

Lecture 4 — Divide and Conquer Continued

Parallel and Sequential Data Structures and Algorithms, 15-210 (Qatar-Spring 2014)

Lectured by Kemal Oflazer — 21 January 2014

Material in this lecture:

- Maximum Contiguous Subsequence Sum (MCSS) problem with two different divide-and-conquer solutions.
- Solving recurrences using substitution
- Euclidean Traveling Salesperson Problem using divide and conquer.

1 Example I: The Maximum Contiguous Subsequence Sum Problem

For a sequence s of n elements, let's write s_i , $0 \leq i < n$, to denote the i 'th element of that sequence and use the “angle bracket” notation for writing the sequence, i.e., $s = \langle s_0, \dots, s_{n-1} \rangle$.

We say that a sequence s' is a *contiguous subsequence* of s if $s' = s_i, s_{i+1}, \dots, s_{i+k}$ for some k .

Example 1.1. For $s = \langle 1, -5, 2, -1, 3 \rangle$, here are some contiguous subsequences:

- $\langle 1 \rangle$,
- $\langle 2, -1, 3 \rangle$, and
- $\langle -5, 2 \rangle$.

The sequence $\langle 1, 2, 3 \rangle$ is not a contiguous subsequence, even though it is a sequence.

As the name suggests, the maximum-contiguous-subsequence problem requires finding the subsequence of a sequence of integers with maximum total sum. We can make this problem precise as follows.

Definition 1.2 (The Maximum Contiguous Subsequence Sum (MCSS) Problem). Given a sequence of numbers, the *maximum contiguous subsequence sum* problem is to find

$$\text{mcss}(s) = \max \left\{ \sum_{k=i}^j s_k : 0 \leq i \leq n-1, 0 \leq j \leq n-1 \right\}.$$

(i.e., the sum of the contiguous subsequence of s that has the largest value).

For an empty sequence, the maximum contiguous subsequence sum is $-\infty$.

[†]Lecture notes by Umut Acar, Guy E Blelloch, Margaret Reid-Miller, and Kanat Tangwongsan, with additional edits by Kemal Oflazer

Example 1.3. For $s = \langle 1, -5, 2, -1, 3 \rangle$, the maximum contiguous subsequence is, $\langle 2, -1, 3 \rangle$. Thus $mcss(s) = 4$.

1.1 Algorithm 1: Using Brute Force

Let's start by using the brute force algorithm to solve this problem.

Question 1.4. To apply brute force, where do we start?

We first start by identifying the structure of the output. In this case, this is just a number. So shall we enumerate all numbers (sums) and check that there is a subsequence that matches that number and we continue until we can no longer find a larger sum?

That is indeed what a literal interpretation of the brute-force technique would suggest.

Question 1.5. Would such an algorithm terminate?

No, because we may never know when to stop unless we know the result a priori, which we don't.

Question 1.6. Can we bound the result to guarantee non-termination?

We can by adding up all positive numbers in the sequence and using that bound but this can still be a very large bound. Furthermore our cost bounds would depend on the elements of the sequence rather than its length.

We thus have already encountered our first challenge. We can tackle this challenge by changing the result type.

Question 1.7. How can we change the result type to something that we can enumerate in finite time?

One natural choice would be to consider the contiguous subsequences directly. Indeed, we can change the result type by reducing this problem to another closely related problem: maximum contiguous subsequence. This problem requires not finding the sum but the sequence itself.

Question 1.8. Can you see how we can solve this problem reducing it to the maximum-contiguous-subsequence problem?

Reducing the problem requires no work because both operate on the same input, a sequence of the same type. If we have a solution to the maximum-contiguous-subsequence problem, then we can solve the maximum-contiguous-subsequence-sum problem, by simply computing the sum.

Question 1.9. *What is the work and span of the reduction?*

All we have to do is compute the sum, which we know by using reduce requires $O(n)$ work and $O(\log n)$ span.

We can now apply the brute-force-technique to the problem. Again, we have to enumerate all possible results.

Question 1.10. *What are all possible results for the maximum-contiguous-subsequence problem? How do we pick the best?*

They are all contiguous subsequences, which can be represented by a pair of integers (i, j) , $0 \leq i \leq j < n$. To pick the best, we simply compute their sum and pick the one with the largest. That is our algorithm for solving the maximum-contiguous-subsequence problem. Since we know how to solve the maximum-contiguous-subsequence-sum problem from this problem, we are done.

Question 1.11. *There is something strange about this algorithm. Do you see what?*

Our algorithm for solving the maximum-contiguous-subsequence problem, already computes the result for the maximum-contiguous-subsequence-sum problem. So the reduction does redundant work.

Question 1.12. *Can you see how we may eliminate this redundancy? (Perhaps by using another technique that we briefly talked about.)*

We can strengthen the problem by requiring it to return the subsequence in addition to the sum. This makes it possible to apply the brute-force technique directly to the problem, without having to go through the reduction.

