

Today we'll be looking at finding approximately-optimal solutions for problems where the difficulty is not that the problem is necessarily *computationally* hard but rather that the algorithm doesn't have all the *information* it needs to solve the problem up front.

Specifically, we will be looking at *online algorithms*, which are algorithms for settings where inputs or data is arriving over time, and we need to make decisions on the fly, without knowing what will happen in the future. This is as opposed to standard problems like sorting where you have all inputs at the start. Data structures are one example of online algorithms (they need to handle sequences of requests, and to do well without knowing in advance which requests will be arriving in the future). We'll talk about some other kinds of examples today.

1 Rent or buy?

Here is a simple online problem that captures a common issue in online decision-making, called the rent-or-buy problem. Say you are just starting to go skiing. You can either rent skis for \$50 or buy them for \$500. You don't know if you'll enjoy skiing, so you decide to rent. Then you decide to go again, and again, and after a while you realize you have shelled out a lot of money renting and you wish you had bought right at the start.¹ The optimal strategy is: if you know you're going to end up skiing more than 10 times, you should buy right at the beginning. If you know you're going to go fewer than 10 times, you should just always rent. (If you know you're going to go exactly 10 then either way is equally good.) But, what if you don't know?

To talk about the quality of an online algorithm, we will look at what's called its *competitive ratio*:

Definition 1 The competitive ratio of an online algorithm ALG is the worst case (i.e., maximum) over possible futures σ of the ratio:

$$\frac{ALG(\sigma)}{OPT(\sigma)},$$

where $ALG(\sigma)$ represents the cost of ALG on σ and $OPT(\sigma)$ is the least possible cost on σ .

E.g., what is competitive ratio of the algorithm that says "buy right away"? The worst case is we only go skiing once. Here the ratio is $500/50 = 10$.

What about the algorithm that says "Rent forever"? Now the worst case is that we keep going skiing. So the competitive ratio of this algorithm is unbounded.

Here's a nice strategy: rent until you realize you should have bought, then buy. (In our case: rent 9 times, then buy). Let's call this algorithm *better-late-than-never*. Formally, if the rental cost is r and the purchase cost is p then the algorithm is to rent $\lceil p/r \rceil - 1$ times and then buy.

Theorem 2 The algorithm *better-late-than-never* has competitive ratio ≤ 2 . If the purchase cost p is an integer multiple of the rental cost r , then the competitive ratio is $2 - r/p$.

Proof: We consider two cases. **Case 1:** if you went skiing fewer than $\lceil p/r \rceil$ times (e.g., 9 or fewer times in the case of $p = 500, r = 50$) then you are optimal. The algorithm never purchased and OPT doesn't purchase either. **Case 2:** If you went skiing $\lceil p/r \rceil$ or more times, then the optimal

¹We are ignoring practical issues such as the type of ski you want depending on your ability level, etc.

solution would have been to buy at the start, so $\text{OPT} = p$. The algorithm paid $r(\lceil p/r \rceil - 1) + p$ (e.g., \$450 + \$500 in our specific case). This is always less than $2p$, and equals $2p - r$ if p is a multiple of r . In Case 1, the ratio of the algorithm's cost to OPT was 1, and in Case 2, the ratio of the algorithm's cost to OPT was less than 2 ($(2p - r)/p = 2 - r/p$ if p was a multiple of r). The worst of these is Case 2, and gives the claimed competitive ratio. ■

Theorem 3 *Algorithm better-late-than-never has the best possible competitive ratio for the ski-rental problem for deterministic algorithms when p is a multiple of r .*

Proof: Consider the event that the day you purchase is the last day you go skiing (this is a legitimate event, since (a) if the algorithm never purchases, we already know its competitive ratio is unbounded, so we may assume a purchase occurs, and (b) the algorithm is deterministic so this occurs after some specific number of rentals). Now, if you rent longer than *better-late-than-never*, then the numerator in Case 2 goes up (the algorithm's cost is larger) but the denominator stays the same, so your ratio is strictly worse. If you rent fewer times (say you rent k fewer times than *better-late-than-never* for some $k \geq 1$), then the numerator in Case 1 goes down by kr but so does the denominator, so again the ratio is worse. ■

2 The elevator problem

You go up to the elevator and press the button. But who knows how long it's going to take to come, if ever? How long should you wait until you give up and take the stairs?

Say it takes time E to get to your floor by elevator (once it comes) and it takes time S by stairs. E.g, maybe $E = 15$ sec, and $S = 45$ sec. How long should you wait until you give up? What strategy has the best competitive ratio?

Answer: wait 30 sec, then take the stairs (in general, wait for $S - E$ time). This is exactly the *better-late-than-never* strategy since we are taking the stairs once we realize we should have taken them at the start. If elevator comes in less than 30 sec, we're optimal. Otherwise, $\text{OPT} = 45$. We took 30+45 sec, so the ratio is $(30 + 45)/45 = 5/3$. Or, in general, the ratio is $(S - E + S)/S = 2 - E/S$.

