

Lecture 17: Online Algorithms & Paging

Rafael Oliveira

University of Waterloo
Cheriton School of Computer Science

rafael.oliveira.teaching@gmail.com

June 19, 2024

Overview

- Part I
 - Why Study Online Algorithms?
 - Competitive Analysis
 - Examples
- Paging & Caching
- Conclusion
- Acknowledgements

Why Study Online Algorithms?

- Online algorithms are important for many applications, when we need to make decisions right when we receive the information.

¹ “Hindsight is 20/20”

Why Study Online Algorithms?

- Online algorithms are important for many applications, when we need to make decisions right when we receive the information.
- Applications in
 - Stock Market
 - Dating
 - Skiing
 - Caching
 - Machine Learning (regret minimization)
 - many more...

¹ “Hindsight is 20/20”

Why Study Online Algorithms?

- Online algorithms are important for many applications, when we need to make decisions right when we receive the information.
- Applications in
 - Stock Market
 - Dating
 - Skiing
 - Caching
 - Machine Learning (regret minimization)
 - many more...
- *Competitive Analysis*: measures performance of our algorithm against best algorithm that could *see into the future* (that is, see the entire input beforehand)¹
 - ① Worst-case analysis

¹ “Hindsight is 20/20”

Different Online Models

We will see other online models in class:

Different Online Models

We will see other online models in class:

- Data Streaming (lecture 19): in this case, we not only receive the input in an online fashion, but we have also *memory constraints*
 - 1 Goal here is to get reasonable (approximate) answers while obeying memory constraints
 - 2 worst-case analysis

Different Online Models

We will see other online models in class:

- Data Streaming (lecture 19): in this case, we not only receive the input in an online fashion, but we have also *memory constraints*
 - 1 Goal here is to get reasonable (approximate) answers while obeying memory constraints
 - 2 worst-case analysis
- Today, we will only see algorithms which must deal with the input as it receives it, *no constraints in memory*.
 - 1 Goal here is to *be competitive* against *any offline algorithm* (that is, algorithms that could see the entire input beforehand)
 - 2 worst-case analysis

Competitive Analysis

- Input is given as a sequence $s = s_1, s_2, \dots, s_n$ of events.

Competitive Analysis

- Input is given as a sequence $s = s_1, s_2, \dots, s_n$ of events.
- Let $C_{opt}(s)$ be the *minimum cost* that *any algorithm* (even one that could look at the *entire input* beforehand) could achieve for input s

Competitive Analysis

- Input is given as a sequence $s = s_1, s_2, \dots, s_n$ of events.
- Let $C_{opt}(s)$ be the *minimum cost* that *any algorithm* (even one that could look at the *entire input* beforehand) could achieve for input s
- Let $C_A(s)$ be the cost of your online algorithm on input s

Competitive Analysis

- Input is given as a sequence $s = s_1, s_2, \dots, s_n$ of events.
- Let $C_{opt}(s)$ be the *minimum cost* that *any algorithm* (even one that could look at the *entire input* beforehand) could achieve for input s
- Let $C_A(s)$ be the cost of your online algorithm on input s

Definition (Deterministic Competitive Ratio)

A deterministic online algorithm A has *competitive ratio* k (aka k -competitive) if for all inputs s , we have:

$$C_A(s) \leq k \cdot C_{opt}(s) + O(1)$$

Competitive Analysis

- Input is given as a sequence $s = s_1, s_2, \dots, s_n$ of events.
- Let $C_{opt}(s)$ be the *minimum cost* that *any algorithm* (even one that could look at the *entire input* beforehand) could achieve for input s
- Let $C_A(s)$ be the cost of your online algorithm on input s

Definition (Deterministic Competitive Ratio)

A deterministic online algorithm A has *competitive ratio* k (aka k -competitive) if for all inputs s , we have:

$$C_A(s) \leq k \cdot C_{opt}(s) + O(1)$$

Definition (Randomized Competitive Ratio)

A randomized online algorithm A has *competitive ratio* k (aka k -competitive) if for all inputs s , we have:

$$\mathbb{E}[C_A(s)] \leq k \cdot C_{opt}(s).$$

- Part I
 - Why Study Online Algorithms?
 - Competitive Analysis
 - Examples
- Paging & Caching
- Conclusion
- Acknowledgements

Ski Rental Problem

- During winter, I am stuck in Canada.