Summary 1.13. *In summary, when trying to apply the brute-force technique, we have encountered a problem, which we solved by first reducing it to another problem. We then realized a redundancy in the resulting algorithm and eliminated that redundancy by strengthening the problem itself. This is a quite common route when designing a good algorithm: we may often find ourselves refining the problem and the solution until it is (close to) perfect.*

Specifying the algorithm and analyzing its cost. The pseudocode for this algorithm using our notation exactly matches the definition of the problem.

Question 1.14. *What is the work and span of the algorithm?*

For each subsequence $i..j$, we can compute its sum by applying a plus reduce. This does $O(j - i)$ work and has $O(\log(j - i))$ span. Furthermore, all the subsequences can be examined independently

in parallel (using, e.g., `tabulate`). This leads to the following bounds:

$$\begin{aligned} W(n) &= 1 + \sum_{1 \leq i \leq j \leq n} W_{\text{reduce}}(j-i) \leq 1 + n^2 \cdot W_{\text{reduce}}(n) = 1 + n^2 \cdot O(n) = O(n^3) \\ S(n) &= 1 + \max_{1 \leq i \leq j \leq n} S_{\text{reduce}}(j-i) \leq 1 + S_{\text{reduce}}(n) = O(\log n) \end{aligned}$$

These are cost bounds for enumerating over all possible subsequences and computing their sums. The final step of the brute-force solution is to find the max over these $O(n^2)$ combinations. Since max reduce for this step has $O(n^2)$ work and $O(\log n)$ span¹, the cost of the final step is subsumed by other costs analyzed above. Overall, we have an $O(n^3)$ -work $O(\log n)$ -span algorithm.

As you might notice already, this algorithm is clearly inefficient. We'll apply divide and conquer to come up with a more efficient solution.

1.2 Algorithm 2: Refining Brute Force with a Reduction

The brute-force algorithm does a lot of redundant work.

Question 1.15. *Can you see where the redundancy is?*

The algorithm repeats the same work many times. To see this let's consider the subsequences that start at some location, for example in the middle. For each position the algorithm considers many ending positions that differ by one, i.e., the sequences are all contained in each other and thus can potentially be computed much more efficiently, avoiding the redundancy.

Question 1.16. *Can you think of an algorithm for computing the maximum contiguous subsequence that starts at some particular position, say, i .*

Simply scan to the right as we compute the sum (point A). Continue as long as the sum increases. When it starts decreasing remember the sum, and continue, until another increase (in the sum) takes place. If it does update your sum and repeat (from point A).

Question 1.17. *What is the work and span of your algorithm?*

The algorithm requires linear work and span.

Question 1.18. *Can you now see a way to improve the brute-force algorithm by reducing our original problem to the problem of finding the problem “maximum contiguous subsequence with a given start”?*

¹Note that it takes the maximum over $\binom{n}{2} \leq n^2$ values, but since $\log n^a = a \log n$, this is simply $O(\log n)$

Yes, we can reduce the original problem to the “with-start” problem by considering all possible start positions and picking the best. The algorithm results in a linear reduction in work.

Exercise 1. Complete this algorithm and analyze its work and span carefully.

Question 1.19. Can this algorithm be optimal?

No, because we have only eliminated the aforementioned redundancy in one dimension, when solving the problem that starts at a particular position. There is still redundancy when solving for, for example, neighboring positions.

1.3 Algorithm 3: Divide And Conquer

To apply the divide-and-conquer technique, we first need to figure out how to divide up the input.

Question 1.20. Can you think of ways of dividing the input?

There are many possibilities. But dividing in the middle often yields the best bounds because it reduces the problem instance of all subproblems.

Question 1.21. Why should we care about reducing the size of all subproblems?

Reducing the problem size for all subproblems is important because the work and span is defined in terms of the sum and the maximum of all subproblems.

So let us divide the sequence in half and recursively solve the problem on both halves, as illustrated by the picture below.

$$\begin{array}{c} \langle \text{--- } s_l \text{ --- } \parallel \text{ --- } s_r \text{ --- } \rangle \\ \Downarrow \\ s_l = \underbrace{\langle \dots \rangle}_{\text{mcSS}(s_l)} \quad s_r = \underbrace{\langle \dots \rangle}_{\text{mcSS}(s_r)} \end{array}$$

Example 1.22. Let $s = \langle -2, 2, -2, -2, 3, 2 \rangle$. By using the approach, we divide the sequence into

$$s_l = \langle -2, 2, -2 \rangle$$

and

$$s_r = \langle -2, 3, 2 \rangle.$$

We can now solve each half to obtain 2 and 5 as the two solutions.

Now that we have the solutions from both halves, can we combine them to generate the solution for the whole problem?

