Divide and Conquer Algorithms

T. M. Murali

March 13 and 15, 2017

T. M. Murali

March 13 and 15, 2017

CS 4104: Divide and Conquer Algorithms

Divide and Conquer Algorithms

- Study three divide and conquer algorithms:
 - Counting inversions.
 - Finding the closest pair of points.
 - Integer multiplication.
- First two problems use clever conquer strategies.
- Third problem uses a clever divide strategy.

Motivation

Inspired by your shopping trends



More top picks for you



- Collaborative filtering: match one user's preferences to those of other users, e.g., purchases, books, music.
- Meta-search engines: merge results of multiple search engines into a better search result.

Fundamental Question

- How do we compare a pair of rankings?
 - ▶ My ranking of songs: ordered list of integers from 1 to *n*.
 - Your ranking of songs: a_1, a_2, \ldots, a_n , a permutation of the integers from 1 to n.



Comparing Rankings





• Suggestion: two rankings of songs are very similar if they have few inversions.

Comparing Rankings



- Suggestion: two rankings of songs are very similar if they have few inversions.
 - ► The second ranking has an *inversion* if there exist *i*, *j* such that *i* < *j* but *a_i* > *a_j*.
 - ► The number of inversions *s* is a measure of the difference between the rankings.
- Question also arises in statistics: Kendall's rank correlation of two lists of numbers is 1 2s/(n(n-1)).

Counting Inversions

Count Inversions

INSTANCE: A list $L = x_1, x_2, ..., x_n$ of distinct integers between 1 and *n*.

SOLUTION: The number of pairs $(i, j), 1 \le i < j \le n$ such $x_i > x_j$.

Counting Inversions

COUNT INVERSIONS **INSTANCE:** A list $L = x_1, x_2, ..., x_n$ of distinct integers between 1 and *n*. **SOLUTION:** The number of pairs $(i, j), 1 \le i < j \le n$ such

Solution: The number of pairs $(I, j), 1 \le I < j \le n$ sucr $x_i > x_j$.



Counting Inversions

COUNT INVERSIONS

INSTANCE: A list $L = x_1, x_2, ..., x_n$ of distinct integers between 1 and *n*.

SOLUTION: The number of pairs $(i, j), 1 \le i < j \le n$ such

 $x_i > x_j$.



• How many inversions can be there in a list of *n* numbers?



• How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.



- How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.
- Sorting removes all inversions in $O(n \log n)$ time. Can we modify the Mergesort algorithm to count inversions?



- How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.
- Sorting removes all inversions in $O(n \log n)$ time. Can we modify the Mergesort algorithm to count inversions?
- Candidate algorithm:
 - **1** Partition L into two lists A and B of size n/2 each.
 - 2 Recursively count the number of inversions in A.
 - **③** Recursively count the number of inversions in B.
 - Gount the number of inversions involving one element in A and one element in B.



- How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.
- Sorting removes all inversions in $O(n \log n)$ time. Can we modify the Mergesort algorithm to count inversions?
- Candidate algorithm:
 - **1** Partition L into two lists A and B of size n/2 each.
 - 2 Recursively count the number of inversions in A.
 - **③** Recursively count the number of inversions in B.
 - Count the number of inversions involving one element in A and one element in B.



- How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.
- Sorting removes all inversions in $O(n \log n)$ time. Can we modify the Mergesort algorithm to count inversions?
- Candidate algorithm:
 - **1** Partition L into two lists A and B of size n/2 each.
 - 2 Recursively count the number of inversions in A.
 - **③** Recursively count the number of inversions in B.
 - Count the number of inversions involving one element in A and one element in B.



- How many inversions can be there in a list of *n* numbers? $\Omega(n^2)$. We cannot afford to compute each inversion explicitly.
- Sorting removes all inversions in $O(n \log n)$ time. Can we modify the Mergesort algorithm to count inversions?
- Candidate algorithm:
 - **1** Partition L into two lists A and B of size n/2 each.
 - 2 Recursively count the number of inversions in A.
 - **③** Recursively count the number of inversions in B.
 - Count the number of inversions involving one element in A and one element in B.





• Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.



- Given lists $A = a_1, a_2, \ldots, a_m$ and $B = b_1, b_2, \ldots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!

Counting Inversions: Conquer Step



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE procedure:
 - Maintain a *current* pointer for each list.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - 2 Append the smaller of the two to the output list.

4 Advance current in the list containing the smaller element.

- Solution Append the rest of the non-empty list to the output.
 - Return the merged list.

Counting Inversions: Conquer Step



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - 2 Append the smaller of the two to the output list.
 - **(3)** If $b_i < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance *current* in the list containing the smaller element.
 - Solution Append the rest of the non-empty list to the output.
 - Return *count* and the merged list.

Counting Inversions: Conquer Step



- Given lists $A = a_1, a_2, \ldots, a_m$ and $B = b_1, b_2, \ldots, b_m$, compute the number of pairs a_i and b_i such $a_i > b_i$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **1** Let a_i and b_i be the elements pointed to by the *current* pointers.
 - 2 Append the smaller of the two to the output list.
 - **3** If $b_i < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - 6 Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - 2 Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **3** If $b_i < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **3** If $b_i < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **(3)** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance *current* in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **③** If $b_j < a_i$, increment *count* by the number of elements remaining in A.
 - 4 Advance current in the list containing the smaller element.
 - Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).



- Given lists $A = a_1, a_2, \dots, a_m$ and $B = b_1, b_2, \dots, b_m$, compute the number of pairs a_i and b_j such $a_i > b_j$.
- Key idea: problem is much easier if A and B are sorted!
- MERGE-AND-COUNT procedure:
 - 1 Maintain a *current* pointer for each list.
 - Maintain a variable count initialised to 0.
 - Initialise each pointer to the front of the list.
 - While both lists are nonempty:
 - **()** Let a_i and b_j be the elements pointed to by the *current* pointers.
 - Append the smaller of the two to the output list.
 - **3** If $b_i < a_i$, increment *count* by the number of elements remaining in A.
 - Advance current in the list containing the smaller element.
 - Solution Append the rest of the non-empty list to the output.
 - Return count and the merged list.
- Running time of this algorithm is O(m).

```
Sort-and-Count(L)
If the list has one element then
    there are no inversions
Else
    Divide the list into two halves:
       A contains the first \lceil n/2 \rceil elements
       B contains the remaining |n/2| elements
    (r_A, A) = Sort-and-Count(A)
    (r_B, B) = \text{Sort-and-Count}(B)
    (r, L) = Merge-and-Count(A, B)
 Endif
 Return r = r_A + r_B + r, and the sorted list L
```

```
Sort-and-Count(L)
If the list has one element then
    there are no inversions
Else
    Divide the list into two halves:
       A contains the first \lceil n/2 \rceil elements
       B contains the remaining |n/2| elements
    (r_A, A) = Sort-and-Count(A)
    (r_B, B) = \text{Sort-and-Count}(B)
    (r, L) = Merge-and-Count(A, B)
 Endif
 Return r = r_A + r_B + r, and the sorted list L
```

• Running time T(n) of the algorithm is $O(n \log n)$ because $T(n) \le 2T(n/2) + O(n)$.

T. M. Murali

March 13 and 15, 2017

CS 4104: Divide and Conquer Algorithms

Counting Inversions: Correctness of Sort-and-Count

• Prove by induction. Strategy: every inversion in the data is counted exactly once.

Counting Inversions: Correctness of Sort-and-Count

- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but $x_k > x_l$. When is this inversion counted by the algorithm?

