

Objectives

- Divide and conquer algorithms
 - Counting inversions
 - Closest pair of points

Mar 2, 2011

CSCI211 - Sprenkle

1

Review

- Describe the template for divide and conquer solutions
- How can you compute D&C running times?
 - Describe first step
 - 2 ways to solve
- What are you looking for when unrolling the recurrence?

Mar 2, 2011

CSCI211 - Sprenkle

2

Review: Divide-and-Conquer

- Divide-and-conquer process
 - Break up problem into several parts
 - Solve each part recursively
 - Combine solutions to sub-problems into overall solution
- Define a **recurrence relation** that describes the running time

Divide et impera.
Veni, vidi, vici.
- Julius Caesar

Mar 2, 2011

CSCI211 - Sprenkle

3

Review: Recurrence Relations

- Use recurrences to analyze/determine the run time of divide and conquer algorithms
 - Number of sub problems
 - Size of sub problems
 - Number of times divided (number of levels)
 - Cost of merging problems
- How to solve
 - Unrolling
 - Substitution

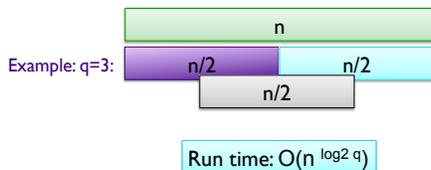
Mar 2, 2011

CSCI211 - Sprenkle

4

Review Example: $q > 2$

- Recurrence relation:
 - For some constant c ,
 - $T(n) \leq q T(n/2) + cn$ when $n > 2$
 - $T(2) \leq c$



Mar 2, 2011

CSCI211 - Sprenkle

5

COUNTING INVERSIONS

Mar 2, 2011

CSCI211 - Sprenkle

6

Problem Context

- Movie site tries to match your movie preferences with others
 - You rank n movies
 - Movie site consults database to find people with similar tastes
 - Collaborative filtering
- Meta-search tools
 - Do same query on several search engines
 - Synthesize results by looking for similarities and differences in search engines' results rankings

Comparing Rankings

- To determine similarity of rankings, need a metric
- Similarity metric: number of inversions between two rankings
 - My rank: $1, 2, \dots, n$
 - Your rank: a_1, a_2, \dots, a_n
 - Movies i and j *inverted* if $i < j$ but $a_i > a_j$

Discuss pros and cons of this metric

	Movies				
	A	B	C	D	E
Me	1	2	3	4	5
You	1	3	4	2	5

What are the inversions?

Comparing Rankings

- To determine similarity of rankings, need a metric
- Similarity metric: number of inversions between two rankings
 - My rank: $1, 2, \dots, n$
 - Your rank: a_1, a_2, \dots, a_n
 - Movies i and j *inverted* if $i < j$ but $a_i > a_j$

Naïve/Brute force solution?

	Movies				
	A	B	C	D	E
Me	1	2	3	4	5
You	1	3	4	2	5

Inversions:
3-2, 4-2

Brute Force Solution

- Look at every pair (i,j) and determine if they are an inversion
- Requires $\Theta(n^2)$ time
 - Note: Already an efficient algorithm but trying to improve upon runtime

Towards a Better Solution...
• Can't look at each inversion individually

Applications

- Voting theory
- Collaborative filtering
- Measuring the "sortedness" of an array
- Sensitivity analysis of Google's ranking function
- Rank aggregation for meta-searching on the Web
- Nonparametric statistics (e.g., Kendall's Tau distance)

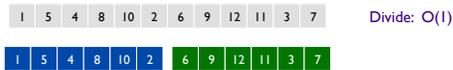
Counting Inversions: Divide-and-Conquer

Assume number represents where item should be in the list

1	5	4	8	10	2	6	9	12	11	3	7
---	---	---	---	----	---	---	---	----	----	---	---

Counting Inversions: Divide-and-Conquer

- **Divide:** separate list into two pieces



What are the inversions in blue and green halves?