Question 1.23. *Can you think of a way to solve the problem by using the solution for the two halves?*


As a starting point, we can simply calculate the maximum of the two halves, e.g., $\max(\text{mcss}(s_l), \text{mcss}(s_r))$. This answer is incorrect, unfortunately. We need a better understanding of the problem to devise a correct combine step.

Notice that the subsequence we're looking for has one of the three forms:

1. the maximum sum lies completely in the left subproblem
2. the maximum sum lies completely in the right subproblem
3. the maximum sum spans across the divide point

The first two cases are easy and have already been taken care of by the recursive calls. The more interesting case is when the largest sum goes between the two subproblems. How do we tackle this case?

It is not hard to see that the maximum sum going across the divide is the sum of the largest sum of a suffix on the left and the largest sum of a prefix on the right, as depicted below:



Question 1.24. *Can you find an algorithm for finding the subsequence with the largest sum that cuts across the divide? Hint: try the problem-reduction technique to reduce the problem to another one that we know.*

We can reduce this problem to the problem of “maximum contiguous sum with a start” by noticing that the maximum sum going across the divide is the largest sum of a suffix on the left and the largest sum of a prefix on the right.

Question 1.25. *Can you specify the reduction?*

The prefix of the right half is easy as it directly maps to the solution at position 0. For the left half, we have to reverse the sequence and ask for the position 0.

Having figured out a way to solve the problem of finding the best contiguous subsequence that cuts across the divide, we can write the pseudocode for our algorithm as follows.

```

1  fun mcss(s) =
2    case (showt s)
3    of EMPTY =  $-\infty$ 
4       | ELT( $x$ ) =  $x$ 
5       | NODE( $L, R$ ) =
6         let ( $m_L, m_R$ ) = (mcss( $L$ ) || mcss( $R$ ))
7            $m_A$  = bestAcross( $L, R$ )
8         in max{ $m_L, m_R, m_A$ }
9    end

```

Is this algorithm correct? Does it always find the maximum contiguous subsequence sum? Before we show a proof of correctness, what do we consider a proof that an algorithm is correct.

Question 1.26. *What technique can we use to show that the algorithm is correct?*

As we briefly talked about in a previous class, we can use the technique of strong induction, which enables us to assume that the theorem that we are trying to prove stands correct for all smaller subproblems.

Proofs of correctness. As was the case in 15-150, you are familiar with writing detailed proofs that reason about essentially every step down to the most elementary operations. You would prove your ML code was correct line by line. Although proving you code is correct is still important, in this class we will step up a level of abstraction and prove that the algorithms are correct. We still expect your proof to be rigorous. But we are more interested in seeing the critical steps highlighted and the standard or obvious steps summarized, with the idea being that if probed, you can easily provide detail on demand. The idea is that we want to make key ideas in an algorithm stand out as much as we can. It will be difficult for us to specify exactly how detailed we expect the proof to be, but you will pick it up by example.

We'll now prove that the mcss algorithm above computes the maximum contiguous subsequence sum. Specifically, we're proving the following theorem:

Theorem 1.27. *Let s be a sequence. The algorithm $mcss(s)$ returns the maximum contiguous subsequence sum in s —and returns $-\infty$ if s is empty.*

Proof. The proof will be by (strong) induction on length. We have two base cases: one when the sequence is empty and one when it has one element. On the empty sequence, it returns $-\infty$ as we stated. On any singleton sequence $\langle x \rangle$, the MCSS is x , for which

$$\max \left\{ \sum_{k=i}^j s_k : 0 \leq i < 1, 0 \leq j < 1 \right\} = \sum_{k=0}^0 s_0 = s_0 = x.$$

For the inductive step, let s be a sequence of length $n \geq 1$, and assume inductively that for any sequence s' of length $n' < n$, $\text{mcss}(s')$ correctly computes the maximum contiguous subsequence sum. Now consider the sequence s and let L and R denote the left and right subsequences resulted from dividing s in half (i.e., $\text{NODE}(L, R) = \text{showt } s$). Furthermore, let $s_{i..j}$ be any contiguous subsequence of s that has the largest sum, and this value is v . Note that the proof has to account for the possibility that there may be many other subsequences with equal sum. Every contiguous subsequence must start somewhere and end after it. We consider the following 3 possibilities corresponding to how the sequence $s_{i..j}$ lies with respect to L and R :

- If the sequence $s_{i..j}$ starts in L and ends in R . Then its sum equals its part in L (a suffix of L) and its part in R (a prefix of R). If we take the maximum of all suffixes in R and prefixes in L and add them this must equal the maximum of all contiguous sequences bridging the two since $\max\{a + b : a \in A, b \in B\} = \max\{a \in A\} + \max\{b \in B\}$. By assumption this equals the sum of $s_{i..j}$ which is v . Furthermore by induction m_L and m_R are sums of other subsequences so they cannot be any larger than v and hence $\max\{m_L, m_R, m_A\} = v$.
- If $s_{i..j}$ lies entirely in L , then it follows from our inductive hypothesis that $m_L = v$. Furthermore m_R and m_A correspond to the maximum sum of other subsequences, which cannot be larger than v . So again $\max\{m_L, m_R, m_A\} = v$.
- Similarly, if $s_{i..j}$ lies entirely in R , then it follows from our inductive hypothesis that $m_R = \max\{m_L, m_R, m_A\} = v$.

We conclude that in all cases, we return $\max\{m_L, m_R, m_A\} = v$, as claimed. \square

Cost analysis. What is the work and span of this algorithm? Before we analyze the cost, let's first remark that it turns out that we can compute the max prefix and suffix sums in parallel by using a primitive called *scan*. For now, we will take it for granted that they can be done in $O(n)$ work and $O(\log n)$ span. dividing takes $O(\log n)$ work and span. This yields the following recurrences:

$$\begin{aligned} W(n) &= 2W(n/2) + O(n) \\ S(n) &= S(n/2) + O(\log n) \end{aligned}$$

Using the definition of big- O , we know that

$$W(n) \leq 2W(n/2) + k_1 \cdot n + k_2,$$

where k_1 and k_2 are constants.

We have solved this recurrence using the recursion tree method. We can also arrive at the same answer by mathematical induction. If you want to go via this route (and you don't know the answer a priori), you'll need to guess the answer first and check it. This is often called the "substitution method." Since this technique relies on guessing an answer, you can sometimes fool yourself by giving a false proof. The following are some tips:

1. Spell out the constants. Do not use big- O —we need to be precise about constants, so big- O makes it super easy to fool ourselves.

2. Be careful that the inequalities always go in the right direction.
3. Add additional lower-order terms, if necessary, to make the induction go through.

Let's now redo the recurrences above using this method. Specifically, we'll prove the following theorem using (strong) induction on n .

Theorem 1.28. *Let a constant $k > 0$ be given. If $W(n) \leq 2W(n/2) + k \cdot n$ for $n > 1$ and $W(1) \leq k$ for $n \leq 1$, then we can find constants κ_1 and κ_2 such that*

$$W(n) \leq \kappa_1 \cdot n \log n + \kappa_2.$$

Proof. Let $\kappa_1 = 2k$ and $\kappa_2 = k$. For the base case ($n = 1$), we check that $W(1) = k \leq \kappa_2$. For the inductive step ($n > 1$), we assume that

$$W(n/2) \leq \kappa_1 \cdot \frac{n}{2} \log\left(\frac{n}{2}\right) + \kappa_2,$$

And we'll show that $W(n) \leq \kappa_1 \cdot n \log n + \kappa_2$. To show this, we substitute an upper bound for $W(n/2)$ from our assumption into the recurrence, yielding

$$\begin{aligned} W(n) &\leq 2W(n/2) + k \cdot n \\ &\leq 2(\kappa_1 \cdot \frac{n}{2} \log\left(\frac{n}{2}\right) + \kappa_2) + k \cdot n \\ &= \kappa_1 n (\log n - 1) + 2\kappa_2 + k \cdot n \\ &= \kappa_1 n \log n + \kappa_2 + (k \cdot n + \kappa_2 - \kappa_1 \cdot n) \\ &\leq \kappa_1 n \log n + \kappa_2, \end{aligned}$$

where the final step follows because $k \cdot n + \kappa_2 - \kappa_1 \cdot n \leq 0$ as long as $n > 1$. □

Question 1.29. *Using divide and conquer, we were able to reduce work to $O(n \log n)$. Can you see where the savings came from by comparing this algorithm to the refined brute-force algorithm that we have considered?*

1.4 Algorithm 3: Divide And Conquer — Version 2.0

We have progressively improved our algorithm from cubic to near linear work. Is it possible to do better than $O(n \log n)$ work using divide and conquer?

Question 1.30. *In general, how do we know that we have a good algorithm?*

We can determine whether we have made enough progress or not by comparing the cost of our algorithm to a lower bound.

Question 1.31. *What is a lower bound for this problem?*

We cannot do less than linear work because we have to inspect each element of the sequence at least once to determine whether it would contribute to the result or not.

Question 1.32. *Can you see a way that we might be able to reduce the work further? Is there some redundancy in our divide-and-conquer algorithm?*

Our divide-and-conquer algorithm has another redundancy: the maximum prefix and maximum suffix must have been computed recursively to solve the subproblems. Thus, we should be avoid re-computing them.

Question 1.33. *How can we avoid re-computing the maximum prefix and suffix?*

Since these should be computed as part of solving the subproblems, we should be able to return them from the recursive calls.

Question 1.34. *Let's take a step back and think about what we did. We now changed the problem slightly so that we return more information. What is the name of this technique?*

This is another application of the strengthening technique same. Now that we have figured out how to strengthen the problem let's go and solve it.

Question 1.35. *Can you see how we can update our divide and conquer algorithm to return also the maximum prefix and suffix in addition to maximum contiguous subsequence.*

So, we need to return a total of three values: the max subsequence sum, the max prefix sum, and the max suffix sum.

Having this information from the subproblems, we can now produce a similar answer tuple for all levels up, in constant work and span per level. More specifically, we strengthen our problem to return a 3-tuple (mcss, max-prefix, max-suffix), and if the recursive calls return (m_1, p_1, s_1, t_1) and (m_2, p_2, s_2) , then we return

$$(\max(s_1 + p_2, m_1, m_2), p_1, s_2).$$

Let's check that our answer makes sense:

Question 1.36. *Don't we have consider the case when s_1 or p_2 is the maximum?*

No, because that case is included in m_1 and m_2 .

Question 1.37. *Are our prefixes and suffixes correct? Can we not have a bigger prefix that contains all of the first sequence?*

Yes, we can. Indeed, we have a bug in our algorithm. We are not returning the prefix and the suffix correctly.


Question 1.38. *Can you think of a way to fix this problem?*

We also need to return the total for each subsequence so that we can compute the maximum prefix and suffix correctly. Thus, we need to return a total of four values: the max subsequence sum, the max prefix sum, the max suffix sum, and the overall sum. Having this information from the subproblems is enough to produce a similar answer tuple for all levels up, in constant work and span per level. More precisely, we strengthen our problem to return a 4-tuple ($mcss$, $max\text{-}prefix$, $max\text{-}suffix$, $total$), and if the recursive calls return (m_1, p_1, s_1, t_1) and (m_2, p_2, s_2, t_2) , then we return

$$(\max(s_1 + p_2, m_1, m_2), \max(p_1, t_1 + p_2), \max(s_1 + t_2, s_2), t_1 + t_2)$$


The following figures clarify these a bit more:

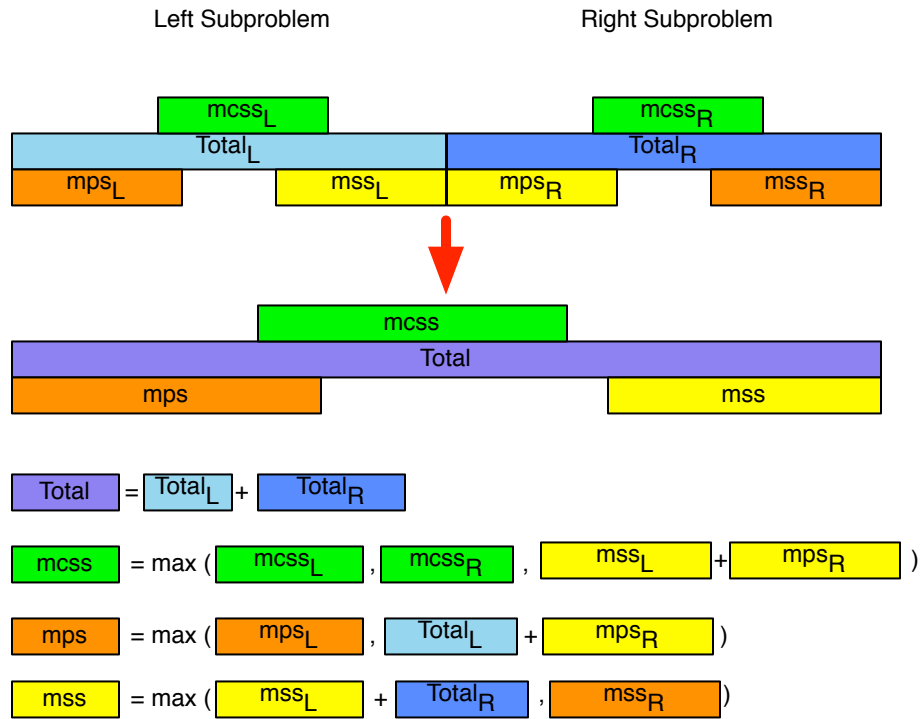
Left Subproblem



Left Subproblem

Right Subproblem





This gives the following pseudocode:

```

1 fun mcss(a) =
2 let
3   fun mcss'(a)
4     case (showt a)
5       of EMPTY =  $(-\infty, -\infty, -\infty, 0)$ 
6         | ELT(x) =  $(x, x, x, x)$ 
7         | NODE(L,R) =
8           let
9              $((m_1, p_1, s_1, t_1), (m_2, p_2, s_2, t_2)) = (mcss'(L) \parallel mcss'(R))$ 
10            in
11               $(\max(s_1 + p_2, m_1, m_2), \max(p_1, t_1 + p_2), \max(s_1 + t_2, s_2), t_1 + t_2)$ 
12            end
13      (m,p,s,t) = mcss'(a)
14 in m end

```

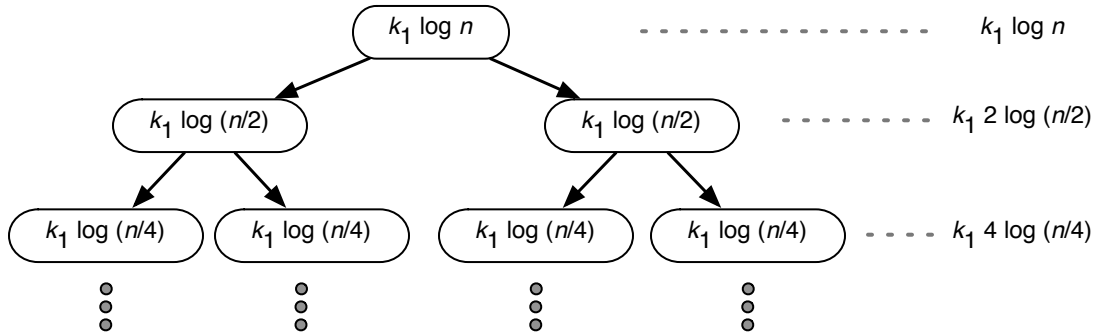
You should verify the base cases are doing the right thing. The SML code for this algorithm is at the end of these notes.

Cost Analysis. Assuming `showt` takes $O(\log n)$ work and span, we have the recurrences

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

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

Note that the span is the same as before, so we'll focus on analyzing the work. Using the tree method, we have



Therefore, the total work is upper-bounded by

$$W(n) \leq \sum_{i=0}^{\log n} k_1 2^i \log(n/2^i)$$

It is not so obvious to what this sum evaluates. The substitution method seems to be more convenient. We'll make a guess that $W(n) \leq \kappa_1 n - \kappa_2 \log n - \kappa_3$. More formally, we'll prove the following theorem:

Theorem 1.39. *Let $k > 0$ be given. If $W(n) \leq 2W(n/2) + k \cdot \log n$ for $n > 1$ and $W(n) \leq k$ for $n \leq 1$, then we can find constants κ_1 , κ_2 , and κ_3 such that*

$$W(n) \leq \kappa_1 \cdot n - \kappa_2 \cdot \log n - \kappa_3.$$

Proof. Let $\kappa_1 = 3k$, $\kappa_2 = k$, $\kappa_3 = 2k$. We begin with the base case. Clearly, $W(1) = k \leq \kappa_1 - \kappa_3 = 3k - 2k = k$. For the inductive step, we substitute the inductive hypothesis into the recurrence and obtain

$$\begin{aligned} W(n) &\leq 2W(n/2) + k \cdot \log n \\ &\leq 2(\kappa_1 \frac{n}{2} - \kappa_2 \log(n/2) - \kappa_3) + k \cdot \log n \\ &= \kappa_1 n - 2\kappa_2(\log n - 1) - 2\kappa_3 + k \cdot \log n \\ &= (\kappa_1 n - \kappa_2 \log n - \kappa_3) + (k \log n - \kappa_2 \log n + 2\kappa_2 - \kappa_3) \\ &\leq \kappa_1 n - \kappa_2 \log n - \kappa_3, \end{aligned}$$

where the final step uses the fact that $(k \log n - \kappa_2 \log n + 2\kappa_2 - \kappa_3) = (k \log n - k \log n + 2k - 2k) = 0 \leq 0$ by our choice of κ 's. \square

Finishing the tree method. It is possible to solve the recurrence directly by evaluating the sum we established using the tree method. We didn't cover this in lecture, but for the curious, here's how you can "tame" it.

$$\begin{aligned}
 W(n) &\leq \sum_{i=0}^{\log n} k_1 2^i \log(n/2^i) \\
 &= \sum_{i=0}^{\log n} k_1 (2^i \log n - i \cdot 2^i) \\
 &= k_1 \left(\sum_{i=0}^{\log n} 2^i \right) \log n - k_1 \sum_{i=0}^{\log n} i \cdot 2^i \\
 &= k_1 (2n - 1) \log n - k_1 \sum_{i=0}^{\log n} i \cdot 2^i.
 \end{aligned}$$

We're left with evaluating $s = \sum_{i=0}^{\log n} i \cdot 2^i$. Observe that if we multiply s by 2, we have

$$2s = \sum_{i=0}^{\log n} i \cdot 2^{i+1} = \sum_{i=1}^{1+\log n} (i-1)2^i,$$

so then

$$\begin{aligned}
 s &= 2s - s = \sum_{i=1}^{1+\log n} (i-1)2^i - \sum_{i=0}^{\log n} i \cdot 2^i \\
 &= ((1 + \log n) - 1)2^{1+\log n} - \sum_{i=1}^{\log n} 2^i \\
 &= 2n \log n - (2n - 2).
 \end{aligned}$$

Substituting this back into the expression we derived earlier, we have $W(n) \leq k_1(2n - 1) \log n - 2k_1(n \log n - n + 1) \in O(n)$ because the $n \log n$ terms cancel.

