

Computability and Complexity (CoCo) 2024

Lecturers

Jakob Nordström (main instructor)

Srikanth Srinivasan

Amir Yehudayoff

Algorithms & Complexity Section

Department of Computer Science (DIKU)

Teaching assistant

Shuo Pang

Course webpage

<https://jakobnordstrom.se/teaching/CoCo24/>

The course webpage contains all practical information about the course.

Announcements, problem sets submissions, questions, and discussions will be dealt with on Absalon.

Course material

Computational Complexity: A Modern Approach by Sanjeev Arora & Boaz Barak
Plus probably some other material towards the end of the course.

Advice: Prepare for lectures by skimming the chapter(s) specified in advance.

Problem sets

4 problem sets will be posted on the course webpage and on Absalon.

See course webpage for (tentative) deadlines and problem set rules.

Collaborations in groups of 2-3 encouraged.

But all solutions should be written **individually** and **from scratch** without sharing any material (or generating it from the Internet).

The best way to learn this type of material is not just by reading, but by “getting your hands dirty” working on problems.

Expect to have to think for quite some time about some of the problems.

COMPUTATIONAL COMPLEXITY THEORY

What is efficiently computable in practice
(given that resources are limited)

Computation

- Informal understanding since ancient times
"write down symbols following certain rules"
- 1st half of 20th century: precise, mathematical definition
- Invention of (electronic) computers
- Now computers omnipresent (laptops, cellphones, Internet)
- But computation is about much more than computers - happens in many other ways
 - o biology (e.g., DNA)
 - o chemistry (e.g., chemical reactions)
 - o neuroscience
 - o physics
 - o economics (e.g., markets computing equilibrium prices)
 - o social networks

All of this seems to be captured by
one computational model
(spoiler alert: TURING MACHINE)

Interesting question: What is computable
in this model?

Answer: Not everything
Exist perfectly well-defined
mathematical functions that
cannot be computed

Even more interesting question: What is
EFFICIENTLY computable?

What this course is about
As we will see, many fascinating
open problems (and some answers)

SOME QUESTIONS

- ① Is solving a problem harder than checking a solution (avoiding exhaustive search)?
- ② Can randomness help solve more problems faster? (If computers can flip fair coins)
- ③ Can any time-efficient algorithm be optimized to use tiny amounts of memory?
- ④ Can every sequential algorithm be efficiently parallelized? (Time T , P processors \Rightarrow speed-up T/P ?)
- ⑤ Can hard problems become significantly easier if we are willing to accept suboptimal, approximate solutions?
- ⑥ Can quantum mechanics be used to build faster computers?
- ⑦ Can computationally hard problems be useful to solve computational problems efficiently?
- ⑧ Can complicated mathematical proofs be verified by quick, random sampling of just a few bits?
- ⑨ Is it possible to provide proofs that reveal absolutely nothing other than the truth of the statement proven?
- ⑩ Is it possible to compute solutions to problems so large that we don't have time to read them or space to store them?

SOME ATTEMPTS AT ANSWERS

- ① Probably yes, but (biggest) open problem in TCS (and all of math) P vs NP
Assume "yes" as axiom? (cf. gravity)
- ② Non-interactively, probably no, — don't know for sure, but strange things would happen otherwise P vs BPP
Interactive computation: Randomness very helpful, as we will see
- ③ Probably no, but big open problem P vs L
- ④ Again, probably no, but big open problem P vs NC
- ⑤ Assuming $P \neq NP$:
 - Sometimes a lot easier
 - Sometimes not at all
- ⑥ Theoretical models seem to predict yes
Not clear if physically realizable
(And recent "quantum supremacy" experiments debated)
- ⑦ Definitely YES! Almost all of modern cryptography builds on this
(and also has connections to ②)

- ⑧ Amazingly, YES! Very connected to ⑤
- ⑨ Again amazingly, YES! Connections to modern cryptography
- ⑩ Yes, sometimes. Buzzwords:
- sublinear-time algorithms
 - streaming algorithms

For many of these questions, there is scientific consensus...

