NP-Completeness

~~To be~~ ~~or not~~ to ~~be~~ ....

Carlos Moreno
`cmoreno@uwaterloo.ca`
EIT-4103

`https://ece.uwaterloo.ca/~cmoreno/ece250`

These slides, the course material, and course web site are based on work by Douglas W. Harder

NP-Completeness

Standard reminder to set phones to silent/vibrate mode, please!

NP-Completeness

- During today's lesson:
  - Look at some important categorizations of algorithms based on their run times.
    - Including one important distinction: polynomial time vs. exponential time.
  - Look into the notion of decision problems, and their relationship to an associated computation or optimization problem.
  - Introduce the sets P and NP.
  - Introduce the notions of NP-hard and NP-complete problems.

NP-Completeness

- For the purpose of this discussion, we will be mostly interested in the distinction of polynomial time algorithms vs. non-polynomial time algorithms (exponential and above):

NP-Completeness

- For the purpose of this discussion, we will be mostly interested in the distinction of polynomial time algorithms vs. non-polynomial time algorithms (exponential and above):
  - Polynomial time: Algorithms with run time $\mathrm{O}\left(n^{\alpha}\right)$ for some $\alpha > 0$.

NP-Completeness

- For the purpose of this discussion, we will be mostly interested in the distinction of polynomial time algorithms vs. non-polynomial time algorithms (exponential and above):
  - Polynomial time: Algorithms with run time $\mathrm{O}\left(n^{\alpha}\right)$ for some $\alpha > 0$.
  - Non-polynomial time: Algorithms with run time $\Omega\left(a^{n}\right)$ for some $a > 1$.

NP-Completeness

- Examples of polynomial time:
  - Constant time
  - Logarithmic
  - Linear time
  - $n \log n$
  - Things like n^2, n^3, $n^3 \log n$
  - Etc.

NP-Completeness

- Examples of non-polynomial time:
  - Exponentials (of any base): 2^n, 10^n, etc.
  - $n\,2^n$
  - $n^2\,2^n$
  - $n!$
  - n^n
  - Etc.

NP-Completeness

- Tractable vs. intractable problems:
  - A problem for which a polynomial time algorithm exists that solves the problem is called a *tractable* problem.
  - Otherwise, we refer to them as *intractable* problems.

NP-Completeness

- Tractable vs. intractable problems:
  - A problem for which a polynomial time algorithm exists that solves the problem is called a *tractable* problem.
  - Otherwise, we refer to them as *intractable* problems.
  - Notice the subtlety — a polynomial time algorithm exists vs. a polynomial time algorithm is known.

NP-Completeness

- Tractable vs. intractable problems:
  - A problem for which a polynomial time algorithm exists that solves the problem is called a *tractable* problem.
  - Otherwise, we refer to them as *intractable* problems.
  - Notice the subtlety — a polynomial time algorithm exists vs. a polynomial time algorithm is known.
    - We often see statements such as "this problem is believed to be intractable" or "is considered intractable", meaning that no polynomial time algorithm is known, and it is not known for sure that none exist, but it is *believed* that none exist.

NP-Completeness

- An important notion related to these is that of *polynomial time reduction*.

NP-Completeness

- An important notion related to these is that of *polynomial time reduction*.
  - We talked about reductions — solving problem A using an algorithm that solves problem B.
    - In this case, we say that A reduces to B.

NP-Completeness

- An important notion related to these is that of *polynomial time reduction*.
  - We talked about reductions — solving problem A using an algorithm that solves problem B.
    - In this case, we say that A reduces to B.
    - Notice that we're applying the notion to the problems, and not necessarily to the algorithms — solving problem A reduces to being able to solve problem B.

NP-Completeness

- The reduction takes some time, since we need to transform an instance of a problem into another problem
  - This typically refers to transforming the input of algorithm A into a valid input for algorithm B, then invoking algorithm B, capture its output and derive (transform) the output for algorithm A.

NP-Completeness

- Relating to today's topic, since we're interested in the distinction between polynomial time vs. non-polynomial time algorithms, then the relevant reductions are those with polynomial run times:
  - They will allow us to prove statements about an algorithm having polynomial time or not.

NP-Completeness

- For example, given:

$$T_A(n) = T_R(n) + T_B(n)$$

- If we know that $T_A(n)$ is non-polynomial and we're trying to prove that $T_B(n)$ is also non-polynomial, then we need $T_R(n)$ to be polynomial — otherwise, the statement (the equation) would say absolutely nothing about $T_B(n)$ being or not polynomial.

NP-Completeness

- For example, given:

$$T_A(n) = T_R(n) + T_B(n)$$

- On the other hand, if $T_B(n)$ is polynomial and we're trying to show that we can find an algorithm A with polynomial time $T_A(n)$ by reducing A to B, then in this case we also need $T_R(n)$ to be polynomial — otherwise, the above equation says that $T_A(n)$ is non-polynomial, so we've failed to find such polynomial time algorithm A.

NP-Completeness

- Some times we refer to polynomial time reduction as an efficient reduction.
  - In this context, polynomial time is efficient, non-polynomial time is inefficient.