You may have noticed this is really the *same* as rent-or-buy where stairs=buy, waiting for E time steps is like renting, and the elevator arriving is like the last time you ever ski. So, this algorithm is optimal for the same reason. Other problems like this: whether it's worth optimizing code, when your laptop should stop spinning the disk between accesses, and many others.

When $E \ll S$, this is very close to being 2-competitive. As you saw in HW#4, you can do better by randomization. Indeed, even in the case where $E \ll S$, you can get close to $\frac{e}{e-1}$ -competitiveness using the zero-sum game approach from the problem of waiting for the bus.

3 An aside

Interesting article in NYT Sept 29, 2007: Talking about a book by Jason Zweig on how people's emotions affect their investing, called "Your money and your brain":

"There is a story in the book about Harry Markowitz, Mr. Zweig said the other day. He was referring to the renowned economist who shared a Nobel for helping found modern portfolio theory and proving the importance of diversification.... Mr. Markowitz was then working at the RAND Corporation and trying to figure out how to allocate his retirement account. He knew what he should do: I should have computed the historical

co-variances of the asset classes and drawn an efficient frontier. (That's efficient-market talk for draining as much risk as possible out of his portfolio.)

But, he said, I visualized my grief if the stock market went way up and I wasn't in it or if it went way down and I was completely in it. So I split my contributions 50/50 between stocks and bonds. As Mr. Zweig notes dryly, Mr. Markowitz had proved incapable of applying his breakthrough theory to his own money."

So, he wasn't applying his own theory but he *was* using competitive analysis: 50/50 guarantees you end up with at least half as much as if you had known in advance what would happen, which is best possible Competitive Ratio you can achieve.

4 List Update

This is a nice problem that illustrates some of the ideas one can use to analyze online algorithms. Here are the ground rules for the problem:

- There's a list of n items. (For simplicity we fix n). The positions in the list are numbered $1, 2, \dots, n$. Position 1 is the front of the list. The initial ordering of the items in the list is the same for any algorithm.
- An item x can be accessed. The operation is called **Access**(x). The cost of the operation is the position i of x in the list.
- After doing an access, the algorithm is allowed to rearrange the list by doing swaps of adjacent elements. The cost of a swap is 1.

So an on-line algorithm is specified by describing which swaps are done after an element is accessed. (Without loss of generality we can give off-line optimum algorithm the power to do its swaps any time it wants, and not associate them with any particular access.)

The goal is to devise and analyze an on-line algorithm for doing all the accesses **Access**(σ_1), **Access**(σ_2), **Access**(σ_3), $\dots$ with a small competitive factor.

Here are several algorithms to consider.

- **Do Nothing:** Don't reorder the list.
- **Single Exchange:** After accessing x , if x is not at the front of the list, swap it with its neighbor toward the front.
- **Frequency Count:** Maintain a frequency of access for each item. Keep the list ordered by non-increasing frequency from front to back.
- **Move To Front (MTF):** After an access to an element x , do a series of swaps to move x to the front of the list.

It's easy to construct sample sequences to show that Do Nothing, Single Exchange and Frequency Count have competitive factors that are $\Omega(n)$.

Theorem 4 *MTF is a 4-competitive algorithm for the list-update problem.*

Proof: We'll use the potential function method. There will be a potential function that depends on the state of the MTF algorithm and the state of the “opponent” algorithm B , which can be any algorithm, even one which can see the future. Using this potential, we'll show that the amortized cost to MTF of an access is at most 4 times the cost of that access to B .

What is the potential function Φ ? Define

$$\Phi_t = 2 \cdot (\text{The number of inversions between } B\text{'s list and MTF's list at time } t)$$

Recall that an inversion is a pair of distinct elements (x, y) that appear in one order in B 's list and in a different order in MTF's list. It's a measure of the similarity of the lists.


We can first analyze the amortized cost to MTF of $\text{Access}(x)$ (where it pays for the list traversal and its swaps, but B only does its access). Then we separately analyze the amortized cost to MTF that is incurred when B does any swaps. (Note that in the latter case MTF incurs zero cost, but it will have a non-zero amortized cost, since the potential function may change. To be complete the analysis must take this into account.). In each case we'll show that the amortized cost to MTF (which is the actual cost, plus the increase in the potential) is at most 4 times the cost to B .

For any particular step, let C_{MTF} and C_B be the actual costs of MTF and B on this step, and $AC_{MTF} = C_{MTF} + \Delta\Phi$ be the amortized cost. Here $\Delta\Phi = \Phi_{new} - \Phi_{old}$ is the increase in Φ . Hence observe that $\Delta\Phi$ may be negative, and the amortized cost may be less than the actual cost. We want to show that

$$AC_{MTF} \leq 4 \cdot C_B$$

We can then sum the amortized costs, which would equal the actual cost of the entire sequence of accesses to MTF plus the final potential (non-negative) minus the initial potential (zero). This would be the four times total cost of B , which would give the result.

