

Notes for Lecture 1

This course assumes CS170, or equivalent, as a prerequisite. We will assume that the reader is familiar with the notions of algorithm and running time, as well as with basic notions of discrete math and probability. We will occasionally refer to Turing machines, especially in this lecture.

A main objective of theoretical computer science is to understand the amount of resources (time, memory, communication, randomness, ...) needed to solve computational problems that we care about. While the design and analysis of algorithms puts upper bounds on such amounts, computational complexity theory is mostly concerned with lower bounds; that is we look for *negative results* showing that certain problems require a lot of time, memory, etc., to be solved. In particular, we are interested in *infeasible* problems, that is computational problems that require impossibly large resources to be solved, even on instances of moderate size. It is very hard to show that a particular problem is infeasible, and in fact for a lot of interesting problems the question of their feasibility is still open. Another major line of work in complexity is in understanding the relations between different computational problems and between different “modes” of computation. For example what is the relative power of algorithms using randomness and deterministic algorithms, what is the relation between worst-case and average-case complexity, how easier can we make an optimization problem if we only look for approximate solutions, and so on. It is in this direction that we find the most beautiful, and often surprising, known results in complexity theory.

Before going any further, let us be more precise in saying what a computational problem is, and let us define some important classes of computational problems. Then we will see a particular incarnation of the notion of “reduction,” the main tool in complexity theory, and we will introduce **NP**-completeness, one of the great success stories of complexity theory. We conclude by demonstrating the use of diagonalization to show some separations between complexity classes. It is unlikely that such techniques will help solving the **P** versus **NP** problem.

1 Computational Problems

In a *computational problem*, we are given an *input* that, without loss of generality, we assume to be encoded over the alphabet $\{0, 1\}$, and we want to return as *output* a solution satisfying some property: a computational problem is then described by the property that the output has to satisfy given the input.

In this course we will deal with four types of computational problems: *decision* problems, *search* problems, *optimization* problems, and *counting* problems.¹ For the moment, we will discuss decision and search problem.

¹This distinction is useful and natural, but it is also arbitrary: in fact every problem can be seen as a search problem

In a *decision* problem, given an input $x \in \{0,1\}^*$, we are required to give a YES/NO answer. That is, in a decision problem we are only asked to verify whether the input satisfies a certain property. An example of decision problem is the 3-coloring problem: given an undirected graph, determine whether there is a way to assign a “color” chosen from $\{1,2,3\}$ to each vertex in such a way that no two adjacent vertices have the same color.

A convenient way to *specify* a decision problem is to give the set $L \subseteq \{0,1\}^*$ of inputs for which the answer is YES. A subset of $\{0,1\}^*$ is also called a *language*, so, with the previous convention, every decision problem can be specified using a language (and every language specifies a decision problem). For example, if we call 3COL the subset of $\{0,1\}^*$ containing (descriptions of) 3-colorable graphs, then 3COL is the language that specifies the 3-coloring problem. From now on, we will talk about decision problems and languages interchangeably.

In a *search* problem, given an input $x \in \{0,1\}^*$ we want to compute some answer $y \in \{0,1\}^*$ that is in some relation to x , if such a y exists. Thus, a search problem is specified by a relation $R \subseteq \{0,1\}^* \times \{0,1\}^*$, where $(x,y) \in R$ if and only if y is an admissible answer given x .

Consider for example the search version of the 3-coloring problem: here given an undirected graph $G = (V,E)$ we want to find, if it exists, a coloring $c : V \rightarrow \{1,2,3\}$ of the vertices, such that for every $(u,v) \in E$ we have $c(u) \neq c(v)$. This is different (and more demanding) than the decision version, because beyond being asked to determine whether such a c exists, we are also asked to construct it, if it exists. Formally, the 3-coloring problem is specified by the relation $R_{3\text{COL}}$ that contains all the pairs (G,c) where G is a 3-colorable graph and c is a valid 3-coloring of G .

2 P and NP

In most of this course, we will study the *asymptotic* complexity of problems. Instead of considering, say, the time required to solve 3-coloring on graphs with 10,000 nodes on some particular model of computation, we will ask what is the best asymptotic running time of an algorithm that solves 3-coloring on all instances. In fact, we will be much less ambitious, and we will just ask whether there is a “feasible” asymptotic algorithm for 3-coloring. Here feasible refers more to the rate of growth than to the running time of specific instances of reasonable size.

