Finding the Closest Pair of Points

Problem: Given n ordered pairs $(x_1, y_1), (x_2, y_2), ..., (x_n, y_n)$, find the distance between the two points in the set that are closest together.

The brute force algorithm is as follows:

Iterate through all possible pairs of points, calculating the distance between each of these pairs. Any time you see a distance shorter than the shortest distance seen, update the shortest distance seen.

Since computing the distance between two points takes $O(1)$ time, and there are a total of $\frac{n(n-1)}{2} = \theta(n^2)$ distinct pairs of points, it follows that the running time of this algorithm is $\theta(n^2)$.

Can we do better?

Here's the idea:

1) Split the set of n points into two halves by a vertical line. (We can do this by sorting all the points by their x-coordinate and then picking the middle point and drawing a vertical line just to the left or right of it.)
2) Recursively call the function to solve the problem on both sets of points.
3) Return the smaller of the two values.

What's the problem with this idea?

The problem is that the actual shortest distance between any two of the original points MIGHT BE between a point from the first set and a point in the second set! Consider this situation:

```
                                |
*                         *     | *                         *
                                |
            A                            B
```

Here we have four points separated into groups A and B. Let the points be labeled 1, 2, 3 and 4 from left to right. Notice that the minimum of the distances between points 1 and 2 and then points 3 and 4 is NOT the actual shortest possible distance between points, since points 2 and 3 are much closer to one another.

Thus, we must adapt our approach:

In step 3, we can "save" the smaller of the two values, (we'll call this δ), then, we have to check to see if there are points (one in each group) that are closer than δ apart.

Do we need to search through all possible pairs of points from the two different sides?

Probably not. We must only consider points that are within a distance of δ to our dividing line.

```
                                        δ
                                |  ---------------------------
*                         *     | *                         *
                                |
            A                            B
        ------------------------------------------------
                    δ                       δ
```

Still, however, one could construct a case where ALL the points on each side are within δ of the vertical line:

```
* | *
* | *
  |
* | *
* | *
```

So, technically speaking, this case is as bad as our original idea where we'd have to compare each pair of points to one another from the different groups.

But, wait, is it really necessary to compare each point on one side with every other point on every other side???

Consider the following rectangle around the dividing line that is constructed by eight δ/2 x δ/2 squares. Note that the diagonal of each square is $\delta/\sqrt{2}$, which is less than δ. Since each square lies on a single side of the dividing line, at most one point lies in each box, because if two points were within a single box the distance between those two points would be less than δ, contradicting the definition of δ. This means, that there are at MOST 7 other points that could possibly be a distance of less than δ apart from a given point, that have a greater y coordinate than that point. (We assume that our point is on the bottom row of this grid; we draw the grid that way.)

Now, we come to the issue of how do we know WHICH 7 points to compare a given point with???

The idea is as follows: as you are processing the points recursively, SORT them based on the y-coordinate. Then, for a given point within the strip, you only need to compare with the next seven points!

Here is a pseudocode version of the algorithm. (I have simplified the pseudocode from the book so that it's easier to get an overall understanding of the flow of the algorithm.)

```
closest_pair(p) {
  mergesort(p, 1, n) // n is number of points
  return rec_cl_pair(p, 1, 2)
}

rec_cl_pair(p, i, j) {

  if (j - i < 3)  { \\ If there are three points or less...
     mergesort(p, i, j)
     return shortest_distance(p[i], p[i+1], p[i+2])
  }

  xval = p[(i+j)/2].x
  deltaL = rec_cl_pair(p, i, (i+j)/2)
  deltaR = rec_cl_pair(p, (i+j)/2+1, j)
  delta = min(deltaL, deltaR)
  merge(p, i, k, j)

  v = vert_strip(p, xval, delta)

  for k=1 to size(v)-1
    for s = (k+1) to min(t, k+7)
      delta = min(delta, dist(v[k], v[s]))
  return delta
}
```

Let's trace through an example with the following 9 points:

(0, 0), (1, 6), (2, 8), (2, 3), (3, 4), (5, 1), (6, 7), (7, 4) and (8, 0).

These are already sorted these by the x-coordinate. We will split the points in half, with the first four in the first half and the last five in the last half.

So, we must recursively call rec_cl_pair on the set

{(0, 0), (1, 6), (2, 8), (2, 3)}

This results in two more recursive calls, to rec_cl_pair on the sets {(0, 0), (1, 6)} and {(2, 8), (2, 3)}. The first call simply returns the distance $\sqrt{37}$ and sorts the points by y-coordinate which leaves them unchanged. The second call returns 5 and swaps the order of the points (2, 3), (2, 8) in the array p.

Then, delta is set to 5, and the points are all merged by y coordinate: (0, 0), (2, 3), (1, 6), (2, 8). (All of these are then copied into the v array since all points are within a distance of 5 of the line x=1.)

Since we have less than 7 points, all points are compared and it is discovered that (1, 6) and (2, 8) are a distance of $\sqrt{5}$ apart.

