

## Topic Notes: Divide and Conquer

*Divide-and-Conquer* is a very common and very powerful algorithm design technique. The general idea:

1. Divide the complete instance of problem into two (sometimes more) subproblems that are smaller instances of the original.
2. Solve the subproblems (recursively).
3. Combine the subproblem solutions into a solution to the complete (original) instance.

While the most common case is that the problem of size  $n$  is divided into 2 subproblems of size  $\frac{n}{2}$ . But in general, we can divide the problem into  $b$  subproblems of size  $\frac{n}{b}$ , where  $a$  of those subproblems need to be solved.

This leads to a general recurrence for divide-and-conquer problems:

$$T(n) = aT(n/b) + f(n), \text{ where } f(n) \in \Theta(n^d), d \geq 0.$$

When we encounter a recurrence of this form, we can use the *master theorem* to determine the efficiency class of  $T$ :

$$T(n) = \begin{cases} \Theta(n^d) & \text{if } a < b^d \\ \Theta(n^d \log n) & \text{if } a = b^d \\ \Theta(n^{\log_b a}) & \text{if } a > b^d \end{cases}$$

Application of this theorem will often allow us to do a quick analysis of many divide-and-conquer algorithms without having to solve the recurrence in detail. We will leave the mathematical proof of the master theorem to the mathematicians.

---

## Mergesort

Each sorting procedure we have considered so far is an “in-place” sort. They require only  $\Theta(1)$  extra space for temporary storage.

Next, we consider *merge sort*, a divide-and-conquer procedure that uses  $\Theta(n)$  extra space in the form of a second array.

It's based on the idea that if you're given two sorted arrays, you can merge them into a third in  $\Theta(n)$  time. Each comparison will lead to one more item being placed into its final location, limiting the number of comparisons to  $n - 1$ .

In the general case, however, this doesn't do anything for our efforts to sort the original array. We have completely unsorted data, not two sorted arrays to merge.

But we can create two arrays to merge if we split the array in half, sort each half independently, and then merge them together (hence the need for the extra  $\Theta(n)$  space).

If we keep doing this recursively, we can reduce the “sort half of the array” problem to the trivial cases.

This approach was invented by John von Neumann in 1945.

How many splits will it take?  $\Theta(\log n)$

Then we will have  $\Theta(\log n)$  merge steps, each of which involves sub-arrays totaling in size to  $n$ , so each merge (which will be  $k$  independent merges into  $\frac{n}{k}$ -element arrays) step has  $\Theta(n)$  operations.

This suggests an overall complexity of  $\Theta(n \log n)$ .

Let's look at pseudocode for this:

```
ALGORITHM MERGESORT(A)
  //Input: an array A[0..n - 1]
  if n > 1 then
    copy first half of array A into a temp array B
    copy second half of array A into a temp array C
    MergeSort(B)
    MergeSort(C)
    Merge(B, C, A)
```

where the merge operation is

```
ALGORITHM MERGE(B, C, A)
  //Input: a sorted array B[0..p - 1]
  //Input: a sorted array C[0..q - 1]
  //Output: a sorted array A[0..(p + q - 1)]
  i ← 0
  j ← 0
  k ← 0
  // take smallest item from B and C and put into A
  while i < p and j < q do
    if B[i] < C[j] then
      A[k] ← B[i]
      i ← i + 1
    else
      A[k] ← C[j]
      j ← j + 1
    k ← k + 1
  // if items remain in B, put them in A
  while i < p do
    A[k] ← B[i]
```

```

    i ← i + 1
    k ← k + 1
// if items remain in C, put them in A
while j < q do
    A[k] ← C[j]
    j ← j + 1
    k ← k + 1

```

Let's do a more formal analysis of mergesort. To keep things simple, we will assume that  $n = 2^k$ .

Our basic operation will be the number of comparisons that need to be made.

