

Objectives

- Divide and conquer
 - Closest pair of points
 - Integer multiplication
 - Matrix multiplication

Mar 7, 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 7, 2011

CSCI211 - Sprenkle

2

Reviewing Closest Pair of Points

Mar 7, 2011

CSCI211 - Sprenkle

3

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

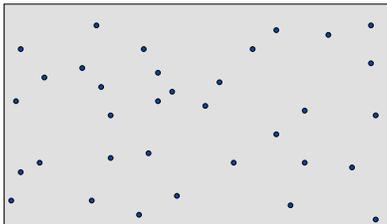
Mar 7, 2011

CSCI211 - Sprenkle

4

Closest Pair of Points

- Recall the approach?



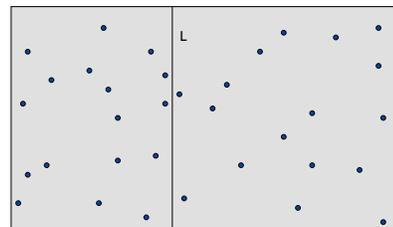
Mar 7, 2011

CSCI211 - Sprenkle

5

Closest Pair of Points

- **Divide:** draw vertical line L so that roughly $\frac{1}{2}n$ points on each side



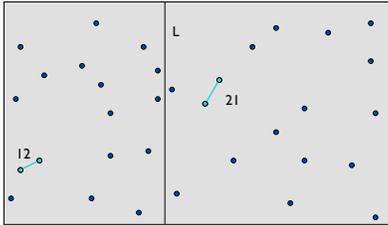
Mar 7, 2011

CSCI211 - Sprenkle

6

Closest Pair of Points

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



Mar 7, 2011

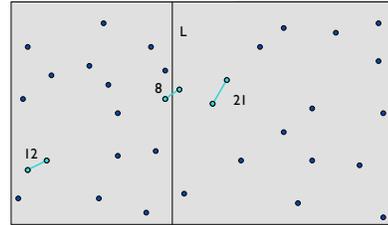
CSCI211 - Sprenkle

7

Closest Pair of Points

- **Divide:** draw vertical line L so that roughly $\frac{1}{2}n$ 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?



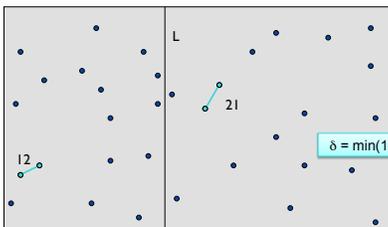
Mar 7, 2011

CSCI211 - Sprenkle

8

Closest Pair of Points

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



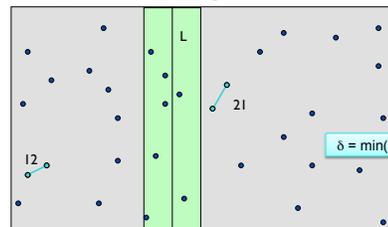
Mar 7, 2011

CSCI211 - Sprenkle

9

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.



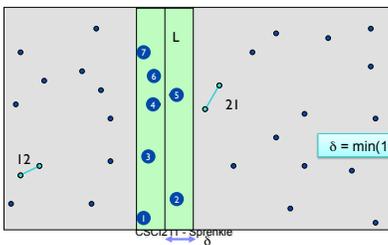
Mar 7, 2011

CSCI211 - Sprenkle

10

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



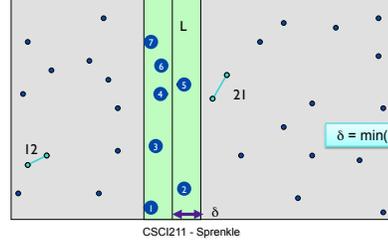
Mar 7, 2011

CSCI211 - Sprenkle

11

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
 • Only checks distances of those within 11 positions in sorted list!



Mar 7, 2011

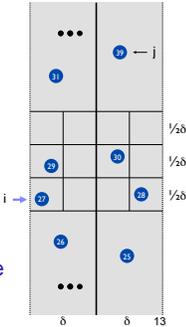
CSCI211 - Sprenkle

12

Analyzing Cost of Combining

Prepare minds to be blown...