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter
- Winters in Canada are veerrry long... so I may go a bunch of times...

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter
- Winters in Canada are veeerrry long... so I may go a bunch of times...
- Having never done this before, I have to decide whether to buy all the equipment or to rent it at the resort.

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter
- Winters in Canada are veeerrry long... so I may go a bunch of times...
- Having never done this before, I have to decide whether to buy all the equipment or to rent it at the resort.
- Buying the equipment costs me 1k CAD. Renting at the resort costs 100 CAD per day.

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter
- Winters in Canada are veeerrry long... so I may go a bunch of times...
- Having never done this before, I have to decide whether to buy all the equipment or to rent it at the resort.
- Buying the equipment costs me 1k CAD. Renting at the resort costs 100 CAD per day.
- Buy or rent?

Ski Rental Problem

- During winter, I am stuck in Canada.
- So I decided to go Skiing this past winter
- Winters in Canada are veeerrry long... so I may go a bunch of times...
- Having never done this before, I have to decide whether to buy all the equipment or to rent it at the resort.
- Buying the equipment costs me 1k CAD. Renting at the resort costs 100 CAD per day.
- Buy or rent?
- Depends on how many times we will go skiing...

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*
 - ② If we go skiing at least 11 times (and surprise ourselves that we can withstand the cold) then clearly *better to buy*

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*
 - ② If we go skiing at least 11 times (and surprise ourselves that we can withstand the cold) then clearly *better to buy*
 - ③ If we go 10 times, it doesn't matter which way it goes...

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*
 - ② If we go skiing at least 11 times (and surprise ourselves that we can withstand the cold) then clearly *better to buy*
 - ③ If we go 10 times, it doesn't matter which way it goes...
- How is this an online algorithm?

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*
 - ② If we go skiing at least 11 times (and surprise ourselves that we can withstand the cold) then clearly *better to buy*
 - ③ If we go 10 times, it doesn't matter which way it goes...
- How is this an online algorithm?
- Each time we go skiing, we have to decide whether to buy or rent (unless we bought it beforehand)

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- Should we buy or rent?
- Depends on how many times we will go skiing...
 - ① If we go skiing 9 times or less (and we see that we are made for beaches and tropical islands), then clearly *better to rent*
 - ② If we go skiing at least 11 times (and surprise ourselves that we can withstand the cold) then clearly *better to buy*
 - ③ If we go 10 times, it doesn't matter which way it goes...
- How is this an online algorithm?
- Each time we go skiing, we have to decide whether to buy or rent (unless we bought it beforehand)
- Algorithm has to decide *when to buy*, knowing only that we have gone skiing t times

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- A 1.9-competitive algorithm:
 - If $t \leq 9$, then rent
 - When $t = 10$, buy

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- A 1.9-competitive algorithm:
 - If $t \leq 9$, then rent
 - When $t = 10$, buy
- Analysis:
 - If $t \leq 9$, then best strategy is to rent: so cost is:

$$\frac{C_A}{C_{opt}} = \frac{100 \cdot t}{100 \cdot t} = 1$$

Ski Rental Problem

- Buying the equipment costs us 1k CAD. Renting at the resort costs 100 CAD per day.
- A 1.9-competitive algorithm:
 - If $t \leq 9$, then rent
 - When $t = 10$, buy
- Analysis:
 - If $t \leq 9$, then best strategy is to rent: so cost is:

$$\frac{C_A}{C_{opt}} = \frac{100 \cdot t}{100 \cdot t} = 1$$

- If $t \geq 10$, we buy at the 10th time, so cost is:

$$\frac{C_A}{C_{opt}} = \frac{100 \cdot 9 + 1000}{1000} = 1.9$$

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...
- One way to do it: go out with all of them at the same time,⁴ and figure out which one is the best!

