ICS 254 - Discrete Structures II

A Lecture Note on Theory of Computation (LNTC 4.3)

Sultan Almuhammadi

September 2018

Introduction

Each branch of Computer Sciences has its own objects of study, and it is time we turned more specifically to those of the theory of computation. The *theory of computation* is the mathematical study of computing machines and their capabilities. We must therefore develop a model for the data that computers manipulate. This is a big area in the theoretical computer science. Unfortunately, we are not going to discuss lots of interesting problems and theorems here since they are out of the scope of this course. However, we will just touch few issues related to the area of modeling computation to get the feeling of studying this topic. We adopt the mathematically expedient choice of representing data by strings of symbols. The best way to start this topic is by introducing the notion of alphabet and languages.

The format of this lecture note is as follows. All concepts and issues (like definitions, theorems, notations, and special notes) are sequentially numbered in brackets [like these] so that the student may need to study and understand them in sequence. Some exercises are given at the end of each section to ensure understanding of the material.

§1 Alphabet and Languages

There are many languages in real life. A language consists of words. Each word consists of letters which are taken from a set of symbols called the alphabet. For Example, the alphabet of English consists of 26 symbols (letters). A sequence of letters forms a string that can be a word in the language. For example, the strings *apple* and *book* are two words in English, but the string *aaabbbccc* is not, even though they are strings on the same alphabet. The binary language (or code) is defined on the alphabet {0, 1}. For any given machine language, some strings on {0, 1} are recognized as valid input (words) and some are not.

In this section, we will use the same analogy to define languages in order to study the computing machines and their capabilities. Let us start a formal approach by introducing the following definitions.

[1] An *alphabet* is a finite set of symbols, denoted by Σ.

Examples: $\Sigma = \{a, b, c\}$ // has 3 symbols
$\Sigma = \{a, <, 3, \#, @, 6\}$ // any object can be a symbol
$\Sigma = \{0, 1\}$ // has 2 symbols, for binary strings
$\Sigma = \{a, b, c, \ldots, z\}$ // has 26 symbols, for English
$\Sigma = \emptyset$ // could be empty, has zero symbols

[2] A *string on* Σ is a finite sequence of symbols from Σ.
Examples: If $\Sigma = \{a, b, c\}$, then $abbbc$ is a string on Σ, and aa is another string.
If $\Sigma = \{0, 1\}$, then 001 and all binary codes are strings on Σ.

[3] The *length* of a string s, denoted by $|s|$, is the number of symbols in s.
Examples: The length of $abbc$ is 4, and the length of 00011 is 5.
$|a| = 1$.
If $s = aaa$, then $|s| = 3$.

[4] The *empty string* (or *null*) is a string of length zero. Thus, it has no symbols. The empty string is defined over every alphabet and it is denoted by λ (Lambda). Therefore, by definition, $|\lambda| = 0$.

[5] **Note:** Although any object can be a symbol, we usually use common characters like letters and numerals as symbols. The alphabets $\{a, b\}$ is commonly used instead of the binary alphabet $\{0, 1\}$. Thus, if Σ is not given explicitly and cannot be determined from the context, then $\Sigma = \{a, b\}$ is assumed by default.

[6] *Concatenation* of strings: let x and y be two strings on the same alphabet. The concatenation of x and y forms a third string denoted by $x \cdot y$ or simply: xy
Example: Let $x = aaab$ and $y = ba$, then $xy = aaabba$
Note: for any string x over some Σ, $x \cdot \lambda = x = \lambda \cdot x$

[7] **Notations:** s^n denotes the *repeated concatenation* of s, which is the string obtained by concatenating s to itself n times. So, $s^2 = s \cdot s$ and $s^3 = s \cdot s \cdot s$. In other words, s^n can be define recursively as follows:

$$s^n = s^{n-1} \cdot s \quad \text{and} \quad s^0 = \lambda$$

Similarly, for any symbol a in some alphabet, a^n denotes the string made of the symbol a repeated n times.
Example: Let $s = abba$, then $s^2 = abbaabba$, and $s^3 = abbaabbaabba$.
The string $a^3 = aaa$, and $b^4 = bbbb$.

[8] In general, the *repeated concatenation* can be applied on any set of symbols, A, and it means the set of all strings made of concatenating n symbols from A possibly with repetition. Thus, $A^n = \{s \mid s \text{ is a string on } A \text{ of length } n\}$.
Examples: $\{a, b\}^3 = \{aaa, aab, aba, abb, baa, bab, bba, bbb\}$
$\{a, b, c\}^2 = \{aa, ab, ac, ba, bb, bc, ca, cb, cc\}$
Σ^n is the set of all strings of length n on Σ.

