

Lecture 5: NP-complete problems and Space complexity

Anup Rao

April 17, 2020

IN THE LAST CLASS, we introduced the concept of NP and NP-completeness. In this lecture, we begin by showing that many problems are NP-complete.

3SAT

A *boolean formula* is an expression of the form

$$(x_1 \wedge \neg x_2) \vee (x_7 \wedge \neg(x_6 \vee \neg x_2)).$$

Formally: it is a circuit where the only allowed gates are \vee, \wedge, \neg , and every gate has fan-out at most 1. Input gates are allowed to repeat. As usual, size of the gates is number of gates, and the fan-in is allowed to be at most 2. The formula is said to be in conjunctive normal form (CNF) if it is an AND of OR's. Similarly, it is said to be in disjunctive normal form (DNF) if it is an OR of ANDS. For example

$$(x_1 \vee \neg x_2) \wedge (\neg x_7 \vee x_9 \vee \neg x_1)$$

is a CNF.

We have the following lemma:

Lemma 1. *Every function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ can be computed by a CNF (resp. DNF) of size $\ell 2^\ell$.*

Proof For each input z such that $f(z) = 0$, we add the literal x_i to the clause if $z_i = 0$ and $\neg z_i$ otherwise. So for example, if $f(0, 1, 0) = 0$, we add the clause $(x_1 \vee \neg x_2 \vee x_3)$. Then note that each clause is 0 on exactly one input, and all inputs x for which $f(x) = 0$ make some clause 0. Every other input evaluates to 1. So, the CNF computes f . The resulting formula is of size $\ell 2^\ell$. The case of DNF's is symmetric. ■

We define $\text{SAT} : \{0, 1\}^* \rightarrow \{0, 1\}$ to be the function that takes as input a boolean formula F , and outputs 1 if and only if there is an x such that $F(x) = 1$. A 3-CNF formula is a CNF where every clause has at most 3 variables. For example:

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

3SAT : $\{0,1\}^* \rightarrow \{0,1\}$ is the function that takes as input 3-CNF and outputs 1 if and only if the formula is satisfiable. Next we show that even this function is NP-complete

Theorem 2. 3SAT is NP-complete.

Proof 3SAT \in NP is easy enough to check. The witness is a satisfying assignment to the formula. The verifier simply evaluates the formula on the given witness, and outputs the results of the evaluation.

Since we have already shown that CKT – SAT is NP-hard, it will be enough to show that CKT – SAT \leq_p SAT.

Given a circuit, we shall output a CNF formula that is satisfiable if and only if the circuit accepts some input. Introduce a new variable y_g for each internal gate g of the circuit. If the internal gate g has inputs h, q , let F_g be the CNF formula on variables y_g, y_h, y_q that is 1 if and only if $y_g = g(y_q, y_h)$. By Lemma 1, this formula is a 3-CNF of constant size. If the output gate is v , the final formula is

$$y_v \wedge \bigwedge_g F_g,$$

which is satisfied if and only if the circuit has a satisfying assignment.

Every clause of this formula has at most 3 variables. To make sure it has *exactly* 3 variables, we replace each clause with less than 3 variables with a 3-CNF that by adding dummy variables. For example, we can replace y_v by a 3-CNF on the variables y_v, z_1, z_2 that computes the same function as y_v :

$$(y_v \vee z_1 \vee z_2) \wedge (y_v \vee \neg z_1 \vee z_2) \wedge (y_v \vee \neg z_1 \vee \neg z_2) \wedge (y_v \vee z_1 \vee \neg z_2).$$

■

We now consider several interesting graph problems and show that they are NP-complete.

3 Coloring

We say that an undirected graph is 3 colorable if one can color every vertex with one of 3 colors so that every edge gets two colors.

$$3\text{COL}(G) = \begin{cases} 1 & \text{if } G \text{ is 3 colorable} \\ 0 & \text{otherwise.} \end{cases}$$

Is the same true for 2SAT? We do not know. There are polynomial time algorithms for 2SAT, so if you found a reduction to 2SAT, you would prove $\mathbf{P} = \mathbf{NP}$. The algorithm works by viewing every clause $(x \vee y)$ as an implication $\neg x \Rightarrow y$ as well as the implication $\neg y \Rightarrow x$. This defines a directed graph where all the vertices correspond to variables and their negations, and the edges correspond to implications. You can show that the formula is satisfiable if and only if there is no path that leads from a variable to its negation.

