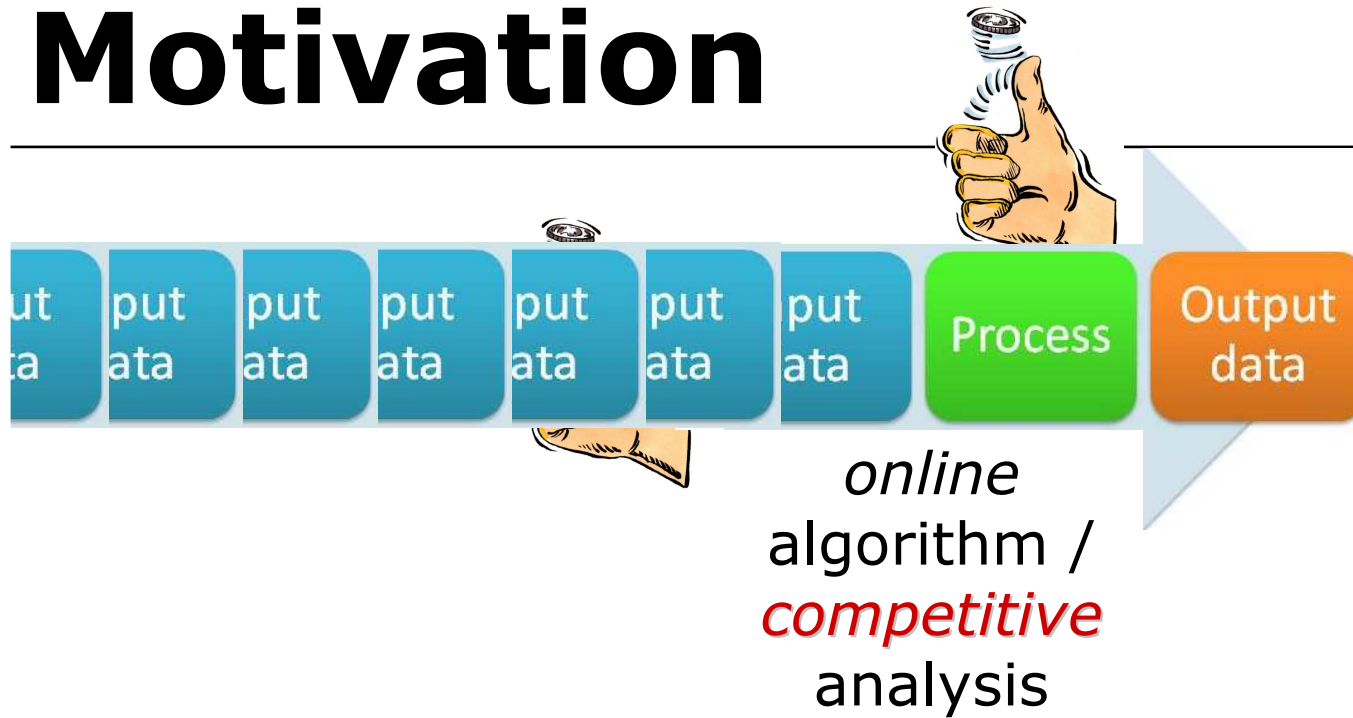


§6 Competitive Analysis of Online Algorithms

- Motivation: Ski Rental
 - *Break-Even* Algorithm
 - is 2-competitive; optimality
- Online Paging
 - *Least-Recently Used* is k -competitive
 - *Least-Frequently Used* is not competitive
 - LRU is optimal among deterministic online
- Randomization and expected competitiveness
 - 1.84-competitive randomized Ski Rental

Motivation



How much better could *omniscient* **offline** algorithm perform?



Ski Rental Problem

Design & Analysis
of Algorithms
Martin Ziegler

Day by day: weather, decision:
(i) **Rent** at \$1 for another day
or (ii) **buy once** for $\$D > 1$

Offline algorithm "knows":
Season will last for L days.
Rent if $L \leq D$, **buy** if $L > D$.
Offline cost = $\min(L, D)$.

Fix **any** ^{deterministic} *online* algorithm \mathcal{A} .

Run on season $L \rightarrow \infty$. **Let** season last $X+1$ days; rerun.
 \mathcal{A} eventually **buys**, $X := \# \text{days rented}$.
$$\frac{\text{Online cost}}{\text{Offline cost}} = \frac{X+D}{\min(X+1, D)} \geq 2-1/D$$

$$\frac{\min(L, 2D-1)}{\min(L, D)} \leq 2-1/D \quad (\text{proof?})$$

Breakeven is
 $(2-1/D)$ -competitive!