But for most of them we don't know how to prove what we believe

... And we could be wrong
(This has happened before)

Fascinating and exciting questions with implications far outside of computer science

Two approaches:

- (A) Concrete, unconditional results for bounded computational models
"low-level" computational complexity
- (B) Connections between computational problems and notions - e.g., assume "yes" to ①, i.e., $P \neq NP$, and see what follows
"high-level" computational complexity

In order to do rigorous study of computation, need formal mathematical setting defining

- what is a computer?
- what does it mean to solve a problem efficiently?
- what is a computational problem?

COMPUTATIONAL PROBLEM

Represent objects as strings in some alphabet Σ

Think of Σ as


- 0 or 1
- a-z A-Z 0-9 _ - () [] < >
- ASCII or UTF-8

Choose Σ as convenient. Does not matter as long as size finite $|\Sigma| < \infty$

STRING $s \in \Sigma^*$ sequence of 0 or more characters from Σ

Some example problems:

- ① Given graph G , vertices s, t , find path in G from s to t (or report none exist)



- ② Given integer N , find prime factors $N = 25957$

- ③ Given propositional logic formula, find satisfying assignment (or report none exist)

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

- ④ Given integers A_1, \dots, A_n and target T , find subset $S \subseteq \{1, \dots, n\}$ such that $\sum_{i \in S} A_i = T$ (or report none exist) $\{2, 3, 5, 7\}$, target 11

Encoding issues

- 1) We can clearly encode all these problems as strings in some way
- 2) From now on, assume we've agreed on some reasonable encoding
 - Details of course matter in practice
 - Not really important for our discussion
 - Avoid silly encodings, e.g., unary (5 encoded as "11111")

SEARCH PROBLEM

Given $x \in \Sigma^*$ encoding problem instance, compute solution $y \in \Sigma^*$

Ex some path; some satisfying assignment

FUNCTION PROBLEM

If every problem has a unique answer, this defines function $f: \Sigma^* \rightarrow \Sigma^*$

Ex prime factors sorted in increasing order; lexicographically smallest path from s to t viewed as string in Σ^*

We will simplify even more

DECISION PROBLEM

Consider functions $f: \Sigma^* \rightarrow \{0, 1\}$

Think of

0	= "no"
1	= "yes"

Examples


- ① Is there a path in G from s to t ?
- ② Is there a prime factor of N of size at most U ?
- ③ Is the given formula satisfiable?
- ④ Is there a subset of $\{A_1, \dots, A_n\}$ summing to the target T ?

Much cleaner to work with mathematically

Often does not matter - efficient algorithm for decision problem will give efficient algorithm for search problems

(Can you work this out for the problems above?)

DECISION PROBLEMS AND LANGUAGES



$f(x) = 1 = \text{yes}$ x is "yes instance" $f(x) = 0 = \text{no}$ x is "no instance"	Historical terminology Rich literature No way of avoiding it...
--	---

We say that an algorithm that computes f DECIDES L

What do we do with strings that are not valid encodings?

Ex If a graph is a list of the edges, is there a path from s to t in

$(s, a), (a, b), (c, b, a), (b, t)$?

"Syntax error" - Define $f(s) = \text{no}$ for strings s that are not valid encodings

EFFICIENT COMPUTATION

What can be computed within reasonable limits on resources such as

- computation time
- computation memory
- computation energy
- network communication

Most important measure for us: TIME

Measure how the number of operations needed scale with the size of the input

- ignore constants depending on low-level details
- look at asymptotic behaviour as input size grows

ASYMPTOTIC NOTATION

$T(n)$ is $O(g(n))$ if exist positive constants c, N
s.t. for $n \geq N$ it holds that $T(n) \leq c \cdot g(n)$

$T(n)$ is $\Omega(g(n))$ if exist positive constants c, N
s.t. for $n \geq N$ it holds that $T(n) \geq c \cdot g(n)$

$T(n)$ is $\Theta(g(n))$ if $T(n)$ is $O(g(n))$ and $\Omega(g(n))$

$T(n)$ is $o(g(n))$ if for all $\epsilon > 0$ exists N s.t.
for $n \geq N$ $T(n) \leq \epsilon \cdot g(n)$

$T(n)$ is $\omega(g(n))$ if for all $K > 0$ exists N s.t.
for $n \geq N$ $T(n) \geq K \cdot g(n)$

But what is our computational model?