NP-Completeness

- Notation:

NP-Completeness

- Notation:
  - If problem A reduces to problem B, that means that solving A can not be harder than solving B.

NP-Completeness

- Notation:
  - If problem A reduces to problem B, that means that solving A can not be harder than solving B.
  - This is denoted by: $A \leqslant_P B$

NP-Completeness

- Notation:
  - If problem A reduces to problem B, that means that solving A can not be harder than solving B.
  - This is denoted by: $A \leqslant_P B$
  - We read it as: A polynomial-time-reduces-to B (or A reduces to B in polynomial time)

NP-Completeness

- Notice that if $A \leqslant_P B$:
  - We can not know whether A is easier (more efficient) than B or equally as hard — the only thing that we know is that A is, *at most*, as hard as B, since we can always use B to solve A.

NP-Completeness

- Notice that if $A \leqslant_P B$:
  - We can not know whether A is easier (more efficient) than B or equally as hard — the only thing that we know is that A is, *at most*, as hard as B, since we can always use B to solve A.
  - But there could always be some other way of solving A which is more efficient — the fact that we find a reduction from A to B says nothing about that.

NP-Completeness

- A fascinating class of questions on algorithmic theory arises when we consider *decision problems* and their complexity.

NP-Completeness

- A fascinating class of questions on algorithmic theory arises when we consider *decision problems* and their complexity.
- A *decision problem* is simply a problem with a yes/no answer
  - An algorithm that solves a decision problem outputs a boolean value.

NP-Completeness

- A fascinating class of questions on algorithmic theory arises when we consider *decision problems* and their complexity.
- A *decision problem* is simply a problem with a yes/no answer
  - An algorithm that solves a decision problem outputs a boolean value.
  - Notice the interesting detail: when reducing a decision problem to another decision problem, we only need to do a transformation for the input — the output for both algorithms is already compatible (at most, we may need to negate the output)

NP-Completeness

- Many (perhaps most) of the decision problems of interest are related to a computational or optimization problem.

NP-Completeness

- Many (perhaps most) of the decision problems of interest are related to a computational or optimization problem.
  - Examples:
  - Decision problem: Determine whether a given graph has a Hamiltonian cycle

NP-Completeness

- Many (perhaps most) of the decision problems of interest are related to a computational or optimization problem.
  - Examples:
  - Decision problem: Determine whether a given graph has a Hamiltonian cycle
    - Related computational problem: Find a Hamiltonian cycle (if there is one) in a given graph.

NP-Completeness

- Many (perhaps most) of the decision problems of interest are related to a computational or optimization problem.
  - Examples:
  - Decision problem: Determine whether a given graph has a Hamiltonian cycle
    - Related computational problem: Find a Hamiltonian cycle (if there is one) in a given graph.
  - Decision problem: Given two vertices in a given graph, is there a path between them shorter than a given value *k*?
    - Optimization problem: Find the shortest path.

NP-Completeness

- The decision problem always reduces to its associated optimization or computational problem

NP-Completeness

- The decision problem always reduces to its associated optimization or computational problem — examples:
  - We can solve the problem of determining whether there is a path shorter than *k* if we have an algorithm that finds the shortest path — simply invoke that algorithm, look at its output and check whether the shortest path's length is less than *k*.

NP-Completeness

- The decision problem always reduces to its associated optimization or computational problem — examples:
  - We can solve the problem of determining whether there is a path shorter than *k* if we have an algorithm that finds the shortest path — simply invoke that algorithm, look at its output and check whether the shortest path's length is less than *k*.
  - We can determine whether a graph has a Hamiltonian cycle if we have an algorithm that finds a Hamiltonian cycle in a graph.

NP-Completeness

- The converse is often true as well — the optimization or computational problem can often be reduced to the decision problem.

NP-Completeness

- The converse is often true as well — the optimization or computational problem can often be reduced to the decision problem. Example:
  - If we have an algorithm that determines whether there is a path shorter than a given value, we can determine the length of the shortest path by using that algorithm:

NP-Completeness

- The converse is often true as well — the optimization or computational problem can often be reduced to the decision problem. Example:
  - If we have an algorithm that determines whether there is a path shorter than a given value, we can determine the length of the shortest path by using that algorithm:
    - Call the decision algorithm many times, until we "narrow down" the length of the shortest path:
      - The algorithm outputs NO when asked whether there is a path shorter than 10, and outputs YES when asked whether there is a path shorter than 11 — we have our answer !

NP-Completeness

- This justifies the use of decision problems in the theoretical study of algorithms and complexities:

NP-Completeness

- This justifies the use of decision problems in the theoretical study of algorithms and complexities:
  - If we show that a decision problem is hard (intractable), then that automatically shows that the associated optimization or computational problem is also hard (since the decision problem reduces to the other one)

NP-Completeness

- Next, let's look at the sets, or *complexity classes*, P and NP

NP-Completeness

- The formal definitions are quite involved, so we'll skip those and go for the *somewhat* practical definitions...