A standard convention is to call an algorithm “feasible” if it runs in polynomial time, i.e. if there is some polynomial p such that the algorithm runs in time at most $p(n)$ on inputs of length n .

We denote by **P** the class of decision problems that are solvable in polynomial time.

We say that a search problem defined by a relation R is a **NP** search problem if the relation is efficiently computable and such that solutions, if they exist, are short. Formally, R is an **NP** search problem if there is a polynomial time algorithm that, given x and y , decides whether $(x,y) \in R$, and if there is a polynomial p such that if $(x,y) \in R$ then $|y| \leq p(|x|)$.

We say that a decision problem L is an **NP** decision problem if there is some **NP** relation R such that $x \in L$ if and only if there is a y such that $(x,y) \in R$. Equivalently, a decision

problem L is an **NP** decision problem if there is a polynomial time algorithm $V(\cdot, \cdot)$ and a polynomial p such that $x \in L$ if and only if there is a y , $|y| \leq p(|x|)$ such that $V(x, y)$ accepts.

We denote by **NP** the class of **NP** decision problems.

Equivalently, **NP** can be defined as the set of decision problems that are solvable in polynomial time by a non-deterministic Turing machine. Suppose that L is solvable in polynomial time by a non-deterministic Turing machine M : then we can define the relation R such that $(x, t) \in R$ if and only if t is a transcript of an accepting computation of M on input x and it's easy to prove that R is an **NP** relation and that L is in **NP** according to our first definition. Suppose that L is in **NP** according to our first definition and that R is the corresponding **NP** relation. Then, on input x , a non-deterministic Turing machine can guess a string y of length less than $p(|x|)$ and then accept if and only if $(x, y) \in R$. Such a machine can be implemented to run in non-deterministic polynomial time and it decides L .

For a function $t : \mathbb{N} \rightarrow \mathbb{N}$, we define by **DTIME**($t(n)$) the set of decision problems that are solvable by a deterministic Turing machine within time $t(n)$ on inputs of length n , and by **NTIME**($t(n)$) the set of decision problems that are solvable by a non-deterministic Turing machine within time $t(n)$ on inputs of length n . Therefore, $\mathbf{P} = \bigcup_k \mathbf{DTIME}(O(n^k))$ and $\mathbf{NP} = \bigcup_k \mathbf{DTIME}(O(n^k))$.

3 NP-completeness

3.1 Reductions

Let A and B be two decision problems. We say that A reduces to B , denoted $A \leq B$, if there is a polynomial time computable function f such that $x \in A$ if and only if $f(x) \in B$.

Two immediate observations: if $A \leq B$ and B is in \mathbf{P} , then also $A \in \mathbf{P}$ (conversely, if $A \leq B$, and $A \notin \mathbf{P}$ then also $B \notin \mathbf{P}$); if $A \leq B$ and $B \leq C$, then also $A \leq C$.

3.2 NP-completeness

A decision problem A is **NP-hard** if for every problem $L \in \mathbf{NP}$ we have $L \leq A$. A decision problem A is **NP-complete** if it is **NP-hard** and it belongs to **NP**.

It is a simple observation that if A is **NP-complete**, then A is solvable in polynomial time if and only if $\mathbf{P} = \mathbf{NP}$.

3.3 An NP-complete problem

Consider the following decision problem, that we call U : we are given in input (M, x, t, l) where M is a Turing machine, $x \in \{0, 1\}^*$ is a possible input, and t and l are integers encoded in unary², and the problem is to determine whether there is a $y \in \{0, 1\}^*$, $|y| \leq l$, such that $M(x, y)$ accepts in $\leq t$ steps.

It is immediate to see that U is in **NP**. One can define a procedure V_U that on input (M, x, t, l) and y accepts if and only if $|y| \leq l$, and $M(x, y)$ accepts in at most t steps.

²The “unary” encoding of an integer n is a sequence of n ones.

Let L be an **NP** decision problem. Then there are algorithm V_L , and polynomials T_L and p_L , such that $x \in L$ if and only if there is y , $|y| \leq p_L(|x|)$ such that $V_L(x, y)$ accepts; furthermore V_L runs in time at most $T_L(|x| + |y|)$. We give a reduction from L to U . The reduction maps x into the instance $f(x) = (V_L, x, T_L(|x| + p_L(|x|)), p_L(|x|))$. Just by applying the definitions, we can see that $x \in L$ if and only if $f(x) \in U$.