The recurrence for the number of comparisons for a mergesort of a problem of size  $n = 2^k$  is

$$C(n) = 2C(n/2) + C_{merge}(n) \quad \text{for } n > 1, C(1) = 0.$$

The best, worst, and average cases depend on how long we are in the main while loop in the merge operation before one of the arrays is empty (as the remaining elements are then taken from the other array with no comparisons needed). Let's consider the worst case, where the merge will take  $n - 1$  comparisons (one array becomes empty only when the other has a single element remaining). This leads to the recurrence:

$$C_{worst}(n) = 2C_{worst}(n/2) + n - 1 \quad \text{for } n > 1, C_{worst}(1) = 0.$$

The master theorem, gives is that  $C_{worst}(n) \in \Theta(n \log n)$ .

---

## Quicksort

Another very popular divide and conquer sorting algorithm is the *quicksort*. This was developed by C. A. R. Hoare in 1962.

Unlike merge sort, quicksort is an in-place sort.

While merge sort divided the array in half at each step, sorted each half, and then merged (where all work is in the merge), quicksort works in the opposite order.

That is, quicksort splits the array (which takes lots of work) into parts consisting of the “smaller” elements and of the “larger” elements, sorts each part, and then puts them back together (trivially).

It proceeds by picking a *pivot* element, moving all elements to the correct side of the pivot, resulting in the pivot being in its final location, and two subproblems remaining that can be solved recursively.

A common and simple choice is to use the leftmost element for the pivot. We put it into its correct position and put all elements on their correct side of the pivot.

Pseudocode for a quicksort:

**ALGORITHM** QUICKSORT( $A, lt, rt$ )

```

//Input: an array  $A[0..n - 1]$ 
//Input: lower  $lt$  and upper  $rt$  bounds of the subarray to sort
//The initial call would be with  $lt = 0$  and  $rt = n - 1$ 
if  $lt < rt$  then
     $s \leftarrow Partition(A, lt, rt)$ 
    //  $s$  is pivot element location
     $QuickSort(A, lt, s - 1)$ 
     $QuickSort(A, s + 1, rt)$ 

```

where the partition is accomplished by:

```

ALGORITHM PARTITION( $A, lt, rt$ )
//Input: an array  $A[0..n - 1]$ 
//Input: lower  $lt$  and upper  $rt$  bounds of the subarray
// to partition
 $p \leftarrow A[lt]$  // select pivot
 $i \leftarrow lt$ 
 $j \leftarrow rt + 1$ 
repeat
    repeat
         $i \leftarrow i + 1$ 
    until  $i = rt$  or  $A[i] \geq p$ 
    repeat
         $j \leftarrow j - 1$ 
    until  $j = lt$  or  $A[j] \leq p$ 
     $swap(A[i], A[j])$ 
until  $i \geq j$ 
 $swap(A[i], A[j])$  // undo last swap
 $swap(A[lt], A[j])$  // place pivot
return  $j$ 

```

Note: we always make a recursive call on a smaller array (but it's easy to make a coding mistake where it doesn't, would would mean the sort never terminates).

The complexity of quicksort is harder to evaluate than merge sort because the pivot will not always wind up in the middle of the array: in the worst case, the pivot is the largest or smallest element.

Again, the basic operation will be the comparisons that take place in the partition.

The `partition` method is clearly  $\Theta(n)$  because every comparison results in left or right moving toward the other and quit when they cross.

In the best case, the pivot element always ends up right in the middle.

This would lead to a number of comparisons according to the recurrence:

$$C_{best}(n) = 2C_{best}(n/2) + n \quad \text{for } n > 1, \quad C_{best}(1) = 0.$$

By solving the recurrence or applying the Master Theorem, we find that  $C_{best}(n) \in \Theta(n \log n)$ , exactly like merge sort.

In the worst case the pivot is at one of the ends and quicksort behaves like a selection sort. This occurs with already-sorted input or reverse-sorted input. To analyze this case, think of the first pass through the full  $n$ -element array. The first element,  $A[0]$ , is chosen as the pivot. The left-to-right inner loop will terminate after one comparison (since  $A[0]$  is the smallest element). The right-to-left inner loop will perform comparisons with  $A[n-1]$ ,  $A[n-2]$ , ... all the way down to  $A[0]$  since we need to “move” the pivot item to its own position. That’s  $n+1$  comparisons for the partition step. In the process, the problem size is decreased by 1, so there will be  $n$  comparisons in the next step,  $n-1$  in the third, and so on. We stop after processing the two-element case (requiring 3 comparisons), so the total comparisons is given by:

$$C_{worst}(n) = (n+1) + n + \dots + 3 = \frac{(n+1)(n+2)}{2} - 3 \in \Theta(n^2)$$