2 Example II: The Euclidean Traveling Salesperson Problem

We'll now turn to another example of divide and conquer. In this example, we will apply it to devise a heuristic method for an **NP**-hard problem. The problem we are concerned with is a variant of the traveling salesperson problem (TSP) from Lecture 1. This variant is known as the Euclidean traveling salesperson (eTSP) problem because in this problem, the points (e.g., cities, nodes, and vertices) lie in a Euclidean space and the distance measure is the Euclidean measure. More specifically, we are interested in the planar version of the eTSP problem, defined as follows:


Definition 2.1 (The Planar Euclidean Traveling Salesperson Problem). Given a set of points P in the 2-d plane, the *planar Euclidean traveling salesperson* (eTSP) problem is to find a tour of minimum total distance that visits all points in P exactly once, where the distance between points (x_1, y_1) and (x_2, y_2) is the Euclidean (i.e. ℓ_2) distance $(\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2})$.

As with the TSP, it is **NP**-hard, but this problem is easier to approximate.² Here is a heuristic divide-and-conquer algorithm that does quite well in practice. In a few weeks, we will see another algorithm based on Minimum Spanning Trees (MST) that gives a constant-approximation guarantee. This divide-and-conquer is more interesting than the ones we have done so far *because it does work both before and after the recursive calls*. Also, as we will see, the recurrence it generates is root dominated.

The basic idea is to split the points by a cut in the plane, solve the TSP on the two halves, and then somehow merge the solutions. For the cut, we can pick a cut that is orthogonal to the coordinate lines. In particular, we can find in which of the two dimensions the points have a larger spread, and then find the median point along that dimension. We'll split just below that point.

To merge the solutions we join the two cycles by making an edge swap

²Unlike the TSP problem, which only has constant approximations, it is known how to approximate this problem to an arbitrary but fixed constant accuracy ϵ in polynomial time (the exponent of n has $1/\epsilon$ dependency). That is, such an algorithm is capable of producing a solution that has length at most $(1 + \epsilon)$ times the length of the best tour.



Consider the figure above: We can remove the edges $e_\ell = (u_\ell, v_\ell)$ and $e_r = (u_r, v_r)$ and either add a new pair of edges (u_ℓ, u_r) and (v_ℓ, v_r) (upper right) or add a new pair of edges (u_ℓ, v_r) and (v_ℓ, u_r) (lower right). In the first case, the increase in the tour cost would be

$$\|u_\ell - u_r\| + \|v_\ell - v_r\| - \|u_\ell - v_\ell\| - \|u_r - v_r\|$$

and in the second case, the increase in the tour cost would be

$$\|u_\ell - v_r\| + \|v_\ell - u_r\| - \|u_\ell - v_\ell\| - \|u_r - v_r\|$$

where $\|u - v\|$ is the Euclidean distance between points u and v . Naturally we are interested in the minimum increase so we only need to consider one or the other. So we define the swap cost as

$$\text{swapCost}((u_\ell, v_\ell), (u_r, v_r)) = \min(\|u_\ell - u_r\| + \|v_\ell - v_r\|, \|u_\ell - v_r\| + \|v_\ell - u_r\|) - \|u_\ell - v_\ell\| - \|u_r - v_r\|$$

Here is the pseudocode for the algorithm

```

1  fun eTSP(P) =
2    case (|P|)
3      of 0,1  $\Rightarrow$  raise TooSmall
4      | 2  $\Rightarrow$  {(P[0], P[1]), (P[1], P[0])}
5      | n  $\Rightarrow$  let
6        (Pℓ, Pr) = splitLongestDim(P)
7        (L, R) = (eTSP(Pℓ) || eTSP(Pr))
8        (c, (e'ℓ, e'r)) = minval <#1 {(swapCost(eℓ, er), (eℓ, er)) : eℓ ∈ L, er ∈ R}
9      in
10       swapEdges(append(L, R), e'ℓ, e'r)
11     end

```


The function $\text{swapEdges}(E, e_\ell, e_r)$ finds the edges e_ℓ and e_r in E and swaps the endpoints (as explained above there are two ways to swap, so the minimum is picked).

Now let's analyze the cost of this algorithm in terms of work and span. We have

$$\begin{aligned} W(n) &= 2W(n/2) + O(n^2) \\ S(n) &= S(n/2) + O(\log n) \end{aligned}$$

We have already seen the recurrence $S(n) = S(n/2) + O(\log n)$, which solves to $O(\log^2 n)$. Here we'll focus on solving the work recurrence.