3.4 The Problem SAT

In SAT (that stands for *CNF-satisfiability*) we are given Boolean variables x_1, x_2, \dots, x_n and a Boolean formula ϕ involving such variables; the formula is given in a particular format called *conjunctive normal form*, that we will explain in a moment. The question is whether there is a way to assign Boolean (TRUE / FALSE) values to the variables so that the formula is satisfied.

To complete the description of the problem we need to explain what is a Boolean formula in conjunctive normal form. First of all, Boolean formulas are constructed starting from variables and applying the operators \vee (that stands for OR), \wedge (that stands for AND) and \neg (that stands for NOT).

The operators work in the way that one expects: $\neg x$ is TRUE if and only if x is FALSE; $x \wedge y$ is TRUE if and only if both x and y are TRUE; $x \vee y$ is TRUE if and only at least one of x or y is TRUE.

So, for example, the expression $\neg x \wedge (x \vee y)$ can be satisfied by setting x to FALSE and y to TRUE, while the expression $x \wedge (\neg x \vee y) \wedge \neg y$ is impossible to satisfy.

A *literal* is a variable or the negation of a variable, so for example $\neg x_7$ is a literal and so is x_3 . A *clause* is formed by taking one or more literals and connecting them with a OR, so for example $(x_2 \vee \neg x_4 \vee x_5)$ is a clause, and so is (x_3) . A *formula in conjunctive normal form* is the AND of clauses. For example

$$(x_3 \vee \neg x_4) \wedge (x_1) \wedge (\neg x_3 \vee x_2)$$

is a formula in conjunctive normal form (from now on, we will just say “CNF formula” or “formula”). Note that the above formula is satisfiable, and, for example, it is satisfied by setting all the variables to TRUE (there are also other possible assignments of values to the variables that would satisfy the formula).

On the other hand, the formula

$$x \wedge (\neg x \vee y) \wedge \neg y$$

is not satisfiable, as it has already been observed.

Theorem 1 (Cook) *SAT is NP-complete.*

We will give a proof of Cook’s Theorem later in the course.

4 Diagonalization

Diagonalization is essentially the only way we know of proving separations between complexity classes. The basic principle is the same as in Cantor’s proof that the set of real

numbers is not countable. First note that if the set of real numbers r in the range $[0, 1)$ is countable then the set of infinite binary sequences is countable: we can identify a real number r in $[0, 1)$ with its binary expansion $r = \sum_{j=1}^{\infty} 2^{-j}r[j]$. If we had an enumeration of real numbers, then we would also have an enumeration of infinite binary string. (The only thing to watch for is that some real numbers may have two possible binary representations, like $0.01000\dots$ and $0.001111\dots$.)

So, suppose towards a contradiction that the set of infinite binary sequences were countable, and let $B_i[j]$ be the j -th bit of the i -th infinite binary sequence. Then define the sequence B whose j -th bits is $1 - B_j[j]$. This is a well-defined sequence but there can be no i such that $B = B_i$, because B differs from B_i in the i -th bit.

Similarly, we can prove that the Halting problem is undecidable by considering the following decision problem D : on input $\langle M \rangle$, the description of a Turing machine, answer NO if $M(\langle M \rangle)$ halts and accepts and YES otherwise. The above problem is decidable if the Halting problem is decidable. However, suppose D were decidable and let T be a Turing machine that solves D , then $T(\langle T \rangle)$ halts and accepts if and only if $T(\langle T \rangle)$ does not halt and accept, which is a contradiction.