[9] *Kleene star* operation: For any symbol $a \in \Sigma$, a^* is the set of all strings made of the symbol a repeated zero or more times.
Thus, $a^* = \{s \mid s = a^n \text{ for all } n \geq 0\} = \{\lambda, a, aa, aaa, aaaa, \dots\}$

[10] In general, the *Kleene start* operator can be applied on any set of symbols, A, and it means the set of all strings made of concatenating zero or more symbols from A.
Thus, $A^* = \{s \mid s \text{ is a string on } A\}$
Examples: Σ^* is the set of all strings on Σ.
$\{a, c\}^* = \{\lambda, a, c, aa, ac, ca, cc, aaa, aac, aca, caa, acc, ccc, aaaa, \dots\}$

[11] *Substring*: a string y is a *substring* of w if and only if there are two strings $x, z \in \Sigma^*$ such that $w = xyz$. (Notice that x and/or z can be null). If x is null, then the substring y is called a *prefix* of w; and if z is null, then y is a *suffix* of w.
Example: aa and bb are two substring of $abbaa$, aa is also a suffix. λ is a substring of every string.

[12] The *reverse* of a string s, denoted by s^R, is the same string spelled backwards. For example, $(abac)^R = caba$. **Note:** the upper case R is reserved for the reversal operator. (Do not confuse with the notation of repeated concatenation s^n in [7])

[13] **Theorem 1.1:** For any two strings $x, y \in \Sigma^*$, $(x \cdot y)^R = y^R \cdot x^R$

[14] A *language*, L, is a set of strings on an alphabet Σ. Thus, $L \subseteq \Sigma^*$
Examples: $L_1 = \{a, aaa, abba\}$ is a language defined on $\Sigma = \{a, b\}$.
$L_2 = \{aaa\}$ is another example of a language on $\{a, b\}$, not using b though.
$L_3 = a^*$ is also a language on $\{a, b\}$ not using b.
$L_4 = \Sigma^*$ is a language that has all strings.
$L_5 = \emptyset$ is the empty language.

[15] *Word*: the strings in a language are also called *words*.
Examples: In the above example, L_1 has three words and L_2 has one word.
The word aaa is a common word in the languages, L_1, L_2 and L_3.
L_3 has infinitely many words.
L_5 has zero words.

[16] The *string-function*: The string w can be considered as a function defined as follows:

$$w : \{1, 2, \ldots, |w|\} \to \Sigma, \text{ where } w(i) \text{ is the } i^{th} \text{ symbol in } w$$

Example: For $w = abbc$, we have $w(1) = a, w(2) = w(3) = b$ and $w(4) = c$.

[17] The *language of palindromes*: $L_p = \{w \mid w = w^R\}$ is the language of all palindrome words, such as: $aba, aa, abbbba, a, b$ and λ. L_p is infinite, i.e. it has infinitely many words.

[18] **Note:** most languages of interest are infinite, so that listing all the strings is not possible. Thus we can define infinite languages, as we define sets in general, by the following schemes:

1. $L = \{w \in \Sigma^* \mid w \text{ has some property}\}$ (set definition)
2. $L = L_1 \cup L_2$ (union of two languages)
3. $L = L_1 \cap L_2$ (intersection of two languages)
4. by string operations, like concatenation and Kleene star. See [19] and [20].

We will study more ways to define infinite languages in the following sections.

[19] *Language concatenation*: $L_1 \cdot L_2 = \{w \mid w = x \cdot y \text{ where } x \in L_1 \text{ and } y \in L_2\}$
Notation: L^n is the set of all strings made of concatenating exactly n strings from L.
Examples: $\{ab, bb\} \cdot \{a, b, \lambda\} = \{aba, abb, ab, bba, bbb, bb\}$
$\{ab, b\}^3 = \{ababab, ababb, abbab, babab, abbb, babb, bbab, bbb\}$
$\{a, b, \lambda\}^5 = \{s \mid s \text{ is a string on } \{a, b\} \text{ and } |s| \leq 5\}$

[20] The *closure operation* (or Kleene star): L^* is the set of all strings obtained by concatenating zero or more strings from L. (The concatenation of *zero* strings is λ).
Example: $\{ab, bb\}^* = \{\lambda, ab, bb, abab, abbb, bbbb, bbab, ababab, \ldots\}$.

[21] **Notation:** L^+ is the set of all strings obtained by concatenating one or more strings from L . thus, $L^+ = L \cdot L^*$. (note: $\lambda \notin L^+$ unless $\lambda \in L$).

[22] **Theorem 1.2:** For any two language A and B , if $A \subseteq B$ then $A^* \subseteq B^*$.

Example 1.1: Let $A = \{a, ab, b, bba\}$ be a language defined over $\Sigma = \{a, b\}$. Find A^*.
Solution: Notice that $\Sigma = \{a, b\} \subseteq A$. Therefore $\Sigma^* = \{a, b\}^* \subseteq A^*$ (by Theorem 1.2). But Σ^* is the set of all strings on Σ, which implies, $A^* \subseteq \Sigma^*$. Hence, $A^* = \Sigma^*$.

