

Theory of Computation Lecture Notes

Prof. Yuh-Dauh Lyuu
Dept. Computer Science & Information Engineering
and
Department of Finance
National Taiwan University

Class Information

- Papadimitriou. *Computational Complexity*. 2nd printing. Addison-Wesley. 1995.
 - We more or less follow the topics of the book.
 - More “advanced” materials may be added.
- You may want to review discrete mathematics.

Class Information (concluded)

- More information and lecture notes can be found at
`www.csie.ntu.edu.tw/~lyuu/complexity.html`
 - Homeworks, exams, solutions and teaching assistants will be announced there.
- Please ask many questions in class.
 - The best way for me to remember you in a large class.^a

^a “[A] science concentrator [...] said that in his eighth semester of [Harvard] college, there was not a single science professor who could identify him by name.” (*New York Times*, September 3, 2003.)

Grading

- Homeworks.
 - Do not copy others' homeworks.
 - Do not give your homeworks for others to copy.
- Two to three exams.
- You must show up for the exams in person.
- If you cannot make it to an exam, please email me or a TA beforehand (unless there is a legitimate reason).
- Missing the final exam will earn a “fail” grade.

Problems and Algorithms

I have never done anything “useful.”
— Godfrey Harold Hardy (1877–1947),
A Mathematician’s Apology (1940)

What This Course Is All About

Computation: What is computation?

Computability: What can be computed?

- There are *well-defined* problems that cannot be computed.
- In fact, “most” problems cannot be computed.

What This Course Is All About (concluded)

Complexity: What is a computable problem's inherent complexity?

- Some computable problems require at least exponential time and/or space.
 - They are said to be **intractable**.
- Some practical problems require superpolynomial resources unless certain conjectures are disproved.
- Other resources besides time?
 - Space, circuit size, program size, number of random bits, VLSI layout area, etc.

Tractability and Intractability

- Polynomial in terms of the input size n defines tractability.
 - $n, n \log n, n^2, n^{90}$.
 - Time, space, and circuit size.
- It results in a fruitful and practical theory of complexity.
- Few practical, tractable problems require a large degree.
- Exponential-time or superpolynomial-time algorithms are usually impractical.
 - $n^{\log n}, 2^{\sqrt{n}}$,^a $2^n, n! \sim \sqrt{2\pi n} (n/e)^n$.

^aSize of depth-3 circuits to compute the majority function (Wolfowitz (2006)).

Growth of Factorials

| n | $n!$ | n | $n!$ |
|-----|-------|-----|--------------------|
| 1 | 1 | 9 | 362,880 |
| 2 | 2 | 10 | 3,628,800 |
| 3 | 6 | 11 | 39,916,800 |
| 4 | 24 | 12 | 479,001,600 |
| 5 | 120 | 13 | 6,227,020,800 |
| 6 | 720 | 14 | 87,178,291,200 |
| 7 | 5040 | 15 | 1,307,674,368,000 |
| 8 | 40320 | 16 | 20,922,789,888,000 |

Growth of *E. Coli*^a

- Under ideal conditions, *E. Coli* bacteria divide every 20 minutes.
- In two days, a single *E. Coli* bacterium would become 2^{144} bacteria.
- They would weigh 2,664 times the Earth!

^aNick Lane, *Power, Sex, Suicide: Mitochondria and the Meaning of Life* (2005).

Turing Machines

Alan Turing (1912–1954)



What Is Computation?

- That can be coded in an **algorithm**.^a
- An algorithm is a detailed step-by-step method for solving a problem.
 - The Euclidean algorithm for the greatest common divisor is an algorithm.
 - “Let s be the least upper bound of compact set A ” is not an algorithm.
 - “Let s be a smallest element of a finite-sized array” can be solved by an algorithm.

^aMuhammad ibn Mūsā Al-Khwārizmī (780–850).

Turing Machines^a

- A Turing machine (TM) is a quadruple $M = (K, \Sigma, \delta, s)$.
- K is a finite set of **states**.
- $s \in K$ is the **initial state**.
- Σ is a finite set of **symbols** (disjoint from K).
 - Σ includes \sqcup (blank) and \triangleright (first symbol).
- $\delta : K \times \Sigma \rightarrow (K \cup \{h, \text{“yes”}, \text{“no”}\}) \times \Sigma \times \{\leftarrow, \rightarrow, -\}$ is a **transition function**.
 - \leftarrow (left), \rightarrow (right), and $-$ (stay) signify cursor movements.

^aTuring (1936).