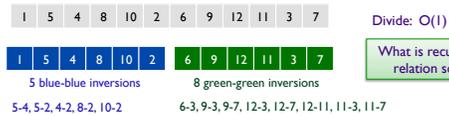
Mar 2, 2011

CSCI211 - Sprenkle

13

Counting Inversions: Divide-and-Conquer

- **Divide:** separate list into two pieces
- **Conquer:** recursively count inversions in each half



What is recurrence relation so far?

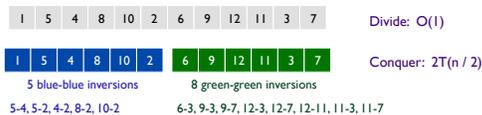
Mar 2, 2011

CSCI211 - Sprenkle

14

Counting Inversions: Divide-and-Conquer

- **Divide:** separate list into two pieces
- **Conquer:** recursively count inversions in each half



What do we need to do next?

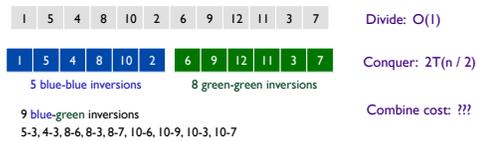
Mar 2, 2011

CSCI211 - Sprenkle

15

Counting Inversions: Divide-and-Conquer

- **Divide:** separate list into two pieces
- **Conquer:** recursively count inversions in each half
- **Combine:** count inversions where a_i and a_j are in different halves, and return sum of three quantities



Combine cost: ???

Total = 5 + 8 + 9 = 22

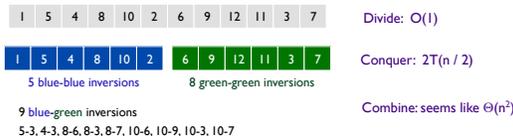
Mar 2, 2011

CSCI211 - Sprenkle

16

Counting Inversions: Divide-and-Conquer

- **Divide:** separate list into two pieces
- **Conquer:** recursively count inversions in each half
- **Combine:** count inversions where a_i and a_j are in different halves, and return sum of three quantities



Total = 5 + 8 + 9 = 22

What would make figuring out blue-green inversions easier?

Mar 2, 2011

CSCI211 - Sprenkle

17

Counting Inversions: Combine

Combine: count blue-green inversions

- > Assume each half is sorted
- > Count inversions where a_i and a_j are in different halves
- > Merge two sorted halves into sorted whole

to maintain sorted invariant



- What does sorting do for us?
- What is our algorithm for counting the inversions and merging?

Mar 2, 2011

CSCI211 - Sprenkle

18

Counting Inversions: Combine

Combine: count blue-green inversions

- Assume each half is sorted
- Count inversions where a_i and a_j are in different halves
- Merge two sorted halves into sorted whole

to maintain sorted invariant



Count: $O(n)$

13 blue-green inversions: $6 + 3 + 2 + 2 + 0 + 0$



Merge: $O(n)$

We'll run through an example in a bit...

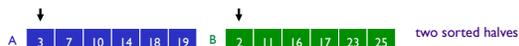
Merge and Count

```

Merge-and-Count(A,B):
  i=0 (front of list A)
  j=0 (front of list B)
  inversions = 0
  output = []
  while A not empty and B not empty:
    output.append( min(A[i], B[j]) )
    if B[j] < A[i]:
      inversions += A.size - i (remaining elements in A)
    update i or j (whichever had smaller element)
    
```

Merge and Count Step

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



two sorted halves



Output array

Total:

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



two sorted halves

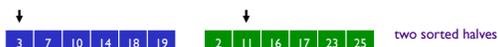


Output array

Total: 6

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



two sorted halves

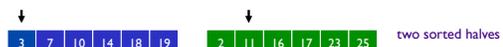


Output array

Total: 6

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



two sorted halves

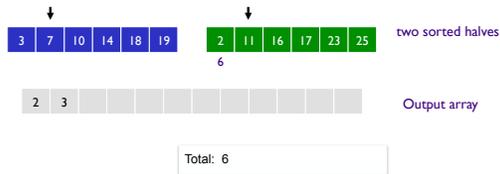


Output array

Total: 6

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



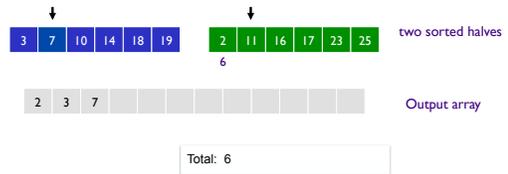
Mar 2, 2011

CSCI211 - Sprenkle

25

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



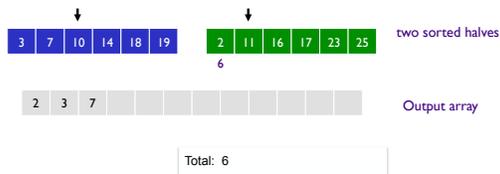
Mar 2, 2011

CSCI211 - Sprenkle

26

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



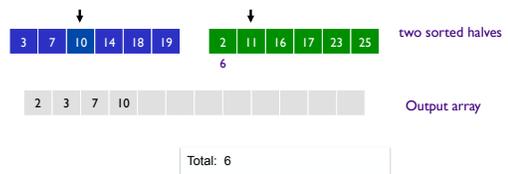
Mar 2, 2011

CSCI211 - Sprenkle

27

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



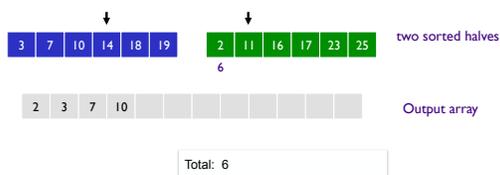
Mar 2, 2011

CSCI211 - Sprenkle

28

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



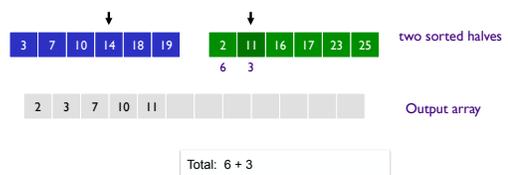
Mar 2, 2011

CSCI211 - Sprenkle

29

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



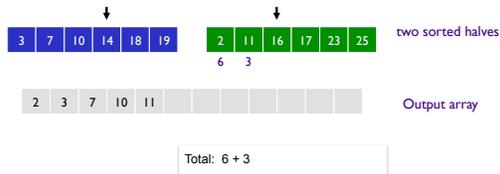
Mar 2, 2011

CSCI211 - Sprenkle

30

Merge and Count

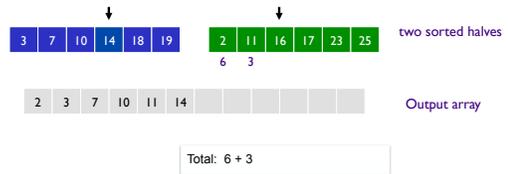
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 31

Merge and Count

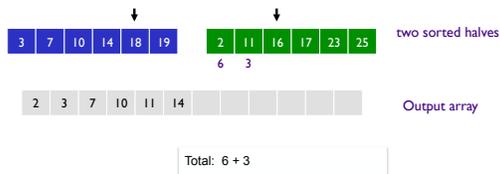
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 32

Merge and Count

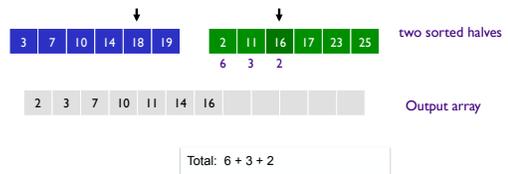
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 33

Merge and Count

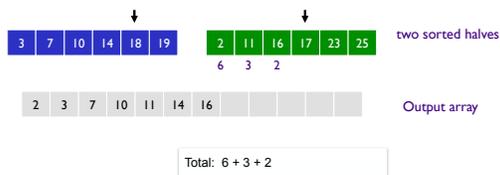
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 34