Although there is an easy polynomial time algorithm for 2-coloring a graph (greedily color the first vertex blue, all its neighbors red, all their neighbors blue and so on), we know of no such algorithm for 3-coloring a graph.

Theorem 3. 3COL is NP-complete.

Sketch of Proof The coloring serves as a witness that can be verified in polynomial time, so $3\text{COL} \in \text{NP}$. Next we show how to reduce 3SAT to 3COL in polynomial time.




Figure 1: Ensuring that a coloring corresponds to a truth assignment

We would like to construct the graph in a way that allows every coloring to be decoded to an assignment to the variables. To this end, we shall have three vertices named T, F, B and $2n$ vertices named $x_1, \bar{x}_1, x_2, \bar{x}_2, \dots, x_n, \bar{x}_n$ that correspond to the variables and their negations. We shall connect every pair of T, F, B so that these three must be given a distinct color. We also connect each x_i and \bar{x}_i to B , so x_i and \bar{x}_i must be given the same as color as T or F . In addition, connect each x_i and \bar{x}_i to ensure that they are assigned the same color. (See Figure 1). Thus any coloring corresponds to an assignment of truth values to the variables.




Figure 2: Ensuring that the assignment satisfies each clause

Next we need to encode each clause of the formula. The idea here is generate a part of the graph that can be colored if and only if the clause is satisfied by the assignment to the corresponding variables. This is shown in Figure 2. We connect the gadget shown there to the

My initial drawing in class had an error!

variables that correspond to the clause we are interested in. If any one of the variables is set to T, then one can color the corresponding vertex in the top row F. This allows us to color the bottom row.

On the other hand, if all variables in the clause are false, then the top row must be colored B, and the bottom row cannot be colored correctly.

■

Independent Set

Given an undirected graph G , an *independent set* in the graph is a set of vertices such that no edge is contained in the set. The goal is find an independent set of maximum size in the graph. We can encode this problem using the following boolean function:

$$\text{ISET}(G, k) = \begin{cases} 1 & \text{if } G \text{ has an independent set of size } k, \\ 0 & \text{otherwise.} \end{cases}$$

If you can compute ISET in polynomial time, then you can find the largest independent set in polynomial time (how?). If on the other hand you can find the largest independent set, then you can also compute ISET. Here we prove:

Theorem 4. ISET is NP-complete.

Proof ISET is in NP, since the independent set of largest size is itself a witness which can be verified in polynomial time. Thus it only remains to show that ISET is NP-hard. To do this, we show how to reduce 3SAT to ISET in polynomial time.

Given a 3SAT instance with m clauses and n variables, we construct a graph with $3m$ variables. Each clause C_i corresponds to 3 vertices, which are all connected to each other. Thus the graph contains m disjoint triangles. In each triangle, we label each of the three vertices with the three literals that occur in the clause. Thus the clause $(a \vee \neg b \vee c)$ leads to the three vertices being labeled $a, \neg b, c$. Finally, for every variable a , we connect every vertex labeled a to every vertex labeled $\neg a$ using an edge.

We claim that the above graph has an independent set of size m if and only if the given 3 CNF is satisfiable. Indeed, suppose the 3 CNF is satisfiable using the assignment to the variables x . Then x must satisfy every clause, so in each clause, some literal must be true. Pick a single vertex from each of the triangles in such a way that we always pick a true vertex. By the construction, every edge must connect a true vertex to a false vertex, so the resulting set is




Figure 3: An example of the input to ISET produced when the input formula is $(a \vee b \vee \neg c) \wedge (\neg d \vee c \vee e) \wedge (\neg f \vee \neg a \vee \neg b) \wedge (b \vee d \vee \neg e) \wedge (e \vee a \vee f)$.

independent. There cannot be a larger independent set in the graph, since every triangle can contain only one vertex.

