Computational Complexity Theory

Lecture I: Intro; Turing machines; Class P

Department of Computer Science, Indian Institute of Science

 Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational **problems** come in various flavors:

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational **problems** come in various flavors:
 a. Decision problem

Example: Is vertex t reachable from vertex s in graph G? Is n a prime number?

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational **problems** come in various flavors:

a. Decision problemb. Search problem

Example: Find a satisfying assignment for a Boolean formula. Find a prime between n and 2n.

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational **problems** come in various flavors:
 - a. Decision problem
 - b. Search problem
 - c. Counting problem
 - Example: Count the number of cycles in a graph.
 - Count the number of perfect matchings in a graph.

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational **problems** come in various flavors:
 - a. Decision problem
 - b. Search problem
 - c. Counting problem
 - d. Optimization problem

Example: Find a minimum size vertex cover in a graph

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Algorithms are <u>methods</u> for solving problems; they are studied using formal <u>models of computation</u>, like Turing machines.
 - a memory with head (like a RAM)
 - a finite control (like a processor)

(...more later in this lecture)

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational resources (required by models of computation) can be:
 - Time (bit operations)
 - Space (memory cells)

- Computational complexity attempts to classify computational problems based on the amount of resources required by algorithms to solve them.
- Computational resources (required by models of computation) can be:
 - Time (bit operations)
 - Space (memory cells)
 - Random bits (magic bits: 0 w. p $\frac{1}{2}$ and 1 w.p $\frac{1}{2}$)
 - Communication (bit exchanges)

Topics to be covered in this course



Structural Complexity

- Classes P, NP, co-NP... NP-completeness.
 - How hard is it to check satisfiability of a Boolean formula?
 - What if the formula has exactly one or no satisfying assignment?

Structural Complexity

- Classes P, NP, co-NP... NP-completeness.
- Space bounded computation.
 - How much space is required to check s-t connectivity?

Structural Complexity

- Classes P, NP, co-NP... NP-completeness.
- Space bounded computation.
- Polynomial Hierarchy (PH).



- How hard is it to check if the largest independent set in **G** has size **k** ?
- How hard is it to check if there is a circuit of size k that computes the same Boolean function as a given Boolean circuit C ?

Circuit Complexity

- The internal workings of an algorithm can be viewed as a Boolean circuit -- a nice combinatorial model of computation that is closely related to Turing Machines.
- The <u>size</u>, <u>depth</u> & <u>width</u> of a circuit correspond to the <u>sequential</u>, <u>parallel</u> & <u>space</u> complexity, respectively, of the algorithm that it represents.

Circuit Complexity

- The internal workings of an algorithm can be viewed as a Boolean circuit -- a nice combinatorial model of computation that is closely related to Turing Machines.
- The <u>size</u>, <u>depth</u> & <u>width</u> of a circuit correspond to the <u>sequential</u>, <u>parallel</u> & <u>space</u> complexity, respectively, of the algorithm that it represents.
- Proving P ≠ NP reduces to showing circuit lower bounds.
 - We will see lower bounds for restricted classes of circuits.

Randomness in Computation

- Probabilistic complexity classes (BPP, RP, co-RP).
 - Does randomization help in improving efficiency?
 - Quicksort has O(n log n) expected time but O(n^2) worst case time.
 - Can SAT be solved in polynomial time using randomness?

Theorem (Schoening, 1999): 3SAT can be solved in randomized $O((4/3)^n)$ time.

Counting Complexity

- Counting complexity classes (class #P).
 - How hard is it to count the number of perfect matchings in a graph?
 - How hard is it to count the number of cycles in a graph?
 - Is counting much harder than the corresponding decision problem?

Hardness of Approximation

• Probabilistically Checkable Proofs (PCPs).

Hardness of approximation results.

Theorem (Hastad, 1997): If there's a poly-time algorithm to compute an assignment that satisfies at least $7/8 + \varepsilon$ fraction of the clauses of an input 3SAT, for any constant $\varepsilon > 0$, then P = NP.