Exercises:

1. Let $\Sigma = \{0, 1\}$, which of the followings are strings on Σ?
   (a) 0
   (b) 012
   (c) 0^n for some $n \geq 1$
   (d) $0^n \cdot 1^n$ for some $n \geq 1$
   (e) $0101010101\ldots$ (endless string of alternating 0's and 1's)
   (f) λ
2. Let A and B be two arbitrary languages on $\Sigma = \{a, b\}$. True or False:
   (a) $A \subseteq \Sigma$
   (b) $A^* = \Sigma^*$
   (c) $(A \cup B)^* \subseteq \Sigma^*$
   (d) $A^* \cup B^* = (A \cup B)^*$
   (e) $A^* \cap B^* = (A \cap B)^*$
3. True or false? And justify your answer.
   (a) $0^* \cup 1^* = \{0, 1\}^*$
   (b) $0^* \cap 1^* = \emptyset$
4. Prove by mathematical induction that $(x^R)^n = (x^n)^R$ for all $n \geq 1$.
5. Let $\Sigma = \{a, b\}$. For each of the following languages on Σ, find, if possible, two strings $x, y \in \Sigma^*$ such that x is a word in the given language, and y is not.
   (a) $\{\lambda, a, b, ab, ba, aa, bb, aaa, aab, aba, abb, baa, bab, bba, bbb\}$
   (b) $\{ab, ba, aa, bb\}^*$
   (c) $\{a, bb, ba\}^*$
   (d) $(a^* \cdot b^*) \cup (b^* \cdot a^*)$
6. Let $\Sigma = \emptyset$. How many different languages can be defined on Σ?

§2 Regular Expressions and Language Representation

A central issue in the theory of computation is the representation of languages by finite specifications. Naturally, any finite language is amenable to finite representation by exhaustive enumeration of all the strings in the language. The issue becomes challenging only when infinite languages are considered.

Let us be more precise about the notion of *finite representation* of a language. The first point to be made is that any such representation must itself be a string, a finite sequence of symbols over some alphabet. Second, we certainly want different languages to have different representations, otherwise the term representation could hardly be considered appropriate. We will see in the following sections that these two requirements already imply that the possibilities for finite representation are severely limited.

The first representation method we will study is called *regular expression* (RE). Since any language, L, by definition is nothing but a (possibly infinite) set of strings on some Σ, we would like to know which strings are there in L and which are not. We need precise expressions to define languages besides the ordinary set definitions and operations. So, let us start by a formal definition of regular expressions, and then we see some examples of how such expressions can be used to define languages.

[1] A *regular expression* (RE) over an alphabet Σ is a string on the alphabet Δ, where $\Delta = \{*, +, (,), \emptyset\} \cup \Sigma$. Thus, Δ contains all the symbols in Σ plus the following symbols: $*, +, (,)$, and $\emptyset$, such that the following rules hold:

1. the elements in Σ and the empty set $\emptyset$ are the basic regular expressions,
2. if α is a regular expression, then so is α^* (star),
3. if α and β are regular expressions, then so is $\alpha\beta$ (concatenation), and
4. if α and β are regular expressions, then so is $\alpha + \beta$ (union).

Note: nothing else is a regular expression unless it follows from rules 1 through 4 with some parentheses if needed to control the precedence (see [2])

Example:
Let $\Sigma = \{a, b\}$, then we have 3 basic RE's (by rule 1), each represents a language, namely:

$\alpha_1 = a$ which represents the language $\{a\}$,
$\alpha_2 = b$ which represents the language $\{b\}$, and
$\alpha_3 = \emptyset$ which represents the language $\emptyset$.

By rule 2, we can build three more RE's:

a^* - for the language $\{\lambda, a, aa, aaa, aaaa, \dots\}$,
b^* - for the language $\{\lambda, b, bb, bbb, bbbb, \dots\}$, and
$\emptyset^*$ - for the language $\{\lambda\}$.

By rule 3, we can build more RE's, like:

aa - for the language $\{aa\}$
$(aa)^*$ - for all strings of even number of a's (and no b's),
$a^*(b^*)$ - for $\{\lambda, a, b, ab, aab, aabb, abb, abbb, aaaabbb \dots\}$, and so on.

By combining rules 2, 3 and 4, we can build RE's for more complicated languages, like:

$(a + b)^*$ - for Σ^*
$a(a + b)^*bb$ - all strings that start in a and end in bb

[2] Precedence rules: the star operation ($*$) comes first, then concatenation, and then union operation ($+$) comes third. Examples:

$ab^*a = a(b^*)a$
$a + ba = a + (ba)$
$(a + b^*)a^* = (a + (b^*))(a^*)$
$b^*(ab^*ab^*)^* = (b^*)(a(b^*)a(b^*))^*$

Example 2.1: Which of the following regular expressions represent a language that contains the word aa?

ab^*a
$a + ba$
$(a + b)a$
$b^*(ab^*ab^*)^*$

Solution:

$aa \in L(ab^*a)$	because it is $a(b^*)a$ and b^* can be λ.
$aa \notin L(a + ba)$	because concatenation has higher precedence.
$aa \in L((a + b)a)$	because of the parenthesis.
$aa \in L(b^*(ab^*ab^*)^*)$	because L is the set of all strings with even number of a's.