A careful analysis can show that quicksort is  $\Theta(n \log n)$  in the average case (under reasonable assumptions on distribution of elements of array). We can proceed by assuming that the partition can occur at any position with the same probability ( $\frac{1}{n}$ ). This leads to a more complex recurrence:

$$C_{avg}(n) = \frac{1}{n} \sum_{s=0}^{n-1} [(n+1) + C_{avg}(s) + C_{avg}(n-1-s)] \quad \text{for } n > 1, C_{avg}(0) = 0, C_{avg}(1) = 0.$$

We will not solve this in detail, but it works out to:

$$C_{avg}(n) \approx 2n \ln n \approx 1.38n \log_2 n \in \Theta(n \log n).$$

Clearly, the efficiency of quicksort depends on the selection of a good pivot. Improving our chances to select a good pivot will ensure quicker progress. One way to do this is to consider three candidates (often the leftmost, rightmost, and middle elements) and take the median of these three as the pivot value.

Other improvements commonly used to improve the overall efficiency of quicksort include (i) switching to a simpler ( $\Theta(n^2)$ ) sort once the subproblems get below a certain threshold size, and (ii) implementing quicksort iteratively rather than recursively.

Quicksort is often the method of choice for general purpose sorting with large data sizes.

## Computational Geometry

We return to two familiar problems from computational geometry to explore divide-and-conquer solutions that are more efficient than the brute force approaches considered previously.

### Divide-and-Conquer Closest Pairs

Our problem is to find among a set of points in the plane the two points that are closest together.

We begin by assuming that the points in the set are ordered by increasing  $x$  coordinate values. If this is not the case, the points can certainly be sorted in  $\Theta(n \log n)$  time as a preprocessing step. As long as the rest of our procedure is at most  $\Theta(n \log n)$ , it will not increase the overall complexity.

We then divide the points into two subsets,  $S_1$  and  $S_2$ , each of which contains  $\frac{n}{2}$  points (which is easy since the points are sorted by  $x$  coordinate).

We then recursively find the closest pair of points in each subset  $S_1$  and  $S_2$ . If the distance between the closest pair in  $S_1 = d_1$  and the distance between the closest pair in  $S_2 = d_2$ . We then know that  $d = \min\{d_1, d_2\}$  is an upper bound on the minimum distance between any pair, but we still need to make sure we check for shorter distances where one point is in  $S_1$  and the other is in  $S_2$ .

The only points which might be closer together than distance  $d$  are those within a strip of width  $d$  from the dividing line between the subsets. For each point within that strip and within one subset, we potentially need to consider all points from within the strip within the other subset. That still sounds like a lot of work. The key observation is that for each point on one side, we have to consider points whose  $y$  coordinates are within  $d$ . This will mean at most 6 points from the other side, since if there are more points than that within  $d$  in the  $y$  coordinate on the other side, at least one pair from among that point would be closer than distance  $d$  to each other.

So how do we find those points to check? They can be found quickly if we also keep the points sorted in order by  $y$  coordinate. Still, this seems difficult but it can be done efficiently (see the text's description for details). Bottom line, we have a constant bound on the number of points to check (we can fairly easily show that each point need only be compared to up to 7 others) for each point in the boundary strip, so the "combine" operation ends up being linear time.

We end up with a recurrence:

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

which given an overall time of  $T(n) \in \Theta(n \log n)$ .

---

## Quickhull

The other computational geometry problem discussed in the text is called *quickhull* – an algorithm for finding the convex hull of a set of points. We will discuss the ideas (but not the details) and go through an example in class, but it is worth reading the section for all of the details.

---

## Binary Trees

No discussion of divide and conquer can proceed without a discussion of *binary trees* and algorithms on them.

You learned about tree data structures and lots of tree-related terminology in data structures. For now, we will just consider a few examples.

We can make analysis convenient by defining a binary tree as either the empty set, or a tree node plus two other binary trees: a left subtree and a right subtree.