Merge and Count

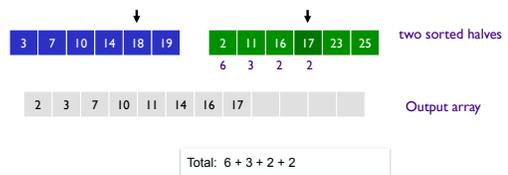
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 35

Merge and Count

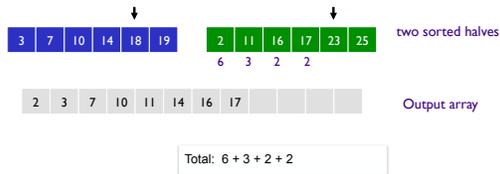
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSC1211 - Sprenkle 36

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



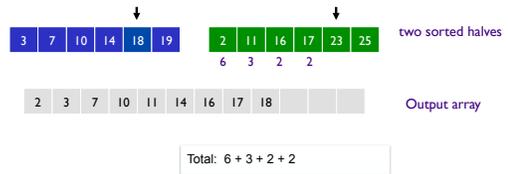
Mar 2, 2011

CSCI211 - Sprenkle

37

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



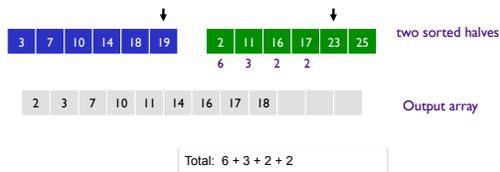
Mar 2, 2011

CSCI211 - Sprenkle

38

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



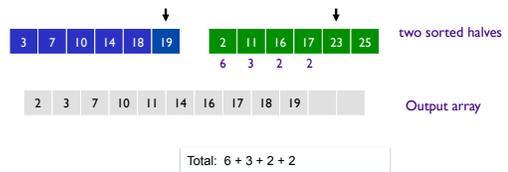
Mar 2, 2011

CSCI211 - Sprenkle

39

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011

CSCI211 - Sprenkle

40

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



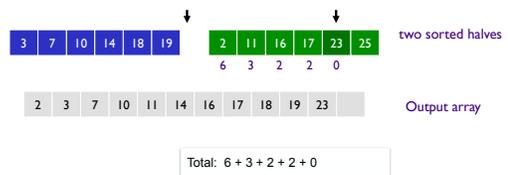
Mar 2, 2011

CSCI211 - Sprenkle

41

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011

CSCI211 - Sprenkle

42

Merge and Count

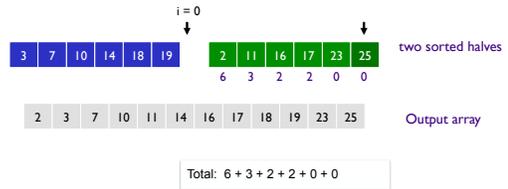
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSCI211 - Sprenkle 43

Merge and Count

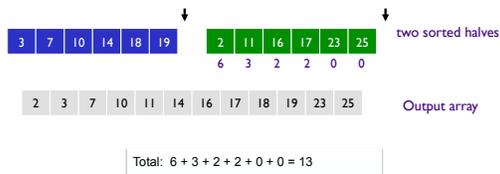
- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSCI211 - Sprenkle 44

Merge and Count

- Given two sorted halves, count number of inversions where a_i and a_j are in different halves
- Combine two sorted halves into sorted whole



Mar 2, 2011 CSCI211 - Sprenkle 45

Counting Inversions: Implementation

- Merge-and-Count Pre-condition. A and B are sorted.
- Sort-and-Count Post-condition. L is sorted.

```
Sort-and-Count(L)
  if list L has one element
    return 0 and the list L
  Divide the list into two halves A and B
  (i_A, A) = Sort-and-Count(A)
  (i_B, B) = Sort-and-Count(B)
  (i, L) = Merge-and-Count(A, B)
  return i = i_A + i_B + i and the sorted list L
```

Recurrence relation?