[3] **Notation:** in Rule 4 of [1], the union operation '+' means 'or' and some books use different notations like: '|' or '$\cup$' instead. We will use '+' in this course unless '+' is a symbol in Σ.

[4] **Notation:** if α is an RE, then $L(\alpha)$ denotes the language represented by α. Since every RE represents a language, we can say that two RE's α and β are equivalent (denoted by $\alpha \equiv \beta$) if and only if they represent the same language. Thus,

$$\alpha \equiv \beta \iff L(\alpha) = L(\beta)$$

For simplicity, we may use the equal sign '=' for both the equivalency of RE's and their languages. For example, we may write: $\alpha = L(\alpha)$ when we talk about the set represented by the RE. To avoid any possible confusion, we shall avoid writing $\alpha \neq \beta$ for different RE's unless we know that $L(\alpha) \neq L(\beta)$. For example, $\alpha = aa^*b$ and $\beta = a^*ab$ are clearly different expressions, but we will avoid writing $\alpha \neq \beta$ since they represent the same language.

[5] *Regular Languages*: L is *regular* if it can be represented by a regular expression. Thus, L is regular if there is a regular expression α such that $L = L(\alpha)$.
Examples: $L = \{aa, ab, bb, ba\}$ is regular for $\alpha = aa + ab + ba + bb$
$L = \{w \mid w \text{ has no } a\}$ is regular for $\alpha = b^*$
$L = \{w \mid w = a^n b^n \text{ for } n \geq 1\}$ is not regular. (See Exercise 4)
L_p (the language of palindromes) is not regular.

[6] **Theorem 2.1:** Regular languages are closed under union, intersection, complement, concatenation and star operations.

[7] **Theorem 2.2:** Any finite language is regular.

Exercises:

1. Let $\Sigma = \{a, b\}$. Give a regular expression that represents each of the following languages.

    (a) All strings that end in a

    (b) All strings that end in a or b

    (c) All strings that start in a or end in b

    (d) All strings of even length

    (e) All strings that contain the substring bbb

    (f) All strings that do not contain the substring bbb

2. Let $\Sigma = \{a, b\}$. For each of the following regular languages on Σ, find two strings $x, y \in \Sigma^*$ such that x is a word in the given language, and y is not. (If there is no such x or y, write "none")

    (a) $a + ab$

    (b) $(ab^* + b)^*$

    (c) $(a^*b^*)^*$

    (d) $a^*(ba + b + aa)^*$

    (e) $\emptyset^*$

    (f) $a^*(bb + aa)^*(ba^*)^*$

    (g) $b^*(a^*b^*ab)^*b^*a^*$

3. Prove Theorem 2.2.

4. Argue why a^*b^* does not represent $L = \{w \mid w = a^n b^n \text{ for } n \geq 1\}$.

§3 Finite-State Automata

Finite state automata (or finite automata for short) are good models for computers with an extremely limited amount of memory. What can a computer do with such a small memory? Many useful things!

In fact, we interact with such computers all the time, as they lie at the heart of various electromechanical devices. Finite automata and their probabilistic counterpart chains are useful tools when we are attempting to recognize patterns in data. These devices are used in speech processing and in optical character recognition. We will now take a close look at finite automata from a mathematical perspective.

We will develop a precise definition of a finite automaton, terminology for describing and manipulating finite automata, and theoretical results that describe their power and limitations. The theoretical development gives us a clearer understanding of what finite automata are and what they can and cannot do; and allows us to practice and become more comfortable with mathematical definitions, theorems, and proofs in a relatively simple setting.

[1] **State Diagrams:**
In beginning to describe the mathematical theory of finite automata, we do so in the abstract, without reference to any particular application. The following figure depicts a finite automaton called M_1.

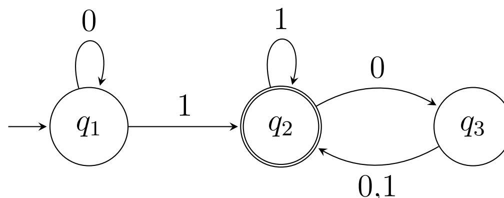

Figure 1: A finite automaton M_1 with three states

Figure 1 is called the *state diagram* of M_1. It has three states, labeled q_l, q_2, and q_3. The *start state*, q_1, is indicated by the arrow pointing at it from nowhere. The *accept state*, q_2 is the one with a double circle. The arrows going from one state to another are called *transitions*. Each transition is labeled with an input symbol.

When this automaton receives an input string such as 1101, it processes that string and produces an output. The output is either *accept or reject*. We will consider only this yes/no type of output to keep things simple.