•
$$k, l \leq \lfloor n/2 \rfloor$$
:

•
$$k, l \ge \lceil n/2 \rceil$$
:

• $k \leq \lfloor n/2 \rfloor, l \geq \lceil n/2 \rceil$:
- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but $x_k > x_l$. When is this inversion counted by the algorithm?
 - ▶ $k, l \leq \lfloor n/2 \rfloor$: $x_k, x_l \in A$, counted in r_A .
 - $k, l \ge \lceil n/2 \rceil$: $x_k, x_l \in B$, counted in r_B .
 - $k \leq \lfloor n/2 \rfloor, l \geq \lceil n/2 \rceil$:



- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but x_k > x_l.
 When is this inversion counted by the algorithm?
 - ▶ $k, l \leq \lfloor n/2 \rfloor$: $x_k, x_l \in A$, counted in r_A .
 - $k, l \ge \lceil n/2 \rceil$: $x_k, x_l \in B$, counted in r_B .
 - ▶ $k \leq \lfloor n/2 \rfloor, l \geq \lceil n/2 \rceil$: $x_k \in A, x_l \in B$. Is this inversion counted by MERGE-AND-COUNT?



- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but x_k > x_l.
 When is this inversion counted by the algorithm?
 - ▶ $k, l \leq \lfloor n/2 \rfloor$: $x_k, x_l \in A$, counted in r_A .
 - $k, l \ge \lceil n/2 \rceil$: $x_k, x_l \in B$, counted in r_B .
 - ▶ $k \leq \lfloor n/2 \rfloor, l \geq \lceil n/2 \rceil$: $x_k \in A, x_l \in B$. Is this inversion counted by MERGE-AND-COUNT? Yes, when x_l is output.



- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but $x_k > x_l$. When is this inversion counted by the algorithm?
 - ▶ $k, l \leq \lfloor n/2 \rfloor$: $x_k, x_l \in A$, counted in r_A .
 - $k, l \ge \lceil n/2 \rceil$: $x_k, x_l \in B$, counted in r_B .
 - ▶ $k \leq \lfloor n/2 \rfloor, l \geq \lceil n/2 \rceil$: $x_k \in A, x_l \in B$. Is this inversion counted by MERGE-AND-COUNT? Yes, when x_l is output.
 - Why is no non-inversion counted, i.e., Why does every pair counted correspond to an inversion?



- Prove by induction. Strategy: every inversion in the data is counted exactly once.
- Base case: n = 1.
- Inductive hypothesis: Algorithm counts number of inversions correctly for all sets of n 1 or fewer numbers.
- Inductive step: Pick an arbitrary k and l such that k < l but x_k > x_l.
 When is this inversion counted by the algorithm?
 - ▶ $k, l \leq \lfloor n/2 \rfloor$: $x_k, x_l \in A$, counted in r_A .
 - $k, l \ge \lceil n/2 \rceil$: $x_k, x_l \in B$, counted in r_B .
 - ▶ $k \leq \lfloor n/2 \rfloor$, $l \geq \lceil n/2 \rceil$: $x_k \in A, x_l \in B$. Is this inversion counted by MERGE-AND-COUNT? Yes, when x_l is output.
 - ▶ Why is no non-inversion counted, i.e., Why does every pair counted correspond to an inversion? When x_l is output, it is smaller than all remaining elements in A, since A is sorted.



T. M. Murali

MULTIPLY INTEGERS **INSTANCE:** Two *n*-digit binary integers *x* and *y* **SOLUTION:** The product *xy*

MULTIPLY INTEGERS **INSTANCE:** Two *n*-digit binary integers *x* and *y* **SOLUTION:** The product *xy*

• Multiply two *n*-digit integers.

MULTIPLY INTEGERS **INSTANCE:** Two *n*-digit binary integers *x* and *y* **SOLUTION:** The product *xy*

- Multiply two *n*-digit integers.
- Result has at most 2*n* digits.