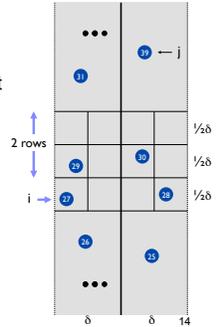
- Def. Let s_i be the point in the 2δ -strip, with the i^{th} smallest y-coordinate
- Claim. If $|i - j| \geq 12$, then the distance between s_i and s_j is at least δ
 - What is the distance of the box?
 - How many points can be in a box?
 - When do we know that points are $> \delta$ apart?



Analyzing Cost of Combining

- Def. Let s_i be the point in the 2δ -strip, with the i^{th} smallest y-coordinate
- Claim. If $|i - j| \geq 12$, then the distance between s_i and s_j is at least δ
- Pf.
 - No two points lie in same $\frac{1}{2}\delta$ -by- $\frac{1}{2}\delta$ box
 - Two points at least 2 rows apart have distance $\geq 2(\frac{1}{2}\delta)$.
- Fact. Still true if we replace 12 with 7.

Cost of combining is therefore...?



Closest Pair Algorithm

```

Closest-Pair( $p_1, \dots, p_n$ )
  Compute separation line L such that half the points
  are on one side and half on the other side.

   $\delta_1 = \text{Closest-Pair}(\text{left half})$ 
   $\delta_2 = \text{Closest-Pair}(\text{right half})$ 
   $\delta = \min(\delta_1, \delta_2)$ 

  Delete all points further than  $\delta$  from separation
  line L

  Sort remaining points by y-coordinate.

  Scan points in y-order and compare distance between
  each point and next 7 neighbors. If any of these
  distances is less than  $\delta$ , update  $\delta$ .

  return  $\delta$ 
    
```

Closest Pair Algorithm

```

Closest-Pair( $p_1, \dots, p_n$ )
  Compute separation line L such that half the points
  are on one side and half on the other side.  $O(n \log n)$ 

   $\delta_1 = \text{Closest-Pair}(\text{left half})$   $2T(n/2)$ 
   $\delta_2 = \text{Closest-Pair}(\text{right half})$ 
   $\delta = \min(\delta_1, \delta_2)$ 

  Delete all points further than  $\delta$  from separation
  line L  $O(n)$ 

  Sort remaining points by y-coordinate.  $O(n \log n)$ 

  Scan points in y-order and compare distance between
  each point and next 7 neighbors. If any of these
  distances is less than  $\delta$ , update  $\delta$ .  $O(n)$ 

  return  $\delta$ 
    
```

$T(n) = 2T(n/2) + O(n \log n)$

Closest Pair of Points: Analysis

- Running time. Solved in 5.2
 - $T(n) \leq 2T(n/2) + O(n \log n) \rightarrow T(n) = O(n \log^2 n)$
- Can we achieve $O(n \log n)$?
 - $T(n) \leq 2T(n/2) + O(n) \rightarrow T(n) = O(n \log n)$
- Yes. Don't sort points in strip from scratch each time.
 - Each recursive returns two lists: all points sorted by y coordinate, and all points sorted by x coordinate
 - Sort by merging two pre-sorted lists

INTEGER AND MATRIX MULTIPLICATION

Integer Arithmetic

- **Add.** Given 2 n -digit integers a and b , compute $a + b$.
 - Algorithm?
 - Runtime?

```

1 1 1 1 1 1 0 1
+ 0 1 1 1 1 1 0 1
-----
1 0 1 0 1 0 0 1 0
    
```

Integer Arithmetic

- **Add.** Given 2 n -digit integers a and b , compute $a + b$.
 - Algorithm?
 - Runtime?

```

1 1 1 1 1 1 0 1
+ 0 1 1 1 1 1 0 1
-----
1 0 1 0 1 0 0 1 0
    
```

