max}$ by claim

$$T_G(W) = t_1 + t_2 \leq \underbrace{\frac{1}{n} \sum_{i=1}^k p_i}_{\leq T_{OPT}(W) \text{ (B1)}} + \underbrace{p_{max}}_{\leq T_{OPT}(W) \text{ (B2)}} \leq 2T_{OPT}(W)$$

as desired.



Worked example: LRU paging

least-recently-used (LRU) : cache w/ k slots stores the most recently requested k pages

hit if user requests for a page in the cache, fault, otherwise

claim: LRU is k -competitive
(i.e., OPT faults ≥ 1 times every time LRU faults k times)

Key insight: (pigeonhole principle)

LRU has k faults
 $\Rightarrow \geq k+1$ distinct pages requested
 \Rightarrow OPT has at least 1 fault