MULTIPLY INTEGERS **INSTANCE:** Two *n*-digit binary integers *x* and *y* **SOLUTION:** The product *xy*

- Multiply two *n*-digit integers.
- Result has at most 2*n* digits.
- Algorithm we learnt in school takes

	1100
	$\times 1101$
12	1100
$\times 13$	0000
36	1100
12	1100
156	10011100
(a)	(b)

Figure 5.8 The elementary-school algorithm for multiplying two integers, in (a) decimal and (b) binary representation.

MULTIPLY INTEGERS **INSTANCE:** Two *n*-digit binary integers *x* and *y* **SOLUTION:** The product *xy*

- Multiply two *n*-digit integers.
- Result has at most 2*n* digits.
- Algorithm we learnt in school takes $O(n^2)$ operations. Size of the input is not 2 but 2n,

	1100
	$\times 1101$
12	1100
$\times 13$	0000
36	1100
12	1100
156	10011100
(a)	(b)

Figure 5.8 The elementary-school algorithm for multiplying two integers, in (a) decimal and (b) binary representation.

• Let us use divide and conquer

- Let us use divide and conquer by splitting each number into first n/2 bits and last n/2 bits.
- Let x be split into x_0 (lower-order bits) and x_1 (higher-order bits) and y into y_0 (lower-order bits) and y_1 (higher-order bits).

- Let us use divide and conquer by splitting each number into first n/2 bits and last n/2 bits.
- Let x be split into x_0 (lower-order bits) and x_1 (higher-order bits) and y into y_0 (lower-order bits) and y_1 (higher-order bits).

$$xy = (x_1 2^{n/2} + x_0)(y_1 2^{n/2} + y_0)$$

=

- Let us use divide and conquer by splitting each number into first n/2 bits and last n/2 bits.
- Let x be split into x_0 (lower-order bits) and x_1 (higher-order bits) and y into y_0 (lower-order bits) and y_1 (higher-order bits).







- Algorithm:
 - **Ompute** x_1y_1 , x_1y_0 , x_0y_1 , and x_0y_0 recursively.
 - Ø Merge the answers, i.e,.,
 - Multiple x_1y_1 by 2^n
 - 2 Add x_1y_0 and x_0y_1 and multiple this sum by $2^{n/2}$
 - **3** Add these two numbers to x_0y_0



- Algorithm:
 - **Ompute** x_1y_1 , x_1y_0 , x_0y_1 , and x_0y_0 recursively.
 - Ø Merge the answers, i.e,.,
 - Multiple x_1y_1 by 2^n
 - 2 Add x_1y_0 and x_0y_1 and multiple this sum by $2^{n/2}$
 - **3** Add these two numbers to x_0y_0
- What is the running time of the conquer step?



- Algorithm:
 - **Ompute** x_1y_1 , x_1y_0 , x_0y_1 , and x_0y_0 recursively.
 - Ø Merge the answers, i.e,.,
 - Multiple x_1y_1 by 2^n
 - 2 Add x_1y_0 and x_0y_1 and multiple this sum by $2^{n/2}$
 - **3** Add these two numbers to x_0y_0
- What is the running time of the conquer step?
 - ► Each of x₁, x₀, y₁, y₀ has n/2 bits, so we can add their products in O(n) time.



- Algorithm:
 - **Orginal** Compute x_1y_1 , x_1y_0 , x_0y_1 , and x_0y_0 recursively.
 - 2 Merge the answers, i.e,.,
 - Multiple x_1y_1 by 2^n
 - 2 Add x_1y_0 and x_0y_1 and multiple this sum by $2^{n/2}$
 - **3** Add these two numbers to $x_0 y_0$
- What is the running time of the conquer step?
 - ► Each of x₁, x₀, y₁, y₀ has n/2 bits, so we can add their products in O(n) time.
- What is the running time T(n)?