Mar 2, 2011 CSCI211 - Sprenkle 46

Analysis

Recurrence Relation:

$$T(n) \leq T(n/2) + T(n/2) + O(n)$$

$$\rightarrow T(n) \in O(n \log n)$$

```
Sort-and-Count(L)
  if list L has one element
    return 0 and the list L
  Divide the list into two halves A and B
  (i_A, A) = Sort-and-Count(A) T(n/2)
  (i_B, B) = Sort-and-Count(B) T(n/2)
  (i, L) = Merge-and-Count(A, B) O(n)
  return i = i_A + i_B + i and the sorted list L
```

Mar 2, 2011 CSCI211 - Sprenkle 47

CLOSEST PAIR OF POINTS

Mar 2, 2011 CSCI211 - Sprenkle 48

Computational Geometry

- Algorithms and data structures for geometrical objects
 - Points, line segments, polygons, etc.
 - Common motivator: large data sets → efficiency
- Some Applications
 - Graphics
 - Robotics
 - motion planning and visibility problems
 - Geographic information systems (GIS)
 - geometrical location and search, route planning

Mar 2, 2011

CSCI211 - Spenkle

49

Closest Pair of Points

- Closest pair.** Given n points in the plane, find a pair with smallest Euclidean distance between them.
 - Special case of nearest neighbor, Euclidean MST, Voronoi
- Brute force?

fast closest pair inspired fast algorithms for these problems

Mar 2, 2011

CSCI211 - Spenkle

50

Closest Pair of Points

- Closest pair.** Given n points in the plane, find a pair with smallest Euclidean distance between them.
 - Special case of nearest neighbor, Euclidean MST, Voronoi.
- Brute force.** Check all pairs of points p and q with $\Theta(n^2)$ comparisons

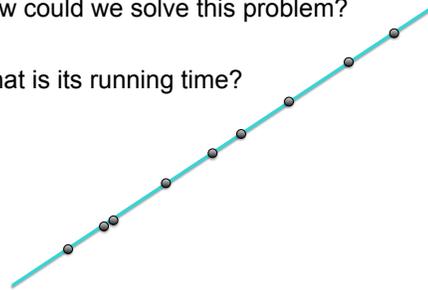
Mar 2, 2011

CSCI211 - Spenkle

51

Simplify: All Points on a Line

- How could we solve this problem?
- What is its running time?



Mar 2, 2011

CSCI211 - Spenkle

52

Simplify: All Points on a Line

- How could we solve this problem?
 - Sort the points
 - Monotonically increasing x/y coordinates
 - No closer points than neighbors in sorted list
 - Step through, looking at the distances between each pair
- What is its running time?
 - $O(n \log n)$

Why won't this work for 2D?

Mar 2, 2011

CSCI211 - Spenkle

53

Closest Pair of Points

- Closest pair.** Given n points in the plane, find a pair with smallest Euclidean distance between them.
 - Special case of nearest neighbor, Euclidean MST, Voronoi.
- Brute force.** Check all pairs of points p and q with $\Theta(n^2)$ comparisons
- 1-D version.** $O(n \log n)$
 - Easy if points are on a line
- Assumption.** No two points have same x coordinate
to make presentation cleaner

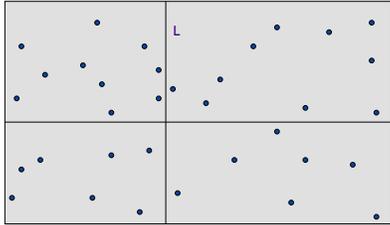
Mar 2, 2011

CSCI211 - Spenkle

54

Closest Pair of Points: First Attempt

- **Divide.** Sub-divide region into 4 quadrants



Any problems with implementing this approach?