It is easy to do something similar with time-bounded computations. Instead of the general halting problem, we can define a “bounded” halting problem, for example we can define the problem BH where given a Turing machine $\langle M \rangle$ and a string x we answer YES if M accepts x within $|x|^3$ steps and NO otherwise. Then we can define the problem D where on input $\langle M \rangle$ and x we answer YES if $M(\langle M \rangle, x)$ rejects within n^3 steps, where n is the length of $(\langle M \rangle, x)$, and NO otherwise. Clearly, there cannot be a machine that solves D in time less than n^3 on inputs of length n , and, similarly, we can deduce that BH cannot be solved in time $o(n^3)$, because an algorithm that solves BH in time $t(n)$ can be easily modified to solve D in time $O(t(2n))$. Finally, we note that BH can be solved in time polynomial in n , which shows $\mathbf{DTIME}(n^3) \neq \mathbf{P}$. Just by replacing the bound n^3 with other bounds, one can show that, for every k , $\mathbf{DTIME}(O(n^k)) \neq \mathbf{P}$, or that $\mathbf{DTIME}(2^n) \neq \mathbf{DTIME}(2^{2n})$ and so on.

If we want to show tighter separations, however, such as $\mathbf{DTIME}(n^3) \neq \mathbf{DTIME}(n^{3.1})$, we need to be more careful in the definition of our “diagonal” problem D . If D is defined as the problem of recognizing pairs $(\langle M \rangle, x)$ such that $M(\langle M \rangle, x)$ rejects within n^3 steps, where n is the length of $(\langle M \rangle, x)$, then it is possible then, in an input of length n for D , $n/2$ bits, say, are devoted to the description of $\langle M \rangle$. To simulate one step of the computation of M , then, a universal Turing machine must scan the entire description of M , and so each step of the simulation will take at least $\Omega(n)$ time, and the total running time for deciding D will be $\Omega(n^4)$. We can overcome this problem if we define D not in terms of the running time of M , but rather in terms of the running time of a fixed universal Turing machine that simulates M . We first state a result about efficient universal Turing machines.

Lemma 2 (Efficient Universal Turing Machine) *There is a Turing machine U that, on input the description $\langle M \rangle$ of a Turing machine M and a string x , behaves like M on input x , that is, if $M(x)$ accepts then $U(\langle M \rangle, x)$ accepts, if $M(x)$ rejects then $U(\langle M \rangle, x)$ rejects, and if $M(x)$ does not halt then $U(\langle M \rangle, x)$ does not halt. Furthermore, if $M(x)$ halts within t steps then $U(\langle M \rangle, x)$ halts within $O(|\langle M \rangle|^{O(1)} \cdot t)$ steps.*

We can now make an argument for a very tight “hierarchy” theorem.

Theorem 3 $\text{DTIME}(o(n^3)) \not\subseteq \text{DTIME}(O(n^3 \log n))$.

PROOF: Consider the following decision problem D : on input $(\langle M \rangle, x)$ answer YES if U rejects (M, x) within $|x|^3$ steps, and NO otherwise.

The problem can be solved in $O(n^3(\log n))$ time. We just need to modify U so that it initializes a counter at $|x|^3$, decreases its value after each step, and, if it hasn't halted already, it halts when the counter reaches zero. With some care, maintaining and updating the counter can be done in time $O(\log n)$ per step even on a one-tape Turing machine.

Suppose by contradiction that D is solvable in time $t(n) = o(n^3)$ on inputs of length n by a machine T . Then, for every x of length n , $U(\langle T \rangle, x)$ also halts in $o(n^3)$ time. Let x be sufficiently long so that $U(\langle T \rangle, x)$ halts in less than $|x|^3$ time. Then, if x is a YES instance of D , it means that U rejects $(\langle T \rangle, x)$ within $|x|^3$ time, which means that T rejects x , which means that T is incorrect on input x . Similarly, if x is a NO instance of D , then U does not reject $(\langle T \rangle, x)$ within $|x|^3$ time, but U halts within $|x|^3$ time, and so it follows that U accepts $(\langle T \rangle, x)$, and so T accepts x , incorrectly. \square

See the homeworks for generalizations of the above proof.

We would like to do the same for non-deterministic time, but we run into the problem that we cannot ask a non-deterministic machine to reject if and only if a non-deterministic machine of comparable running time accepts. If we could do so, then we would be able to prove $\text{NP} = \text{coNP}$. A considerably subtler argument must be used instead, which uses the following simple fact.

Theorem 4 *On input the description $\langle M \rangle$ of a non-deterministic Turing machine M , a string x and an integer $t > n$, the problem of deciding whether M accepts x within t steps is solvable in deterministic $|\langle M \rangle|^{O(1)} 2^{O(t)}$ time.*

We will also need a theorem about efficient universal non-deterministic Turing machines.