NP-Completeness

- The formal definitions are quite involved, so we'll skip those and go for the *somewhat* practical definitions...
  - The complexity class P is the set of (decision) problems that can be decided in polynomial time.

NP-Completeness

- The formal definitions are quite involved, so we'll skip those and go for the *somewhat* practical definitions...
  - The complexity class P is the set of (decision) problems that can be decided in polynomial time.
    - That is, the problems for which a polynomial time algorithm exists that solves (decides) the problem.

NP-Completeness

- The formal definitions are quite involved, so we'll skip those and go for the *somewhat* practical definitions...
  - The complexity class P is the set of (decision) problems that can be decided in polynomial time.
    - That is, the problems for which a polynomial time algorithm exists that solves (decides) the problem.
  - Straightforward enough: P stands for Polynomial time (so, we're talking about the set of Polynomial time decidable problems)

NP-Completeness

- The complexity class NP is a bit trickier:
  - NP stands for Nondeterministic Polynomial time

NP-Completeness

- The complexity class NP is a bit trickier:
  - NP stands for Nondeterministic Polynomial time
    - The set of (decision) problems that can be decided in polynomial time using a *nondeterministic* machine!

NP-Completeness

- The complexity class NP is a bit trickier:
  - NP stands for Nondeterministic Polynomial time
    - The set of (decision) problems that can be decided in polynomial time using a *nondeterministic* machine!
  - Now, we don't want to deal with nondeterministic (as in, random) machines, so here is the tricky part:

NP-Completeness

- The complexity class NP is a bit trickier:
  - NP stands for Nondeterministic Polynomial time
    - The set of (decision) problems that can be decided in polynomial time using a *nondeterministic* machine!
  - Now, we don't want to deal with nondeterministic (as in, random) machines, so here is the tricky part:
  - It turns out that the above definition is equivalent to the following, more convenient one:
    - NP is the set of decision problems that can be decided in polynomial time if provided with a certificate, typically corresponding to a solution to the associated optimization or computational problem.

NP-Completeness

- The key detail here is counting on an efficient (i.e., polynomial time) *verification* algorithm.

NP-Completeness

- The key detail here is counting on an efficient (i.e., polynomial time) *verification* algorithm.
- For example: Determining whether a given graph has a Hamiltonian cycle is an NP problem:
  - It can be efficiently *verified* — that is, if we're given a certificate claiming to be a Hamiltonian cycle for the graph, we can verify that claim in polynomial time.

NP-Completeness

- The key detail here is counting on an efficient (i.e., polynomial time) *verification* algorithm.
- For example: Determining whether a given graph has a Hamiltonian cycle is an NP problem:
  - It can be efficiently *verified* — that is, if we're given a certificate claiming to be a Hamiltonian cycle for the graph, we can verify that claim in polynomial time (we only need to verify that all the vertices are visited, that all edges correspond to existing edges, and that no vertex is visited more than once).

NP-Completeness

- The key detail here is counting on an efficient (i.e., polynomial time) *verification* algorithm.
- The subset-sum problem is in NP as well:
  - Given a certificate corresponding to the subset of values that add to 0 (or the given parameter, in the more general form of the problem), we can verify in linear time that they indeed add to 0.

NP-Completeness

- Here's another tricky detail: The set P $\subseteq$ NP. That is, if a problem can be decided in polynomial time, then it can be verified in polynomial time.

NP-Completeness

- Here's another tricky detail: The set P $\subseteq$ NP. That is, if a problem can be decided in polynomial time, then it can be verified in polynomial time.
  - Just ignore the certificate — solve (decide) the problem in polynomial time, which can be done without the certificate, since the problem is in the set P.

NP-Completeness

- Here's a big (tricky?) question: Is NP $\subseteq$ P?

NP-Completeness

- Here's a big (tricky?) question: Is NP $\subseteq$ P?
  - That is: is every problem that can be efficiently verified also efficiently solved?

NP-Completeness

- Here's a big (tricky?) question: Is NP $\subseteq$ P?
  - That is: is every problem that can be efficiently verified also efficiently solved?
  - Basic intuition obviously says NO:

NP-Completeness

- Here's a big (tricky?) question: Is NP $\subseteq$ P?
  - That is: is every problem that can be efficiently verified also efficiently solved?
  - Basic intuition obviously says NO:
    - Think of finding a needle in a haystack vs. verifying that there is a needle at a location that we're told.

NP-Completeness

- Here's a big (tricky?) question: Is NP $\subseteq$ P?
  - That is: is every problem that can be efficiently verified also efficiently solved?
  - Basic intuition obviously says NO:
    - Think of finding a needle in a haystack vs. verifying that there is a needle at a location that we're told.
    - Find the factorization of a number vs. verify that a set of numbers is the factorization of a given number.

NP-Completeness

- We're used to this asymmetry between solving a problem and verifying a given solution.
- Some technologies even *rely* on this!