- Course no.: E0 224 Credits: 3:1
- Instructor: Chandan Saha
- Lecture time: M,W 2-3:30 pm. Venue: GMeet
- Link: meet.google.com/aab-vqzy-weu
- Course homepage:

https://www.csa.iisc.ac.in/~chandan/courses/complexity21/home.html

- **Prerequisites**: Basic familiarity with algorithms; Mathematical maturity.
- Primary reference: Computational Complexity A Modern Approach by Sanjeev Arora and Boaz Barak.
- Lectures: Slides will be posted on the course homepage. Recordings will also be shared.
- Number of lectures: ~25.

- **Communications:** All course related communications will happen over a <u>Google group</u>. I'll share the group-id after <u>Aug 16</u> (course registration deadline).
- **Google form:** Need to fill in a Google form. Link will be shared in class after <u>Aug 16</u>. Requires sign-in to Gmail.
- Till Aug 16, links to the recordings will be provided on the course homepage.
- Subsequent recordings will be shared with the Google group. Lecture videos will appear in your Google drive.

Grading policy: Three assignments - 45%
 One presentation - 25%
 Final exam - 30%

Assignments

- First assignment: Will posted on <u>Aug 31</u>; due date will be <u>Sep 14</u>.
- Second assignment: Will posted on <u>Sep 30</u>; due date will be <u>Oct 14</u>.
- Third assignment: Will posted on <u>Oct 31</u>; due date will be <u>Nov 14</u>.
- Mode: Assignments will be posted on the course homepage. You need to e-mail me your assignment as a pdf file (use Latex). (Will use Moodle if required.)

Presentations

- A group of 2 students would present a paper/result.
- Duration of a presentation: ~45 mins
- Mode: Pre-recorded talks using slides.
- I will start giving topics to present from Sep 15.All topics will be handed out by Oct 15.
- You get 4 weeks to prepare a presentation.

Final exam

- Would be a one-on-one ~I hr long oral exam, if there aren't many students crediting the course.
- Else, it would be a one-day long take-home exam.



- An algorithm is a set of instructions or rules.
- To understand the performance of an algorithm we need a <u>model of computation</u>. Turing machine is one such *natural* model (introduced by Alan Turing in 1936).

- An algorithm is a set of instructions or rules.
- To understand the performance of an algorithm we need a <u>model of computation</u>. Turing machine is one such *natural* model (introduced by Alan Turing in 1936).
- ATM consists of:
 - Memory tape(s)
 - A finite set of rules

- An algorithm is a set of instructions or rules.
- To understand the performance of an algorithm we need a <u>model of computation</u>. Turing machine is one such *natural* model (introduced by Alan Turing in 1936).
- ATM consists of:
 - Memory tape(s)
 - A finite set of rules
- Turing machines A mathematical way to describe algorithms.

- An algorithm is a set of instructions or rules.
- To understand the performance of an algorithm we need a <u>model of computation</u>. Turing machine is one such *natural* model (introduced by Alan Turing in 1936).
- ATM consists of:
 - Memory tape(s)
 - A finite set of rules

(e.g. of a physical realization of a TM is a simple adder)

• Definition. A k-tape Turing Machine M is described by a tuple (Γ, Q, δ) such that

- Definition. A k-tape Turing Machine M is described by a tuple (Γ, Q, δ) such that
- M has k memory tapes (input/work/output tapes) with *heads*;
- Fis a finite set of alphabets. (Every memory cell contains an element of F)

- Definition. A k-tape Turing Machine M is described by a tuple (Γ, Q, δ) such that
- M has k memory tapes (input/work/output tapes) with *heads*;
- Fis a finite set of alphabets. (Every memory cell contains an element of

has a <mark>blank</mark> symbol

- Definition. A k-tape Turing Machine M is described by a tuple (Γ, Q, δ) such that
- M has k memory tapes (input/work/output tapes) with heads:
- Fis a finite set of alphabets. (Every memory cell contains an element of Γ)
- Q is a finite set of states. (special states: q_{start} , q_{halt}) δ is a function from Q x Γ^{k} to Q x Γ^{k} x {L,S,R}

- Definition. A k-tape Turing Machine M is described by a tuple (Γ, Q, δ) such that
- M has k memory tapes (input/work/output tapes) with heads;
- Fis a finite set of alphabets. (Every memory cell contains an element of Γ)
- Q is a finite set of states. (special states: q_{start} , q_{halt}) δ is a function from Q x Γ^{k} to Q x Γ^{k} x {L,S,R}

known as transition function; it captures the dynamics of M

Turing Machines: Computation

• Start configuration.

