

Lecture 20: Online Algorithms & k -server

Rafael Oliveira

University of Waterloo
Cheriton School of Computer Science

rafael.oliveira.teaching@gmail.com

July 20, 2021

Overview

- Online Algorithms: Randomized Lower Bounds
- k -server on a line
- Conclusion
- Acknowledgements

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:

worst-case \rightarrow *expectation over the randomness of your algorithm*

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

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).

Randomized Online Algorithms & Game Theory


- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests
- Entry (A, s) of payoff matrix: $C_A(s)$



Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests
- Entry (A, s) of payoff matrix: $C_A(s)$
- Algorithm wants to minimize cost $- C_A(s)$
- Adversary wants to maximize it $C_A(s)$

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests
- Entry (A, s) of payoff matrix: $C_A(s)$
- Algorithm wants to minimize cost
- Adversary wants to maximize it
 - Randomized algorithm \Leftrightarrow mixed strategies!

} pure strategies

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests
- Entry (A, s) of payoff matrix: $C_A(s)$
- Algorithm wants to minimize cost
- Adversary wants to maximize it
 - Randomized algorithm \Leftrightarrow mixed strategies!
- As we showed in lecture 12, if one player is using *mixed strategy*, the other player has as best response a *pure strategy*

Randomized Online Algorithms & Game Theory

- Think of online algorithms as being a zero-sum, two-player game between you (the algorithm) and an adversary (the entity choosing the sequence of requests).
- Each of your strategies is a different *deterministic algorithm*
- Each of adversary's strategies is a sequence of requests
- Entry (A, s) of payoff matrix: $C_A(s)$
- Algorithm wants to minimize cost
- Adversary wants to maximize it
 - Randomized algorithm \Leftrightarrow mixed strategies!
- As we showed in lecture 12, if one player is using *mixed strategy*, the other player has as best response a *pure strategy*

LP duality applied to games

Theorem (Yao's minimax principle)

If for some input distribution, no deterministic algorithm is k -competitive, then no randomized algorithm is k -competitive!

randomized algorithm \leftrightarrow mixed strategy for you

probability distribution over the possible sequences of requests \leftrightarrow mixed strategy for adversary
(how to choose sequence of requests at random)

From Lect. 12

if you use randomized algorithm then adversary mixed strat. } competitiveness for rand. alg.

can (and will) use pure strategy as best response worst possible req.

= if adversary uses prob. distribution over requests then you mixed strat. can use det. algorithm

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*
 - Expected number of requests per fault: $k + 1$

Overall $E[C_A(\sigma)] = \frac{n}{k+1}$

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*
 - Expected number of requests per fault: $k + 1$
- 5 Offline Algorithm (OPT)
 - OPT can see the whole input beforehand (still use Farthest in Future)

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*
 - Expected number of requests per fault: $k + 1$
- 5 Offline Algorithm (OPT)
 - OPT can see the whole input beforehand (still use Farthest in Future)
 - Farthest in Future faults only after $k + 1$ distinct pages seen

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution*
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*
 - Expected number of requests per fault: $k + 1$
- 5 Offline Algorithm (OPT)
 - OPT can see the whole input beforehand (still use Farthest in Future)
 - Farthest in Future faults only after $k + 1$ distinct pages seen
 - Expected number of requests per fault:¹ $\Theta(k \log k)$ (see reference)

$$\mathbb{E} [C_{\text{opt}}(s)] = \frac{n}{\Theta(k \log k)}$$

¹Here expectation is over the choice of input.

Lower Bound - Randomized Paging Algorithms

- 1 Setting: $k + 1$ distinct pages, cache of size k , n requests
- 2 Distribution of inputs: *uniform distribution* $\frac{E(CA)}{E(OPT)} = \text{competitive ratio}$
- 3 Equivalently: each page has probability $\frac{1}{k+1}$ of being chosen
- 4 Online Algorithm
 - No matter what our (*fixed*) deterministic algorithm A does, only k pages in cache, with probability $\frac{1}{k+1}$ requested page *not in memory*
 - Expected number of requests per fault: $k + 1$
- 5 Offline Algorithm (OPT)
 - OPT can see the whole input beforehand (still use Farthest in Future)
 - Farthest in Future faults only after $k + 1$ distinct pages seen
 - Expected number of requests per fault:¹ $\Theta(k \log k)$ (see reference)

Theorem

Any randomized algorithm for paging with k pages is $\Omega(\log k)$ -competitive!

¹Here expectation is over the choice of input.