$O(n)$ operations

Integer Arithmetic

- **Multiply.** Given 2 n -digit integers a and b , compute $a \times b$.
 - Algorithm?
 - Runtime?

```

1 1 0 1 0 1 0 1
* 0 1 1 1 1 0 1
    
```

Integer Arithmetic

- **Multiply.** Given 2 n -digit integers a and b , compute $a \times b$.
 - Brute force solution: $\Theta(n^2)$ bit operations

Goal: Faster algorithm

```

      1 1 0 1 0 1 0 1
      * 0 1 1 1 1 0 1
      -----
      1 1 0 1 0 1 0 1 0
    0 0 0 0 0 0 0 0 0
  1 1 0 1 0 1 0 1 0
 1 1 0 1 0 1 0 1 0
1 1 0 1 0 1 0 1 0
1 1 0 1 0 1 0 1 0
1 1 0 1 0 1 0 1 0
0 0 0 0 0 0 0 0 0
0 1 1 0 1 0 0 0 0 0 0 0 0 0 0 0 1 0
    
```

Divide-and-Conquer Multiplication: Warmup

- To multiply 2 n -digit integers:
 - Multiply 4 $\frac{1}{2} n$ -digit integers
 - Add 2 $\frac{1}{2} n$ -digit integers and shift to obtain result

Higher order bits Lower order bits

Shift

$$\begin{aligned}
 x &= 2^{n/2} \cdot x_1 + x_0 \\
 y &= 2^{n/2} \cdot y_1 + y_0 \\
 xy &= (2^{n/2} \cdot x_1 + x_0)(2^{n/2} \cdot y_1 + y_0) = 2^n \cdot x_1 y_1 + 2^{n/2} \cdot (x_1 y_0 + x_0 y_1) + x_0 y_0
 \end{aligned}$$

A B C D

What is the recurrence relation?

- How many subproblems?
- What is merge cost?
- What is its runtime?

Divide-and-Conquer Multiplication: Warmup

- To multiply two n -digit integers:
 - Multiply 4 $\frac{1}{2} n$ -digit integers
 - Add 2 $\frac{1}{2} n$ -digit integers and shift to obtain result

Higher order bits Lower order bits

Shift

$$\begin{aligned}
 x &= 2^{n/2} \cdot x_1 + x_0 \\
 y &= 2^{n/2} \cdot y_1 + y_0 \\
 xy &= (2^{n/2} \cdot x_1 + x_0)(2^{n/2} \cdot y_1 + y_0) = 2^n \cdot x_1 y_1 + 2^{n/2} \cdot (x_1 y_0 + x_0 y_1) + x_0 y_0
 \end{aligned}$$

A B C D

$$T(n) = 4T(n/2) + \Theta(n) \Rightarrow T(n) = \Theta(n^2)$$

↑
assumes n is a power of 2

Not an improvement over brute force

Karatsuba Multiplication

- To multiply two n-digit integers:
 - Add 2 1/2n digit integers
 - Multiply 3 1/2n-digit integers
 - Add, subtract, and shift 1/2n-digit integers to obtain result

$$\begin{aligned}
 x &= 2^{n/2} \cdot x_1 + x_0 \\
 y &= 2^{n/2} \cdot y_1 + y_0 \\
 xy &= 2^n \cdot x_1 y_1 + 2^{n/2} \cdot (x_1 y_0 + x_0 y_1) + x_0 y_0 \\
 &= 2^n \cdot x_1 y_1 + 2^{n/2} \cdot (x_1 + x_0)(y_1 + y_0) - x_1 y_1 - x_0 y_0 + x_0 y_0
 \end{aligned}$$

A
B
A
C
C

What is the recurrence relation? Runtime?