> All tapes other than the input tape contain blank symbols.

> The input tape contains the input string.

> All the head positions are at the start of the tapes.

 \geq The machine is in the start state q_{start} .

Turing Machines: Computation

• Start configuration.

> All tapes other than the input tape contain blank symbols.

> The input tape contains the input string.

- > All the head positions are at the start of the tapes.
- \geq The machine is in the start state q_{start} .

• Computation.

> A step of computation is performed by applying δ .

• Halting.

 \geq Once the machine enters q_{halt} it stops computation.

Turing Machines: Running time

- Let f: {0,1}* → {0,1}* and T: N → N and M be a Turing machine.
- Definition. We say M computes f if on every x in {0,1}*, M halts with f(x) on its output tape beginning from the start configuration with x on its input tape.

Turing Machines: Running time

- Let f: {0,1}* → {0,1}* and T: N → N and M be a Turing machine.
- Definition. We say M computes f if on every x in {0,1}*, M halts with f(x) on its output tape beginning from the start configuration with x on its input tape.
- Definition. M computes f in T(|x|) time, if for every x in {0,1}*, M halts within T(|x|) steps of computation and outputs f(x).

• In this course, we would be dealing with

Turing machines that <u>halt on every input</u>.
 Computational problems that can be solved by Turing machines.

- In this course, we would be dealing with
 - Turing machines that <u>halt on every input</u>.
 Computational problems that can be solved by Turing machines.
- Can every computational problem be solved using Turing machines?

- There are problems for which there exists *no* TM that halts on every input instances of the problem and outputs the correct answer.
 - Input: A system of polynomial equations in many variables with integer coefficients.
 - Output: Check if the system has integer solutions.
 - Question: Is there an algorithm to solve this problem?

• There are problems for which there exists *no* TM that halts on every input instances of the problem and outputs the correct answer.

> A typical input instance:

 $x^{2}y + 5y^{3} = 3$ $x^{2} + z^{5} - 3y^{2} = 0$ Integer solutions for x, y, z? $y^{2} - 4z^{6} = 0$

- There are problems for which there exists *no* TM that halts on every input instances of the problem and outputs the correct answer.
 - Input: A system of polynomial equations in many variables with integer coefficients.
 - Output: Check if the system has integer solutions.
 - Question: Is there an algorithm to solve this problem?

(Hilbert's tenth problem, 1900)

- There are problems for which there exists *no* TM that halts on every input instances of the problem and outputs the correct answer.
 - Input: A system of polynomial equations in many variables with integer coefficients.
 - Output: Check if the system has integer solutions.
 - Question: Is there an algorithm to solve this problem?
- Theorem. There doesn't exist any algorithm (realizable by a TM) to solve this problem. (Davis, Putnam, Robinson, Matiyasevich 1970)

Why Turing Machines?

- TMs are natural and intuitive.
- Church-Turing thesis: "Every physically realizable computation device – whether it's based on silicon, DNA, neurons or some other alien technology – can be simulated by a Turing machine".

--- [quoted from Arora-Barak's book]

Why Turing Machines?

- TMs are natural and intuitive.
- Church-Turing thesis: "Every physically realizable computation device – whether it's based on silicon, DNA, neurons or some other alien technology – can be simulated by a Turing machine".

--- [quoted from Arora-Barak's book]

Several other computational models can be simulated by TMs.

Why Turing Machines?

- TMs are natural and intuitive.
- Church-Turing thesis: "Every physically realizable computation device – whether it's based on silicon, DNA, neurons or some other alien technology – can be simulated by a Turing machine".