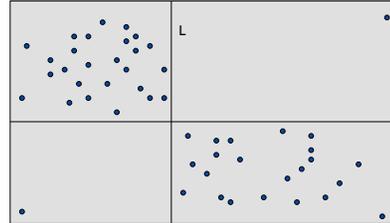
Mar 2, 2011

CSCI211 - Sprenkle

55

Closest Pair of Points: First Attempt

- **Divide.** Sub-divide region into 4 quadrants
- **Obstacle.** Impossible to ensure $n/4$ points in each piece



Mar 2, 2011

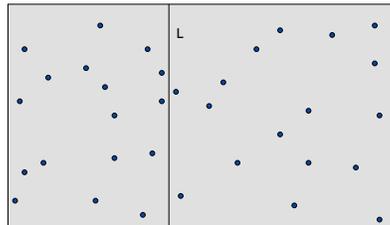
CSCI211 - Sprenkle

56

Closest Pair of Points

- **Divide:** draw vertical line L so that roughly $1/2n$ points on each side

How do we implement this?



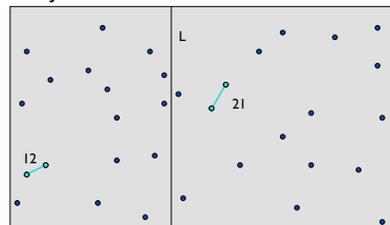
Mar 2, 2011

CSCI211 - Sprenkle

57

Closest Pair of Points

- **Divide:** draw vertical line L so that roughly $1/2n$ points on each side
- **Conquer:** find closest pair in each side recursively



Mar 2, 2011

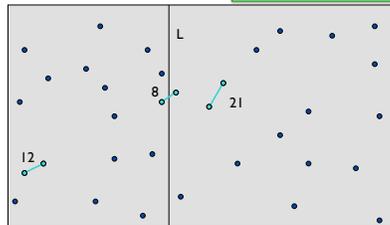
CSCI211 - Sprenkle

58

Closest Pair of Points

- **Divide:** draw vertical line L so that roughly $1/2n$ points on each side
- **Conquer:** find closest pair in each side recursively
- **Combine:** find closest pair with one point in each side *seems like $\Theta(n^2)$*
- Return best of 3 solutions

Do we need to check all pairs?



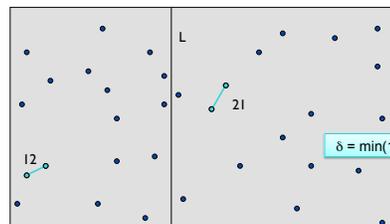
Mar 2, 2011

CSCI211 - Sprenkle

59

Closest Pair of Points

- Find closest pair with one point in each side, assuming that distance $< \delta$
where $\delta = \min(\text{left_min_dist}, \text{right_min_dist})$



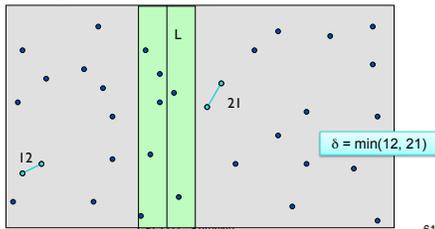
Mar 2, 2011

CSCI211 - Sprenkle

60

Closest Pair of Points

- Find closest pair with one point in each side, assuming that distance $< \delta$.
 - Observation: only need to consider points within δ of line L .



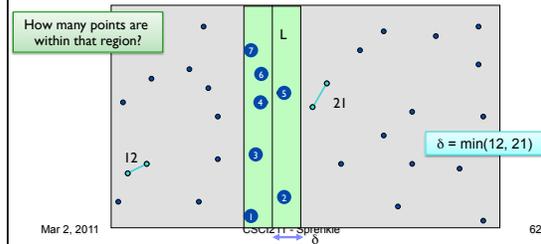
Mar 2, 2011

CSCI211 - Sprenkle

61

Closest Pair of Points

- Find closest pair w/ 1 point in each side, assuming that distance $< \delta$.
 - Observation: only consider points within δ of line L
 - Sort points in 2δ -strip by their y coordinate



Mar 2, 2011

CSCI211 - Sprenkle

62

Assignments

- Continue reading Chapter 5
- PS5 due Friday
- SSA: opportunities for extra credit
 - If you want to attend "Truth Values", let me know so I can give you a ticket (FREE!)
- PS6 due next Friday

Mar 2, 2011

CSCI211 - Sprenkle

63