To find the height of a tree:

```
ALGORITHM height(T)
  if (T is empty) return -1
  else return max(height(T.left), height(T.right)) + 1
```

To analyze this, we first note that the size of the problem is the number of nodes in our tree  $T$ , denoted  $n(T)$ .

The basic operation is either the comparison needed to find the max or the addition of 1 once we find the max. A recurrence for either:

$$A(n(T)) = A(n(T.left)) + A(n(T.right)) + 1 \quad \text{for } n(T) > 0, A(0) = 0$$

Note, however, that there are in fact more checks to see if the tree is empty than there are additions or comparisons. That check happens at the base case only.

The number of additions/max comparisons:  $A(n) = n$ , while the number of checks for an empty tree:  $C(n) = 2n + 1$ .

The other key divide and conquer operations are the recursively defined tree traversals, which you have seen previously:

1. *preorder*: visit the root, then visit the left subtree, then visit the right subtree.
2. *in-order* visit the left subtree, then visit the root, then visit the right subtree.
3. *postorder*: visit the left subtree, then visit the right subtree, then visit the root.

Pseudocode is straightforward, for example:

```
ALGORITHM inorder(T)
  if (T is not empty)
    inorder(T.left)
    visit(T.root)
    inorder(T.right)
```

Analysis is similar to that of height for this and for preorder and postorder. Each node is visited once.

## Strassen's Matrix Multiplication

Our text describes a divide and conquer multiplication of large numbers, but since you likely have seen or will see that in a discrete math course, we will move along to consider Strassen's matrix-matrix multiplication.

This algorithm improves upon the standard matrix-matrix multiplication by observing that the product of two  $2 \times 2$  matrices can be performed using 7 multiplications instead of the usual 8.

This in itself does not seem that significant, especially when we consider that the standard algorithm requires 4 additions, while Strassen's algorithm requires 18 additions/subtractions.

Strassen's method computes the product

$$C = A \cdot B$$

for  $2 \times 2$  matrices:

$$\begin{aligned} \begin{bmatrix} C_{00} & C_{01} \\ C_{10} & C_{11} \end{bmatrix} &= \begin{bmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{bmatrix} \cdot \begin{bmatrix} B_{00} & B_{01} \\ B_{10} & B_{11} \end{bmatrix} \\ &= \begin{bmatrix} M_1 + M_4 - M_5 + M_7 & M_3 + M_5 \\ M_2 + M_4 & M_1 + M_3 - M_2 + M_6 \end{bmatrix} \end{aligned}$$

where

$$\begin{aligned} M_1 &= (A_{00} + A_{11}) \cdot (B_{00} + B_{11}) \\ M_2 &= (A_{10} + A_{11}) \cdot B_{00} \\ M_3 &= A_{00} \cdot (B_{01} - B_{11}) \\ M_4 &= A_{11} \cdot (B_{10} - B_{00}) \\ M_5 &= (A_{00} + A_{01}) \cdot B_{11} \\ M_6 &= (A_{10} - A_{00}) \cdot (B_{00} + B_{01}) \\ M_7 &= (A_{01} - A_{11}) \cdot (B_{10} + B_{11}) \end{aligned}$$

The real benefit comes when we apply this idea to get a divide-and-conquer algorithm for multiplying matrices. To multiply  $2n \times n$  matrices, we break it into  $4 \frac{n}{2} \times \frac{n}{2}$  matrices and use Strassen's algorithm to multiply them.

Our recurrence for the number of multiplications with this approach:

$$M(n) = 7M(n/2) \quad \text{for } n > 1, \quad M(1) = 1.$$

From which the Master Theorem (or a backward substitution) will yield:



$$M(n) \in \Theta(n^{\log_2 7}) \approx n^{2.807}.$$

Which is definitely a slower rate of growth than  $\Theta(n^3)$ .

An analysis of the (larger) number of additions/subtractions results in the same efficiency class:  
 $A(n) \in \Theta(n^{\log_2 7})$ .

Many other approaches have been invented with even smaller rates of growth than Strassen's algorithm, but most have constant factors that make them impractical for real usage.

For us, it is interesting that an algorithm can exist for matrix-matrix multiplication that requires fewer than  $\Theta(n^3)$  element multiplications.