Karatsuba Multiplication

- Theorem.** [Karatsuba-Ofman, 1962] Can multiply two n-digit integers in $O(n^{1.585})$ bit operations

$$\begin{aligned}
 x &= 2^{n/2} \cdot x_1 + x_0 \\
 y &= 2^{n/2} \cdot y_1 + y_0 \\
 xy &= 2^n \cdot x_1 y_1 + 2^{n/2} \cdot (x_1 y_0 + x_0 y_1) + x_0 y_0 \\
 &= 2^n \cdot x_1 y_1 + 2^{n/2} \cdot ((x_1 + x_0)(y_1 + y_0) - x_1 y_1 - x_0 y_0) + x_0 y_0
 \end{aligned}$$

A
B
A
C
C

$$\begin{aligned}
 T(n) &\leq T(\lfloor n/2 \rfloor) + T(\lfloor n/2 \rfloor) + T(\lfloor 1 + \lfloor n/2 \rfloor \rfloor) + \Theta(n) \\
 &\Rightarrow T(n) = O(n^{\log_2 3}) = O(n^{1.585})
 \end{aligned}$$

recursive calls add, subtract, shift

MATRIX MULTIPLICATION

Matrix Multiplication

- Given 2 n-by-n matrices A and B, compute $C = AB$

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$

➢ Ex: $c_{12} = a_{11} b_{12} + a_{12} b_{22} + a_{13} b_{32} + \dots + a_{1n} b_{n2}$

Solve using brute force ...

Matrix Multiplication

- Given 2 n-by-n matrices A and B, compute $C = AB$

$$c_{ij} = \sum_{k=1}^n a_{ik} b_{kj}$$

➢ Ex: $c_{12} = a_{11} b_{12} + a_{12} b_{22} + a_{13} b_{32} + \dots + a_{1n} b_{n2}$

- Brute force.** $\Theta(n^3)$ arithmetic operations
- Fundamental question:** Can we improve upon brute force?

Matrix Multiplication: Warmup

- Divide:** partition A and B into 1/2n-by-1/2n blocks
- Conquer:** multiply 8 1/2n-by-1/2n recursively
- Combine:** add appropriate products using 4 matrix additions

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \times \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}$$

$$\begin{aligned}
 C_{11} &= (A_{11} \times B_{11}) + (A_{12} \times B_{21}) \\
 C_{12} &= (A_{11} \times B_{12}) + (A_{12} \times B_{22}) \\
 C_{21} &= (A_{21} \times B_{11}) + (A_{22} \times B_{21}) \\
 C_{22} &= (A_{21} \times B_{12}) + (A_{22} \times B_{22})
 \end{aligned}$$

Recurrence relation? Runtime?

Matrix Multiplication: Warmup

- **Divide:** partition A and B into 1/2n-by-1/2n blocks
- **Conquer:** multiply 8 1/2n-by-1/2n recursively
- **Combine:** add appropriate products using 4 matrix additions

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \times \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}$$

$$\begin{aligned} C_{11} &= (A_{11} \times B_{11}) + (A_{12} \times B_{21}) \\ C_{12} &= (A_{11} \times B_{12}) + (A_{12} \times B_{22}) \\ C_{21} &= (A_{21} \times B_{11}) + (A_{22} \times B_{21}) \\ C_{22} &= (A_{21} \times B_{12}) + (A_{22} \times B_{22}) \end{aligned}$$

$$T(n) = \underbrace{8T(n/2)}_{\text{recursive calls}} + \underbrace{\Theta(n^2)}_{\text{add. form submatrices}} \Rightarrow T(n) = \Theta(n^3)$$

Mar 7, 2011

CSCI211 - Sprenkle

31

Matrix Multiplication: Key Idea

Trading expensive multiplication for less expensive addition/subtraction

- Multiply 2-by-2 block matrices with only 7 multiplications and 15 additions

$$\begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \times \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}$$

$$\begin{aligned} C_{11} &= P_3 + P_4 - P_2 + P_6 \\ C_{12} &= P_1 + P_2 \\ C_{21} &= P_3 + P_4 \\ C_{22} &= P_3 + P_1 - P_3 - P_7 \end{aligned}$$

$$\begin{aligned} P_1 &= A_{11} \times (B_{12} - B_{22}) \\ P_2 &= (A_{11} + A_{12}) \times B_{22} \\ P_3 &= (A_{21} + A_{22}) \times B_{11} \\ P_4 &= A_{22} \times (B_{21} - B_{11}) \\ P_5 &= (A_{11} + A_{22}) \times (B_{11} + B_{22}) \\ P_6 &= (A_{12} - A_{22}) \times (B_{21} + B_{22}) \\ P_7 &= (A_{11} - A_{21}) \times (B_{11} + B_{12}) \end{aligned}$$