NP-Completeness

- We're used to this asymmetry between solving a problem and verifying a given solution.
- Some technologies even *rely* on this! Modern cryptography, for example, relies on the asymmetry of discrete logarithm problem:

NP-Completeness

- We're used to this asymmetry between solving a problem and verifying a given solution.
- Some technologies even *rely* on this! Modern cryptography, for example, relies on the asymmetry of discrete logarithm problem:

  given a^x mod m, determine the value of x — the asymmetry being: given x, it is trivial to obtain a^x mod m, but given a^x mod m, it is *believed* that finding x is an intractable problem)

NP-Completeness

- Notice, BTW, that NP $\subseteq$ P would imply that P = NP (i.e., that the two sets are the same), since P $\subseteq$ NP

NP-Completeness

- Despite this intuition, as obvious as it may seem, whether P = NP is an unanswered question; not only that: it is considered the single most important open problem in theoretical computer science!

NP-Completeness

- Despite this intuition, as obvious as it may seem, whether P = NP is an unanswered question; not only that: it is considered the single most important open problem in theoretical computer science!
  - Perhaps the puzzling aspect is: if it seems so obvious that P ≠ NP, why has no-one been able to mathematically prove it?

NP-Completeness

- Despite this intuition, as obvious as it may seem, whether P = NP is an unanswered question; not only that: it is considered the single most important open problem in theoretical computer science!
  - Perhaps the puzzling aspect is: if it seems so obvious that P ≠ NP, why has no-one been able to mathematically prove it?
    - Corollary: if no-one, despite a lot of brilliant people trying really hard, has been able to prove it (that P ≠ NP), could it be because it is not true that P ≠ NP??

NP-Completeness

- NP-complete problems:
  - In short, a set of problems in NP that all reduce to each other!

NP-Completeness

- NP-complete problems:
  - In short, a set of problems in NP that all reduce to each other!
  - Huh?? How could this happen???

NP-Completeness

- NP-complete problems:
  - In short, a set of problems in NP that all reduce to each other!
  - Huh?? How could this happen??? (back to this in a minute)

NP-Completeness

- NP-complete problems:
  - In short, a set of problems in NP that all reduce to each other!
  - One important implication is that these problems are all equivalent:
    - If someone ever finds a polynomial time solution to one of them, then *all* of them have polynomial time solutions.
    - Conversely, if someone ever proves that one of these problems is intractable, then we know that *all* of them are intractable.
      - Corollary: this would prove that $P \neq NP$!!!

NP-Completeness

- NP-complete problems:
  - How do we prove that a given problem is NP-complete?

NP-Completeness

- NP-complete problems:
  - How do we prove that a given problem is NP-complete?
    - If we can reduce any other problem X_{NPC} known to be NP-complete, then we're done — every other NP-complete problem reduces in polynomial time to X_{NPC}, and therefore to our given problem

NP-Completeness

- NP-complete problems:
  - How do we prove that a given problem is NP-complete?
    - If we can reduce any other problem X_{NPC} known to be NP-complete, then we're done — every other NP-complete problem reduces in polynomial time to X_{NPC}, and therefore to our given problem
    - Wait ... That shows that every other NP-complete problem reduces to ours — how do we know that ours reduce to either one of the other NP-complete problems?

NP-Completeness

- NP-complete problems:
  - How do we prove that a given problem is NP-complete?
    - Turns out that this detail has already been taken care of (that's how we know about this class of NP-complete problems to begin with!)
    - So, we really only need to show that some problem known to be NP-complete reduces to ours, and that proves that ours is also NP-complete!

NP-Completeness

- NP-complete problems:
  - It all goes back to the first problem discovered to be NP-complete — the Circuit satisfiability problem:
    - Given a combinational circuit (with logic gates AND, OR, and NOT) with *n* inputs and a single output (a boolean output), is there any combination of input values that produces an output of 1? (a set of input values that *satisfies* the logic given by the circuit)

NP-Completeness

- NP-complete problems:
  - It all goes back to the first problem discovered to be NP-complete — the Circuit satisfiability problem:
    - Given a combinational circuit (with logic gates AND, OR, and NOT) with *n* inputs and a single output (a boolean output), is there any combination of input values that produces an output of 1? (a set of input values that *satisfies* the logic given by the circuit)
      - We observe that the problem can be trivially solved in exponential time — try all 2^n possible combinations of n binary values, checking whether the output is 1.
      - It is not known, however, whether a polynomial time solution exists (none is known, but no-one has been able to prove that none exist)

NP-Completeness

- NP-complete problems:
  - It all goes back to the first problem discovered to be NP-complete — the Circuit satisfiability problem:
    - It is also straightforward to see that this problem is NP
      - If given a certificate consisting of a combination of 1 and 0's that satisfies the circuit, we can verify it in polynomial time (there is the implicit assumption that the number of logic gates is polynomial with respect to the number of inputs)

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?
    - Well, you guys know how a digital computer works, right? (you guys saw ECE-124 and are now near completion of ECE-222, so for you, this proof will be a piece of cake!)