Analysis of $\text{Access}(x)$. First look at what happens when MTS accesses x and brings it to the front of its list. Say the picture looks like this:



Look at the elements that lie before x in MTF's list, and partition them as follows:

$$S = \{y \mid y \text{ is before } x \text{ in MTF and } y \text{ is before } x \text{ in } B\}$$

$$T = \{y \mid y \text{ is before } x \text{ in MTF and } y \text{ is after } x \text{ in } B\}$$

What is the cost of the access to MTF in terms of these sets?

$$C_{MTF} = 1 + \underbrace{|S| + |T|}_{\text{find cost}} + \underbrace{|S| + |T|}_{\text{swap cost}} = 1 + 2(|S| + |T|).$$

On the other hand, since all of S lies before x in B , the cost of the algorithm B is at least

$$C_B \geq |S| + 1.$$

What happens to the potential as a result of this operation? Well, here's MTF after the operation:

requests whichever page the algorithm *doesn't* have in its fast memory. By design, this will cause the algorithm to have a page fault every time. However, if we knew the future, every time we had a page fault we could always throw out the item whose next request is farthest in the future. Since there are k pages in our fast memory, for one of them, this next request has to be at least k time steps in the future, and since $N = k + 1$, this means we won't have a page fault for at least $k - 1$ more steps (until that one is requested). So, the algorithm that knows the future has a page fault at most once every k steps, and the ratio is k .


Here is a neat *randomized* algorithm with a competitive ratio of $O(\log k)$.

Marking Algorithm:

- The initial state is pages $1, \dots, k$ in fast memory. Start with all pages unmarked.
- When a page is requested,
 - if it's in fast memory already, mark it.
 - if it's not, then throw out a random unmarked page. (If all pages in fast memory are marked, unmark everything first. For analysis purposes, call this the end of a "phase"). Then bring in the page and mark it.

We can think of this as a 1-bit randomized LRU, where marks represent "recently used" vs "not recently used".

The figure below illustrates a phase of the algorithm for $N = 5, k = 4$. Each page (shown as a box) contains the probability that that page is in fast memory. A blue dot indicates that that page is marked. The column on the left shows requested page, and the column on the right shows the expected cost of that access. The phase is comprised of accesses to k distinct pages. (Access to pages that are in the cache with probability 1, i.e., that have probability 0 of incurring a page fault, are not shown.) At the end of the phase, all marks are erased, and the next phase begins.



Theorem 5 When the marking algorithm is run on a sequence σ of accesses, we have:

$$\frac{\mathbb{E}[\text{MARKING}(\sigma)]}{\text{OPT}(\sigma)} \leq \begin{cases} H_k & \text{if } N = k + 1 \\ 2H_k & \text{if } N > k + 1 \end{cases}$$

Note that $H_k := 1 + 1/2 + \dots + 1/k$ is the k th harmonic number. Recall that $H_k \leq 1 + \ln k$.


Proof: We will only show the proof for the special case of $N = k + 1$. For general N , the proof follows similar lines but just is a bit more complicated.

First off all, the expected cost in each phase incurred by the marking algorithm is H_k . This should be clear from the figure above.

To complete the proof we will show the following lemma.

Lemma 6 The cost of the first m phases for OPT is at least m .

Proof: The following figure illustrates a phase. At the beginning (and end) of a phase MARKING is in a known state (no probabilities). And the set of pages that MARKING has in fast memory is either the same or different from that of OPT .



We'll define a potential function which is 1 if the two algorithms have the same pages in fast memory and 0 if these pages differ. This will allow us to do an amortized analysis of OPT during a phase.

The following chart shows the cost of OPT , the change in potential, and the amortized cost of OPT as a function of the equality predicate at the beginning and end of a phase.

start	end	cost to OPT is at least	$\Delta\Phi$	amortized cost to OPT is at least
=	=	1	0	1
=	≠	2	-1	1
≠	=	0	1	1
≠	≠	1	0	1

The fourth and fifth columns of this chart follow from the third. So let's verify that the third column is correct.

If OPT and MARKING start and end in the same state then the very first access costs OPT 1. (All the other accesses can be free if it immediately puts itself in MARKING 's final state.)

Suppose OPT and MARKING start in the same state and end in different states. There's no way that OPT can do this for a cost of only 1. The very first access costs OPT 1. Now if it goes immediately to a state that differs from MARKING 's final state, it will be hit by a costly access at some point during the phase. (The remaining k pages are all accessed during the phase.)

Suppose OPT and MARKING start in different states and end in the same state. This could cost OPT 0 if it starts in MARKING 's final state.

Finally suppose OPT and MARKING start and end in different states. The only way that OPT might possibly get away with 0 cost in this case is if it never moves. But the only page that is not accessed during this phase is the one omitted by MARKING's final state. So the one that OPT attempted to squeak by omitting is hit during the phase and OPT incurs a cost of at least 1.

Recall the usual amortized cost equation:

$$\sum_{1 \leq i \leq m} c_i = \sum_{1 \leq i \leq m} ac_i + \Phi_{initial} - \Phi_{final}$$

where c_i denotes the cost of OPT in the i th phase, and ac_i is its amortized cost. First note $ac_i = 1$ so $\sum_{1 \leq i \leq m} ac_i = m$. Now the initial potential is 1 and the final potential is at most 1. Therefore we have

$$\sum c_i \geq m$$

■
■