Mar 7, 2011

CSCI211 - Sprenkle

32

Fast Matrix Multiplication

[Strassen, 1969]

- **Divide:** partition A and B into 1/2n-by-1/2n blocks
- **Compute:** 14 1/2n-by-1/2n matrices via 10 matrix additions
- **Conquer:** multiply 7 1/2n-by-1/2n matrices recursively
- **Combine:** 7 products into 4 terms using 8 matrix additions
- **Analysis.**

$$T(n) = \underbrace{7T(n/2)}_{\text{recursive calls}} + \underbrace{\Theta(n^2)}_{\text{add. submatr.}} \Rightarrow T(n) = \Theta(n^{\log_2 7}) = O(n^{2.81})$$

- Assume n is a power of 2.
- T(n) = # arithmetic operations.

Mar 7, 2011

CSCI211 - Sprenkle

33

Fast Matrix Multiplication in Practice

- Implementation issues: problems with putting theory into practice
 - Sparsity
 - Caching effects
 - Numerical stability
 - Theoretically correct but possible problems with round off errors, etc
 - Odd matrix dimensions
 - Crossover to classical algorithm around n = 128

Mar 7, 2011

CSCI211 - Sprenkle

34

Fast Matrix Multiplication in Practice

- Common misperception: "Strassen is only a theoretical curiosity."
 - Advanced Computation Group at Apple Computer reports **8x** speedup on G4 Velocity Engine when n ~ 2,500
 - Range of instances where it's useful is a subject of controversy
- Can "Strassenize" Ax=b, determinant, eigenvalues, and other matrix ops

Mar 7, 2011

CSCI211 - Sprenkle

35

Fast Matrix Multiplication in Theory

- Q. Multiply two 2-by-2 matrices with only 7 scalar multiplications?
- A. Yes! [Strassen, 1969] $\Theta(n^{\log_2 7}) = O(n^{2.81})$
- Q. Multiply two 2-by-2 matrices with only 6 scalar multiplications?
- A. Impossible [Hopcroft and Kerr, 1971] $\Theta(n^{\log_2 6}) = O(n^{2.59})$
- Q. Two 3-by-3 matrices with only 21 scalar multiplications?
- A. Also impossible $\Theta(n^{\log_3 21}) = O(n^{2.77})$
- Q. Two 70-by-70 matrices with only 143,640 scalar multiplications?
- A. Yes! [Pan, 1980] $\Theta(n^{\log_{70} 143640}) = O(n^{2.80})$
- **Decimal wars.**
 - December, 1979: $O(n^{2.521813})$
 - January, 1980: $O(n^{2.521801})$

Mar 7, 2011

CSCI211 - Sprenkle

36

Fast Matrix Multiplication in Theory

- **Best known.** $O(n^{2.376})$
[Coppersmith-Winograd, 1987]
 - But *really* large constant
- **Conjecture.** $O(n^{2+\epsilon})$ for any $\epsilon > 0$.
- **Caveat.** Theoretical improvements to Strassen are progressively less practical.

Mar 7, 2011

CSCI211 - Sprenkle

37

PS4 Feedback

- Whenever you develop an algorithm, **analyze** its running time (e.g., Prob 4)
- Be explicit
 - Explicitly state the metric, who is the pronoun (say which algorithm), etc.
 - Creative in solution, not in explanation
 - Follow template for proofs for stays ahead or exchange argument

Mar 7, 2011

CSCI211 - Sprenkle

38

Assignments

- Wiki for 5.1-5.4 due Wednesday
- Continue reading Chapter 5, start Chapter 6
- PS6 due Friday
 - May want to try to implement problems 2 and 3 (to some extent) to help ensure that your algorithm is correct

Mar 7, 2011

CSCI211 - Sprenkle

39