A TM Schema

δ

▷1000110000111001110001110□□□□

More about δ

- The program has the **halting state** (h), the **accepting state** (“yes”), and the **rejecting state** (“no”).
- Given current state $q \in K$ and current symbol $\sigma \in \Sigma$,

$$\delta(q, \sigma) = (p, \rho, D).$$

- It specifies the next state p , the symbol ρ to be written over σ , and the direction D the cursor will move *afterwards*.
- We require $\delta(q, \triangleright) = (p, \triangleright, \rightarrow)$ so that the cursor never falls off the left end of the string.

More about δ (concluded)

- Think of the program as lines of codes:

$$\delta(q_1, \sigma_1) = (p_1, \rho_1, D_1),$$

$$\delta(q_2, \sigma_2) = (p_2, \rho_2, D_2),$$

\vdots

$$\delta(q_n, \sigma_n) = (p_n, \rho_n, D_n).$$

- Given the state q and the symbol under the cursor σ , the machine finds the line that matches (q, σ) .
- This line of code is then executed.

The Operations of TMs

- Initially the state is s .
- The string on the tape is initialized to a \triangleright , followed by a *finite-length* string $x \in (\Sigma - \{\sqcup\})^*$.
- x is the **input** of the TM.
 - The input must not contain \sqcup s (why?)!
- The cursor is pointing to the first symbol, always a \triangleright .
- The TM takes each step according to δ .
- The cursor may overwrite \sqcup to make the string longer during the computation.

“Physical” Interpretations

- The tape: computer memory and registers.
- δ : program.
- K : instruction numbers.
- s : “main()” in C.
- Σ : **alphabet** much like the ASCII code.

Program Count

- A program has a *finite* size.
- Recall that
$$\delta : K \times \Sigma \rightarrow (K \cup \{h, \text{“yes”}, \text{“no”}\}) \times \Sigma \times \{\leftarrow, \rightarrow, -\}.$$
- So $|K| \times |\Sigma|$ “lines” suffice to specify a program, one line per pair from $K \times \Sigma$ ($|x|$ denotes the length of x).
- Given K and Σ , there are

$$((|K| + 3) \times |\Sigma| \times 3)^{|K| \times |\Sigma|}$$

possible δ 's (see next page).

- This is a constant—albeit large.

K Σ

| | | |
|--|--|--|
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |

$(|K| + 3) \times |\Sigma| \times 3$
possibilities

The Halting of a TM

- A TM M may **halt** in three cases.
 - “yes”: M **accepts** its input x , and $M(x) = \text{“yes”}$.
 - “no”: M **rejects** its input x , and $M(x) = \text{“no”}$.
 - h : $M(x) = y$ means the string (tape) consists of a \triangleright , followed by a finite string y , whose last symbol is not \sqcup , followed by a string of \sqcup s.
 - y is the **output** of the computation.
 - y may be empty denoted by ϵ .
- If M never halts on x , then write $M(x) = \nearrow$.

Why TMs?

- Because of the simplicity of the TM, the model has the advantage when it comes to complexity issues.
- One can conceivably develop a complexity theory based on something similar to C++ or Java, say.
- But the added complexity does not yield additional fundamental insights.
- We will describe TMs in pseudocode.

Remarks

- A problem is computable if there is a TM that halts with the correct answer.
 - If a TM (i.e., program) does not always halt, it does not solve the problem, assuming the problem is computable.^a

^aContributed by Ms. Amy Liu (J94922016) on May 15, 2006. Control-C is not a legitimate way to halt a program.

Remarks (concluded)

- Any computation model must be physically realizable.
 - A model that requires nearly infinite precision to build is not physically realizable.
 - For example, if the TM required a voltage of exactly 100 to work, it would not be considered a successful model for computation.
- Although a TM requires a tape of infinite length, which is not realizable, it is not a major conceptual problem.^a

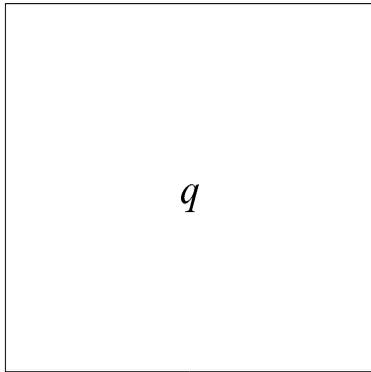
^aThanks to a lively discussion on September 20, 2006.

The Concept of Configuration

- A **configuration** is a complete description of the current state of the computation.
- The specification of a configuration is sufficient for the computation to continue as if it had not been stopped.
 - What does your PC save before it sleeps?
 - Enough for it to resume work later.
- Similar to the concept of state in Markov process.

Configurations (concluded)

- A configuration is a triple (q, w, u) :
 - $q \in K$.
 - $w \in \Sigma^*$ is the string to the left of the cursor (inclusive).
 - $u \in \Sigma^*$ is the string to the right of the cursor.
- Note that (w, u) describes both the string and the cursor position.