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...
- One way to do it: go out with all of them at the same time,⁴ and figure out which one is the best!
- Not possible, due to time constraints and society's value system

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...
- One way to do it: go out with all of them at the same time,⁴ and figure out which one is the best!
- Not possible, due to time constraints and society's value system
- So we have to go out with one of them at a time, and decide whether we want to stay with them or date another person, in which case we must break up

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...
- One way to do it: go out with all of them at the same time,⁴ and figure out which one is the best!
- Not possible, due to time constraints and society's value system
- So we have to go out with one of them at a time, and decide whether we want to stay with them or date another person, in which case we must break up
- Clearly *online setting* (pun intended)

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- In the high-tech life, you decide to join a dating site...
- There are n people that you are interested in dating, and you would like to date the best person² out there.³
- But you don't know who is the best person in advance...
- One way to do it: go out with all of them at the same time,⁴ and figure out which one is the best!
- Not possible, due to time constraints and society's value system
- So we have to go out with one of them at a time, and decide whether we want to stay with them or date another person, in which case we must break up
- Clearly *online setting* (pun intended)
- **Goal:** maximize probability of dating the best person

²Assumptions: people are comparable AND we know how to do it

³Go big or go home lonely!

⁴Also assuming they will all want to date us...

Secretary Dating Problem

- Consider the following algorithm:

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - ① Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - ① Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - ② Pick random order of the n people: call it π

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - 1 Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - 2 Pick random order of the n people: call it π
 - 3 Go out with n/e of them and reject them⁵

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - 1 Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - 2 Pick random order of the n people: call it π
 - 3 Go out with n/e of them and reject them⁵
 - 4 After first n/e dates, you will decide to settle if the person you found is *better* than *anyone else you have dated before*

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - ① Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - ② Pick random order of the n people: call it π
 - ③ Go out with n/e of them and reject them⁵
 - ④ After first n/e dates, you will decide to settle if the person you found is *better* than *anyone else you have dated before*
- This algorithm picks the best person (i.e., the one ranked 1) with probability $\approx 1/e$

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - ① Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - ② Pick random order of the n people: call it π
 - ③ Go out with n/e of them and reject them⁵
 - ④ After first n/e dates, you will decide to settle if the person you found is *better* than *anyone else you have dated before*
- This algorithm picks the best person (i.e., the one ranked 1) with probability $\approx 1/e$
- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- Consider the following algorithm:
 - ① Let's assume that all people you want to date are ranked and associate them with their rank: $1, 2, \dots, n$
 - ② Pick random order of the n people: call it π
 - ③ Go out with n/e of them and reject them⁵
 - ④ After first n/e dates, you will decide to settle if the person you found is *better* than *anyone else you have dated before*
- This algorithm picks the best person (i.e., the one ranked 1) with probability $\approx 1/e$
- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.
- What is the probability that we pick the number 1 in our list?

⁵It's not about them, it's about you... you haven't seen enough, too young to commit, etc.

Secretary Dating Problem

- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.
- What is the probability that we pick the number 1 in our list?