Now, let's call rec_cl_pair on the set

{(3, 4), (5, 1), (6, 7), (7, 4), (8, 0)}

We first make two recursive calls to the sets: {(3, 4), (5, 1)} and {(6, 7), (7, 4), (8, 0)}. The first will result in deltaL being set to $\sqrt{13}$ and the points getting sorted by y: {(5, 1), (3, 4)}.

The second results in delta being set to $\sqrt{10}$ and the points getting sorted by y: {(8, 0), (7, 4), (6, 7)}. We then merge all of these points: {(8, 0), (5, 1), (3, 4), (7, 4), (6, 7)}. All of these points get copied into v.

Then, we discover that NO pair of points beats the original deltaR of $\sqrt{10}$. This is then the closest distance between any pair of points within the second set of data.

Finally, we must merge ALL the points together by y-coordinate:

(0, 0), (8, 0), (5, 1), (2, 3), (3, 4), (7, 4) , (1, 6), (6, 7), (2, 8)

this time, we only pick those points that are within $\sqrt{10}$ of the line x=2 to copy into v. These points are:

(0, 0), (5, 1), (2, 3), (3, 4), (1, 6), (2, 8)

Now, we scan through all pairs to discover that the shortest distance between any of the two points is $\sqrt{2}$.

Strassen's algorithm:Matrix multiplication

The standard method of matrix multiplication of two n x n matrices takes T(n) = O(n³).

The following algorithm multiplies nxn matrices A and B:

```
// Initialize C.
for i = 1 to n
  for j = 1 to n
    for k = 1 to n
        C [i, j] += A[i, k] * B[k, j];
```

Stassen's algorithm is a Divide-and-Conquer algorithm that beat the bound. The usual multiplication of two n x n matrices takes

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \qquad B = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \qquad C = \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix}$$

if $C = AB$ **, then we have the following:**

$$c_{11} = a_{11}b_{11} + a_{12}b_{21}$$
$$c_{12} = a_{11}b_{12} + a_{12}b_{22}$$
$$c_{21} = a_{21}b_{11} + a_{22}b_{21}$$
$$c_{22} = a_{21}b_{12} + a_{22}b_{22}$$

8   n/2 * n/2 matrix multiples plus
4   n/2 * n/2 matrix additions

T(n) = 8T(n/2) + O(n²)
Plug in   a = 8, b = 2, k = 2   → log_b^a =3   →   T(n)= O(n³)

Strassen showed how two matrices can be multiplied using only 7 multiplications and 18 additions:

Consider calculating the following 7 products:

$q_1 = (a_{11} + a_{22}) * (b_{11} + b_{22})$
$q_2 = (a_{21} + a_{22}) * b_{11}$
$q_3 = a_{11} * (b_{12} - b_{22})$
$q_4 = a_{22} * (b_{21} - b_{11})$
$q_5 = (a_{11} + a_{12}) * b_{22}$
$q_6 = (a_{21} - a_{11}) * (b_{11} + b_{12})$
$q_7 = (a_{12} - a_{22}) * (b_{21} + b_{22})$

It turns out that

$c_{11} = q_1 + q_4 - q_5 + q_7$
$c_{12} = q_3 + q_5$
$c_{21} = q_2 + q_4$
$c_{22} = q_1 + q_3 - q_2 + q_6$

Note: I haven't actually looked at Strassen's paper. This is just from the text. In going through other books, I saw a different set of products than these. (I also saw these in other books too though.) The different set of products I saw required 24 additions instead of 18, so this is a slight improvement (though not asymptotic) over the other algorithm I saw in a different book.

Let's verify a couple of these:

$$q_1 + q_4 - q_5 + q_7 = (a_{11} + a_{22}) * (b_{11} + b_{22}) + a_{22} * (b_{21} - b_{11})$$
$$- (a_{11} + a_{12}) * b_{22} + (a_{12} - a_{22}) * (b_{21} + b_{22})$$

$$= a_{11}b_{11} + a_{11}b_{22} + a_{22}b_{11} + a_{22}b_{22} + a_{22}b_{21} - a_{22}b_{11} - a_{11}b_{22} - a_{12}b_{22}$$
$$+ a_{12}b_{21} + a_{12}b_{22} - a_{22}b_{21} - a_{22}b_{22}$$

$$= a_{11}b_{11} + a_{12}b_{21} \text{ , since everything else cancels out.}$$

$$q_3 + q_5 = a_{11}*(b_{12} - b_{22}) + (a_{11} + a_{12}) * b_{22}$$
$$= a_{11}b_{12} - a_{11}b_{22} + a_{11}b_{22} + a_{12}b_{22}$$
$$= a_{11}b_{12} + a_{12}b_{22}$$

I have no idea how Strassen came up with these combinations. He probably realized that he wanted to determine each element in the product using less than 8 multiplications. From there, he probably just played around with it. In the algorithms text by Cormen, Leiserson and Rivest, they show how one could derive these products.

If we let T(n) be the running time of Strassen's algorithm, then it satisfies the following recurrence relation:

$$T(n) = 7T(n/2) + O(n^2)$$

It's important to note that the hidden constant in the $O(n^2)$ term is larger than the corresponding constant for the standard divide and conquer algorithm for this problem. However, for large matrices this algorithm yields an improvement over the standard one with respect to time.