The processing begins in M_1's start state. The automaton receives the symbols from the input string one by one from left to right. After reading each symbol, M_1 moves from one state to another along the transition that has that symbol as its label. When it reads the last symbol, M_1 produces its output. The output is either *accept* if M_1 ends up in an accept state after reading the last symbol, or *reject* otherwise.

Example 3.1: Does M_1 accept the string 1101?

Solution:
When we feed the input string 1101 to the machine M_1 in Figure 1, the processing proceeds as follows:

1. start in state q_1;
2. read 1, follow the transition from q_1 to q_2;
3. read 1, follow the transition from q_2 to q_2;
4. read 0, follow the transition from q_2 to q_3;
5. read 1, follow the transition from q_3 to q_2;
6. accept because M_1 is in an accept state q_2 at the end of the input.

Example 3.2: Test M_1 on the following input strings: 1, 01, 11, 0101010101, 100, 0100, 110000, 0101000000, 0, 10, and 101000.

Solution:
Experimenting with this machine on a variety of input strings reveals that it accepts the strings 1, 01, 11, and 0101010101. In fact, M_1 accepts any string that ends with 1, as it goes to its accept state q_2 whenever it reads the symbol 1. In addition, it accepts strings 100, 0100, 110000, and 0101000000, and any string that ends in an even number of zeros following the last 1. It rejects other strings, such as 0, 10, and 101000. Can you describe the language consisting of all strings that M_1 accepts? (Exercise 2.a)

[2]*Deterministic* machines: In the above example, M_1 is *deterministic*. This means that at every state q, when the machine M1 reads an input symbol, it knows (or determines) where to go next according to the input symbol and the current state. We call this type of machines Deterministic Finite Automata (DFA). We will see another type of machines (in Section 4) where the next state is not precisely determined. Hence, they are called Nondeterministic Finite Automata (NFA).

Now we define deterministic finite automata formally. Though state diagrams are easier to grasp intuitively, we need the formal definition too, for two specific reasons. First, a formal definition is precise. It resolves any uncertainties about what is allowed in a finite automaton. If you were uncertain about whether finite automata were allowed to have 0 accept states or whether they must have exactly one transition exiting every state for each possible input symbol, you could consult the formal definition and verify that the answer is yes in both cases. Second, a formal definition provides notation. Good notation helps you think and express your thoughts clearly.

[3] A deterministic finite automaton (DFA) is a quintuple $(Q, \Sigma, \delta, s, F)$, where

1. Q is a finite set called the *states*,
2. Σ is a finite set called the *alphabet*,
3. $\delta : Q \times \Sigma \to Q$ is the *transition function*,
4. $s \in Q$ is the *start state*, and
5. $F \subseteq Q$ is the set of *accept states* (also called *final states*).

The formal definition precisely describes what we mean by a finite automaton. For example, returning to the earlier question of whether 0 accept states is allowable, you can see that setting F to be the empty set $\emptyset$ yields 0 accept states, which is allowable. Furthermore, the transition function δ specifies exactly one next state for each possible combination of a state and an input symbol. That answers our other question affirmatively, showing that exactly one transition arrow exits every state for each possible input symbol.

Example 3.3: Describe M_1 formally.
Solution:
$M_1 = (Q, \Sigma, \delta, q_1, F)$ where

1. $Q = \{q_1, q_2, q_3\}$,
2. $\Sigma = \{0, 1\}$,
3. δ is the transaction table described as follows:

δ	0	1
q_1	q_1	q_2
q_2	q_3	q_2
q_3	q_2	q_2

4. q_1 is the start state, and
5. $F = \{q_2\}$

[4] The *language of a machine* M, denoted by $L(M)$, is the set of all strings accepted by M. If $L(M) = A$, we say that M recognizes or accepts the language A.

Example 3.4: Design a DFA that recognizes the language

$$L = \{w \mid w \text{ does not contain the substring } bbb\}.$$

Solution:
The idea is to accept after 0, 1 or 2 consecutive b's but reject if you ever reach three consecutive b's. Any other combinations of a's and b's is accepted. We can draw the DFA, M_2, as shown in Figure 2. The machine will go to state q_3 after reading three consecutive b's, but then it will never leave, and eventually rejected. Therefore q_3 here is called a *trap state.*

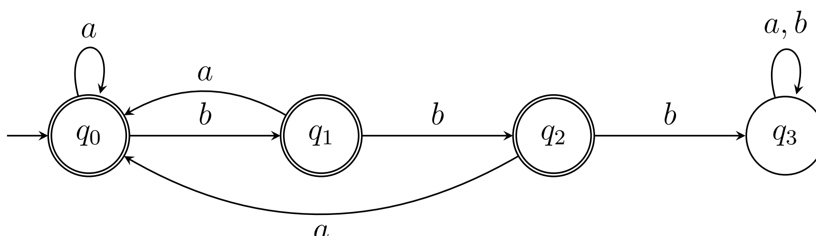

Figure 2: DFA M_2 for $L = \{w \mid w \text{ does not contain the substring } bbb\}$

$M_2 = (Q, \Sigma, \delta, q_0, \{q_0, q_1, q_2\})$, with $Q = \{q_0, q_1, q_2, q_3\}$, $\Sigma = \{a, b\}$ and the transition function δ given in the table below, does indeed accept the specified language.