Secretary Dating Problem

- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.
- What is the probability that we pick the number 1 in our list?
- Suppose we pick a person at time k , then want to compute probability

$$P_k = \Pr[\pi(k) = 1 \text{ and we pick person at time } k]$$

Secretary Dating Problem

- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.
- What is the probability that we pick the number 1 in our list?
- Suppose we pick a person at time k , then want to compute probability

$$P_k = \Pr[\pi(k) = 1 \text{ and we pick person at time } k]$$

- Then our final success probability will be $P = \sum_{k>t}^n P_k$

Secretary Dating Problem

- More general algorithm: given a time t , go on t dates and from date $t + 1$ onwards you decide to settle with a person who is better than the previous ones.
- What is the probability that we pick the number 1 in our list?
- Suppose we pick a person at time k , then want to compute probability

$$P_k = \Pr[\pi(k) = 1 \text{ and we pick person at time } k]$$

- Then our final success probability will be $P = \sum_{k>t}^n P_k$
- If $\pi(k) = 1$, then $1 - P_k$ is the probability that we picked a person between $[t + 1, k - 1]$, which means someone in this range better than the first t people.

$$P_k = \Pr[\pi(k) = 1 \text{ and } \min \pi(1), \dots, \pi(k - 1) \text{ is in } \{\pi(1), \dots, \pi(t)\}]$$

Secretary Dating Problem

- Final success probability will be $P = \sum_{k>t}^n P_k$

Secretary Dating Problem

- Final success probability will be $P = \sum_{k>t}^n P_k$
- From previous slide

$$\begin{aligned} P_k &= \Pr[\pi(k) = 1 \text{ and } \min \pi(1), \dots, \pi(k-1) \text{ is in } \{\pi(1), \dots, \pi(t)\}] \\ &= \frac{1}{n} \cdot \frac{t}{k-1} \end{aligned}$$

Secretary Dating Problem

- Final success probability will be $P = \sum_{k>t}^n P_k$
- From previous slide

$$\begin{aligned} P_k &= \Pr[\pi(k) = 1 \text{ and } \min \pi(1), \dots, \pi(k-1) \text{ is in } \{\pi(1), \dots, \pi(t)\}] \\ &= \frac{1}{n} \cdot \frac{t}{k-1} \end{aligned}$$

- We get

$$P = \sum_{k>t}^n \frac{1}{n} \cdot \frac{t}{k-1} = \frac{t}{n} \cdot \sum_{k>t}^n \frac{1}{k-1} \approx \frac{t}{n} \cdot (\ln n - \ln t) = \frac{t}{n} \cdot \ln(n/t)$$

Secretary Dating Problem

- Final success probability will be $P = \sum_{k>t}^n P_k$
- From previous slide

$$\begin{aligned} P_k &= \Pr[\pi(k) = 1 \text{ and } \min \pi(1), \dots, \pi(k-1) \text{ is in } \{\pi(1), \dots, \pi(t)\}] \\ &= \frac{1}{n} \cdot \frac{t}{k-1} \end{aligned}$$

- We get

$$P = \sum_{k>t}^n \frac{1}{n} \cdot \frac{t}{k-1} = \frac{t}{n} \cdot \sum_{k>t}^n \frac{1}{k-1} \approx \frac{t}{n} \cdot (\ln n - \ln t) = \frac{t}{n} \cdot \ln(n/t)$$

- Optimizing we get that we should set $t = n/e$, which gives us $1/e$ probability.

Secretary Dating Problem

- Final success probability will be $P = \sum_{k>t}^n P_k$
- From previous slide

$$\begin{aligned} P_k &= \Pr[\pi(k) = 1 \text{ and } \min \pi(1), \dots, \pi(k-1) \text{ is in } \{\pi(1), \dots, \pi(t)\}] \\ &= \frac{1}{n} \cdot \frac{t}{k-1} \end{aligned}$$

- We get

$$P = \sum_{k>t}^n \frac{1}{n} \cdot \frac{t}{k-1} = \frac{t}{n} \cdot \sum_{k>t}^n \frac{1}{k-1} \approx \frac{t}{n} \cdot (\ln n - \ln t) = \frac{t}{n} \cdot \ln(n/t)$$

- Optimizing we get that we should set $t = n/e$, which gives us $1/e$ probability.
- Wait a second, where is the competitive analysis?

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list
 - Expected rank of our life-long partner is $\Omega(n)$

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list
 - Expected rank of our life-long partner is $\Omega(n)$
- Can we do better?

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list
 - Expected rank of our life-long partner is $\Omega(n)$
- Can we do better?
- Yes! There is an algorithm that picks person of average rank $O(1)$, which is therefore $O(1)$ -competitive.

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list
 - Expected rank of our life-long partner is $\Omega(n)$
- Can we do better?
- Yes! There is an algorithm that picks person of average rank $O(1)$, which is therefore $O(1)$ -competitive.
- Complicated algorithm, based on computing time steps $t_0 \leq t_1 \leq \dots$ and between timesteps t_k and t_{k+1} we are willing to pick person who is $\leq k + 1$ best from our current list.

Making Dating Competitive

- To make the dating problem competitive, we would have to modify it a little bit.
 - We can simply minimize the rank.
 - Say we always want to end up with someone (loneliness has a cost of $-\infty$, after all nobody wants to be alone)
 - Previous algorithm would then either pick the best person, or the last person in the order.
 - With constant probability, rank of the last person is $\Omega(n)$, so we either date the best, or we date someone in the “bottom percentile” of our list
 - Expected rank of our life-long partner is $\Omega(n)$
- Can we do better?
- Yes! There is an algorithm that picks person of average rank $O(1)$, which is therefore $O(1)$ -competitive.
- Complicated algorithm, based on computing time steps $t_0 \leq t_1 \leq \dots$ and between timesteps t_k and t_{k+1} we are willing to pick person who is $\leq k + 1$ best from our current list.
- That is, as we get older, we become more desperate to find someone and lower our expectations...

- Part I
 - Why Study Online Algorithms?
 - Competitive Analysis
 - Examples

- Paging & Caching

- Conclusion

- Acknowledgements

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \leftrightarrow request takes negligible time

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \leftrightarrow request takes negligible time
 - Otherwise we have *miss* \leftrightarrow need to fetch data from slower memory
 - In negligible extra time, can also copy new data & location to cache

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \leftrightarrow request takes negligible time
 - Otherwise we have *miss* \leftrightarrow need to fetch data from slower memory
 - In negligible extra time, can also copy new data & location to cache
 - If cache full, must delete an old entry before copying new data

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \leftrightarrow request takes negligible time
 - Otherwise we have *miss* \leftrightarrow need to fetch data from slower memory
 - In negligible extra time, can also copy new data & location to cache
 - If cache full, must delete an old entry before copying new data
- Main question: which entry of the cache to delete?

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \leftrightarrow request takes negligible time
 - Otherwise we have *miss* \leftrightarrow need to fetch data from slower memory
 - In negligible extra time, can also copy new data & location to cache
 - If cache full, must delete an old entry before copying new data
- Main question: which entry of the cache to delete?
- Cost function: *number of cache misses*

Online Paging Problem

- Computer memory is hierarchical: cache \rightarrow L1 \rightarrow L2 \rightarrow main memory
- Memory can be modelled in the following way:
 - Each layer of memory is an array with certain number of pages (hence the name)
 - Page stores the content of the item and its location in main memory
 - When we get a request (\Leftrightarrow event in online jargon), we first look up in cache, then L1, then L2, then main memory
 - If request is in cache, we have a *hit* \Leftrightarrow request takes negligible time
 - Otherwise we have *miss* \Leftrightarrow need to fetch data from slower memory
 - In negligible extra time, can also copy new data & location to cache
 - If cache full, must delete an old entry before copying new data
- Main question: which entry of the cache to delete?
- Cost function: *number of cache misses*
- Simplification: assume we only have cache and main memory.

Common Heuristics

- 1 **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past

Common Heuristics

- 1 **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past
- 2 **Random**: delete random page.