- Algorithm:
 - **1** Compute x_1y_1 , x_1y_0 , x_0y_1 , and x_0y_0 recursively.
 - 2 Merge the answers, i.e,.,
 - Multiple x_1y_1 by 2^n
 - 2 Add x_1y_0 and x_0y_1 and multiple this sum by $2^{n/2}$
 - **3** Add these two numbers to $x_0 y_0$
- What is the running time of the conquer step?
 - ► Each of x₁, x₀, y₁, y₀ has n/2 bits, so we can add their products in O(n) time.
- What is the running time T(n)?

$$T(n) \leq 4T(n/2) + cn \leq O(n^2)$$

Improving the Algorithm

- Four sub-problems lead to an $O(n^2)$ algorithm.
- How can we reduce the number of sub-problems?

Improving the Algorithm

- Four sub-problems lead to an $O(n^2)$ algorithm.
- How can we reduce the number of sub-problems?
 - ► No need to compute x₁y₀ and x₀y₁ independently; we just need their sum.



- Compute x₁y₁, x₀y₀ and (x₀ + x₁)(y₀ + y₁) recursively and then compute (x₁y₀ + x₀y₁) by subtraction.
- Strategy: simple arithmetic manipulations.

Final Algorithm

Recursive-Multiply(x,y): Write $x = x_1 \cdot 2^{n/2} + x_0$ $y = y_1 \cdot 2^{n/2} + y_0$ Compute $x_1 + x_0$ and $y_1 + y_0$ p = Recursive-Multiply($x_1 + x_0$, $y_1 + y_0$) x_1y_1 = Recursive-Multiply(x_1 , y_1) x_0y_0 = Recursive-Multiply(x_0 , y_0) Return $x_1y_1 \cdot 2^n + (p - x_1y_1 - x_0y_0) \cdot 2^{n/2} + x_0y_0$

Final Algorithm

Recursive-Multiply(x,y): Write $x = x_1 \cdot 2^{n/2} + x_0$ $y = y_1 \cdot 2^{n/2} + y_0$ Compute $x_1 + x_0$ and $y_1 + y_0$ p = Recursive-Multiply($x_1 + x_0$, $y_1 + y_0$) x_1y_1 = Recursive-Multiply(x_1 , y_1) x_0y_0 = Recursive-Multiply(x_0 , y_0) Return $x_1y_1 \cdot 2^n + (p - x_1y_1 - x_0y_0) \cdot 2^{n/2} + x_0y_0$

- We have three sub-problems of size n/2.
- What is the running time T(n)?

$$T(n) \leq 3T(n/2) + cn$$

Final Algorithm

Recursive-Multiply(x,y): Write $x = x_1 \cdot 2^{n/2} + x_0$ $y = y_1 \cdot 2^{n/2} + y_0$ Compute $x_1 + x_0$ and $y_1 + y_0$ p = Recursive-Multiply($x_1 + x_0$, $y_1 + y_0$) x_1y_1 = Recursive-Multiply(x_1 , y_1) x_0y_0 = Recursive-Multiply(x_0 , y_0) Return $x_1y_1 \cdot 2^n + (p - x_1y_1 - x_0y_0) \cdot 2^{n/2} + x_0y_0$

- We have three sub-problems of size n/2.
- What is the running time T(n)?

$$\begin{array}{rcl} T(n) & \leq & 3T(n/2) + cn \\ & \leq & O(n^{\log_2 3}) = O(n^{1.59}) \end{array}$$

Computational Geometry

- Algorithms for geometric objects: points, lines, segments, triangles, spheres, polyhedra, ldots.
- Started in 1975 by Shamos and Hoey.
- Problems studied have applications in a vast number of fields: ecology, molecular biology, statistics, computational finance, computer graphics, computer vision, ...

Computational Geometry

- Algorithms for geometric objects: points, lines, segments, triangles, spheres, polyhedra, ldots.
- Started in 1975 by Shamos and Hoey.
- Problems studied have applications in a vast number of fields: ecology, molecular biology, statistics, computational finance, computer graphics, computer vision, ...