δ	a	b
q_0	q_0	q_1
q_1	q_0	q_2
q_2	q_0	q_3
q_3	q_3	q_3

[5] *The extended transition function*: When the transition function δ is applied on a state $q_1 \in Q$ and a symbol $a \in \Sigma$, it yields the next state $q_2 \in Q$. This is denoted by $\delta(q_1, a) = q_0$. For example, in the DFA M_2, $\delta(q_1, a) = q_0$, and $\delta(q_1, b) = q_2$. It is convenient to introduce the *extended transition function* $\delta^*\colon Q \times \Sigma^* \to Q$. This function takes an initial state and a string as inputs, and it outputs the state that the automaton will end up in after reading the string. For example, $\delta^*(q_2, aaab) = q_1$ in M_2.

[6] **Theorem 3.1:** The languages recognized by DFA are regular. Thus, every DFA has an equivalent regular expression that represents the same language.

[7] **Theorem 3.2:** Every regular expression has an equivalent DFA. Thus, if α is a regular expression, then there is a DFA, M, such that $L(M) = L(\alpha)$.

Exercises:

1. Let $\Sigma = \{a, b\}$. Draw a DFA that recognizes each of the following languages.
   (a) Σ^*
   (b) $\emptyset$
   (c) $\{\lambda\}$
   (d) All strings of even length
   (e) All strings that start in a and end in bb
   (f) ab^*
   (g) All strings that contain the substring aba
   (h) All strings that do not contain the substring aba
   (i) All strings that have even number of a's and even number of b's
   (j) All strings that have even number of a's or odd number of b's
   (k) All strings that start and end in the same symbol
2. Write a regular expression for the language of the following DFA.
   (a) M_1 (shown in Figure 1)
   (b) M_2 (shown in Figure 2)
3. Using Theorems 3.1 and 3.2, prove Theorem 2.1.

§4 Nondeterministic Finite Automata

In this section we add a powerful feature to finite automata. This feature is called *nondeterminism*, and is essentially the ability to change states in a way that is only partially determined by the current state and input symbol. That is, we shall now permit several possible "next states" for a given combination of current state and input symbol. The automaton, as it reads the input string, may choose at each step to go into any one of these legal next states; the choice is not determined by anything in our model, and is therefore said to be *nondeterministic*. On the other hand, the choice is not wholly unlimited either; only those next states that are legal from a given state with a given input symbol can be chosen. A nondeterministic finite automaton (NFA) can be a much more convenient device to design than a deterministic finite automaton (DFA). We will show some examples.

[1] Two main differences between DFA and NFA:
1. NFA have multiple next states on the same input (zero or more next states)
2. NFA allow λ-transitions

Example 4.1: (multiple next states)
This example explains the first difference. Consider the language $L = (ab + aba)^*$, which is accepted by the deterministic finite automaton illustrated in Figure 3.

Even with the diagram, it takes a few moments to ascertain that a deterministic finite automaton is shown; one must check that there are exactly two arrows leaving each node, one labeled a and one labeled b. Some thought is needed to convince oneself that the language accepted by this fairly complex device is the simple language $L = (ab + aba)^*$.

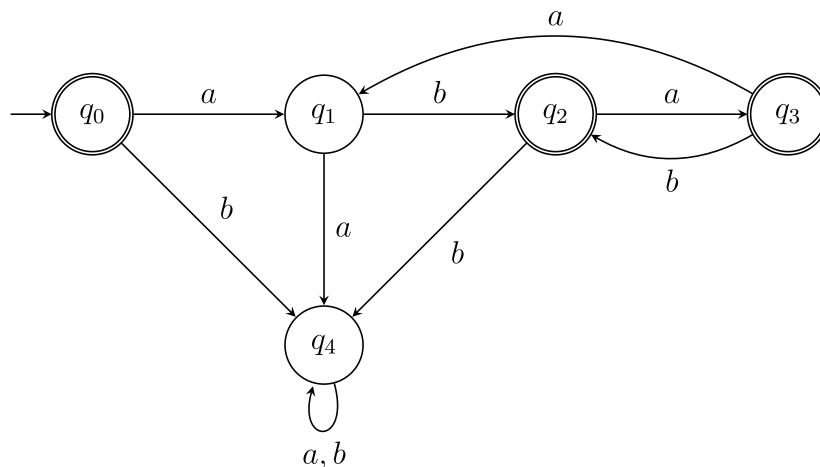

Figure 3: DFA for $L = (ab + aba)^*$

However, L is accepted by the simple nondeterministic device shown in Figure 4. When this device is in state q_1, and the input symbol is b, there are two possible next states, q_0 and q_2. Thus Figure 4 does not represent a deterministic finite automaton. Nevertheless, there is a natural way to interpret the diagram as a device accepting L. A string is accepted if there is some way to get from the initial state (q_0) to a final state (in this case, q_0) while following arrows labeled with the symbols of the string. For example, ab is accepted by going from q_0 to q_l to q_0; and aba is accepted by going from q_0 to q_1 to q_2 to q_0.