Conversely, if the graph has an independent set of size m , then there must be exactly one vertex in every triangle of the construction, or else one of the triangle edges would be included in the set. Now pick the assignment to the variables in such a way that all the vertices of the independent set are labeled with true. There is always a way to do this, since by construction every time we try to set a variable in this process, it has not already been set to a different value by the construction of the graph and the property that the set is independent.

Thus the reduction is to read the input formula and construct the above graph in polynomial time. ■

Hamiltonian Path

Given a directed graph G , a Hamiltonian path is a path that visits every vertex of the graph exactly once. We define the function

$$\text{HPATH}(G) = \begin{cases} 1 & \text{if } G \text{ has a Hamiltonian path} \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 5. HPATH is NP-complete.

Proof Given a path in the graph, one can check in polynomial time whether or not it is a Hamiltonian path. Thus $\text{HPATH} \in \mathbf{NP}$ using the path as a witness. Next we show that you can reduce 3SAT to HPATH, proving that HPATH is NP-hard.

Suppose the formula has n variables and m clauses. We shall construct a graph on $(2m + 2)n + 2$ vertices that encodes assignments to the formulas as follows. We start by constructing a graph that will contain $(2m + 2)n$ vertices named $v_{i,j}$, where $i \in \{1, 2, \dots, n\}$ and $j \in \{1, 2, \dots, m\}$. For every i and $1 \leq j < j + 1 \leq m$, we have the edges $(v_{i,j}, v_{i,j+1})$ and $(v_{i,j+1}, v_{i,j})$. Thus these vertices can be thought of as arranged




Figure 4: An example showing how to generate a directed graph for the Hamiltonian path problem using a single clause from the formula.

in n rows, where in each row the path can go left or right. For every $1 \leq i < i+1 \leq n$, we add the edges

$$(v_{i,1}, v_{i+1,1}), (v_{i,1}, v_{i+1,n}), (v_{i,n}, v_{i+1,1}), (v_{i,n}, v_{i+1,n}).$$

Finally we add two special vertices s, t , with edges

$$(s, v_{1,1}), (s, v_{1,n}), (v_{n,1}, t), (v_{n,n}, t).$$

By construction, every Hamiltonian path in the graph must start at s and end at t , and must traverse each row in order. Each row can be traversed in either left to right or right to left fashion. We shall imagine that traversing the row left to right corresponds to assigning the i 'th variable the value 0, and traversing it the other way corresponds to assigning the value 1.

Next we add some vertices to encode the constraints given by the clauses. Without loss of generality we assume that each clause contains a variable at most once (since we can always reduce the formula to this case). For the j 'th clause C_j , we add the vertex c_j . For every

variable x_i that the clause contains unnegated, we add the edges $(v_{i,2j}, c_j), (c_j, v_{i,2j-1})$. For every variable x_j that is contained in the clause as $\neg x_j$, we add the edges $(v_{i,2j-1}, c_j), (c_j, v_{i,2j})$. By construction, any Hamiltonian path that takes the edge $(v_{i,2j}, c_j)$, must take $(c_j, v_{i,2j-1})$ next, or $v_{i,2j-1}$ will never be visited. Similarly, any Hamiltonian path that takes the edge $(v_{i,2j-1}, c_j)$ must take $(c_j, v_{i,2j})$ next. We claim that the graph has a Hamiltonian path if and only if the formula is satisfiable.

Indeed, if the formula is satisfiable, then traverse each row in the direction corresponding to the satisfying assignment. Since each clause is satisfied by some variable, we can visit the vertex for the clause when we traverse the first variable that satisfies it. Conversely, if there is a Hamiltonian path, then the construction ensures that this path corresponds to an assignment to the variables, and this path must visit every clause vertex, which guarantees that each clause vertex is satisfied by some variable. ■

Subset Sum

In the subset sum problem, the input is a collection of numbers a_1, \dots, a_k , as well as a target number t . The goal is compute whether or not some subset of the numbers a_1, \dots, a_k sums to t .

$$\text{SubSum}(a_1, \dots, a_k, t) = \begin{cases} 1 & \text{if there is a subset } S \subseteq \{1, 2, \dots, n\} \text{ such that } \sum_{i \in S} a_i = t, \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 6. *SubSum is NP-complete.*

We sketch the proof. SubSum is in NP, since there is an obvious polynomial time computable verifier for the problem. The witness is a subset S , and the verifier simply checks that $\sum_{i \in S} a_i = t$, which can be done in polynomial time.