CLOSEST PAIR OF POINTS

INSTANCE: A set P of n points in the plane

SOLUTION: The pair of points in *P* that are the closest to each other.

Computational Geometry

- Algorithms for geometric objects: points, lines, segments, triangles, spheres, polyhedra, ldots.
- Started in 1975 by Shamos and Hoey.
- Problems studied have applications in a vast number of fields: ecology, molecular biology, statistics, computational finance, computer graphics, computer vision, ...

CLOSEST PAIR OF POINTS

INSTANCE: A set P of n points in the plane **SOLUTION:** The pair of points in P that are the closest to each other.

- At first glance, it seems any algorithm must take $\Omega(n^2)$ time.
- Shamos and Hoey figured out an ingenious $O(n \log n)$ divide and conquer algorithm.

- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use $d(p_i, p_j)$ to denote the Euclidean distance between p_i and p_j . For a specific pair of points, can compute $d(p_i, p_j)$ in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.

- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use d(p_i, p_j) to denote the Euclidean distance between p_i and p_j. For a specific pair of points, can compute d(p_i, p_j) in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.
- How do we solve the problem in 1D?



- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use d(p_i, p_j) to denote the Euclidean distance between p_i and p_j. For a specific pair of points, can compute d(p_i, p_j) in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.
- How do we solve the problem in 1D?
 - Sort: closest pair must be adjacent in the sorted order.



- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use d(p_i, p_j) to denote the Euclidean distance between p_i and p_j. For a specific pair of points, can compute d(p_i, p_j) in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.
- How do we solve the problem in 1D?
 - Sort: closest pair must be adjacent in the sorted order.
 - Divide and conquer after sorting: closest pair must be closest of
 - **1** closest pair in left half: distance δ_I .
 - 2 closest pair in right half: distance δ_r .
 - Solution closest among pairs that span the left and right halves and are at most $\min(\delta_l, \delta_r)$ apart. How many such pairs do we need to consider?



- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use d(p_i, p_j) to denote the Euclidean distance between p_i and p_j. For a specific pair of points, can compute d(p_i, p_j) in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.
- How do we solve the problem in 1D?
 - Sort: closest pair must be adjacent in the sorted order.
 - Divide and conquer after sorting: closest pair must be closest of
 - **1** closest pair in left half: distance δ_l .
 - 2 closest pair in right half: distance δ_r .
 - Solution of the set of the se



- Let $P = \{p_1, p_2, ..., p_n\}$ with $p_i = (x_i, y_i)$.
- Use d(p_i, p_j) to denote the Euclidean distance between p_i and p_j. For a specific pair of points, can compute d(p_i, p_j) in O(1) time.
- Goal: find the pair of points p_i and p_j that minimise $d(p_i, p_j)$.
- How do we solve the problem in 1D?
 - Sort: closest pair must be adjacent in the sorted order.
 - Divide and conquer after sorting: closest pair must be closest of
 - **1** closest pair in left half: distance δ_l .
 - 2 closest pair in right half: distance δ_r .
 - Oclosest among pairs that span the left and right halves and are at most min(δ_l, δ_r) apart. How many such pairs do we need to consider? Just one!

• Generalize the second idea to 2D.



Closest Pair: Algorithm Skeleton

- Divide P into two sets Q and R of n/2 points such that each point in Q has x-coordinate less than any point in R.
- **2** Recursively compute closest pair in Q and in R, respectively.



Closest Pair: Algorithm Skeleton

- Divide P into two sets Q and R of n/2 points such that each point in Q has x-coordinate less than any point in R.
- **2** Recursively compute closest pair in Q and in R, respectively.
- Solution Let δ_Q be the distance computed for Q, δ_R be the distance computed for R, and δ = min(δ_Q, δ_R).