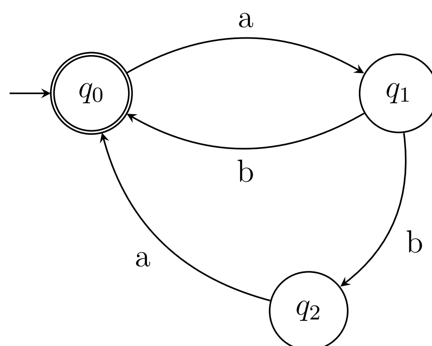

Figure 4: NFA for $L = (ab + aba)^*$

Of course, the device might guess wrong and go from q_0 to q_1 to q_0 to q_1 on input aba, winding up in a non-final state, but this does not matter, since there is *some way* of getting from the initial to the final state with this input.

On the other hand, the input abb is not accepted, since there is no way to get from q_0 back to q_0 while reading this string. Indeed, you notice that from q_0 there is no state to be entered when the input is b.

Example 4.2: (λ-transition)
Sometimes it is convenient to allow state diagrams in which arrows can be labeled either by symbols in Σ or by the empty string λ. We illustrate λ-transitions in this example.

Consider the device in Figure 5. This device accepts the same language $L = (ab + aba)^*$ as in Example 4.1 above. From q_2 this machine can return to q_0 either by reading an a or immediately, without consuming any input.

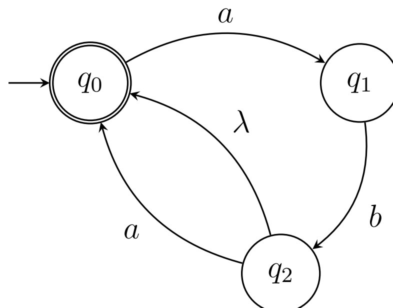

Figure 5: NFA with λ-transition for $L = (ab + aba)^*$

[2] A *nondeterministic finite automata* (NFA) is a quintuple $M = (Q, \Sigma, \delta, s, F)$, where

1. Q is a finite set of states,
2. Σ is a finite alphabet,
3. $\delta : Q \times (\Sigma \cup \{\lambda\}) \to 2^Q$ is the transition function. (2^Q is the power-set of Q)
4. $s \in Q$ is the *start state*, and
5. $F \subseteq Q$ is the set of *final states.*

Note: δ here can be described as a table of the function that maps the current state and an input symbol or λ to a subset of states belongs to the power-set of Q.

Example 4.3: Describe the NFA M_3, given in Figure 6, formally.

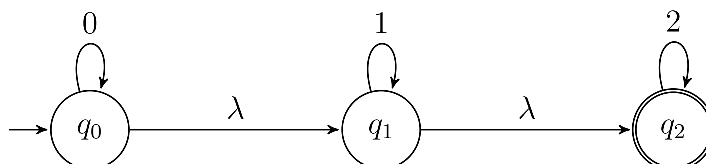

Figure 6: M_3, an example of NFA

Solution:
M_3 is defined formally as follows: $M_3 = (Q,\ \Sigma,\ \delta,\ q_0,\ \{q_2\})$, where $Q = \{q_0,\ q_1,\ q_2\}$, and $\Sigma = \{0, 1\}$. The table of the transitions δ can be described as:

δ	0	1	2	λ
q_0	$\{q_0\}$	$\emptyset$	$\emptyset$	$\{q_1\}$
q_1	$\emptyset$	$\{q_1\}$	$\emptyset$	$\{q_2\}$
q_2	$\emptyset$	$\emptyset$	$\{q2\}$	$\emptyset$

This gives a precise definition of the NFA and the language it recognize. For example, the word 002 belongs to $L(M_3)$ because it is accepted by M_3 through the path q_0, q_0, q_0, q_1, q_2, q_2 with arcs labeled 0, 0, λ, λ, 2.

Example 4.4: Describe the NFA M_4, given in Figure 7, formally.