--- [quoted from Arora-Barak's book]

Several other computational models can be simulated by TMs.

Possible exception: Quantum machines!

Basic facts about TMs

- Time constructible functions. A function T: $N \rightarrow N$ is <u>time constructible</u> if $T(n) \ge n$ and there's a TM that computes the function that maps x to T(|x|) in O(T(|x|)) time.
- Examples: $T(n) = n^2$, or 2^n , or $n \log n$

Turing Machines: Robustness

- Let f: {0,1}* → {0,1}* and T: N → N be a time constructible function.
- Binary alphabets suffice.

If a TM M computes f in T(n) time using Γ as the alphabet set, then there's another TM M' that computes f in time 4.log |Γ|.T(n) using {0, I, blank} as the alphabet set.

Turing Machines: Robustness

- Let f: {0,1}* → {0,1}* and T: N → N be a time constructible function.
- Binary alphabets suffice.

If a TM M computes f in T(n) time using Γ as the alphabet set, then there's another TM M' that computes f in time 4.log |Γ|.T(n) using {0, I, blank} as the alphabet set.

• A single tape suffices.

> If a TM M computes f in T(n) time using k tapes then there's another TM M' that computes f in time $5k \cdot T(n)^2$ using a single tape that is used for input, work and output.

Every TM can be represented by a finite string over {0,1}.

... simply encode the description of the TM.

- Every TM can be represented by a finite string over {0,1}.
- Every string over {0,1} represents some TM.
 ...invalid strings map to a fixed, trivial TM.

- Every TM can be represented by a finite string over {0,1}.
- Every string over {0,1} represents some TM.
- Every TM has infinitely many string representations. ... allow padding with arbitrary number of 0's

- Every TM can be represented by a finite string over {0,1}.
- Every string over {0,1} represents some TM.
- Every TM has infinitely many string representations.



- Every TM can be represented by a finite string over {0,1}.
- Every string over {0,1} represents some TM.
- Every TM has infinitely many string representations.
- ATM (i.e., its string representation) can be given as an input to another TM !!

Universal Turing Machines

- Theorem. There exists a TM U that on every input x, α in {0,1}* outputs $M_{\alpha}(x)$.
- Further, if M_{α} halts within T steps then U halts within C. T. log T steps, where C is a constant that depends only on M_{α} 's <u>alphabet size</u>, <u>number of states</u> and <u>number of tapes</u>.

Universal Turing Machines

- Theorem. There exists a TM U that on every input x, α in {0,1}* outputs $M_{\alpha}(x)$.
- Further, if M_{α} halts within T steps then U halts within C. T. log T steps, where C is a constant that depends only on M_{α} 's alphabet size, number of states and number of tapes.
- Physical realization of UTMs are modern day electronic computers.

Complexity classes P and FP

• In the initial part of this course, we'll focus primarily on decision problems.

- In the initial part of this course, we'll focus primarily on decision problems.
- Decision problems can be naturally identified with Boolean functions, i.e., functions from {0,1}* to {0,1}.

- In the initial part of this course, we'll focus primarily on decision problems.
- Decision problems can be naturally identified with Boolean functions, i.e., functions from {0,1}* to {0,1}.
- Boolean functions can be naturally identified with sets of {0,1} strings, also called languages.



• Definition. We say a TM M <u>decides a language</u> $L \subseteq \{0, I\}^*$ if M computes f_L , where $f_L(x) = I$ if and only if $x \in L$.

The characteristic function of L .

Complexity Class P

- Let T: $N \rightarrow N$ be some function.
- Definition: A language L is in DTIME(T(n)) if there's a TM that decides L in time O(T(n)).
- Definition: Class $P = \bigcup_{c>0} DTIME (n^c)$.

Complexity Class P

- Let T: $N \rightarrow N$ be some function.
- Definition: A language L is in DTIME(T(n)) if there's a TM that decides L in time O(T(n)).

Definition: Class P = U DTIME (n^c).
 Deterministic polynomial-time