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?
    - Well, you guys know how a digital computer works, right? (you guys saw ECE-124 and are now near completion of ECE-222, so for you, this proof will be a piece of cake!)
      - Any problem that is in NP has an efficient verification algorithm.

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?
    - Well, you guys know how a digital computer works, right? (you guys saw ECE-124 and are now near completion of ECE-222, so for you, this proof will be a piece of cake!)
      - Any problem that is in NP has an efficient verification algorithm.
      - That algorithm can be implemented in a digital computer.

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?
    - Well, you guys know how a digital computer works, right? (you guys saw ECE-124 and are now near completion of ECE-222, so for you, this proof will be a piece of cake!)
      - Any problem that is in NP has an efficient verification algorithm.
      - That algorithm can be implemented in a digital computer.
      - Because the verification occurs in polynomial time, then that means that the algorithm will complete execution in a polynomial number of clock cycles.

NP-Completeness

- NP-complete problems:
  - How is it that it was discovered that every problem in NP reduces to the circuit satisfiability problem?
    - Well, you guys know how a digital computer works, right? (you guys saw ECE-124 and are now near completion of ECE-222, so for you, this proof will be a piece of cake!)
      - Any problem that is in NP has an efficient verification algorithm.
      - That algorithm can be implemented in a digital computer.
      - Because the verification occurs in polynomial time, then that means that the algorithm will complete execution in a polynomial number of clock cycles.
      - ... Anyone sees where this is going?

NP-Completeness

- NP-complete problems:
  - A digital computer has a state, represented by all of its internal registers (the ones that are "visible" to the user through the ISA, and the ones that aren't; i.e., the ones that are part of the control circuitry)
  - At each clock cycle, a huge (but fixed-size) set of combinational circuits map the current state to the next state.
    - States are "captured" at each clock edge by the arrays of flip-flops (the registers), and then the outputs of these registers are fed to the combinational circuits to determine the next state.

NP-Completeness

- NP-complete problems:
  - So, I'll ask again: anyone sees where this is going?

NP-Completeness

- NP-complete problems:
  - We could construct a circuit that "executes" the entire algorithm by eliminating the clock and the registers that provide a "state", and instead, simply put multiple copies (stages) of the combinational circuits that determine the next state.

NP-Completeness

- NP-complete problems:
  - We could construct a circuit that “executes” the entire algorithm by eliminating the clock and the registers that provide a “state”, and instead, simply put multiple copies (stages) of the combinational circuits that determine the next state.
    - The output of each stage (representing the next state) is not fed back to registers that will capture the value: instead, that output is directly fed as input to the next stage of combinational circuits.

NP-Completeness

- NP-complete problems:

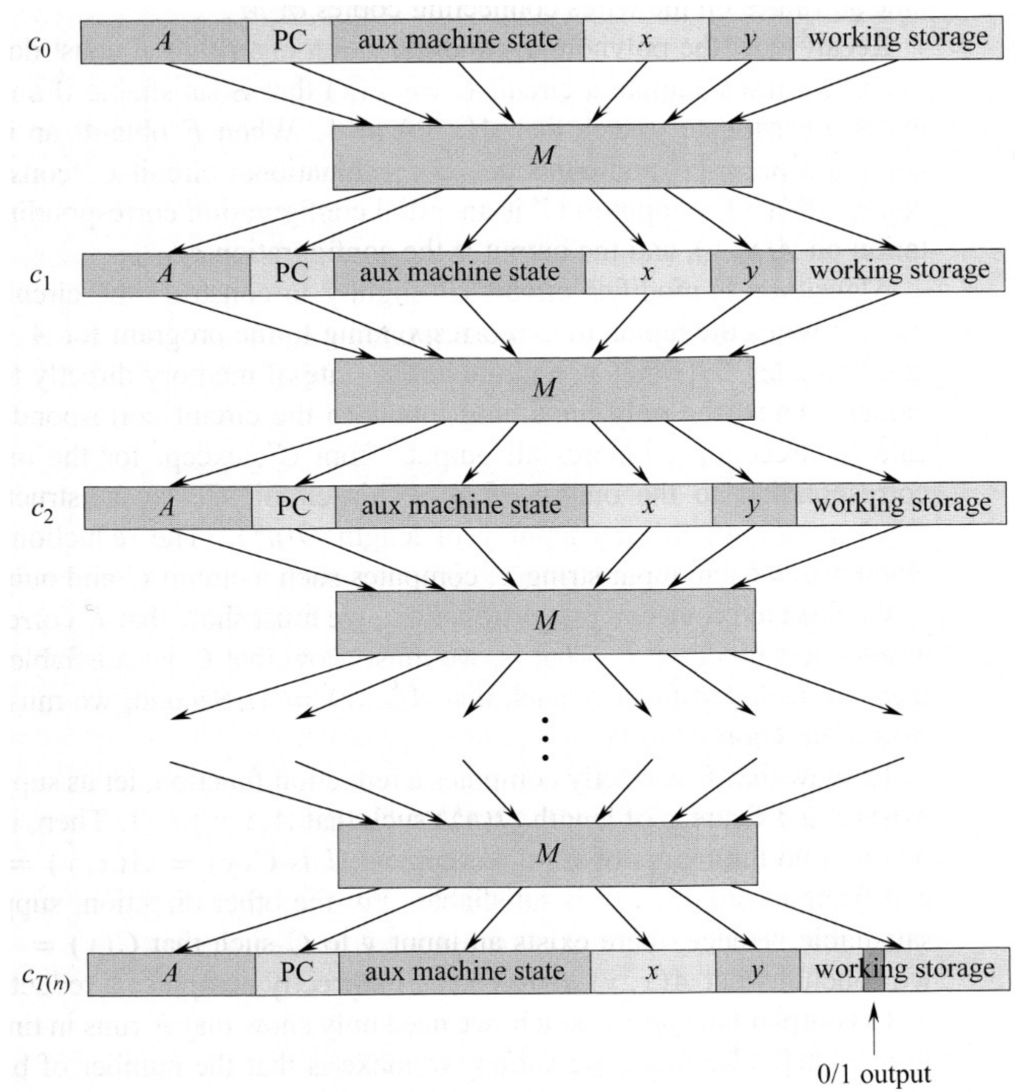

CLRS, Fig. 34.9

NP-Completeness

- NP-complete problems:
  - Bottom line: for any given NP problem, we can implement the the verification procedure, and from there, generate a combinational circuit that "executes" it sort of instantaneously (plus-minus circuits response time and signals propagation time).

NP-Completeness

- NP-complete problems:
  - Bottom line: for any given NP problem, we can implement the the verification procedure, and from there, generate a combinational circuit that “executes” it sort of instantaneously (plus-minus circuits response time and signals propagation time).
  - If the output of that circuit is 1 for a given input, then that input verifies the problem.
    - So, if we could solve the circuit satisfiability problem in polynomial time, then we would be able to *find* a solution to any problem for which a solution can be verified in polynomial time.

NP-Completeness

- NP-complete problems:
  - Recall that a solution to circuit-sat tells you whether there exists a combination of ones and zeros that produces a 1 as the output.
  - But that output is the solution to the other NP problem.

NP-Completeness

- NP-complete problems:
  - Recall that a solution to circuit-sat tells you whether there exists a combination of ones and zeros that produces a 1 as the output.
  - But that output is the solution to the other NP problem.
  - Important detail: the number of stages in that huge circuit is the number of clock cycles that the algorithm takes — since we're talking about NP problems, then that number of cycles is polynomial, and thus the number of logic gates will be polynomial with respect to the size of the input!

NP-Completeness

- NP-complete problems:
  - Bottom line (as in, the *real* bottom line): this construction shows that *every NP problem* reduces to the circuit satisfiability problem!

NP-Completeness

- NP-complete problems:
  - Bottom line (as in, the *real* bottom line): this construction shows that *every NP problem* reduces to the circuit satisfiability problem!
  - That means that this problem is at least as hard as every other problem in NP — thus, it is at least as hard as the hardest problem(s) in NP

NP-Completeness

- NP-complete problems:
  - Bottom line (as in, the *real* bottom line): this construction shows that *every NP problem* reduces to the circuit satisfiability problem!
  - That means that this problem is at least as hard as every other problem in NP — thus, it is at least as hard as the hardest problem(s) in NP
    - This is the definition of an NP-hard problem: A problem such that every problem in NP can be reduced to it.

NP-Completeness

- NP-complete problems:
  - Bottom line (as in, the *real* bottom line): this construction shows that *every NP problem* reduces to the circuit satisfiability problem!
  - That means that this problem is at least as hard as every other problem in NP — thus, it is at least as hard as the hardest problem(s) in NP
    - This is the definition of an NP-hard problem: A problem such that every problem in NP can be reduced to it.
    - Notice that an NP-hard problem does not have to be in NP (it doesn't even need to be a decision problem!)
      - An NP-complete problem is an NP-hard problem that is NP

NP-Completeness

- NP-complete problems — the complete picture:
  - Circuit-sat was the firts problem to be shown to be NP-hard; since it clearly is NP, then that makes it NP-complete.

NP-Completeness

- NP-complete problems — the complete picture:
  - Circuit-sat was the firts problem to be shown to be NP-hard; since it clearly is NP, then that makes it NP-complete.
  - If we can reduce circuit-sat to any other NP problem, that would make that other problem *equivalent* to circuit-sat — since circuit-sat reduces to that problem, but we know that that problem reduces to circuit-sat

NP-Completeness

- NP-complete problems — the complete picture:
  - Circuit-sat was the firts problem to be shown to be NP-hard; since it clearly is NP, then that makes it NP-complete.
  - If we can reduce circuit-sat to any other NP problem, that would make that other problem *equivalent* to circuit-sat — since circuit-sat reduces to that problem, but we know that that problem reduces to circuit-sat
    - So, circuit-sat really “closes the loop” that ties every NP-complete problem to every other NP-complete problem: solving either one solves every one.