To show that SubSum is NP-hard, we shall show that

$$3\text{SAT} \leq_P \text{SubSum}.$$

We describe the polynomial time reduction next. Given a 3-sat formula ϕ , our algorithm needs to output numbers a_1, \dots, a_k and t such that $\text{SubSum}(a_1, \dots, a_k, t) = 1$ if and only if ϕ is satisfiable.

Suppose ϕ has n variables and m clauses. Then, we will have $k = 2n + 2m$, and all of the numbers a_1, \dots, a_k and t will be $n + m$ digit numbers, written in base 10. Moreover, all the digits of a_1, \dots, a_k will be either 0 or 1, and the numbers will be chosen in such a way that adding any subset of a_1, \dots, a_k will never produce a carry.

Example: suppose we are given the formula $(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_2 \vee x_3 \vee x_4) \wedge (\neg x_1, \vee \neg x_3 \vee \neg x_4) \vee (\neg x_2, \vee \neg x_3 \vee x_4)$. There are 4 variables and 4 clauses, so the polynomial time reduction will generate 16 numbers, each with 8-digits, and a target number with 8-digits:

$$\begin{aligned} t_1 &= 10001000 \\ f_1 &= 10000010 \\ t_2 &= 01000000 \\ f_2 &= 01001101 \\ t_3 &= 00101100 \\ f_3 &= 00100011 \\ t_4 &= 00010101 \\ f_4 &= 00010010 \\ b_1 &= 00001000 \\ c_1 &= 00001000 \\ b_2 &= 00000100 \\ c_2 &= 00000100 \\ b_3 &= 00000010 \\ c_3 &= 00000010 \\ b_4 &= 00000001 \\ c_4 &= 00000001 \end{aligned}$$

The target number will be:

$$t = 11113333.$$

For each variable x_i of the formula ϕ , we shall have two numbers: t_i and f_i . The i 'th digit of t_i and f_i will be set to 1 and all of the remaining $n - 1$ digits in the first n digits will be set to 0. Meanwhile, in the target number t , all of the first n digits will be set to 1. This choice ensures that choosing any subset of $t_1, f_1, \dots, t_n, f_n$ that sums to t corresponds to choosing either t_i or f_i to be included in the set, for each i . In other words, a subset of these numbers that sums to t corresponds to a truth assignment to the variables x_1, \dots, x_n .

Next, we need to add more digits to ensure that this truth assignment satisfies all the clauses. For every i, j , if x_i occurs in the j 'th clause, we make the $n + j$ 'th digit of t_i 1. If $\neg x_i$ occurs in the j 'th clause, we make the $n + j$ 'th digit of f_i 1. All other digits (upto the $n + m$ 'th digit) of t_i, f_i are set to 0. This choice ensures that if the subset chosen satisfies the j 'th clause, then the j 'th digit of the sum will be either 1, 2 or 3. Finally, we add two numbers b_j, c_j , which are 0 in all digits, except for the j 'th digit. The j 'th digit of both numbers is 1. This ensures that if the j 'th clause is satisfied by the assignment, then one can pick 0, 1 or 2 elements of $\{b_j, c_j\}$ to add to the subset, so that the sum of the j 'th digits is 3.

Space

NEXT, WE TURN OUR ATTENTION TO SPACE. Recall, that the space complexity of an execution of a Turing machine is defined to be the maximum value attained by the pointer to the work tape during the execution. So, it is just a count of the number of cells used on the work tape during the execution of the algorithm.

The smallest space class that makes sense is $L = DSPACE(\log n)$. This is because even maintaining a pointer to the input takes $\log n$ work space. While we do not necessarily need to maintain such pointers in the work tape, if we want to be able to design algorithms that have the same complexity regardless of the specific choices made for the Turing machine, then we need to maintain such pointers in order to simulate one Turing machine by another.