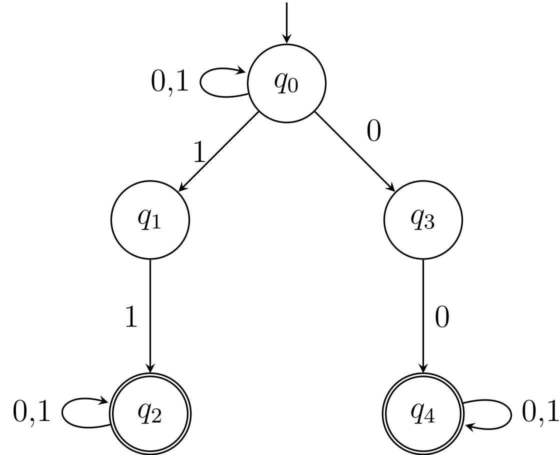

Figure 7: $M4$, another example of NFA

Solution:
$M_4 = (Q, \Sigma, \delta, q_0, \{q_2, q_4\})$, where $Q = \{q_0, q_1, q_2, q_3, q_4\}$, and $\Sigma = \{0, 1\}$. The table of the transitions δ can be described as follows:

δ	0	1	λ
q_0	$\{q_0, q_3\}$	$\{q_0, q_1\}$	$\emptyset$
q_1	$\emptyset$	$\{q_2\}$	$\emptyset$
q_2	$\{q_2\}$	$\{q_2\}$	$\emptyset$
q_3	$\{q_4\}$	$\emptyset$	$\emptyset$
q_4	$\{q_4\}$	$\{q_4\}$	$\emptyset$

Example 4.5: Does M_4 shown in Figure 7 accept the input 01001?

Solution:
We examine the propagation of states of the NFA under the input string 01001. Figure 8 shows that there is a way from q_0 to one of the two final states, q_4. Therefore, 01001 is accepted by M_4.

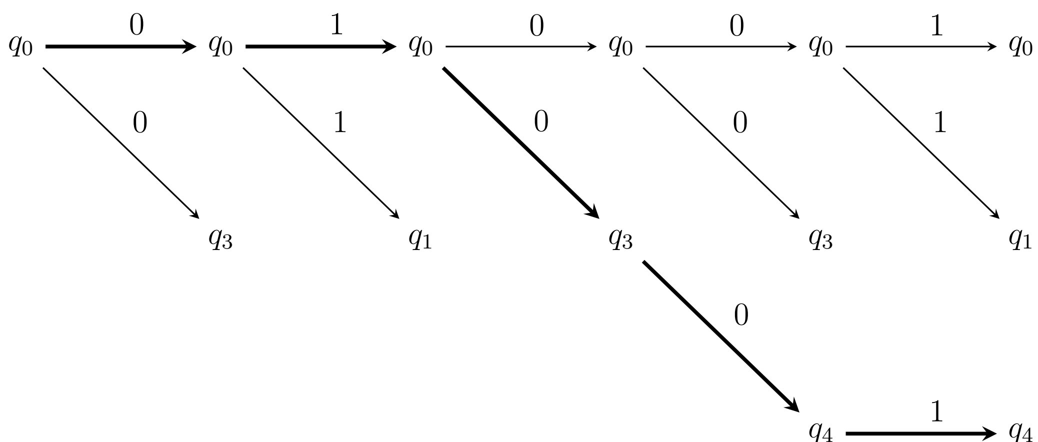

Figure 8: State propagation of M_4 under the input 01001

[3] **Note:** The extended transition function in an NFA, $\delta^*(q, x)$, outputs a set of all states that are reachable from q after reading the string x. In the above example, M_4, $d^*(q_0, 01001) = \{q_0, q_1, q_4\}$

[4] The language of an NFA, $M = (Q, S, d, q_0, F)$, can be formally defined as:

$$L(M) = \{w \mid \delta^*(q_0, w) \cap F \neq \emptyset\}$$

[5] **Note:** a deterministic finite automaton is just a special type of a nondeterministic finite automaton, where the transition function δ is just a function from $(Q \times S)$ to Q. Therefore, every DFA is an NFA by definition.

[6] **Theorem 4.1:** Every nondeterministic finite automaton is equivalent to some deterministic finite automaton. Thus, every NFA has a DFA that recognizes the same language.

[7] **Theorem 4.2:** The languages recognized by NFA are regular. (This follows from Theorem 3.1 and Theorem 4.1)

Exercises:

1. Let $\Sigma = \{a, b\}$. Using as few states as possible, draw an NFA that recognizes:
   (a) $\emptyset$
   (b) All strings that start in a and end in bb
   (c) All strings that start and end in the same symbol
   (d) $(aa + ab + ba)^*$
   (e) All strings of even length that start in a
   (f) All strings that contain the substring $abba$
   (g) All strings that do not contain the substring $abba$
2. Let $\Sigma = \{a, b, c\}$. Formally describe an NFA that recognizes:
   (a) All strings that do not contain a.
   (b) All strings that start in a and have an even number of c's.
   (c) All strings in which the number of a's, b's and c's are even.
3. Write a regular expression equivalent to the following NFA.
   (a) M_3 (shown in Figure 6).
   (b) M_4 (shown in Figure 7).
4. Consider the NFA M_3. Compute:
   (a) $\delta(q_0, 0)$
   (b) $\delta(q_0, 1)$
   (c) $\delta(q_1, 2)$
   (d) $\delta^*(q_1, 012)$
5. Consider the NFA M_4. Compute: $\delta^*(q_0, 1101)$