Lecture 19

NP-Completeness I

19.1 Overview

In the past few lectures we have looked at increasingly more expressive problems that we were able to solve using efficient algorithms. In this lecture we introduce a class of problems that are so expressive — they are able to model *any* problem in an extremely large class called **NP** — that we believe them to be *intrinsically unsolvable by polynomial-time algorithms*. These are the **NP-complete** problems. What is particularly surprising about this class is that they include many problems that at first glance appear to be quite benign. Specific topics in this lecture include:

- Reductions and expressiveness

- Informal definitions and the ESP problem

- Formal definitions: decision problems, P and NP.

- Circuit-SAT and 3-SAT

19.2 Introduction: Reduction and Expressiveness

In the last few lectures have seen a series of increasingly more expressive problems: network flow, min cost max flow, and finally linear programming. These problems have the property that you can code up a lot of different problems in their "language". So, by solving these well, we end up with important tools we can use to solve other problems.

To talk about this a little more precisely, it is helpful to make the following definitions:

Definition 19.1 *We say that an algorithm runs in* **Polynomial Time** *if, for some constant c, its running time is $O(n^c)$, where n is the size of the input.*

In the above definition, "size of input" means "number of bits it takes to write the input down". So, to be precise, when defining a problem and asking whether or not a certain algorithm runs in polynomial time, it is important to say how the input is given. For instance, the basic Ford-Fulkerson algorithm is *not* a polynomial-time algorithm for network flow when edge capacities are written in binary, but both of the Edmonds-Karp algorithms *are* polynomial-time.

Definition 19.2 *A Problem A is* **poly-time reducible** *to problem B (written as $A \leq_p B$) if we can solve problem A in polynomial time given a polynomial time black-box algorithm for problem B. Problem A is* **poly-time equivalent** *to problem B ($A =_p B$) if $A \leq_p B$ and $B \leq_p A$.*

For instance, we gave an efficient algorithm for Bipartite Matching by showing it was poly-time reducible to Max Flow. Notice that it could be that $A \leq_p B$ and yet our fastest algorithm for solving problem A might be slower than our fastest algorithm for solving problem B (because our reduction might involve several calls to the algorithm for problem B, or might involve blowing up the input size by a polynomial but still nontrivial amount).

19.3   Our first NP-Complete Problem: ESP

Many of the problems we would like to solve have the property that if someone handed us a solution, we could at least check if the solution was correct. For instance the TRAVELING SALESMAN PROBLEM asks: "Given a weighted graph G and an integer k, does G have a tour that visits all the vertices and has total length at most k?" We may not know how to find such a tour quickly, but if someone gave such a tour to us, we could easily check if it satisfied the desired conditions (visited all the vertices and had total length at most k). Similarly, for the 3-COLORING problem: "Given a graph G, can vertices be assigned colors red, blue, and green so that no two neighbors have the same color?" we don't know of any polynomial-time algorithms for solving the problem but we could easily check a proposed solution if someone gave one to us. The class of problems of this type — namely, if the answer is YES, then there exists a polynomial-length proof that can be checked in polynomial time — is called **NP**. (we define the class **NP** formally in Section 19.5).

Let's consider now what would be a problem *so expressive* that if we could solve it, we could solve any problem of this kind. Moreover, let's see if we can define the problem so that it is of this kind as well. Here is a natural candidate:

Definition 19.3 Existence of a verifiable Solution Problem (ESP)*: The input to this problem is in three parts. The first part is a program $V(I, X)$, written in some standard programming language, that has two arguments.[1] The second part is a string I intended as a first argument, and the third part is a bound b written in unary (a string of b 1s). Question: does there exist a string X, $|X| \leq b$, such that $V(I, X)$ halts in at most b steps and outputs YES?*

What we will show is that (a) ESP $\in$ **NP** and (b) for any problem $Q \in$ **NP** we have $Q \leq_p$ ESP. (I.e., if you "had ESP" you could solve any problem in **NP**).[2]

Let's begin with (a): why is ESP $\in$ **NP**? This is the reason for the bound b written in unary. If we didn't have b at all, then (since we can't even in general tell if a program is ever going to halt) the ESP question would not even be computable. However, with the bound b, if the answer is YES, then there is a short proof (namely the string X) that we can check in polynomial time (just run $V(I, X)$ for b steps). The reason we ask for b to be written in unary is precisely so that this check counts as being polynomial time: if b were in binary, then this check could take time exponential

---

[1] We use "V" for the program because we will think of it as a solution-verifier.