NP-Completeness

- NP-complete problems — the complete picture:

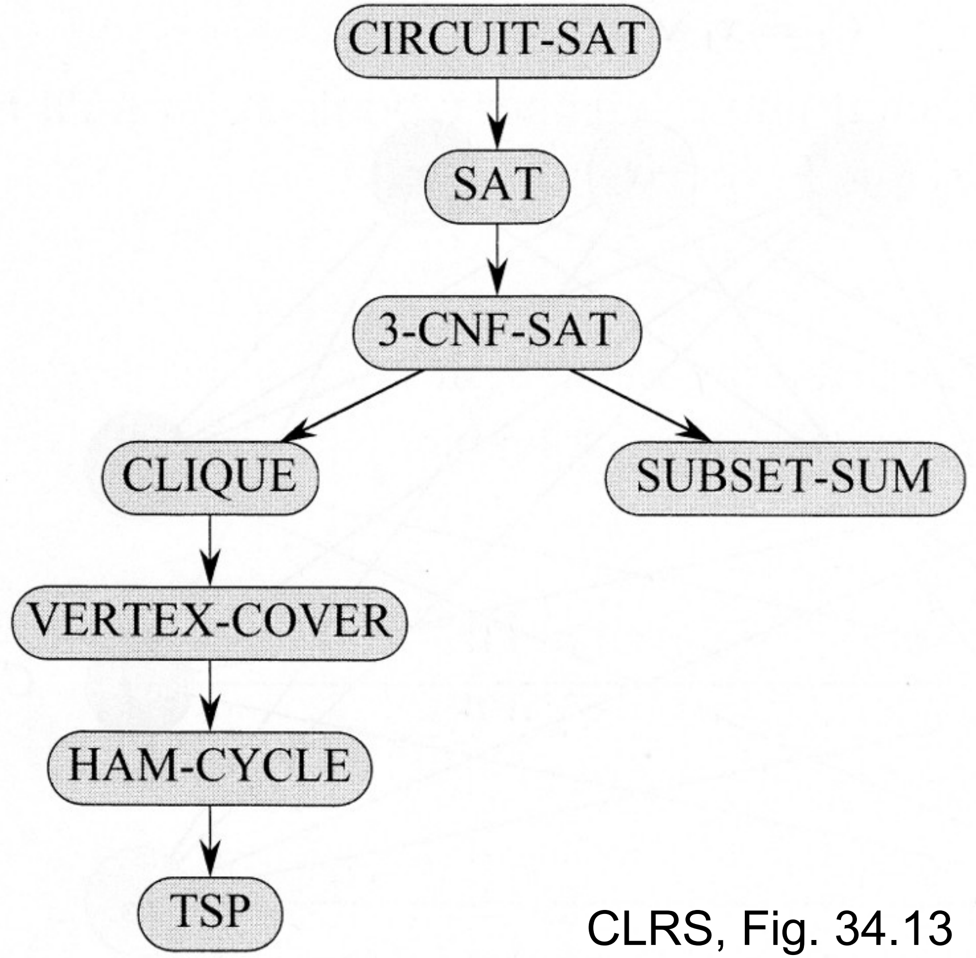

CLRS, Fig. 34.13

NP-Completeness

- For example, the Travelling Salesman Problem (TSP) is a surprisingly ubiquitous problem (it shows up in many optimization problems, in pattern recognition, circuit design, DNA sequencing, economics, etc.)

NP-Completeness

- For example, the Travelling Salesman Problem (TSP) is a surprisingly ubiquitous problem (it shows up in many optimization problems, in pattern recognition, circuit design, DNA sequencing, economics, etc.)
  - Given a complete weighted graph (a graph for which every vertex is adjacent to every other vertex), is there a path of length less than or equal to a given parameter *k* that visits each vertex exactly once returning to the starting vertex?

NP-Completeness

- For example, the Travelling Salesman Problem (TSP) is a surprisingly ubiquitous problem (it shows up in many optimization problems, in pattern recognition, circuit design, DNA sequencing, economics, etc.)
  - Given a complete weighted graph (a graph for which every vertex is adjacent to every other vertex), is there a path of length less than or equal to a given parameter *k* that visits each vertex exactly once returning to the starting vertex?
    - The related optimization problem being: find the shortest path that visits each vertex exactly once returning to the first one.

NP-Completeness

- Side note:
  - In its original formulation, the problem is given *n* cities with pairwise distances, find the shortest itinerary that visits each city exactly once and returns to the starting city. (but this immediately translates into the more abstract formulation with a complete graph)

NP-Completeness

- TSP is obviously NP — if we're given a solution consisting of a path that meets the requirements, we can easily verify that each vertex is visited exactly once, and we can easily compute the length of the path and verify that it is less than or equal to *k*.

NP-Completeness

- Reducing the Hamiltonian cycle problem to TSP is surprisingly straightforward:

NP-Completeness

- Reducing the Hamiltonian cycle problem to TSP is surprisingly straightforward:
  - Given a graph *G*, construct a complete weighted graph *G′* that has the same vertices as *G*, and such that for every pair of vertices (u,v), the edge (u,v) in G' has weight 0 if the edge (u,v) is present in G, and weight 1 otherwise.