Lemma 5 *There is a non-deterministic Turing machine NU that, on input the description $\langle M \rangle$ of a non-deterministic Turing machine M and a string x :*

- *If $M(x)$ has an accepting computation of length t then $NU(\langle M \rangle, x)$ has an accepting computation of length $O(|\langle M \rangle|^{O(1)} \cdot t)$*
- *If no computational paths of $M(x)$ accepts, then no computational paths of $NU(\langle M \rangle, x)$ accepts.*

Finally, we can state and prove a special case of the *non-deterministic hierarchy theorem*.

Theorem 6 $\text{NTIME}(o(n^3)) \not\subseteq \text{NTIME}(O(n^3 \log n))$.

PROOF: Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be defined inductively so that $f(1) = 2$ and $f(k+1) = 2^{k \cdot (f(k))^3}$. Consider the following decision problem D : on input $x = (\langle M \rangle, 1^t)$, where M is a non-deterministic Turing machine,

1. if $t = f(k)$ for some k , then the answer is YES if and only if the simulation of $M(\langle M \rangle, 1^{1+f(k-1)})$ as in Theorem 4 returns NO within t steps,

2. otherwise answer YES if and only if $NU(\langle M \rangle, (\langle M \rangle, 1^{t+1}))$ accepts within t^3 steps.

We first observe that D is solvable by a non-deterministic Turing machine running in $O(n^3 \log n)$ time, where n is the input length.

Suppose that D were decided by a non-deterministic Turing machine T running in time $o(n^3)$, and consider inputs of the form $\langle T \rangle, 1^t$ (which are solved by T in time $o(t^3)$). Pick a sufficiently large k , and consider the behaviour of T on inputs $(\langle T \rangle, 1^t)$ for $f(k-1) < t < f(k)$; since all such inputs fall in case (2), we have that $T(\langle T \rangle, 1^t) = T(\langle T \rangle, 1^{t+1})$ for all such t and, in particular,

$$T(\langle T \rangle, 1^{1+f(k-1)}) = T(\langle T \rangle, 1^{f(k)}) \tag{1}$$

On the other hand, the input $(\langle T \rangle, 1^{f(k)})$ falls in case (2), and since $T(\langle T \rangle, 1^{1+f(k-1)})$ can be simulated deterministically in time $2^{(k-1) \cdot (f(k-1))^3}$, if k is large enough, and so the correct answer on input $(\langle T \rangle, 1^{f(k)})$ is NO if and only if $T(\langle T \rangle, 1^{1+f(k-1)})$ accepts, which is in contradiction to Equation 1. \square

5 References

The time hierarchy theorem is proved in [HS65], which is also the paper that introduced the term “Computational Complexity.” The non-deterministic hierarchy theorem is due to Cook [Coo73]. The notion of NP-completeness is due to Cook [Coo71] and Levin [Lev73], and the recognition of its generality is due to Karp [Kar72].

References

- [Coo71] Stephen A Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing*, pages 151–158, 1971. 7
- [Coo73] Stephen A Cook. A hierarchy for nondeterministic time complexity. *Journal of Computer and System Sciences*, 7(4):343–353, 1973. 7
- [HS65] Juris Hartmanis and Richard E Stearns. On the computational complexity of algorithms. *Transactions of the AMS*, 117:285–306, 1965. 7
- [Kar72] Richard M Karp. Reducibility among combinatorial problems. In R.E. Miller and J.W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972. 7
- [Lev73] Leonid A. Levin. Universal search problems. *Problemi Peredachi Informatsii*, 9:265–266, 1973. 7

Exercises

1. Show that if $\mathbf{P} = \mathbf{NP}$ for decision problems, then every \mathbf{NP} search problem can be solved in polynomial time.
2. Generalize Theorem 3. Say that a monotone non-decreasing function $t : \mathbb{N} \rightarrow \mathbb{N}$ is time-constructible if, given n , we can compute $t(n)$ in $O(t(n))$ time. Show that if $t(n)$ and $t'(n)$ are two time-constructible functions such that $t'(n) > t(n) > n^3$ and $\lim_{n \rightarrow \infty} \frac{t(n) \log t(n)}{t'(n)} = 0$ then $\mathbf{DTIME}(t'(n)) \not\subseteq \mathbf{DTIME}(t(n))$.