In anticipation of recurrences that you'll encounter later in class, we'll attempt to solve a more general form of recurrences. Let $\varepsilon > 0$ be a constant. We'll solve the recurrence

$$W(n) = 2W(n/2) + k \cdot n^{1+\varepsilon}$$

by the substitution method.

Theorem 2.2. *Let $\varepsilon > 0$. If $W(n) \leq 2W(n/2) + k \cdot n^{1+\varepsilon}$ for $n > 1$ and $W(1) \leq k$ for $n \leq 1$, then for some constant κ ,*

$$W(n) \leq \kappa \cdot n^{1+\varepsilon}.$$

Proof. Let $\kappa = \frac{1}{1-1/2^\varepsilon} \cdot k$. The base case is easy: $W(1) = k \leq \kappa_1$ as $\frac{1}{1-1/2^\varepsilon} \geq 1$. For the inductive step, we substitute the inductive hypothesis into the recurrence and obtain

$$\begin{aligned} W(n) &\leq 2W(n/2) + k \cdot n^{1+\varepsilon} \\ &\leq 2\kappa \left(\frac{n}{2}\right)^{1+\varepsilon} + k \cdot n^{1+\varepsilon} \\ &= \kappa \cdot n^{1+\varepsilon} + \left(2\kappa \left(\frac{n}{2}\right)^{1+\varepsilon} + k \cdot n^{1+\varepsilon} - \kappa \cdot n^{1+\varepsilon}\right) \\ &\leq \kappa \cdot n^{1+\varepsilon}, \end{aligned}$$

where in the final step, we argued that

$$\begin{aligned} 2\kappa \left(\frac{n}{2}\right)^{1+\varepsilon} + k \cdot n^{1+\varepsilon} - \kappa \cdot n^{1+\varepsilon} &= \kappa \cdot 2^{-\varepsilon} \cdot n^{1+\varepsilon} + k \cdot n^{1+\varepsilon} - \kappa \cdot n^{1+\varepsilon} \\ &= \kappa \cdot 2^{-\varepsilon} \cdot n^{1+\varepsilon} + (1 - 2^{-\varepsilon})\kappa \cdot n^{1+\varepsilon} - \kappa \cdot n^{1+\varepsilon} \\ &\leq 0. \end{aligned}$$

□

Solving the recurrence directly. Alternatively, we could use the tree method and evaluate the sum directly. As argued before, the recursion tree here has depth $\log n$ and at level i (again, the root is at level 0), we have 2^i nodes, each costing $k \cdot (n/2^i)^{1+\varepsilon}$. Thus, the total cost is

$$\begin{aligned} \sum_{i=0}^{\log n} k \cdot 2^i \cdot \left(\frac{n}{2^i}\right)^{1+\varepsilon} &= k \cdot n^{1+\varepsilon} \cdot \sum_{i=0}^{\log n} 2^{-i \cdot \varepsilon} \\ &\leq k \cdot n^{1+\varepsilon} \cdot \sum_{i=0}^{\infty} 2^{-i \cdot \varepsilon}. \end{aligned}$$

But the infinite sum $\sum_{i=0}^{\infty} 2^{-i \cdot \varepsilon}$ is at most $\frac{1}{1-1/2^\varepsilon}$. Hence, we conclude $W(n) \in O(n^{1+\varepsilon})$.

3 SML Code

Note that the following code for Algorithm 3 for the MCSS problem uses SML *records* to give names to each of the four fields. When passing around multiple values, such records can make your code more clear and less prone to errors than using unnamed tuples. It is easy to forget the order you store things in tuples, and harder for people who are reading your code to know the order. Fields of a record can be extracted using pattern matching. Unlike tuples, you can extract only the fields that you need by including ellipse (...) at the end of the pattern. As with tuples you can access a single field using `#fieldname record`, as can be seen at the end of the code.

```

functor mcssDivConq(Seq : SEQUENCE) =
struct
  fun MCSS (A) =
  let
    fun MCSS' A =
      case Seq.showt A
      of Seq.ELT(v) => {mcss = v, prefix = v, suffix = v, total = v}
       | Seq.NODE(A1,A2) =>
        let
          val (B1, B2) = Primitives.par(fn () => MCSS'(A1),
                                         fn () => MCSS'(A2))
          val {mcss = M1, prefix = P1, suffix = S1, total = T1} = B1
          val {mcss = M2, prefix = P2, suffix = S2, total = T2} = B2
        in
          {mcss = Int.max(S1 + P2, Int.max(M1, M2)),
           prefix = Int.max(P1, T1 + P2),
           suffix = Int.max(S2, S1 + T2),
           total = T1 + T2}
        end
      in
        case Seq.showt A
        of Seq.EMPTY => NONE
         | _ => SOME(#mcss(MCSS' A))
      end
  end
end

```