NP-Completeness

- Reducing the Hamiltonian cycle problem to TSP is surprisingly straightforward:
  - Given a graph *G*, construct a complete weighted graph *G′* that has the same vertices as *G*, and such that for every pair of vertices (u,v), the edge (u,v) in G' has weight 0 if the edge (u,v) is present in G, and weight 1 otherwise.
    - Then, we ask the algorithm that decides TSP whether there is a path of length 0 in *G′* that visits each vertex exactly once.

NP-Completeness

- Reducing the Hamiltonian cycle problem to TSP is surprisingly straightforward:
  - Given a graph *G*, construct a complete weighted graph *G′* that has the same vertices as *G*, and such that for every pair of vertices (u,v), the edge (u,v) in G' has weight 0 if the edge (u,v) is present in G, and weight 1 otherwise.
    - Then, we ask the algorithm that decides TSP whether there is a path of length 0 in *G′* that visits each vertex exactly once... Anyone sees why this works?

NP-Completeness

- Reducing the Hamiltonian cycle problem to TSP is surprisingly straightforward:
  - Given a graph *G*, construct a complete weighted graph *G′* that has the same vertices as *G*, and such that for every pair of vertices (u,v), the edge (u,v) in G' has weight 0 if the edge (u,v) is present in G, and weight 1 otherwise.
    - Then, we ask the algorithm that decides TSP whether there is a path of length 0 in *G′* that visits each vertex exactly once... Anyone sees why this works?
      - If there is, then it means that only edges with weight 0 could have been used — but those are edges that are present in *G*, and so if there is such a path in *G′*, then there is a Hamiltonian cycle in *G*.

NP-Completeness

- The clique and vertex cover are also related by a simple construction:

NP-Completeness

- The clique and vertex cover are also related by a simple construction:
  - Clique problem: a clique is a sub-graph in which every vertex is connected to every other vertex. The clique problem being: given a graph, is there a clique in it? (or, is there a clique of at least *k* vertices?)

NP-Completeness

- The clique and vertex cover are also related by a simple construction:
  - Clique problem: a clique is a sub-graph in which every vertex is connected to every other vertex. The clique problem being: given a graph, is there a clique in it? (or, is there a clique of at least *k* vertices?)
  - Vertex cover: a vertex cover is a sub-graph such that every vertex in the graph is adjacent to some vertex or vertices in that sub-graph (so, that subset of vertices "covers" the entire graph)

NP-Completeness

- The reduction is not as trivial as Hamiltonian cycle to TSP, but it is still remarkably simple:
  - Given a graph *G*, we construct a graph *G′* with the same set of vertices, and where the set of edges is the complement of the set of edges in *G* (every edge present in *G* is missing in *G′* and every edge missing in *G* is present in *G′*).

NP-Completeness

- The reduction is not as trivial as Hamiltonian cycle to TSP, but it is still remarkably simple:
  - Given a graph G, we construct a graph G' with the same set of vertices, and where the set of edges is the complement of the set of edges in G (every edge present in G is missing in G' and every edge missing in G is present in G').
    - It can be shown that G has a clique of size k if and only if G' has a vertex cover of size $n-k$.
    - This property provides the sought reduction.

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?
  - Instances of these problems or problems similar to them tend to show up very often in engineering (not only in software!!).

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?
  - Instances of these problems or problems similar to them tend to show up very often in engineering (not only in software!!).
  - If a problem is NP-complete, we don't know that no efficient solution exists, but no-one knows of an efficient solution for either one of these problems, and it is *believed* that none exist!

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?

  - As engineers, knowing about these problems (and knowing which problems are NP-complete) can save us lots of time — time which could have otherwise wasted looking for an algorithm to solve a problem which is extremely unlikely (virtually impossible) that we would find a solution!

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?
  - As engineers, knowing about these problems (and knowing which problems are NP-complete) can save us lots of time — time which could have otherwise wasted looking for an algorithm to solve a problem which is extremely unlikely (virtually impossible) that we would find a solution!
  - Plus, c'mon — this has to be in the category of *all-time most awesome ideas ... ever* !!

NP-Completeness

- Big picture:
  - Why are these notions relevant for us, engineers?

  - As engineers, knowing about these problems (and knowing which problems are NP-complete) can save us lots of time — time which could have otherwise wasted looking for an algorithm to solve a problem which is extremely unlikely (virtually impossible) that we would find a solution!
  - Plus, c'mon — this has to be in the category of *all-time most awesome ideas ... ever* !!
    - So, who was worried about a shortage of coolness in this course? :-)

Summary

- During today's lesson:
  - Categorized problems as polynomial time (tractable) vs. non-polynomial time (intractable).
  - Looked into the notion of decision problems, and their relationship to an associated computation or optimization problem.
  - Introduced the sets P and NP, and the big, mysterious, open question: Is P = NP?
  - Introduced the notions of NP-hard and NP-complete problems.
  - Argued about the relevance of these notions in a practical setting (for an engineer)