Common Heuristics

- 1 **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past
- 2 **Random**: delete random page.
- 3 **First-in, First-out (FIFO)**: delete page that has been in cache the *longest*

Common Heuristics

- ① **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past
- ② **Random**: delete random page.
- ③ **First-in, First-out (FIFO)**: delete page that has been in cache the *longest*
- ④ **Least Frequently Used (LFU)**: delete page in cache which has been requested *least often*

Common Heuristics

- 1 **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past
- 2 **Random**: delete random page.
- 3 **First-in, First-out (FIFO)**: delete page that has been in cache the *longest*
- 4 **Least Frequently Used (LFU)**: delete page in cache which has been requested *least often*

Today, we will analyze the **Least Recently Used** heuristic. We will assume that *the size of our cache is k pages*.

Common Heuristics

- 1 **Least Recently Used (LRU)**: delete page in cache whose *most recent request* happened *furthest* in the past
- 2 **Random**: delete random page.
- 3 **First-in, First-out (FIFO)**: delete page that has been in cache the *longest*
- 4 **Least Frequently Used (LFU)**: delete page in cache which has been requested *least often*

Today, we will analyze the **Least Recently Used** heuristic. We will assume that *the size of our cache is k pages*.

- 1 **Least Recently Used (LRU)**: k -competitive
- 2 **Random**: k -competitive
- 3 **First-in, First-out (FIFO)**: k -competitive
- 4 **Least Frequently Used (LFU)**: NOT competitive

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

- 1 Upper bound: divide input sequence into phases.
 - First phase starts immediately after our algorithm first faults, ends right after the algorithm faults k more times
 - Second phase starts at end of first phase, ends when algorithm faults for additional k times
 - and so on...

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

- ① Upper bound: divide input sequence into phases.
 - First phase starts immediately after our algorithm first faults, ends right after the algorithm faults k more times
 - Second phase starts at end of first phase, ends when algorithm faults for additional k times
 - and so on...
- ② We will prove that OPT algorithm faults *at least* once per phase

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

- 1 Upper bound: divide input sequence into phases.
 - First phase starts immediately after our algorithm first faults, ends right after the algorithm faults k more times
 - Second phase starts at end of first phase, ends when algorithm faults for additional k times
 - and so on...
- 2 We will prove that OPT algorithm faults *at least* once per phase
- 3 This gives us that $C_A \leq k \cdot C_{opt}$, which is what we want.

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

- 1 Upper bound: divide input sequence into phases.
 - First phase starts immediately after our algorithm first faults, ends right after the algorithm faults k more times
 - Second phase starts at end of first phase, ends when algorithm faults for additional k times
 - and so on...
- 2 We will prove that OPT algorithm faults *at least* once per phase
- 3 This gives us that $C_A \leq k \cdot C_{opt}$, which is what we want.
- 4 Examples of phases, for $k = 3$:
1, 1, 2, 2, 1, 3, 4, 3, 2, 4, 5, 6, 15, 4, 4, 2, 3, 5, 6, 4,5

LRU Analysis

Theorem

For cache of size k , LRU is k -competitive.

- 1 Upper bound: divide input sequence into phases.
 - First phase starts immediately after our algorithm first faults, ends right after the algorithm faults k more times
 - Second phase starts at end of first phase, ends when algorithm faults for additional k times
 - and so on...
- 2 We will prove that OPT algorithm faults *at least* once per phase
- 3 This gives us that $C_A \leq k \cdot C_{opt}$, which is what we want.
- 4 Examples of phases, for $k = 3$:

1, 1, 2, 2, 1, 3, 4, 3, 2, 4, 5, 6, 15, 4, 4, 2, 3, 5, 6, 4, 5
1 (1, 2, 2, 1, 3, 4) (3, 2, 4, 5, 6) (15, 4, 4, 2) (3, 5, 6) (4, 5

LRU Analysis - Example

Examples of phases, for $k = 3$:

1 (1, 2, 2, 1, 3, 4) (3, 2, 4, 5, 6) (15, 4, 4, 2) (3, 5, 6) (4, 5

LRU Analysis - Upper Bound

- Need to prove that OPT will fault at least once per phase.
- If the same page faulted twice in one phase:

LRU Analysis - Upper Bound

- If each page faulted once in a phase.