Optimal ???

adversary
argument

Breakeven algorithm:
Rent $D-1$ days, then **buy**.

$$\text{Online cost: } L \quad \text{if } L < D, \\ 2D-1 \quad \text{if } L \geq D.$$


Competitive ratio = online cost / offline cost

Online Paging

k pages of *fast* memory, caching $K \gg k$ *slow* pages.

For any sequence $\underline{a} = a_1, \dots, a_N \in \{1, \dots, K\}$ of accesses, minimize the number of cache misses/load/evictions.

- LRU
- LFU
- FIFO
- MIN




Input revealed gradually; *online* algorithm must make decisions with partial knowledge.
Analyze *online* algorithm's output in comparison to optimal *offline* algorithm: **competitive ratio**.

Online Paging: LRU

Theorem: LRU has competitive ratio $k = \# \text{pages}$

Proof: Compare LRU to optimal offline algorithm \mathcal{A} , started with *same* initial cache contents.



Divide $1, \dots, N$ into *rounds* $1 < t_0 < t_1 < \dots < t_M = N$ s.t. LRU incurs precisely k faults in $(t_{m-1} \dots t_m]$ and $[1 \dots k]$ faults in $[1 \dots t_1]$.

In each round, $\geq k+1$ pages get accessed;

fault

hence \mathcal{A} incurs at least 1 page fault!

*Whenever a new page is accessed, evict the one **Least Recently Used**.*

Analyze *online* algorithm output in comparison to the *offline* optimum: **competitive ratio**.

Online Paging: LFU

Theorem: LFU has no (finite) competitive ratio!

Proof: Compare LFU to the optimal online algorithm on the following access sequence for $K=k+1$:

repeat $(m+1)$ -times

$2, 3, \dots, k, 2, 3, \dots, k, 2, 3, \dots, k, \dots, 2, 3, \dots, k,$

$1, k+1, 1, k+1, 1, k+1, \dots, 1, k+1$

repeat
 m -times

fault

*Whenever a new page is accessed,
evict the one **Least Frequently Used**.*

Analyze *online* algorithm output in comparison to the *offline* optimum: **competitive ratio**.

Optimality in Online Paging

Theorem: Every deterministic online algorithm \mathcal{A} has competitive ratio $\geq k = \# \text{pages}$

Proof: Let $K > k$. Simulate \mathcal{A} on initial access sequence $\underline{a} = (1, 2, \dots, k)$.


"adversary"

Pigeonhole: choose $a_{n+1} \in \{1, \dots, K\}$ not in \mathcal{A} 's cache after serving (a_1, a_2, \dots, a_n) .

\mathcal{A} faults on every request!

Pigeonhole: Omniscient *offline* algorithm \mathcal{B} (can) choose to evict a page not to be accessed in the next $k-1$ steps.

\mathcal{B} faults only every k -th request!



$\dots a_1, a_2, a_3, \dots, a_k, a_{k+1}, \dots$

Analyze *online* algorithm output in comparison to the *offline* optimum: **competitive ratio**.

Randomized Online Algorithm

Each morning: (i) Rent at \$1 for another day
or (ii) buy once for $\$D > 1$

Break-even is $(2 - 1/D)$ -competitive, and best possible:

Fix any algorithm \mathcal{A} , run on ∞ season, let $X = \#$ days it rents before buying. Restart, abort season on day $\#X+1$.

adversary

Randomized Ski Rental: Flip a fair coin.

Head: Break-even (\approx rent for D days, then buy)

Tail: Rent for $\frac{2}{3}D$ days, then buy.

$$L \geq D: \mathbb{E}[\text{cost}] = \frac{1}{2} \cdot (2D) + \frac{1}{2} \cdot (\frac{2}{3} + 1) \cdot D = \frac{11}{6} \cdot D$$

1.833

$$\frac{2}{3}D \leq L < D: \mathbb{E}[\text{cost}] = \frac{1}{2} \cdot (L) + \frac{1}{2} \cdot (\frac{2}{3} + 1) \cdot D \leq (\frac{1}{2} + \frac{1}{2} \cdot (\frac{2}{3} + 1)^{2/3}) \cdot L$$

$$L < \frac{2}{3}D: \mathbb{E}[\text{cost}] = \frac{1}{2} \cdot (L) + \frac{1}{2} \cdot (L) = L$$

1.75

"Randomization can beat an adversary!"

§6 Summary

- Motivation: Ski Rental
 - *Break-Even* Algorithm
 - is 2-competitive; optimality
- Online Paging
 - *Least-Recently Used* is k -competitive
 - *Least-Frequently Used* is not competitive
 - LRU is optimal among deterministic online
- Randomization and expected competitiveness
 - 1.84-competitive randomized Ski Rental