As usual the non-deterministic version of this class is when the machine can make non-deterministic choices, and is called $NL = NSPACE(\log n)$. There is a subtle issue about the definition of NL : if we allow the machine to remember the non-deterministic choices that it made for free (for example by giving it access to a guess tape that it can read from), then the power of the class changes significantly. Another interesting class is $PSPACE = \bigcup_c DSPACE(n^c)$. The corresponding non-deterministic class is actually equal to $PSPACE$, as we shall prove below.

So far, we have only discussed time complexity.

For example, if we are designing an algorithm to add two n -bit integers a, b , then if a, b are written on two different tapes (or interleaved on one tape), the computation can be carried out with $O(1)$ space. If, on the other hand, the inputs are written on one tape a, b , then we need space $O(\log n)$ in order to correctly maintain counters to allow us to scan between the corresponding bits a_i and b_i .

A very useful fact when composing space bounded computations is the following:

Claim 7. *If it takes space $s_1(n) \geq \log n$ to compute f and space $s_2(n) \geq \log n$ to compute g , then one can compute the composition $f(g(x))$ in space $O(s_1(n) + s_2(n))$.*

The idea is that in the computation of f , every time we need to lookup an output symbol of $g(x)$, we can recompute it. Thus, as long as $s_1(n), s_2(n)$ are enough to store pointers into the output locations, we actually only need to sum the spaces to compute the composition.

Savitch's Algorithm

One of the most interesting small space algorithms is Savitch's graph search algorithm.

Theorem 8 (Savitch). *Given a directed graph G with two special vertices s, t , there is an algorithm that can compute whether or not there is a path from s to t in the graph, using space $O(\log^2 n)$.*

Proof We shall give a recursive algorithm that can compute the values $A(u, v, i)$ as defined below:

$$A(u, v, i) = \begin{cases} 1 & \text{if there is a path from } u \text{ to } v \text{ of length } 2^i, \\ 0 & \text{else.} \end{cases}$$

Note that $A(u, v, i) = 1$ if and only if $\exists z$ such that $A(u, z, i-1) = 1$ and $A(z, v, i-1) = 1$. Thus, to compute $A(u, v, i)$, do

1. For all z , recursively compute $A(u, z, i-1)$ and $A(z, v, i-1)$, and output 1 if both computations result in 1.
2. Otherwise output 0.

If the size of the graph is 2^s , there are $s+1$ recursive calls, where $A(u, v, 0)$ can be computed trivially by looking up the corresponding bit in the input. In each recursive call, the algorithm needs to store only the vertices u, v, z , which takes $O(\log n)$ space. Thus the total space used is $O(\log^2 n)$. ■

One reason Savitch's algorithm is so important is because, in some sense, graph search is a complete problem for small space computation. Let us discuss this point next.




Figure 5: An example of a configuration graph.

Configuration Graphs

Given an input x to a (possibly non-deterministic) Turing machine M , the configuration graph $G_{M,x}$ is the directed graph where there is a distinct vertex for every possible value of the pointers to the input and work tapes, the value of the string written in the work tape and the current line-number of the line of code that is about to be executed in the machine. There is an edge from u to v if and only if the configuration u could possibly become the configuration v after one step of the program is executed.

Lemma 9. *If the machine uses space $s(n) \geq \Omega(\log n)$, then the number of vertices in the configuration graph is at most $2^{O(s(n))}$.*

Proof The number of options for locations of the pointers is at most $n \cdot s(n)$. The number of options for the contents of the work tape is at most $2^{O(s(n))}$. The number of options for the lines of code is $O(1)$. Thus, the number of different vertices in the graph is at most the product of these numbers, which is at most $2^{O(s(n))}$. ■

The number of options for the pointer that points to the input tape is at most n . This is because we do not allow the pointer on the input tape to move past the actual input. As we discussed in class, even if we did not place this restriction, we can prove that if the Turing machine moves the input pointer more than $2^{O(s(n))}$ steps beyond the input, then the machine does not halt. So, even without this restriction, the number of possible values for the input pointer is at most $2^{O(s(n))}$.