LRU Analysis - Upper Bound

- If each page faulted once in a phase.
- **Claim:** in the beginning of each phase, content of OPT and content of our algorithm A intersect in at least one page.
- Proof: Look at last fault page in previous phase.

LRU Analysis - Upper Bound

- If each page faulted once in a phase.
- **Claim:** in the beginning of each phase, content of OPT and content of our algorithm A intersect in at least one page.
- Since OPT and A had a common page, then OPT must have faulted as well (since each page faulted in this phase)

Lower Bound - Deterministic Paging Algorithms

Theorem

Any deterministic algorithm for paging with k pages is at least k -competitive!

- Proof by trolling.⁶ Let's use $k + 1$ pages, and let A be our paging algorithm.

⁶Common lower bound technique for online algorithms, also commonly used online as well :)

Lower Bound - Deterministic Paging Algorithms

Theorem

Any deterministic algorithm for paging with k pages is at least k -competitive!

- Proof by trolling.⁶ Let's use $k + 1$ pages, and let A be our paging algorithm.
- **Input sequence:** at each step, request page that A *doesn't have*.

⁶Common lower bound technique for online algorithms, also commonly used online as well :)

Lower Bound - Deterministic Paging Algorithms

Theorem

Any deterministic algorithm for paging with k pages is at least k -competitive!

- Proof by trolling.⁶ Let's use $k + 1$ pages, and let A be our paging algorithm.
- **Input sequence:** at each step, request page that A *doesn't have*.
- A faults every single time.

⁶Common lower bound technique for online algorithms, also commonly used online as well :)

Lower Bound - Deterministic Paging Algorithms

Theorem

Any deterministic algorithm for paging with k pages is at least k -competitive!

- Proof by trolling.⁶ Let's use $k + 1$ pages, and let A be our paging algorithm.
- **Input sequence:** at each step, request page that A *doesn't have*.
- A faults every single time.
- **Offline Algorithm:** on cache miss, delete page which is requested *furthest in the future*.

⁶Common lower bound technique for online algorithms, also commonly used online as well :)

Lower Bound - Deterministic Paging Algorithms

Theorem

Any deterministic algorithm for paging with k pages is at least k -competitive!

- Proof by trolling.⁶ Let's use $k + 1$ pages, and let A be our paging algorithm.
- **Input sequence:** at each step, request page that A *doesn't have*.
- A faults every single time.
- **Offline Algorithm:** on cache miss, delete page which is requested *furthest in the future*.
- When offline algorithm deletes a page, it's next delete happens after at least k steps.

⁶Common lower bound technique for online algorithms, also commonly used online as well :)

Conclusion

- Online algorithms are important for many applications, when we need to make decisions right when we receive the information.
- Applications in
 - Stock Market
 - Dating
 - Skiing
 - Caching
 - Machine Learning (regret minimization)
 - many more...
- *Competitive Analysis*: measures performance of our algorithm against best algorithm that could *see into the future*

Acknowledgement

- Lecture based largely on:
 - Lecture 17 of Luca's Optimization class
 - Lectures 19 and 20 of Karger's 6.854 Fall 2004 algorithms course
 - [Motwani & Raghavan 2007, Chapter 13]
- See Luca's Lecture 17 notes at
<https://lucatrevisan.github.io/teaching/cs261-11/lecture17.pdf>
- See Karger's Lecture 19 notes at
<http://courses.csail.mit.edu/6.854/06/scribe/s22-online.pdf>
- See Karger's Lecture 20 notes at
<http://courses.csail.mit.edu/6.854/06/scribe/s24-paging.pdf>

References I

-  Motwani, Rajeev and Raghavan, Prabhakar (2007)
Randomized Algorithms