- Online Algorithms: Randomized Lower Bounds
- k -server on a line
- Conclusion
- Acknowledgements

k -server Problem

- **Setup:** we are given a metric space (X, d) .

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.
- Main question: which server to move?

k -server problem

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.
- Main question: which server to move?
- Cost function: *total distance travelled*

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.
- Main question: which server to move?
- Cost function: *total distance travelled*
- Goal: minimize distance travelled

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.
- Main question: which server to move?
- Cost function: *total distance travelled*
- Goal: minimize distance travelled
- Paging is special case of this problem (points of simplex)

Practice problem: prove this!

k -server Problem

- **Setup:** we are given a metric space (X, d) .
- Online algorithm manages k mobile servers, each server is located at a point in X
- A request specifies a point in X , to which a server *must be moved*, unless we already have a server there.
- Main question: which server to move?
- Cost function: *total distance travelled*
- Goal: minimize distance travelled
- Paging is special case of this problem (points of simplex)
- Today's Simplification: assume X is a *line*. Think $X = \mathbb{R}$

Attempt 1: Greedy

- 1 Strategy: just move the server which is closest to the request to it

Attempt 1: Greedy

- ① Strategy: just move the server which is closest to the request to it
- ② Not competitive.

Attempt 1: Greedy

- ① Strategy: just move the server which is closest to the request to it
- ② Not competitive.
- ③ Scenario: two servers A and B , initially located at 0 and 1 respectively

Attempt 1: Greedy


- 1 Strategy: just move the server which is closest to the request to it
- 2 Not competitive.
- 3 Scenario: two servers A and B , initially located at 0 and 1 respectively
- 4 **Requests:** sequence given by $s_{2k-1} = 3/4$, $s_{2k} = 5/4$, for $k \geq 1$

$$s = \frac{3}{4}, \frac{5}{4}, \frac{3}{4}, \frac{5}{4}, \frac{3}{4}, \frac{5}{4}, \dots$$

Attempt 1: Greedy

Over n requests $C_{\text{greedy}}(n) = \frac{1}{4} + \frac{1}{2}(n-1)$

- 1 Strategy: just move the server which is closest to the request to it
- 2 Not competitive.
- 3 Scenario: two servers A and B , initially located at 0 and 1 respectively
- 4 **Requests:** sequence given by $s_{2k-1} = 3/4$, $s_{2k} = 5/4$, for $k \geq 1$
- 5 Only server B will move



Attempt 1: Greedy


- 1 Strategy: just move the server which is closest to the request to it
- 2 Not competitive.
- 3 Scenario: two servers A and B , initially located at 0 and 1 respectively
- 4 **Requests:** sequence given by $s_{2k-1} = 3/4$, $s_{2k} = 5/4$, for $k \geq 1$
- 5 Only server B will move
- 6 Best strategy: put A on $3/4$, B on $5/4$

$$C_{\text{opt}}(s) = \frac{3}{4} + \frac{1}{4} = 1$$

Attempt 2: Double Coverage (DC)


- If request falls between two servers, move both towards request at same rate until one reaches it

(Simplification: never query midpoint between n servers)



Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.



Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

- 1 How to model OPT (offline algorithm)?

Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

- 1 How to model OPT (offline algorithm)?
- 2 Will assume that OPT algorithm moves *exactly one server at a time*.

Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

- 1 How to model OPT (offline algorithm)?
- 2 Will assume that OPT algorithm moves *exactly one server at a time*.
- 3 This is w.l.o.g., because can convert *any offline strategy* into a strategy that moves one server per request, by deferring moves to the future

Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

- 1 How to model OPT (offline algorithm)?
- 2 Will assume that OPT algorithm moves *exactly one server at a time*.
- 3 This is w.l.o.g., because can convert *any offline strategy* into a strategy that moves one server per request, by deferring moves to the future
- 4 How to analyze competitiveness?

Attempt 2: Double Coverage (DC)

- If request falls between two servers, move both towards request at same rate until one reaches it
- Else, just move the closest server to the request.

Theorem

For k servers, Double Coverage is k -competitive.

- 1 How to model OPT (offline algorithm)?
- 2 Will assume that OPT algorithm moves *exactly one server at a time*.
- 3 This is w.l.o.g., because can convert *any offline strategy* into a strategy that moves one server per request, by deferring moves to the future
- 4 How to analyze competitiveness?
- 5 Potential Function:
 - match each server from DC to a server of OPT
 - track changes as requests come

Potential Method - Recap

- In potential method, we have a potential function Φ_t for each time t

Potential Method - Recap

- In potential method, we have a potential function Φ_t for each time t
- Real cost of operation: c_t

Potential Method - Recap

- In potential method, we have a potential function Φ_t for each time t
- Real cost of operation: c_t
- Ammortized cost at time t :

$$\gamma_t = c_t + \underbrace{\Phi_t - \Phi_{t-1}}_{\text{change in potential}}$$

Potential Method - Recap

- In potential method, we have a potential function Φ_t for each time t
- Real cost of operation: c_t
- Ammortized cost at time t :

$$\gamma_t = c_t + \Phi_t - \Phi_{t-1}$$

- Total ammortized cost:

$$\sum_{t=1}^n \gamma_t = \sum_{t=1}^n (c_t + \Phi_t - \Phi_{t-1})$$

$$= \Phi_n - \Phi_0 + \sum_{t=1}^n c_t$$

final potential

initial potential

total true cost

Potential Method - Recap

- In potential method, we have a potential function Φ_t for each time t
- Real cost of operation: c_t
- Ammortized cost at time t :

$$\gamma_t = c_t + \Phi_t - \Phi_{t-1}$$

- Total ammortized cost:

$$\begin{aligned}\sum_{t=1}^n \gamma_t &= \sum_{t=1}^n c_t + \Phi_t - \Phi_{t-1} \\ &= \cancel{\Phi_n} - \Phi_0 + \sum_{t=1}^n c_t \cong -\Phi_0 + \sum_{t=1}^n c_t\end{aligned}$$

- If potential function is always *non-negative*

$$\sum_{t=1}^n c_t \leq \cancel{\Phi_0} + \sum_{t=1}^n \gamma_t$$

$\Phi_t \geq 0 \quad \forall t$
and if $\Phi_0 = 0$

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

$$\text{if } \delta_t = \alpha \cdot C_{\text{opt}}(\text{time } t)$$

$$e_t = C_A(\text{time } t)$$

$$\sum e_t = C_A(\wedge) \leq \alpha \cdot C_{\text{opt}} \\ \sum \delta_t$$

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t

assume they start at the same state

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t
- Let M_t be cost of minimum cost matching between DC's servers and OPT servers
- Let S_t be sum of pairwise distances of DC's servers

initially have all of them at same pt


$$M_0 = 0, \quad S_0 = 0$$

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t
- Let M_t be cost of minimum cost matching between DC's servers and OPT servers
- Let S_t be sum of pairwise distances of DC's servers
- Our potential function will be

$$\Phi_t = k \cdot M_t + S_t$$



DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t
- Let M_t be cost of minimum cost matching between DC's servers and OPT servers
- Let S_t be sum of pairwise distances of DC's servers
- Our potential function will be

$$\Phi_t = k \cdot M_t + S_t$$

- Note that $\Phi_t \geq 0$ at all times

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t
- Let M_t be cost of minimum cost matching between DC's servers and OPT servers
- Let S_t be sum of pairwise distances of DC's servers
- Our potential function will be

$$\Phi_t = k \cdot M_t + S_t$$

- Note that $\Phi_t \geq 0$ at all times
- Use Amortized Analysis to compute amortized cost of DC

DC Analysis - Potential Function

Main idea: have the *amortized cost* per request be (a multiple of) the cost of OPT, while the actual cost is the cost of DC.

- Consider the state of DC and of OPT at time t
- Let M_t be cost of minimum cost matching between DC's servers and OPT servers
- Let S_t be sum of pairwise distances of DC's servers
- Our potential function will be

$$\Phi_t = k \cdot M_t + S_t$$

} Changes if OPT changes or if DC changes

- Note that $\Phi_t \geq 0$ at all times
- Use Amortized Analysis to compute amortized cost of DC
- Break requests into two parts:
 - First account for OPT move
 - Then account for DC move

DC Analysis - Potential Function

1 OPT moves

DC Analysis - Potential Function

① OPT moves


- If OPT moves a distance d , the distance from the moved server to the matched DC's server increases by d

DC Analysis - Potential Function

① OPT moves

- If OPT moves a distance d , the distance from the moved server to the matched DC's server increases by d

- So $M_{t+1} \leq M_t + d$



M_{t+1} old matching now has cost $\leq \underbrace{\text{old cost}}_{M_t} + d$

\Rightarrow there exists a matching of cost $\leq M_t + d$

M_{t+1} is a minimum cost matching

DC Analysis - Potential Function

1 OPT moves

- If OPT moves a distance d , the distance from the moved server to the matched DC's server increases by d
- So $M_{t+1} \leq M_t + d$
- Thus potential increased (so far) by $\Phi_{t+1} - \Phi_t \leq k \cdot d$

$$\Phi_{t+1} - \Phi_t = k(M_{t+1} - M_t) \leq k \cdot d$$

(we still have not accounted for DC moves)

DC Analysis - Potential Function

① OPT moves


- If OPT moves a distance d , the distance from the moved server to the matched DC's server increases by d
- So $M_{t+1} \leq M_t + d$
- Thus potential increased (so far) by $\Phi_{t+1} - \Phi_t \leq k \cdot d$
- Real cost incurred by DC: $c_{t+1} = 0$

DC Analysis - Potential Function

1 OPT moves

- If OPT moves a distance d , the distance from the moved server to the matched DC's server increases by d
- So $M_{t+1} \leq M_t + d$
- Thus potential increased (so far) by $\Phi_{t+1} - \Phi_t \leq k \cdot d$
- Real cost incurred by DC: $c_{t+1} = 0$
- Ammortized cost of DC: $\gamma_{t+1} \leq k \cdot d$

DC Analysis - Potential Function



② DC moves


- ① The request falls between two servers A and B . Say that B is taken to the location requested.

DC Analysis - Potential Function

② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .

DC Analysis - Potential Function



② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)

$$\sum \text{pairwise distances} = \sum_i (\text{distances from } x_i)$$

$$d(A', B') = d(A, B) - 2\delta \quad \text{this does not get canceled}$$

DC Analysis - Potential Function

② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)

DC Analysis - Potential Function

② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ

$$S_{t+1} = S_t - 2\delta$$


DC Analysis - Potential Function

② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)

is sum from OPT

DC Analysis - Potential Function



2 DC moves

- 1 The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move

DC Analysis - Potential Function

② DC moves

- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move
 - $M_{t+1} \leq M_t$

DC Analysis - Potential Function

2 DC moves

- 1 The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move
 - $M_{t+1} \leq M_t$
 - Potential Change: $\Phi_{t+1} - \Phi_t \leq k \cdot 0 - 2 \cdot \delta = -2 \cdot \delta$

$$\leq \underbrace{k(M_{t+1} - M_t)}_{\leq 0} + (S_{t+1} - S_t) = -2\delta$$

DC Analysis - Potential Function

2 DC moves

- 1 The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move
 - $M_{t+1} \leq M_t$
 - Potential Change: $\Phi_{t+1} - \Phi_t \leq k \cdot 0 - 2 \cdot \delta = -2 \cdot \delta$
 - Real cost incurred by DC: $c_{t+1} = 2\delta$

DC Analysis - Potential Function

$$\begin{aligned}\delta_{t+1} &= c_{t+1} + \Delta\Phi \\ &= -2\delta + \Delta\Phi \leq -2\delta\end{aligned}$$

2 DC moves

- 1 The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move
 - $M_{t+1} \leq M_t$
 - Potential Change: $\Phi_{t+1} - \Phi_t \leq k \cdot 0 - 2 \cdot \delta = -2 \cdot \delta$
 - Real cost incurred by DC: $c_{t+1} = 2\delta$
 - Ammortized cost of DC: $\gamma_{t+1} \leq 2\delta - 2\delta = 0$

DC Analysis - Potential Function

② DC moves


- ① The request falls between two servers A and B . Say that B is taken to the location requested.
 - Both servers move a distance δ .
 - Thus pairwise distances decrease by 2δ (because they are in a line)
 - Changes in other pairwise distances cancel out (because line)
 - Thus S decreases by 2δ
 - B has match at destination (problem constraint)
 - A may be further from its match, but balanced by B 's move
 - $M_{t+1} \leq M_t$
 - Potential Change: $\Phi_{t+1} - \Phi_t \leq k \cdot 0 - 2 \cdot \delta = -2 \cdot \delta$
 - Real cost incurred by DC: $c_{t+1} = 2\delta$
 - Ammortized cost of DC: $\gamma_{t+1} \leq 2\delta - 2\delta = 0$
- ② Only one server moves (request outside the border)

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ



DC Analysis - Potential Function


- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$
 - Each pairwise distance (A, B) (where B is another of DC's servers) increases by δ



DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$
 - Each pairwise distance (A, B) (where B is another of DC's servers) increases by δ
 - Total distance increased: $S_{t+1} - S_t \leq (k - 1) \cdot \delta$

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$
 - Each pairwise distance (A, B) (where B is another of DC's servers) increases by δ
 - Total distance increased: $S_{t+1} - S_t \leq (k - 1) \cdot \delta$
 - Change in potential:

$$\Delta\Phi \leq \underbrace{-k \cdot \delta} + \underbrace{(k - 1) \cdot \delta} = -\delta$$

||

$$k(M_{t+1} - M_t) + (S_{t+1} - S_t)$$

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$
 - Each pairwise distance (A, B) (where B is another of DC's servers) increases by δ
 - Total distance increased: $S_{t+1} - S_t \leq (k - 1) \cdot \delta$
 - Change in potential:

$$\Delta\Phi \leq -k \cdot \delta + (k - 1) \cdot \delta = -\delta$$

- Real cost incurred by DC: $c_{t+1} = \delta$

DC Analysis - Potential Function

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server, say A , moves (request outside the border)
 - Suppose A moved δ
 - A has its match (from OPT's server) at destination
 - $M_{t+1} \leq M_t - \delta$
 - Each pairwise distance (A, B) (where B is another of DC's servers) increases by δ
 - Total distance increased: $S_{t+1} - S_t \leq (k - 1) \cdot \delta$
 - Change in potential:

$$\Delta\Phi \leq -k \cdot \delta + (k - 1) \cdot \delta = -\delta$$

- Real cost incurred by DC: $c_{t+1} = \delta$
- Ammortized cost at this step: $\gamma_{t+1} = \leq \delta - \delta = 0$

$$c_t \sim \geq \Delta\Phi$$
$$c_t + \Delta\Phi$$

DC Analysis - Wrapping Up

- ① OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
- ② DC moves
 - ① The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - ② Only one server moves (request outside the border)
 - Ammortized cost at this step: $\gamma_t \leq \delta - \delta = 0$

$$\delta_t \leq 0$$

Total ^{amortized} cost per request is the sum of
the amortized cost from ① (OPT)
plus " " " from ②

$$\therefore \text{at each request } \gamma_t^{\text{bind}} \leq k \cdot d_{\text{opt},t}$$

DC Analysis - Wrapping Up

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
 - 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server moves (request outside the border)
 - Ammortized cost at this step: $\gamma_t \leq \delta - \delta = 0$
- By our potential function inequality, we have:

$$\underbrace{\sum_{t=1}^n c_t}_{C_{DC}} \leq \Phi_0 + \underbrace{\sum_{t=1}^n \gamma_t}_{\text{total amortized cost}} \leq \sum_{t=1}^n k \cdot d_{opt,t}$$

$k \cdot C_{opt}$
||

DC Analysis - Wrapping Up

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
 - 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server moves (request outside the border)
 - Ammortized cost at this step: $\gamma_t \leq \delta - \delta = 0$
- By our potential function inequality, we have:

$$\sum_{t=1}^n c_t \leq \Phi_0 + \sum_{t=1}^n \gamma_t$$

- Since $\gamma_t \leq k \cdot d$ whenever OPT moves d , and $\gamma_t \leq 0$ when OPT doesn't move, we have that $\sum_t \gamma_t \leq k \cdot C_{opt}$

DC Analysis - Wrapping Up

DC is k -competitive

- 1 OPT moves distance d
 - Ammortized cost of DC: $\gamma_t \leq k \cdot d$
 - 2 DC moves
 - 1 The request falls between two servers.
 - Ammortized cost of DC: $\gamma_t \leq 0$
 - 2 Only one server moves (request outside the border)
 - Ammortized cost at this step: $\gamma_t \leq \delta - \delta = 0$
- By our potential function inequality, we have:

$$\sum_{t=1}^n c_t \leq \Phi_0 + \sum_{t=1}^n \gamma_t$$

- Since $\gamma_t \leq k \cdot d$ whenever OPT moves d , and $\gamma_t \leq 0$ when OPT doesn't move, we have that $\sum_t \gamma_t \leq k \cdot C_{opt}$
- Since Φ_0 is the initial state, we can regard it as constant (even 0, if require that servers start at a certain place)

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*
- Saw how to use *minimax theorem* in *Yao's principle* to prove lower bounds for randomized online algorithms.
- combined amortized analysis in the online setting to solve k-server

Acknowledgement

- Lecture based largely on:
 - Lectures 18 & 20 of Karger's 6.854 Fall 2004 algorithms course
 - [Motwani & Raghavan 2007, Chapter 13]
- See Karger's Lecture 18 notes at
<http://courses.csail.mit.edu/6.854/06/scribe/s23-onlineRandomLb.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