

## Lecture 6

# Concrete models and tight upper/lower bounds

### 6.1 Overview

In this lecture, we will examine some simple, concrete models of computation, each with a precise definition of what counts as a step, and try to get tight upper and lower bounds for a number of problems. Unlike many of the other lectures, in this one we will not be using  $O$ ,  $\Theta$ , and  $\Omega$ , and we will instead try to examine exact quantities as much as possible. Specific models and problems examined in this lecture include:

- The number of exchanges needed to sort an array.
- The number of comparisons needed to find the largest and second-largest elements in an array, and a more precise look at the number of comparisons needed to sort.
- The number of probes into a graph needed to determine if the graph is connected (the evasiveness of connectivity).

### 6.2 Terminology and setup

In this lecture, we will look at (worst-case) upper and lower bounds for a number of problems in several different concrete models. Each model will specify exactly what operations may be performed on the input, and how much they cost. Typically, each model will have some operations that cost 1 step (like performing a comparison, or swapping a pair of elements), some that are free, and some that are not allowed at all.

By an *upper bound* of  $f(n)$  for some problem, we mean that there exists an algorithm that takes at most  $f(n)$  steps on any input of size  $n$ . By a *lower bound* of  $g(n)$ , we mean that for any algorithm there exists an input on which it takes at least  $g(n)$  steps. The reason for this terminology is that if we think of our goal as being to understand the “true complexity” of each problem, measured in terms of the best possible worst-case guarantee achievable by any algorithm, then an upper bound of  $f(n)$  and lower bound of  $g(n)$  means that the true complexity is somewhere between  $g(n)$  and  $f(n)$ .

## 6.3 Sorting in the exchange model

Consider a shelf containing  $n$  unordered books to be arranged alphabetically. In each step, we can swap any pair of books we like. How many swaps do we need to sort all the books? Formally, we are considering the problem of *sorting* in the *exchange model*.

**Definition 6.1** *In the exchange model, an input consists of an array of  $n$  items, and the only operation allowed on the items is to swap a pair of them at a cost of 1 step. All other (planning) work is free: in particular, the items can be examined and compared to each other at no cost.*

**Question:** how many exchanges are necessary (lower bound) and sufficient (upper bound) in the exchange model to sort an array of  $n$  items in the worst case?

**Claim 6.1 (Upper bound)**  $n - 1$  exchanges is sufficient.

**Proof:** To prove an upper bound of  $n - 1$  we just need to give an algorithm. For instance, consider the algorithm that in step 1 puts the smallest item in location 1, swapping it with whatever was originally there. Then in step 2 it swaps the second-smallest item with whatever is currently in location 2, and so on (if in step  $k$ , the  $k$ th-smallest item is already in the correct position then we just do a no-op). No step ever undoes any of the previous work, so after  $n - 1$  steps, the first  $n - 1$  items are in the correct position. This means the  $n$ th item must be in the correct position too. ■

But are  $n - 1$  exchanges necessary in the worst-case? If  $n$  is even, and no book is in its correct location, then  $n/2$  exchanges are clearly necessary to “touch” all books. But can we show a better lower bound than that?

**Claim 6.2 (Lower bound)** *In fact,  $n - 1$  exchanges are necessary, in the worst case.*

**Proof:** Here is how we can see it. Create a graph in which a directed edge  $(i, j)$  means that that the book in location  $i$  must end up at location  $j$ . For instance, consider the example in Figure 6.1. Note that this is a special kind of directed graph: it is a permutation — a set of cycles. In particular, every book points to *some* location, perhaps its own location, and every location is pointed to by exactly one book. Now consider the following points:

1. What is the effect of exchanging any two elements (books) that are in the same cycle?

**Answer:** Suppose the graph had edges  $(i_1, j_1)$  and  $(i_2, j_2)$  and we swap the elements in locations  $i_1$  and  $i_2$ . Then this causes those two edges to be replaced by edges  $(i_2, j_1)$  and  $(i_1, j_2)$  because now it is the element in location  $i_2$  that needs to go to  $j_1$  and the element in  $i_1$  that needs to go to  $j_2$ . This means that if  $i_1$  and  $i_2$  were in the same cycle, that cycle now becomes two disjoint cycles.

2. What is the effect of exchanging any two elements that are in different cycles?

**Answer:** If we swap elements  $i_1$  and  $i_2$  that are in different cycles, then the same argument as above shows that this merges those two cycles into one cycle.




Figure 6.1: Graph for input [6 3 4 5 2 1 7]

3. How many cycles are in the final sorted array?

Answer: The final sorted array has  $n$  cycles.

Putting the above 3 points together, suppose we begin with an array consisting of a single cycle, such as  $[n, 1, 2, 3, 4, \dots, n-1]$ . Each operation at best increases the number of cycles by 1 and in the end we need to have  $n$  cycles. So, this input requires  $n-1$  operations. ■

## 6.4 The comparison model

Let's now go back to the comparison model of computation we looked at last time.

**Definition 6.2** *In the comparison model, we have an input containing  $n$  items, but the only information the algorithm can get about the items is by comparing pairs of them, where each comparison returns YES or NO. Each comparison costs 1 step. But exchanges and moves are free.*

In the last lecture, we looked at sorting in the comparison model, and gave a lower bound of  $\lg(n!)$  on the number of comparisons needed.<sup>1</sup> Let us begin with a simple generalization. Suppose you have some problem where there are  $M$  possible different outputs the algorithm might be required to produce; e.g., for sorting,  $M = n!$ . Then, we have a worst-case lower bound of  $\lg M$ . The reason is that the algorithm needs to find out which of these  $M$  outputs is the right one, and each YES/NO question could be answered in a way that removes at most half of the possibilities remaining from consideration. So, in the worst case, it takes at least  $\lg M$  steps to find the right answer.

Just to get a better handle on what exactly  $\lg(n!)$  looks like, since today's theme is tight bounds, we can use the fact that  $n! \in [(n/e)^n, n^n]$ . So this means that:

$$\begin{aligned} n \lg n - n \lg e &< \lg(n!) < n \lg n \\ n \lg n - 1.433n &< \lg(n!) < n \lg n. \end{aligned}$$

Since  $1.433n$  is a low-order term, sometimes people will write this fact this as:  $\lg(n!) = (n \lg n)(1 - o(1))$ , meaning that the ratio between  $\lg(n!)$  and  $n \lg n$  goes to 1 as  $n$  goes to infinity.

<sup>1</sup>We will use “lg” to mean “log<sub>2</sub>”.

### 6.4.1 Almost-tight upper-bounds for comparison-based sorting

Assume  $n$  is a power of 2 — in fact, let's assume this for the entire rest of today's lecture. Can you think of an algorithm that makes at most  $n \lg n$  comparisons, and so is tight in the leading term? In fact, there are several algorithms, including:

**Binary insertion sort** If we perform insertion-sort, using binary search to insert each new element, then the number of comparisons made is at most  $\sum_{k=2}^n \lceil \lg k \rceil \leq n \lg n$ . Note that insertion-sort spends a lot in moving items in the array to make room for each new element, and so is not especially efficient if we count movement cost as well, but it does well in terms of comparisons.

**Mergesort** Merging two lists of  $n/2$  elements each requires at most  $n - 1$  comparisons. So, unrolling the recurrence we get  $(n - 1) + 2(n/2 - 1) + 4(n/4 - 1) + \dots + n/2(2 - 1) = n \lg n - (n - 1) < n \lg n$ .

### 6.4.2 Finding the maximum of $n$ elements

How many comparisons are necessary and sufficient to find the maximum of  $n$  elements, in the comparison model of computation?

**Claim 6.3 (Upper bound)**  $n - 1$  comparisons are sufficient to find the maximum of  $n$  elements.

**Proof:** Just scan left to right, keeping track of the largest element so far. This makes at most  $n - 1$  comparisons. ■

Now, let's try for a lower bound. One simple lower bound is that since there are  $n$  possible answers for the location of the minimum element, our previous argument gives a lower bound of  $\lg n$ . But clearly this is not at all tight. In fact, we can give a better lower bound of  $n - 1$ .

**Claim 6.4 (Lower bound)**  $n - 1$  comparisons are needed in the worst-case to find the maximum of  $n$  elements.

**Proof:** Suppose some algorithm  $\mathcal{A}$  claims to find the maximum of  $n$  elements using less than  $n - 1$  comparisons. Consider an arbitrary input of  $n$  distinct elements, and construct a graph in which we join two elements by an edge if they are compared by  $\mathcal{A}$ . If fewer than  $n - 1$  comparisons are made, then this graph must have at least two components. Suppose now that algorithm  $\mathcal{A}$  outputs some element  $u$  as the maximum, where  $u$  is in some component  $C_1$ . In that case, pick a different component  $C_2$  and add a large positive number (e.g., the value of  $u$ ) to every element in  $C_2$ . This process does not change the result of any comparison made by  $\mathcal{A}$ , so on this new set of elements, algorithm  $\mathcal{A}$  would still output  $u$ . Yet this now ensures that  $u$  is not the maximum, so  $\mathcal{A}$  must be incorrect. ■

Since the upper and lower bounds are equal, these bounds are tight.

### 6.4.3 Finding the second-largest of $n$ elements

How many comparisons are necessary (lower bound) and sufficient (upper bound) to find the second largest of  $n$  elements? Again, let us assume that all elements are distinct.

**Claim 6.5 (Lower bound)**  $n - 1$  comparisons are needed in the worst-case to find the second-largest of  $n$  elements.

**Proof:** The same argument used in the lower bound for finding the maximum still holds. ■

Let us now work on finding an upper bound. Here is a simple one to start with.


**Claim 6.6 (Upper bound #1)**  $2n - 3$  comparisons are sufficient to find the second-largest of  $n$  elements.

**Proof:** Just find the largest using  $n - 1$  comparisons, and then the largest of the remainder using  $n - 2$  comparisons, for a total of  $2n - 3$  comparisons. ■

We now have a gap:  $n - 1$  versus  $2n - 3$ . It is not a huge gap: both are  $\Theta(n)$ , but remember today's theme is tight bounds. So, which do you think is closer to the truth? It turns out, we can reduce the upper bound quite a bit:

**Claim 6.7 (Upper bound #2)**  $n + \lg n - 2$  comparisons are sufficient to find the second-largest of  $n$  elements.

**Proof:** As a first step, let's find the maximum element using  $n - 1$  comparisons, but in a tennis-tournament or playoff structure. That is, we group elements into pairs, finding the maximum in each pair, and recurse on the maxima. E.g.,



Now, given just what we know from comparisons so far, what can we say about possible locations for the second-highest number (i.e., the second-best player)? The answer is that the second-best must have been directly compared to the best, and lost.<sup>2</sup> This means there are only  $\lg n$  possibilities for the second-highest number, and we can find the maximum of them making only  $\lg(n) - 1$  more comparisons. ■

---

<sup>2</sup>Apparently first person to have pointed this out was Charles Dodgson (better known as Lewis Carroll!), writing about the proper way to award prizes in lawn tennis tournaments.

At this point, we have a lower bound of  $n - 1$  and an upper bound of  $n + \lg(n) - 2$ , so they are nearly tight. It turns out that, in fact, the lower bound can be improved to exactly meet the upper bound.<sup>3</sup>

## 6.5 Query models, and the evasiveness of connectivity

To finish with something totally different, let's look at the query complexity of determining if a graph is connected. Assume we are given the adjacency matrix  $G$  for some  $n$ -node graph. That is,  $G[i, j] = 1$  if there is an edge between  $i$  and  $j$ , and  $G[i, j] = 0$  otherwise. We consider a model in which we can *query* any element of the matrix  $G$  in 1 step. All other computation is free. That is, imagine the graph matrix has values written on little slips of paper, face down. In one step we can turn over any slip of paper. How many slips of paper do we need to turn over to tell if  $G$  is connected?

**Claim 6.8 (Easy upper bound)**  $n(n-1)/2$  queries are sufficient to determine if  $G$  is connected.

**Proof:** This just corresponds to querying every pair  $(i, j)$ . Once we have done that, we know the entire graph and can just compute for free to see if it is connected. ■

Interestingly, it turns out the simple upper-bound of querying every edge is a lower bound too. Because of this, connectivity is called an “evasive” property of graphs.

**Theorem 6.9 (Lower bound)**  $n(n-1)/2$  queries are necessary to determine connectivity in the worst case.

**Proof:** Here is the strategy for the adversary: when the algorithm asks us to flip over a slip of paper, we return the answer 0 *unless* that would force the graph to be disconnected, in which case we answer 1. (It is not important to the argument, but we can figure this out by imagining that all un-turned slips of paper are 1 and seeing if that graph is connected.) Now, here is the key claim:

*Claim:* we maintain the invariant that for any un-asked pair  $(u, v)$ , the graph revealed so far has no path from  $u$  to  $v$ .

*Proof of claim:* If there was, consider the last edge  $(u', v')$  revealed on that path. We could have answered 0 for that and kept the same connectivity in the graph by having an edge  $(u, v)$ . So, that contradicts the definition of our adversary strategy.

Now, to finish the proof: Suppose an algorithm halts without examining every pair. Consider some unasked pair  $(u, v)$ . If the algorithm says “connected,” we reveal all-zeros for the remaining unasked edges and then there is no path from  $u$  to  $v$  (by the key claim) so the algorithm is wrong. If the algorithm says “disconnected,” we reveal all-ones for the remaining edges, and the algorithm is wrong by definition of our adversary strategy. So, the algorithm must ask for all edges. ■

We'll see more arguments like this when we talk about spanning trees later on in the course.

---

<sup>3</sup>First shown by Kislitsyn (1964).