▷1000110000111001110001110□□□□

- $w = \triangleright 1000110000$.
- $u = 111001110001110$.

Yielding

- Fix a TM M .
- Configuration (q, w, u) **yields** configuration (q', w', u') in one step,

$$(q, w, u) \xrightarrow{M} (q', w', u'),$$

if a step of M from configuration (q, w, u) results in configuration (q', w', u') .

- $(q, w, u) \xrightarrow{M^k} (q', w', u')$: Configuration (q, w, u) yields configuration (q', w', u') in $k \in \mathbb{N}$ steps.
- $(q, w, u) \xrightarrow{M^*} (q', w', u')$: Configuration (q, w, u) yields configuration (q', w', u') .

Example: How To Insert a Symbol

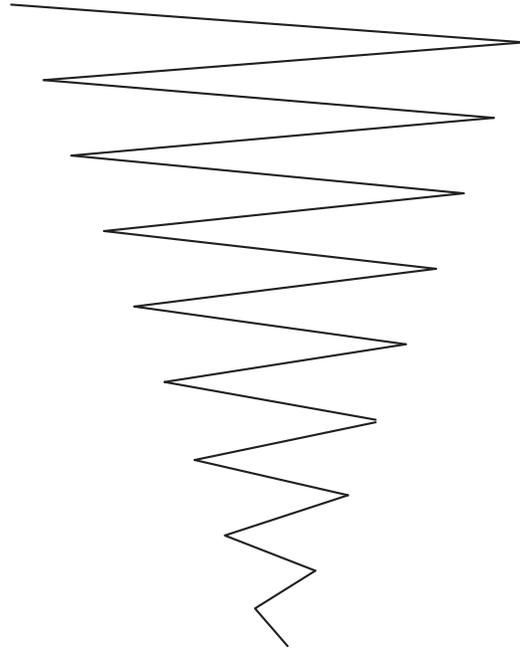
- We want to compute $f(x) = ax$.
 - The TM moves the last symbol of x to the right by one position.
 - It then moves the next to last symbol to the right, and so on.
 - The TM finally writes a in the first position.
- The total number of steps is $O(n)$, where n is the length of x .

Palindromes

- A string is a **palindrome** if it reads the same forwards and backwards (e.g., 001100).
- A TM program can be written to recognize palindromes:
 - It matches the first character with the last character.
 - It matches the second character with the next to last character, etc. (see next page).
 - “yes” for palindromes and “no” for nonpalindromes.
- This program takes $O(n^2)$ steps.
- We cannot do better.^a

^aHennie (1965).

1000110000000100111



Comments on Lower-Bound Proofs

- They are usually difficult.
 - Worthy of a Ph.D. degree.
- An algorithm whose running time matches a lower bound means it is optimal.
 - The simple $O(n^2)$ algorithm for PALINDROME is optimal.
- This happens rarely and is model dependent.
 - Searching, sorting, PALINDROME, matrix-vector multiplication, etc.

Decidability and Recursive Languages

- Let $L \subseteq (\Sigma - \{\square\})^*$ be a **language**, i.e., a set of strings of symbols with a *finite* length.
 - For example, $\{0, 01, 10, 210, 1010, \dots\}$.
- Let M be a TM such that for any string x :
 - If $x \in L$, then $M(x) = \text{“yes.”}$
 - If $x \notin L$, then $M(x) = \text{“no.”}$
- We say M **decides** L .
- If L is decided by some TM, then L is **recursive**.
 - Palindromes over $\{0, 1\}^*$ are recursive.

Acceptability and Recursively Enumerable Languages

- Let $L \subseteq (\Sigma - \{\square\})^*$ be a language.
- Let M be a TM such that for any string x :
 - If $x \in L$, then $M(x) = \text{“yes.”}$
 - If $x \notin L$, then $M(x) = \nearrow$.
- We say M **accepts** L .

Acceptability and Recursively Enumerable Languages (concluded)

- If L is accepted by some TM, then L is called a **recursively enumerable language**.^a
 - A recursively enumerable language can be generated by a TM, thus the name.
 - That is, there is an algorithm such that for every $x \in L$, it will be printed out eventually.

^aPost (1944).

Emil Post (1897–1954)



Recursive and Recursively Enumerable Languages

Proposition 1 *If L is recursive, then it is recursively enumerable.*

- We need to design a TM that accepts L .
- Let TM M decide L .
- We next modify M 's program to obtain M' that accepts L .
- M' is identical to M except that when M is about to halt with a “no” state, M' goes into an infinite loop.
- M' accepts L .