Closest Pair: Algorithm Skeleton

- Divide P into two sets Q and R of n/2 points such that each point in Q has x-coordinate less than any point in R.
- **2** Recursively compute closest pair in Q and in R, respectively.
- Solution Let δ_Q be the distance computed for Q, δ_R be the distance computed for R, and δ = min(δ_Q, δ_R).
- Compute pair (q, r) of points such that $q \in Q$, $r \in R$, $d(q, r) < \delta$ and d(q, r) is the smallest possible.



Closest Pair: Proof Sketch

- Prove by induction: Let (s, t) be the closest pair.
 - (i) both are in Q: computed correctly by recursive call.
 - (ii) both are in R: computed correctly by recursive call.
 - (iii) one is in Q and the other is in R: computed correctly in O(n) time by the procedure we will discuss.
- Strategy: Pairs of points for which we do not compute the distance between cannot be the closest pair.

• Overall running time is $O(n \log n)$.



Closest Pair: Conquer Step

- Line L passes through right-most point in Q.
- Let S be the set of points within distance δ of L. (In image, $\delta = \delta_R$.)



Closest Pair: Conquer Step

- Line L passes through right-most point in Q.
- Let S be the set of points within distance δ of L. (In image, $\delta = \delta_{R}$.)
- Claim: There exist $q \in Q$, $r \in R$ such that $d(q, r) < \delta$ if and only if $q, r \in S$.



Closest Pair: Conquer Step

- Line L passes through right-most point in Q.
- Let S be the set of points within distance δ of L. (In image, $\delta = \delta_{R}$.)
- Claim: There exist $q \in Q$, $r \in R$ such that $d(q, r) < \delta$ if and only if $q, r \in S$.
- Corollary: If $t \in Q S$ or $u \in R S$, then (t, u) cannot be the closest pair.



• Intuition: "too many" points in S that are closer than δ to each other \Rightarrow there must be a pair in Q or in R that are less than δ apart.

- Intuition: "too many" points in S that are closer than δ to each other \Rightarrow there must be a pair in Q or in R that are less than δ apart.
- Let S_y denote the set of points in S sorted by increasing y-coordinate and let s_y denote the y-coordinate of a point s ∈ S.



- Intuition: "too many" points in S that are closer than δ to each other \Rightarrow there must be a pair in Q or in R that are less than δ apart.
- Let S_y denote the set of points in S sorted by increasing y-coordinate and let s_y denote the y-coordinate of a point s ∈ S.
- Claim: If there exist $s, s' \in S$ such that $d(s, s') < \delta$ then s and s' are at most 15 indices apart in S_y .



- Intuition: "too many" points in S that are closer than δ to each other \Rightarrow there must be a pair in Q or in R that are less than δ apart.
- Let S_y denote the set of points in S sorted by increasing y-coordinate and let s_y denote the y-coordinate of a point s ∈ S.
- Claim: If there exist $s, s' \in S$ such that $d(s, s') < \delta$ then s and s' are at most 15 indices apart in S_y .
- Converse of the claim: If there exist s, s' ∈ S such that s' appears 16 or more indices after s in S_y, then s'_y − s_y ≥ δ.



- Intuition: "too many" points in S that are closer than δ to each other \Rightarrow there must be a pair in Q or in R that are less than δ apart.
- Let S_y denote the set of points in S sorted by increasing y-coordinate and let s_y denote the y-coordinate of a point s ∈ S.
- Claim: If there exist $s, s' \in S$ such that $d(s, s') < \delta$ then s and s' are at most 15 indices apart in S_y .
- Converse of the claim: If there exist s, s' ∈ S such that s' appears 16 or more indices after s in S_y, then s'_y − s_y ≥ δ.
- Use the claim in the algorithm: For every point $s \in S_y$, compute distances only to the next 15 points in S_y .
- Other pairs of points cannot be candidates for the closest pair T. M. Murali March 13 and 15, 2017



• Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y - s_y \ge \delta$.



- Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y s_y \ge \delta$.
- Pack the plane with squares of side $\delta/2$.



- Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y s_y \ge \delta$.
- Pack the plane with squares of side $\delta/2$.
- Each square contains at most one point.



- Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y s_y \ge \delta$.
- Pack the plane with squares of side $\delta/2$.
- Each square contains at most one point.
- Let s lie in one of the squares.



- Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y s_y \ge \delta$.
- Pack the plane with squares of side $\delta/2$.
- Each square contains at most one point.
- Let s lie in one of the squares.
- Any point in the third row of the packing below s has a *y*-coordinate at least δ more than s_y.



- Claim: If there exist $s, s' \in S$ such that s' appears 16 or more indices after s in S_y , then $s'_y s_y \ge \delta$.
- Pack the plane with squares of side $\delta/2$.
- Each square contains at most one point.
- Let s lie in one of the squares.
- Any point in the third row of the packing below s has a *y*-coordinate at least δ more than s_y.
- We get a count of 12 or more indices (textbook says 16).



Closest Pair: Final Algorithm

```
Closest-Pair(P)
  Construct P_r and P_r (O(n log n) time)
  (p_{n}^{*}, p_{1}^{*}) = \text{Closest-Pair-Rec}(P_{n}, P_{n})
Closest-Pair-Rec(Pr, Pr)
  If |P| \le 3 then
     find closest pair by measuring all pairwise distances
  Endif
  Construct Q_x, Q_y, R_x, R_y (O(n) time)
  (q_0^*, q_1^*) = \text{Closest-Pair-Rec}(Q_x, Q_y)
  (r_{\alpha}^{*}, r_{1}^{*}) = \text{Closest-Pair-Rec}(R_{x}, R_{y})
  \delta = \min(d(q_0^*, q_1^*), d(r_0^*, r_1^*))
  x^* = maximum x-coordinate of a point in set O
  L = \{(x, y) : x = x^*\}
  S = \text{points in } P \text{ within distance } \delta \text{ of } L.
  Construct S. (O(n) time)
  For each point s \in S_{n}, compute distance from s
      to each of next 15 points in S.
      Let s, s' be pair achieving minimum of these distances
      (O(n) \text{ time})
  If d(s,s') < \delta then
      Return (s,s')
  Else if d(q_{0,*}^*q_1^*) < d(r_{0,*}^*r_1^*) then
      Return (q_{0,*}^*, q_1^*)
  Else
      Return (r_0*, r_1*)
  Endif
```

Closest Pair: Final Algorithm

```
Closest-Pair(P)

Construct P_x and P_y (O(n log n) time)

(p_0^*, p_1^*) = \text{Closest-Pair-Rec}(P_x, P_y)

Closest-Pair-Rec(P_x, P_y)

If |P| \leq 3 then

find closest pair by measuring all pairwise distances

Endif
```

Construct
$$Q_x$$
, Q_y , R_x , R_y (O(n) time,
 $(q_0^*, q_1^*) = \text{Closest-Pair-Rec}(Q_x, Q_y)$
 $(r_0^*, r_1^*) = \text{Closest-Pair-Rec}(R_x, R_y)$

$$\delta = \min(d(q_0^*, q_1^*), d(r_0^*, r_1^*))$$

$$x^* = \max \min x - \text{coordinate of a point in set } Q$$

Closest Pair: Final Algorithm

```
x^* = maximum x-coordinate of a point in set Q
```

```
L = \{(x, y) : x = x^*\}
```

```
S = points in P within distance \delta of L.
```

```
Construct S_y (O(n) time)
For each point s \in S_y, compute distance from s
to each of next 15 points in S_y
Let s, s' be pair achieving minimum of these distances
(O(n) time)
```

```
If d(s,s') < \delta then

Return (s,s')

Else if d(q_0^*,q_1^*) < d(r_0^*,r_1^*) then

Return (q_0^*,q_1^*)

Else

Return (r_0^*,r_1^*)

Endif
```