The TURING MACHINE

- Seems to be able to simulate all physically realizable computational methods with little overhead
- But very simple, so mathematically nice to work with
- But so simple and stupid that they are very annoying to deal with
 - o CLRS completely skips details (for those of you who took AADS)
 - o We will follow Chapter 1 in Arora-Barak, but will be brief and informal

TURING MACHINE (TM)

- Fixed alphabet Σ (of finite size)
- Program Q (or "set of states" of TM)
- Tapes
 - o input tape, contains input, read-only (after input, special EOF/blank symbol)
 - o work/output tape (initialized to EOF/blank symbols)Tapes have a starting position but no end
- Read-write heads positioned on tapes (start in starting position)

(Can have more tapes if we want
Does not really matter — see Arora-Barak)

At each time step the TM:

- (a) reads symbols at current position on all tapes
- (b) writes symbol to work tapes
- (c) move tape head left or right one step (or stand still)
- (d) jump to new state $q' \in Q$

(b) - (d) depend on

- current state q
- ~~positions~~ ^{symbols} read on tapes in (a)

Special state q_{halt} - TM stops
Running time # steps before reaching q_{halt}

To compute a function:

- write value on output tape
- then move to q_{halt}

Ex TM that decides whether # 1s in binary string odd

ALWAYS

~~q_{start}~~ : Read symbol s on input tape, move input head right

q_{start} : If $s = 0$ go to q_{start}

If $s = 1$ go to q_{odd}

If $s = EOF$ write 0 on output and go to q_{halt}

q_{odd} : If $s = 0$ go to q_{odd}

If $s = 1$ go to q_{start}

If $s = EOF$ write 1 on output and go to q_{halt}

FACTS ABOUT TURING MACHINES

- ① Expressive: Can do "normal things" efficiently, so we can allow ourselves to write pseudocode (including subroutine calls)
- ② Robust to tweaks
 - change of alphabet
 - adding more work tapes
- ③ Description of TM can be written as string and given as input to other TM
- ④ There is a UNIVERSAL TURING MACHINE that can simulate any other Turing machine M given its string representation. This is EFFICIENT — if original TM M runs in time T , then simulation runs in time $O(T \log T)$

All of this needs proving, of course, but we do not have time or patience for this now... So take it on faith.

From this it follows that there are UNCOMPUTABLE / UNDECIDABLE PROBLEMS!

Perfectly well-defined mathematical functions that cannot be correctly computed by any algorithm

HALTING PROBLEM

Fix alphabet Σ

Fix some way of encoding Turing machines

Consider the language

$$\text{HALT} = \{ \langle M, x \rangle \mid \text{Turing machine } M \text{ halts on input } x \}$$

(As before, if M or x is not valid, then $\langle M, x \rangle$ is a no instance)

THEOREM The language HALT is not decidable by any Turing machine

Proof By contradiction.

Suppose that H is a TM that decides HALT. We can construct another TM H' that simulates H as a subroutine

Then we can feed H' to H with a suitable input

This is all legitimate, so if we reach a contradiction, then H cannot exist

TM H' with input M

```
if  $H(M, M) = \text{yes}$  then
  while true // infinite loop
  endwhile
else //  $H(M, M) = \text{no}$ 
  halt
```

What does H' do when given input $M = H'$?

- a) H' halts on $H' \Rightarrow$
 $H(H', H') = \text{yes} \Rightarrow$
 H' gets stuck in infinite loop
- b) H' does not halt on $H' \Rightarrow$
 $H(H', H') = \text{no} \Rightarrow$
 H' halts

Contradiction! Hence H does not exist \square

Another example:

Given set of polynomials with integer coefficients, do these equations have a common integral solution?

Undecidable!

Does not mean that no instance of these problems can ever be solved. But no algorithm can:

- always terminate
- always give correct answer