[2] Thanks to Manuel Blum for suggesting the acronym.

in the number of bits in the input (much like Ford-Fulkerson is not a polynomial-time algorithm if the capacities are written in binary).

Now, let's go to (b): why is it the case that for any problem $Q \in \mathbf{NP}$ we have $Q \leq_p$ ESP? Consider some $\mathbf{NP}$ problem we might want to solve like 3-COLORING. We don't know any fast ways of solving that problem, but we can easily write a program V that given inputs $I = G$ and $X =$ an assignment of colors to the vertices, verifies whether X indeed satisfies our requirements (uses at most three colors and gives no two adjacent vertices the same color). Furthermore, this solution-verifier is linear time. So, if we had an algorithm to solve the ESP, we could feed in this V, feed in the graph G, feed in a bound b that is linear in the size of G, and solve the 3-COLORING problem. Similarly, we could do this for the TRAVELING SALESMAN PROBLEM: program V, given inputs $I = (G, k)$ and $X =$ a description of a tour through G, just verifies that the tour indeed has length at most k and visits all the vertices. More generally, we can do this for any problem Q in $\mathbf{NP}$. By definition of $\mathbf{NP}$, YES-instances of Q must have short proofs that can be easily checked: i.e., they must have such a solution-verifier V that we can plug into our magic ESP algorithm.

Thus, we have shown that ESP satisfies both conditions (a) and (b) and therefore is **NP-complete**.

19.4 Search versus Decision

Technically, a polynomial-time algorithm for the ESP just tells us if a solution exists, but doesn't actually produce it. How could we use an algorithm that just answers the YES/NO question of ESP to actually find a solution X? If we can do this, then we can use it to actually *find* the coloring or *find* the tour, not just smugly tell us that there is one. The problem of actually finding a solution is often called the *search* version of the problem, as opposed to the *decision* version that just asks whether or not the solution exists. That is, we are asking: can we reduce the search version of the ESP to the decision version?

It turns out that in fact we can, by essentially performing binary search. In particular, once we know that a solution X exists, we want to ask: "how about a solution whose first bit is 0?" If, say, the answer to that is YES, then we will ask: "how about a solution whose first two bits are 00?" If, say, the answer to that is NO (so there must exist a solution whose first two bits are 01) we will then ask: "how about a solution whose first three bits are 010?" And so on. The key point is that we can do this using a black-box algorithm for the decision version of ESP as follows. Given a string of bits S, we define a new program $V_S(I, X) = V(I, X_S)$ where X_S is the string X whose first $|S|$ bits are replaced by S. We then feed our magic ESP algorithm the program V_S instead of V. This way, using at most b calls to the decision algorithm, we can solve the search problem too.

So, if we had a polynomial-time algorithm for the decision version of ESP, we immediately get a polynomial-time algorithm for the search version of ESP. So, we can *find* the tour or coloring or whatever.

The ESP seems pretty stylized. But we can now show that other simpler-looking problems have the property that if you could solve them in polynomial-time, then you could solve the ESP in polynomial time as well, so they too are **NP**-complete. That is, they are so expressive that if we *could* solve them in polynomial-time, then it would mean that for any problem where we could *check* a proposed solution efficiently, we could also *find* such a solution efficiently. Now, onto formal definitions.

19.5   Formal definitions: P, NP, and NP-Completeness

We will formally be considering decision problems: problems whose answer is YES or NO. E.g., "Does the given network have a flow of value at least k?" or "Does the given graph have a 3-coloring?" For such problems, we can split all possible instances into two categories: YES-instances (whose correct answer is YES) and NO-instances (whose correct answer is NO). We can also put any ill-formed instances into the NO category. We now define the complexity classes **P** and **NP**.

Definition 19.4 P *is the set of decision problems solvable in polynomial time.*

E.g., the decision version of the network flow problem: "Given a network G and a flow value k, does there exist a flow $\geq k$?" belongs to **P**.

Definition 19.5 NP *is the set of decision problems that have polynomial-time* verifiers. *Specifically, problem Q is in* **NP** *if there is a polynomial-time algorithm $V(I, X)$ such that:*

- *If I is a YES-instance, then there exists X such that $V(I, X) = $ YES.*

- *If I is a NO-instance, then for all X, $V(I, X) = $ NO.*

Furthermore, X should have length polynomial in size of I (since we are really only giving V time polynomial in the size of the instance, not the combined size of the instance and solution).

The second input X to the verifier V is often called a *witness*. E.g., for 3-coloring, the witness that an answer is YES is the coloring. For factoring, the witness that N has a factor between 2 and k is a factor. For the TRAVELING SALESMAN PROBLEM: "Given a weighted graph G and an integer k, does G have a tour that visits all the vertices and has total length at most k?" the witness is the tour. All these problems belong to **NP**. Of course, any problem in **P** is also in **NP**, since V could just ignore X and directly solve I. So, **P** $\subseteq$ **NP**.

A huge open question in complexity theory is whether **P** $=$ **NP**. It would be quite strange if they were equal since that would mean that any problem for which a solution can be easily *verified* also has the property that a solution can be easily *found*. So most people believe **P** $\neq$ **NP**. But, it's very hard to prove that a fast algorithm for something does *not* exist. So, it's still an open problem.

Definition 19.6 *Problem Q is* **NP***-complete if:*

1. *Q is in* **NP***, and*

2. *For any other problem Q' in* **NP***, $Q' \leq_p Q$.*

So if Q is **NP**-complete and you could solve Q in polynomial time, you could solve *any* problem in **NP** in polynomial time. If Q just satisfies part (2) of the definition, then it's called **NP**-hard.

As we showed above, the ESP is **NP**-complete: it belongs to **NP** (that was the reason for including the bound b in unary, so that running the verifier for b steps counts as being polynomial-time), and we saw that we could use a polynomial-time algorithm for the ESP to solve any other problem in **NP** in polynomial-time.

19.6 Circuit-SAT and 3-SAT

Though the ESP is **NP**-complete, it is a bit unweildy. We will now develop two more natural problems that also are **NP**-complete: CIRCUIT-SAT and 3-SAT. Both of them will be obviously in **NP**. To show they are **NP**-complete, we will show that ESP $\leq_p$ CIRCUIT-SAT, and then that CIRCUIT-SAT $\leq_p$ 3-SAT. Notice that this is enough: it means that if you had a polynomial-time algorithm for 3-SAT then you would also have a polynomial-time algorithm for CIRCUIT-SAT; and if you had a polynomial-time algorithm for CIRCUIT-SAT, then you would also have a polynomial-time algorithm for ESP; and we already know that if you have a polynomial-time algorithm for ESP, you can solve any problem in **NP** in polynomial-time. In other words, to show that a problem Q is **NP**-complete, we just need to show that $Q' \leq_p Q$ for *some* **NP**-complete problem Q' (plus show that $Q \in$ **NP**).

Definition 19.7 CIRCUIT-SAT: *Given a circuit of NAND gates with a single output and no loops (some of the inputs may be hardwired). Question: is there a setting of the inputs that causes the circuit to output 1?*

Theorem 19.1 *CIRCUIT-SAT is **NP**-complete.*

Proof Sketch: First of all, CIRCUIT-SAT is clearly in **NP**, since you can just guess the input and try it. To show it is **NP**-complete, we need to show that if we could solve this, then we could solve the ESP. Say we are given V, I, and b, and want to tell if there exists X such that $V(I, X)$ halts and outputs YES within at most than b steps. Since we only care about running V for b steps we can assume it uses at most b bits of memory, including the space for its arguments. We will now use the fact that one can construct a RAM with b bits of memory (including its stored program) and a standard instruction set using only $O(b \log b)$ NAND gates and a clock. By unrolling this design for b levels, we can remove loops and create a circuit that simulates what V computes within b time steps. We then hardwire the inputs corresponding to I and feed this into our CIRCUIT-SAT solver. ∎

So, we now have one more **NP**-complete problem. Still, CIRCUIT-SAT looks complicated: we weren't expecting to be able to solve it in polynomial-time. However, now we will show that a much simpler-looking problem, 3-SAT has the property that CIRCUIT-SAT $\leq_p$ 3-SAT.

Definition 19.8 3-SAT: *Given: a CNF formula (AND of ORs) over n variables $x_1, \ldots, x_n$, where each clause has at most 3 variables in it. E.g., $(x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_3) \wedge (x_1 \vee x_3) \wedge \ldots$. Goal: find an assignment to the variables that satisfies the formula if one exists.*

We'll save the proof that 3-SAT is **NP**-complete for the next lecture, but before we end, here is formally how we are going to do our reductions. Say we have some problem A that we know is **NP**-complete. We want to show problem B is **NP**-complete too. Well, first we show B is in **NP** but that is usually the easy part. The main thing we need to do is show that $A \leq_p B$; that is, any polynomial-time algorithm for B would give a polynomial-time algorithm for A. We will do this through the following method called a *many-one* or *Karp* reduction:

Many-one (Karp) reduction from problem A to problem B: To reduce problem A to problem B we want a function f that takes arbitrary instances of A to instances of B such that: