

*Life only avails, not the having lived. Power ceases in the instant of repose;
it resides in the moment of transition from a past to a new state,
in the shooting of the gulf, in the darting to an aim.*

— Ralph Waldo Emerson, “Self Reliance”, *Essays, First Series* (1841)

*O Marvelous! what new configuration will come next?
I am bewildered with multiplicity.*

— William Carlos Williams, “At Dawn” (1914)

3 Finite-State Machines

3.1 Intuition

Suppose we want to determine whether a given string $w[1..n]$ of bits represents a multiple of 5 in binary. After a bit of thought, you might realize that you can read the bits in w one at a time, from left to right, keeping track of the value modulo 5 of the prefix you have read so far.

```

MULTIPLEOF5( $w[1..n]$ ):
   $rem \leftarrow 0$ 
  for  $i \leftarrow 1$  to  $n$ 
     $rem \leftarrow (2 \cdot rem + w[i]) \bmod 5$ 
  if  $rem = 0$ 
    return TRUE
  else
    return FALSE

```

Aside from the loop index i , which we need just to read the entire input string, this algorithm has a single local variable rem , which has only four different values: 0, 1, 2, 3, or 4.

This algorithm already runs in $O(n)$ time, which is the best we can hope for—after all, we have to read every bit in the input—but we can speed up the algorithm *in practice*. Let’s define a **change** or **transition** function $\delta: \{0, 1, 2, 3, 4\} \times \{0, 1\} \rightarrow \{0, 1, 2, 3, 4\}$ as follows:

$$\delta(q, a) = (2q + a) \bmod 5.$$

(Here I’m implicitly converting the symbols **0** and **1** to the corresponding integers 0 and 1.) Since we already know all values of the transition function, we can store them in a precomputed table, and then replace the computation in the main loop of MULTIPLEOF5 with a simple array lookup.

We can also modify the return condition to check for different values modulo 5. To be completely general, we replace the final if-then-else lines with another array lookup, using an array $A[0..4]$ of booleans describing which final mod-5 values are “acceptable”.

After both of these modifications, our algorithm looks like one of the following, depending on whether we want something iterative or recursive (with $q = 0$ in the initial call):

```

DoSOMETHINGCOOL( $w[1..n]$ ):
   $q \leftarrow 0$ 
  for  $i \leftarrow 1$  to  $n$ 
     $q \leftarrow \delta[q, w[i]]$ 
  return  $A[q]$ 

```

```


DoSOMETHINGCOOL( $q, w$ ):
  if  $w = \epsilon$ 
    return  $A[q]$ 
  else
    decompose  $w = a \cdot x$ 
    return DoSOMETHINGCOOL( $\delta(q, a), x$ )

```

If we want to use our new DoSOMETHINGCOOL algorithm to implement MULTIPLEOF5, we simply give the arrays δ and A the following hard-coded values:

q	$\delta[q, 0]$	$\delta[q, 1]$	$A[q]$
0	0	1	TRUE
1	2	3	FALSE
2	4	0	FALSE
3	1	2	FALSE
4	3	4	FALSE

We can also visualize the behavior of DoSOMETHINGCOOL by drawing a directed graph, whose vertices represent possible values of the variable q —the possible *states* of the algorithm—and whose edges are labeled with input symbols to represent transitions between states. Specifically, the graph includes the labeled directed edge $p \xrightarrow{a} q$ if and only if $\delta(p, a) = q$. To indicate the proper return value, we draw the “acceptable” final states using doubled circles. Here is the resulting graph for MULTIPLEOF5:



State-transition graph for MULTIPLEOF5

If we run the MULTIPLEOF5 algorithm on the string 00101110110 (representing the number 374 in binary), the algorithm performs the following sequence of transitions:

$$0 \xrightarrow{0} 0 \xrightarrow{0} 0 \xrightarrow{1} 1 \xrightarrow{0} 2 \xrightarrow{1} 0 \xrightarrow{1} 1 \xrightarrow{1} 3 \xrightarrow{0} 1 \xrightarrow{1} 3 \xrightarrow{1} 2 \xrightarrow{0} 4$$

Because the final state is not the “acceptable” state 0, the algorithm correctly returns FALSE. We can also think of this sequence of transitions as a walk in the graph, which is completely determined by the start state 0 and the sequence of edge labels; the algorithm returns TRUE if and only if this walk ends at an “acceptable” state.

3.2 Formal Definitions

The object we have just described is an example of a *finite-state machine*. A finite-state machine is a formal model of any system/machine/algorithm that can exist in a finite number of *states* and that transitions among those states based on sequence of *input* symbols.

Finite-state machines are also known as *deterministic finite-state automata*, abbreviated *DFA*s. The word “deterministic” means that the behavior of the machine is completely *determined* by the input string; we’ll discuss nondeterministic automata in the next lecture. The word “automaton” (the singular of “automata”) comes from ancient Greek αὐτόματος meaning “self-acting”, from the roots αὐτό- (“self”) and -ματος (“thinking, willing”, the root of Latin *mentus*).

Formally, every finite-state machine consists of five components:

- An arbitrary finite set Σ , called the *input alphabet*.
- Another arbitrary finite set Q , whose elements are called *states*.
- An arbitrary *transition* function $\delta : Q \times \Sigma \rightarrow Q$.
- A *start state* $s \in Q$.
- A subset $A \subseteq Q$ of *accepting states*.

The behavior of a finite-state machine is governed by an *input string* w , which is a finite sequence of symbols from the input alphabet Σ . The machine *reads* the symbols in w one at a time in order (from left to right). At all times, the machine has a *current state* q ; initially q is the machine's start state s . Each time the machine reads a symbol a from the input string, its current state *transitions* from q to $\delta(q, a)$. After all the characters have been read, the machine *accepts* w if the current state is in A and *rejects* w otherwise. In other words, every finite state machine runs the algorithm `DoSomethingCool!`

More formally, we extend the transition function $\delta : Q \times \Sigma \rightarrow Q$ of any finite-state machine to a function $\delta^* : Q \times \Sigma^* \rightarrow Q$ that transitions on *strings* as follows:

$$\delta^*(q, w) := \begin{cases} q & \text{if } w = \epsilon, \\ \delta^*(\delta(q, a), x) & \text{if } w = ax. \end{cases}$$

Finally, a finite-state machine *accepts* a string w if and only if $\delta^*(s, w) \in A$, and *rejects* w otherwise. (Compare this definition with the recursive formulation of `DoSomethingCool!`)

For example, our final `MULTIPLEOF5` algorithm is a DFA with the following components:

- input alphabet: $\Sigma = \{0, 1\}$
- state set: $Q = \{0, 1, 2, 3, 4\}$
- transition function: $\delta(q, a) = (2q + a) \bmod 5$
- start state: $s = 0$
- accepting states: $A = \{0\}$

This machine rejects the string `00101110110`, because

$$\begin{aligned} \delta^*(0, 00101110110) &= \delta^*(\delta(0, 0), 0101110110) \\ &= \delta^*(0, 0101110110) = \delta^*(\delta(0, 0), 101110110) \\ &= \delta^*(0, 101110110) = \delta^*(\delta(0, 1), 01110110) = \dots \\ &\quad \vdots \\ \dots &= \delta^*(1, 110) = \delta^*(\delta(1, 1), 10) \\ &= \delta^*(3, 10) = \delta^*(\delta(3, 1), 0) \\ &= \delta^*(2, 0) = \delta^*(\delta(2, 0), \epsilon) \\ &= \delta^*(4, \epsilon) = 4 \notin A. \end{aligned}$$

We have already seen a more graphical representation of this entire sequence of transitions:

$$0 \xrightarrow{0} 0 \xrightarrow{0} 0 \xrightarrow{1} 1 \xrightarrow{0} 2 \xrightarrow{1} 0 \xrightarrow{1} 1 \xrightarrow{1} 3 \xrightarrow{0} 1 \xrightarrow{1} 3 \xrightarrow{1} 2 \xrightarrow{0} 4$$


The arrow notation is easier to read and write for specific examples, but surprisingly, most people actually find the more formal functional notation easier to use in formal proofs. Try them both!

We can equivalently define a DFA as a directed graph whose vertices are the states Q , whose edges are labeled with symbols from Σ , such that every vertex has exactly one outgoing edge with each label. In our drawings of finite state machines, the start state s is always indicated by an incoming arrow, and the accepting states A are always indicated by doubled circles. By induction, for any string $w \in \Sigma^*$, this graph contains a unique walk that starts at s and whose edges are labeled with the symbols in w in order. The machine accepts w if this walk ends at an accepting state. This graphical formulation of DFAs is incredibly useful for developing intuition and even designing DFAs. For proofs, it's largely a matter of taste whether to write in terms of extended transition functions or labeled graphs, but (as much as I wish otherwise) I actually find it easier to write *correct* proofs using the functional formulation.

3.3 Another Example

The following drawing shows a finite-state machine with input alphabet $\Sigma = \{0, 1\}$, state set $Q = \{s, t\}$, start state s , a single accepting state t , and the transition function

$$\delta(s, 0) = s, \quad \delta(s, 1) = t, \quad \delta(t, 0) = t, \quad \delta(t, 1) = s.$$



A simple finite-state machine.

For example, the two-state machine M at the top of this page accepts the string 00101110100 after the following sequence of transitions:

$$s \xrightarrow{0} s \xrightarrow{0} s \xrightarrow{1} t \xrightarrow{0} t \xrightarrow{1} s \xrightarrow{1} t \xrightarrow{1} s \xrightarrow{0} s \xrightarrow{1} t \xrightarrow{0} t \xrightarrow{0} t.$$

The same machine M rejects the string 11101101 after the following sequence of transitions:

$$s \xrightarrow{1} t \xrightarrow{1} s \xrightarrow{1} t \xrightarrow{0} t \xrightarrow{1} s \xrightarrow{1} t \xrightarrow{0} t \xrightarrow{1} s.$$

Finally, M rejects the empty string, because the start state s is not an accepting state.

From these examples and others, it is easy to conjecture that the language of M is the set of all strings of 0 s and 1 s with an odd number of 1 s. So let's prove it!

Proof (tedious case analysis): Let $\#(a, w)$ denote the number of times symbol a appears in string w . We will prove the following stronger claims by induction, for any string w .

$$\delta^*(s, w) = \begin{cases} s & \text{if } \#(1, w) \text{ is even} \\ t & \text{if } \#(1, w) \text{ is odd} \end{cases} \quad \text{and} \quad \delta^*(t, w) = \begin{cases} t & \text{if } \#(1, w) \text{ is even} \\ s & \text{if } \#(1, w) \text{ is odd} \end{cases}$$

Let's begin. Let w be an arbitrary string. Assume that for any string x that is shorter than w , we have $\delta^*(s, x) = s$ and $\delta^*(t, x) = t$ if x has an even number of 1 s, and $\delta^*(s, x) = t$ and $\delta^*(t, x) = s$ if x has an odd number of 1 s. There are five cases to consider.

- If $w = \epsilon$, then w contains an even number of 1 s and $\delta^*(s, w) = s$ and $\delta^*(t, w) = t$ by definition.

- Suppose $w = \mathbf{1}x$ and $\#(\mathbf{1}, w)$ is even. Then $\#(\mathbf{1}, x)$ is odd, which implies

$$\begin{aligned} \delta^*(s, w) &= \delta^*(\delta(s, \mathbf{1}), x) && \text{by definition of } \delta^* \\ &= \delta^*(t, x) && \text{by definition of } \delta \\ &= s && \text{by the inductive hypothesis} \\ \delta^*(t, w) &= \delta^*(\delta(t, \mathbf{1}), x) && \text{by definition of } \delta^* \\ &= \delta^*(s, x) && \text{by definition of } \delta \\ &= T && \text{by the inductive hypothesis} \end{aligned}$$

Since the remaining cases are similar, I'll omit the line-by-line justification.

- If $w = \mathbf{1}x$ and $\#(\mathbf{1}, w)$ is odd, then $\#(\mathbf{1}, x)$ is even, so the inductive hypothesis implies

$$\begin{aligned} \delta^*(s, w) &= \delta^*(\delta(s, \mathbf{1}), x) = \delta^*(t, x) = t \\ \delta^*(t, w) &= \delta^*(\delta(t, \mathbf{1}), x) = \delta^*(s, x) = s \end{aligned}$$

- If $w = \mathbf{0}x$ and $\#(\mathbf{1}, w)$ is even, then $\#(\mathbf{1}, x)$ is even, so the inductive hypothesis implies

$$\begin{aligned} \delta^*(s, w) &= \delta^*(\delta(s, \mathbf{0}), x) = \delta^*(s, x) = s \\ \delta^*(t, w) &= \delta^*(\delta(t, \mathbf{0}), x) = \delta^*(t, x) = t \end{aligned}$$

- Finally, if $w = \mathbf{0}x$ and $\#(\mathbf{1}, w)$ is odd, then $\#(\mathbf{1}, x)$ is odd, so the inductive hypothesis implies

$$\begin{aligned} \delta^*(s, w) &= \delta^*(\delta(s, \mathbf{0}), x) = \delta^*(s, x) = t \\ \delta^*(t, w) &= \delta^*(\delta(t, \mathbf{0}), x) = \delta^*(t, x) = s \end{aligned} \quad \square$$

Notice that this proof contains $|Q|^2 \cdot |\Sigma| + |Q|$ separate inductive arguments. For every pair of states p and q , we must argue about the language of all strings w such that $\delta^*(p, w) = q$, and we must consider every possible first symbol in w . We must also argue about $\delta(p, \varepsilon)$ for every state p . Each of those arguments is typically straightforward, but it's easy to get lost in the deluge of cases.

For this particular proof, however, we can reduce the number of cases by switching from tail recursion to *head* recursion. The following identity holds for all strings $x \in \Sigma^*$ and symbols $a \in \Sigma$:

$$\boxed{\delta^*(q, xa) = \delta(\delta^*(q, x), a)}$$

We leave the inductive proof of this identity as a straightforward exercise (hint, hint).

Proof (clever renaming, head induction): Let's rename the states with the integers 0 and 1 instead of s and t . Then the transition function can be described concisely as $\delta(q, a) = (q + a) \bmod 2$. We claim that for every string w , we have $\delta^*(0, w) = \#(\mathbf{1}, w) \bmod 2$.

Let w be an arbitrary string, and assume that for any string x that is shorter than w that $\delta^*(0, x) = \#(\mathbf{1}, x) \bmod 2$. There are only two cases to consider: either w is empty or it isn't.

- If $w = \varepsilon$, then $\delta^*(0, w) = 0 = \#(\mathbf{1}, w) \bmod 2$ by definition.


- Otherwise, $w = xa$ for some string x and some symbol a , and we have

$$\begin{aligned}
 \delta^*(0, w) &= \delta(\delta^*(0, x), a) && \text{by definition of } \delta^* \\
 &= \delta(\#(\mathbf{1}, x) \bmod 2, a) && \text{by the inductive hypothesis} \\
 &= (\#(\mathbf{1}, x) \bmod 2 + a) \bmod 2 && \text{by definition of } \delta \\
 &= (\#(\mathbf{1}, x) + a) \bmod 2 && \text{by definition of mod 2} \\
 &= (\#(\mathbf{1}, x) + \#(\mathbf{1}, a)) \bmod 2 && \text{because } \#(\mathbf{1}, \emptyset) = 0 \text{ and } \#(\mathbf{1}, \mathbf{1}) = 1 \\
 &= (\#(\mathbf{1}, xa)) \bmod 2 && \text{by definition of } \# \\
 &= (\#(\mathbf{1}, w)) \bmod 2 && \text{because } w = xa \quad \square
 \end{aligned}$$

Hmmm. This “clever” proof is certainly shorter than the earlier brute-force proof, but is it actually *better*? Simpler? More intuitive? Easier to understand? I’m skeptical. Sometimes brute force really is more effective.

3.4 Yet Another Example

As a more complex example, consider the *Rubik’s cube*, a well-known mechanical puzzle invented independently by Ernő Rubik in Hungary and Terutoshi Ishigi in Japan in the mid-1970s. This puzzle has precisely 519,024,039,293,878,272,000 distinct configurations. In the unique *solved* configuration, each of the six faces of the cube shows exactly one color. We can change the configuration of the cube by rotating one of the six faces of the cube by 90 degrees, either clockwise or counterclockwise. The cube has six faces (front, back, left, right, up, and down), so there are exactly twelve possible turns, typically represented by the symbols $R, L, F, B, U, D, \bar{R}, \bar{L}, \bar{F}, \bar{B}, \bar{U}, \bar{D}$, where the letter indicates which face to turn and the presence or absence of a bar over the letter indicates turning counterclockwise or clockwise, respectively. Thus, we can represent a Rubik’s cube as a finite-state machine with 519,024,039,293,878,272,000 states and an input alphabet with 12 symbols; or equivalently, as a directed graph with 519,024,039,293,878,272,000 vertices, each with 12 outgoing edges. In practice, the number of states is *far* too large for us to actually draw the machine or explicitly specify its transition function; nevertheless, the number of states is still finite. If we let the start state s and the sole accepting state be the solved state, then the language of this finite state machine is the set of all move sequences that leave the cube unchanged.



A complicated finite-state machine.

3.5 Building DFAs

This section describes a few examples of building DFAs that accept particular languages, thereby proving that those languages are automatic. As usual in algorithm design, there is no purely

mechanical recipe—no *automatic* method—no *algorithm*—for building DFAs in general. However, the following examples show several useful design strategies.

3.5.1 Superstrings

Perhaps the simplest rule of thumb is to try to construct an algorithm that looks like MULTIPLEOF5: A simple for-loop through the symbols, using a *constant* number of variables, where each variable (except the loop index) has only a *constant* number of possible values. Here, “constant” means an actual number that is not a function of the input size n . You should be able to compute the number of possible values for each variable *at compile time*.

For example, the following algorithm determines whether a given string in $\Sigma = \{0, 1\}$ contains the substring **11**.

```

CONTAINS11(w[1..n]):
  found ← FALSE
  for i ← 1 to n
    if i = 1
      last2 ← w[1]
    else
      last2 ← w[i - 1] · w[i]
    if last2 = 11
      found ← TRUE
  return found
    
```

Aside from the loop index, this algorithm has exactly two variables.

- A boolean flag *found* indicating whether we have seen the substring **11**. This variable has exactly two possible values: TRUE and FALSE.
- A string *last2* containing the last (up to) three symbols we have read so far. This variable has exactly 7 possible values: ϵ , 0, 1, 00, 01, 10, and 11.

Thus, altogether, the algorithm can be in at most $2 \times 7 = 14$ possible states, one for each possible pair (*found*, *last2*). Thus, we can encode the behavior of CONTAINS11 as a DFA with fourteen states, where the start state is (FALSE, ϵ) and the accepting states are all seven states of the form (TRUE, *). The transition function is described in the following table (split into two parts to save space):

q	$\delta[q, 0]$	$\delta[q, 1]$	q	$\delta[q, 0]$	$\delta[q, 1]$
(FALSE, ϵ)	(FALSE, 0)	(FALSE, 1)	(TRUE, ϵ)	(TRUE, 0)	(TRUE, 1)
(FALSE, 0)	(FALSE, 00)	(FALSE, 01)	(TRUE, 0)	(TRUE, 00)	(TRUE, 01)
(FALSE, 1)	(FALSE, 10)	(TRUE, 11)	(TRUE, 1)	(TRUE, 10)	(TRUE, 11)
(FALSE, 00)	(FALSE, 00)	(FALSE, 01)	(TRUE, 00)	(TRUE, 00)	(TRUE, 01)
(FALSE, 01)	(FALSE, 10)	(TRUE, 11)	(TRUE, 01)	(TRUE, 10)	(TRUE, 11)
(FALSE, 10)	(FALSE, 00)	(FALSE, 01)	(TRUE, 10)	(TRUE, 00)	(TRUE, 01)
(FALSE, 11)	(FALSE, 10)	(TRUE, 11)	(TRUE, 11)	(TRUE, 10)	(TRUE, 11)


For example, given the input string **1001011100**, this DFA performs the following sequence of

transitions and then accepts.

$$\begin{aligned}
 (\text{FALSE}, \varepsilon) &\xrightarrow{1} (\text{FALSE}, 1) \xrightarrow{0} (\text{FALSE}, 10) \xrightarrow{0} (\text{FALSE}, 00) \xrightarrow{1} \\
 &(\text{FALSE}, 01) \xrightarrow{0} (\text{FALSE}, 10) \xrightarrow{1} (\text{FALSE}, 01) \xrightarrow{1} \\
 &(\text{TRUE}, 11) \xrightarrow{1} (\text{TRUE}, 11) \xrightarrow{0} (\text{TRUE}, 10) \xrightarrow{0} (\text{TRUE}, 00)
 \end{aligned}$$

3.5.2 Reducing states


You can probably guess that the brute-force DFA we just constructed has considerably more states than necessary, especially after seeing its transition graph:



Our brute-force DFA for strings containing the substring 11

For example, we don't need actually to remember both of the last two symbols, but only the penultimate symbol, because the last symbol is the one we're currently reading. This observation allows us to reduce the number of states from fourteen to only six. Once the flag part of the state is set to TRUE, we know the machine will eventually accept, so we might as well merge all the accepting states together. Finally, and more subtly, because all transitions out of (FALSE, ε) and (FALSE, 0) are identical, we can merge those two states together as well. In the end, we obtain the following DFA with just three states:

- The start state, which indicates that the machine has not read the substring 11 and did not just read the symbol 1.
- An intermediate state, which indicates that the machine has not read the substring 11 but just read the symbol 1.
- A unique accept state, which indicates that the machine has read the substring 11.



A minimal DFA for superstrings of 11


While it is important not to use an excessive number of states—too many states makes a DFA hard to understand—there is really no point in trying to reduce DFAs *by hand* to the absolute

minimum number of states. Clarity is much more important than brevity (especially in this class), and DFAs with too *few* states can also be hard to understand. At the end of this note, I'll describe an efficient algorithm to transform any given DFA into an equivalent DFA with the fewest possible states.

3.5.3 Every this after that

Suppose we want to accept the set of strings in which every occurrence of the substring 00 occurs after every occurrence of the substring 11. Equivalently, we want to *reject* every string in which some 00 occurs before 11. Often the easiest way to design a DFA to check whether a string is *not* in some set is first to build a DFA that *is* in that set and then invert which states in that machine are accepting.

From the previous example, we know that there is a three-state DFA M_{11} that accepts the set of strings with the substring 11 and a nearly identical DFA M_{00} that accepts the set of strings containing the substring 00. By identifying the accept state of M_{00} with the start state of M_{11} , we obtain a five-state DFA that accepts the set of strings with 00 before 11. Finally, by inverting which states are accepting, we obtain the DFA we want.



Building a DFA for the language of strings in which every 00 is after every 11.


3.6 Combining DFAs: The Product Construction

Now suppose we want to accept all strings that contain both 00 and 11 as substrings, in either order. Intuitively, we'd like to run two of our earlier DFAs in parallel—the DFA M_{00} to detect superstrings of 00 and the DFA M_{11} to detect superstrings of 11—and then accept the input string if and only if both of these DFAs accept. In fact, we can encode precisely this “parallel computation” into a single DFA as follows:

- The states of the new DFA are all ordered pairs (p, q) , where p is a state in M_{00} and q is a state in M_{11} .
- The start state of the new DFA is the pair (s, s') , where s is the start state of M_{00} and s' is the start state of M_{11} .
- The new DFA includes the transition $(p, q) \xrightarrow{a} (p', q')$ if and only if M_{00} contains the transition $p \xrightarrow{a} p'$ and M_{11} contains the transition $q \xrightarrow{a} q'$.

- Finally, (p, q) is an accepting state of the new DFA if and only if p is an accepting state in M_{00} and q is an accepting state in M_{11} .

The resulting nine-state DFA is shown below, with the two factor DFAs M_{00} and M_{11} shown in gray for reference. (The state (a, a) can be removed, because it has no incoming arrows, but let's not worry about that now.)



Building a DFA for the language of strings containing both 00 and 11 .

More generally, let $M_1 = (\Sigma, Q_1, \delta_1, s_1, A_1)$ be an arbitrary DFA that accepts some language L_1 , and let $M_2 = (\Sigma, Q_2, \delta_2, s_2, A_2)$ be an arbitrary DFA that accepts some language L_2 (over the same alphabet Σ). We can construct a third DFA $M = (\Sigma, Q, \delta, s, A)$ that accepts the intersection language $L_1 \cap L_2$ as follows.

$$\begin{aligned}
 Q &:= Q_1 \times Q_2 = \{(p, q) \mid p \in Q_1 \text{ and } q \in Q_2\} \\
 s &:= (s_1, s_2) \\
 A &:= A_1 \times A_2 = \{(p, q) \mid p \in A_1 \text{ and } q \in A_2\} \\
 \delta((p, q), a) &:= (\delta_1(p, a), \delta_2(q, a))
 \end{aligned}$$

To convince ourselves that this product construction is actually correct, let's consider the extended transition function $\delta^*: (Q \times Q') \times \Sigma^* \rightarrow (Q \times Q')$, which acts on strings instead of individual symbols. Recall that this function is defined recursively as follows:

$$\delta^*((p, q), w) := \begin{cases} q & \text{if } w = \varepsilon, \\ \delta^*(\delta((p, q), a), x) & \text{if } w = ax. \end{cases}$$

This function behaves exactly as we should expect:

Lemma 3.1. $\delta^*((p, q), w) = (\delta_1^*(p, w), \delta_2^*(q, w))$ for any string w .

Proof: Let w be an arbitrary string. Assume $\delta^*((p, q), x) = (\delta_1^*(p, x), \delta_2^*(q, x))$ for every string x that is shorter than w . As usual, there are two cases to consider.

- First suppose $w = \varepsilon$:

$$\begin{aligned} \delta^*((p, q), \varepsilon) &= (p, q) && \text{by the definition of } \delta^* \\ &= (\delta_1^*(p, \varepsilon), q) && \text{by the definition of } \delta_1^* \\ &= (\delta_1^*(p, \varepsilon), \delta_2^*(q, \varepsilon)) && \text{by the definition of } \delta_2^* \end{aligned}$$

- Now suppose $w = ax$ for some symbol a and some string x :

$$\begin{aligned} \delta^*((p, q), ax) &= \delta^*(\delta((p, q), a), x) && \text{by the definition of } \delta^* \\ &= \delta^*((\delta_1(p, a), \delta_2(q, a)), x) && \text{by the definition of } \delta \\ &= (\delta_1^*((\delta_1(p, a), x), \delta_2^*(\delta_2(q, a), x))) && \text{by the induction hypothesis} \\ &= (\delta_1^*(p, ax), \delta_2^*(q, ax)) && \text{by the definitions of } \delta_1^* \text{ and } \delta_2^*. \end{aligned}$$

In both cases, we conclude that $\delta^*((p, q), w) = (\delta_1^*(p, w), \delta_2^*(q, w))$. □

An immediate consequence of this lemma is that for every string w , we have $\delta^*(s, w) \in A$ if and only if both $\delta_1^*(s_1, w) \in A_1$ and $\delta_2^*(s_2, w) \in A_2$. In other words, M accepts w if and only if **both** M_1 accepts w **and** M_2 accept w , as required.

As usual, this construction technique does not necessarily yield *minimal* DFAs. For example, in our first example of a product DFA, illustrated above, the central state (a, a) cannot be reached by any other state and is therefore redundant. Whatever.

Similar product constructions can be used to build DFAs that accept any other boolean combination of languages; in fact, the only part of the construction that changes is the choice of accepting states. For example:

- To accept the union $L_1 \cup L_2$, define $A = \{(p, q) \mid p \in A_1 \text{ **or** } q \in A_2\}$.
- To accept the difference $L_1 \setminus L_2$, define $A = \{(p, q) \mid p \in A_1 \text{ **but** } q \notin A_2\}$.
- To accept the symmetric difference $L_1 \oplus L_2$, define $A = \{(p, q) \mid p \in A_1 \text{ **xor** } q \in A_2\}$.

Examples of these constructions are shown on the next page.


Moreover, by cascading this product construction, we can construct DFAs that accept arbitrary boolean combinations of arbitrary finite collections of regular languages.

3.7 Automatic Languages and Closure Properties

The **language** of a finite state machine M , denoted $L(M)$ is the set of all strings in Σ^* that M accepts. We call a language **automatic** if it is the language of some finite state machine. Our product construction examples let us prove that the set of automatic languages is **closed** under simple boolean operations.

Theorem 3.2. *Let L and L' be arbitrary automatic languages over an arbitrary alphabet Σ .*

- $\bar{L} = \Sigma^* \setminus L$ is automatic.
- $L \cup L'$ is automatic.
- $L \cap L'$ is automatic.
- $L \setminus L'$ is automatic.
- $L \oplus L'$ is automatic.



DFA for (a) strings that contain 00 or 11, (b) strings that contain either 00 or 11 but not both, and (c) strings that contain 11 if they contain 00. These DFA are identical except for their choices of accepting states.

Eager students may have noticed that a Google search for the phrase “automatic language” turns up *no* results that are relevant for this class, except perhaps this lecture note. That’s because “automatic” is just a synonym for “regular”! This equivalence was first observed by Stephen Kleene (the inventor of regular expressions) in 1956.

Theorem 3.3 (Kleene). *For any regular expression R , there is a DFA M such that $L(R) = L(M)$. For any DFA M , there is a regular expression R such that $L(M) = L(R)$.*

Unfortunately, we don’t yet have all the tools we need to prove Kleene’s theorem; we’ll return to the proof in the next lecture note, after we have introduced *nondeterministic* finite-state machines. The proof is actually constructive—there are explicit algorithms that transform arbitrary DFA into equivalent regular expressions and vice versa.¹

This equivalence between regular and automatic languages implies that the set of **regular** languages is also closed under simple boolean operations. The union of two regular languages is regular *by definition*, but it’s much less obvious that *every* boolean combination of regular languages can also be described by regular expressions.

Corollary 3.4. *Let L and L' be arbitrary **regular** languages over an arbitrary alphabet Σ .*

- $\bar{L} = \Sigma^* \setminus L$ is regular.
- $L \cap L'$ is regular.
- $L \setminus L'$ is regular.
- $L \oplus L'$ is regular.

Conversely, because concatenations and Kleene closures of regular languages are regular *by definition*, we can immediately conclude that concatenations and Kleene closures of automatic languages are automatic.

Corollary 3.5. *Let L and L' be arbitrary automatic languages.*

- $L \cdot L'$ is automatic.
- L^* is automatic.


¹These conversion algorithms run in exponential time in the worst case, but that’s unavoidable. There are regular languages whose smallest accepting DFA is exponentially larger than their smallest regular expression, and there are regular languages whose smallest regular expression is exponentially larger than their smallest accepting DFA.

These results give us several options to prove that a given languages is regular or automatic. We can either (1) build a regular expression that describes the language, (2) build a DFA that accepts the language, or (3) build the language from simpler pieces from other regular/automatic languages. (Later we'll see a fourth option, and possibly even a fifth.)

3.8 Fooling Sets

But how would we prove that a language L is *not* regular? By Theorem 3.3, it suffices to prove that there is no finite-state automaton that accepts L . Equivalently, we need to prove that any automaton that accepts L requires infinitely many states. That may sound tricky, what with the “infinitely many”, but there’s actually a fairly simple technique to prove exactly that.

Perhaps the single most important feature of DFAs is that they have no memory other than the current state. Once a DFA enters a particular state, all *future* transitions depend only on that state and *future* input symbols; *past* input symbols are simply forgotten. For example, consider our very first DFA, which accepts the binary representations of integers divisible by 5.



DFA accepting binary multiples of 5.

The strings **0010** and **11011** both lead this DFA to state 2, although they follow different transitions to get there. Thus, for any string z , the strings **0010** z and **11011** z also lead to the same state in this DFA. In particular, **0010** z leads to the accepting state if and only if **11011** z leads to the accepting state. It follows that **0010** z is divisible by 5 if and only if **11011** z is divisible by 5.

More generally, for any DFA $M = (\Sigma, Q, s, A, \delta)$, we say that two strings x and y are **M -equivalent**, written $x \equiv_M y$, if $\delta^*(s, x) = \delta^*(s, y)$, that is, if x and y lead M to the same state. Two strings x and y are **L -equivalent**, written $x \equiv_L y$, if they are M -equivalent for any DFA M that recognizes.



Swap these two paragraphs. Let x and y be arbitrary strings.

If there is a DFA M that accepts L , such that
 x and y lead to the same state in M ,
 Then for every string $z \in \Sigma^*$
 $xz \in L \iff yz \in L$.

Now consider the contrapositive of this statement.

If there is a string $z \in \Sigma^*$ such that
 $xz \in L$ xor $yz \in L$,
 Then for every DFA M that accepts L
 x and y lead to different states in M .

Fix an arbitrary language L over an arbitrary alphabet Σ . For any strings $x, y, z \in \Sigma^*$, we say that z **distinguishes x and y** if exactly one of the strings xz and yz is in L . If no suffix distinguishes x and y , we say that x and y are **L -equivalent** and write $x \equiv_L y$. In other words,

$x \equiv_L y \iff$ For every string $z \in \Sigma^*$, we have $xz \in L$ if and only if $yz \in L$.

For example, let L_{eo} denote the language of strings over $\{0, 1\}$ with an even number of 0s and an odd number of 1s. Then the strings $x = 01$ and $y = 0011$ are distinguished by the string $z = 100$, because

$$\begin{aligned}xz &= 01 \cdot 100 = 01100 \in L_{eo} \\yz &= 0011 \cdot 100 = 0011100 \notin L_{eo}.\end{aligned}$$

On the other hand, it is quite easy to prove (hint, hint) that the strings 0001 and 1011 are L_{eo} -equivalent.

Now let M be an arbitrary DFA for our arbitrary language L , and let x be y be arbitrary strings (not necessarily in L). If x and y lead to the same state in M —that is, if $\delta^*(s, x) = \delta^*(s, y)$ —then we have

$$\delta^*(s, xz) = \delta^*(\delta^*(s, x), z) = \delta^*(\delta^*(s, y), z) = \delta^*(s, yz)$$

for any string z . In particular, either M accepts both x and y , or M rejects both x and y , and therefore $x \equiv_L y$. It follows that if x and y are not L -equivalent, then *any* DFA that accepts L has at least two distinct states $\delta^*(s, x) \neq \delta^*(s, y)$.

Finally, a **fooling set** for L is a set F of strings such that every pair of strings in F has a distinguishing suffix. For example, $F = \{01, 101, 010, 1010\}$ is a fooling set for the language L_{eo} of strings with an even number of 0s and an odd number of 1s, because each pair of strings in F has a distinguishing suffix:

- 0 distinguishes 01 and 101;
- 0 distinguishes 01 and 010;
- 0 distinguishes 01 and 1010;
- 10 distinguishes 101 and 010;
- 1 distinguishes 101 and 1010;
- 1 distinguishes 010 and 1010.

The pigeonhole principle now implies that for any integer k , if language L is accepted by a DFA with k states, then every fooling set for L contains at most k strings. This simple observation has two immediate corollaries.

First, for any integer k , if L has a fooling set of size k , then every DFA that accepts L has at least k states. For example, the fooling set $\{01, 101, 010, 1010\}$ proves that any DFA for L_{eo} has at least four states. Thus, we can use fooling sets to prove that certain DFAs are as small as possible.

Second, and more interestingly, if a language L is accepted by any DFA, then every fooling set for L must be finite. Equivalently:

If L has an infinite fooling set, then L is not accepted by any DFA.

This is arguably both the simplest and most powerful method for proving that a language is non-regular. Here are a few canonical examples of the fooling-set technique in action.

Lemma 3.6. *The language $L = \{0^n 1^n \mid n \geq 0\}$ is not regular.*

Proof: Consider the set $F = \{0^n \mid n \geq 0\}$, or more simply $F = 0^*$. Let x and y be arbitrary distinct strings in F . Then we must have $x = 0^i$ and $y = 0^j$ for some integers $i \neq j$. The suffix $z = 1^i$ distinguishes x and y , because $xz = 0^i 1^i \in L$, but $yz = 0^i 1^j \notin L$. We conclude that F is a fooling set for L . Because F is infinite, L cannot be regular. \square

Lemma 3.7. *The language $L = \{ww^R \mid w \in \Sigma^*\}$ of even-length palindromes is not regular.*

Proof: Let x and y be arbitrary distinct strings in 0^*1 . Then we must have $x = 0^i 1$ and $y = 0^j 1$ for some integers $i \neq j$. The suffix $z = 10^i$ distinguishes x and y , because $xz = 0^i 110^i \in L$, but $yz = 0^i 110^j \notin L$. We conclude that 0^*1 is a fooling set for L . Because 0^*1 is infinite, L cannot be regular. \square

Lemma 3.8. *The language $L = \{0^{2^n} \mid n \geq 0\}$ is not regular.*

Proof: Let x and y be arbitrary distinct strings in L . Then we must have $x = 0^{2^i}$ and $y = 0^{2^j}$ for some integers $i \neq j$. The suffix $z = 0^{2^i}$ distinguishes x and y , because $xz = 0^{2^i+2^i} = 0^{2^{i+1}} \in L$, but $yz = 0^{2^i+2^j} \notin L$. We conclude that L itself is a fooling set for L . Because L is infinite, L cannot be regular. \square

Lemma 3.9. *The language $L = \{0^p \mid p \text{ is prime}\}$ is not regular.*

Proof: Again, we use 0^* as our fooling set, but the actual argument is somewhat more complicated than in our earlier examples.

Let x and y be arbitrary distinct strings in 0^* . Then we must have $x = 0^i$ and $y = 0^j$ for some integers $i \neq j$. Without loss of generality, assume that $i < j$. Let p be any prime number larger than i . Because $p + 0(j-i)$ is prime and $p + p(j-i) > p$ is not, there must be a positive integer $k \leq p$ such that $p + (k-1)(j-i)$ is prime but $p + k(j-i)$ is not. Then the suffix $0^{p+(k-1)j-ki}$ distinguishes x and y :

$$\begin{aligned} xz &= 0^i 0^{p+(k-1)j-ki} = 0^{p+(k-1)(j-i)} \in L && \text{because } p + (k-1)(j-i) \text{ is prime;} \\ yz &= 0^j 0^{p+(k-1)j-ki} = 0^{p+k(j-i)} \notin L && \text{because } p + k(j-i) \text{ is not prime.} \end{aligned}$$

(Because $i < j$ and $i < p$, the suffix $0^{p+(k-1)j-ki} = 0^{(p-i)+(k-1)(j-i)}$ has positive length and therefore *actually exists!*) We conclude that 0^* is indeed a fooling set for L , which implies that L is not regular. \square

One natural question that many students ask is “How did you come up with that fooling set?” Perhaps the simplest rule of thumb is that for most languages L —in particular, for almost all languages that students are asked to prove non-regular on homeworks or exams—either some simple regular language like 0^* or 10^*1 is a fooling set, or the language L itself is a fooling set. (Of course, there are well-engineered counterexamples.)

*3.9 The Myhill-Nerode Theorem

The fooling set technique implies a *necessary* condition for a language to be accepted by a DFA—the language must have no infinite fooling sets. In fact, this condition is also *sufficient*.

The following powerful theorem was first proved by Anil Nerode in 1958, strengthening a 1957 result of John Myhill.²

The Myhill-Nerode Theorem. *For any language L , the following are equal:*

- (a) *the minimum number of states in a DFA that accepts L ,*
- (b) *the maximum size of a fooling set for L , and*
- (c) *the number of equivalence classes of \equiv_L .*

In particular, L is accepted by a DFA if and only if every fooling set for L is finite.

Proof: Let L be an arbitrary language.

We have already proved that the size of any fooling set for L is at most the number of states in any DFA that accepts L , so (a) \leq (b). It also follows directly from the definitions that $F \subseteq \Sigma^*$ is a fooling set for L if and only if F contains at most one string in each equivalence class of \equiv_L ; thus, (b) = (c). We complete the proof by showing that (a) \geq (c).

We have already proved that if \equiv_L has an infinite number of equivalence classes, there is no DFA that accepts L , so assume that the number of equivalence classes is finite. For any string w , let $[w]$ denote its equivalence class. We define a DFA $M_{\equiv} = (\Sigma, Q, s, A, \delta)$ as follows:

$$\begin{aligned} Q &:= \{[w] \mid w \in \Sigma^*\} \\ s &:= [\varepsilon] \\ A &:= \{[w] \mid w \in L\} \\ \delta([w], a) &:= [w \cdot a] \end{aligned}$$

We claim that this DFA accepts the language L ; this claim completes the proof of the theorem.

But before we can prove anything about this DFA, we first need to verify that it is actually well-defined. Let x and y be two strings such that $[x] = [y]$. By definition of L -equivalence, for any string z , we have $xz \in L$ if and only if $yz \in L$. It immediately follows that for any symbol $a \in \Sigma$ and any string z' , we have $xaz' \in L$ if and only if $yaz' \in L$. Thus, by definition of L -equivalence, we have $[xa] = [ya]$ for every symbol $a \in \Sigma$. We conclude that the function δ is indeed well-defined.

An easy inductive proof implies that $\delta^*([\varepsilon], x) = [x]$ for every string x . Thus, M accepts string x if and only if $[x] = [w]$ for some string $w \in L$. But if $[x] = [w]$, then by definition (setting $z = \varepsilon$), we have $x \in L$ if and only if $w \in L$. So M accepts x if and only if $x \in L$. In other words, M accepts L , as claimed, so the proof is complete. \square

*3.10 Minimal Automata

Given a DFA $M = (\Sigma, Q, s, A, \delta)$, suppose we want to find another DFA $M' = (\Sigma, Q', s', A', \delta')$ with the fewest possible states that accepts the same language. In this final section, we describe an efficient algorithm to minimize DFAs, first described (in slightly different form) by Edward Moore in 1956. We analyze the running time of Moore's in terms of two parameters: $n = |Q|$ and $\sigma = |\Sigma|$.

²Myhill considered the finer equivalence relation $x \sim_L y$, meaning $wxz \in L$ if and only if $wyz \in L$ for all strings w and z , and proved that L is regular if and only if \sim_L defines a finite number of equivalence classes. Like most of Myhill's early automata research, this result appears in an unpublished Air Force technical report. The modern Myhill-Nerode theorem appears (in an even more general form) as a minor lemma in Nerode's 1958 paper, which (not surprisingly) does not cite Myhill.

In the preprocessing phase, we find and remove any states that cannot be reached from the start state s ; this filtering can be performed in $O(n\sigma)$ time using any graph traversal algorithm. So from now on we assume that all states are reachable from s .

Now we recursively define two states p and q in the trimmed DFA to be *distinguishable*, written $p \not\sim q$, if at least one of the following conditions holds:

- $p \in A$ and $q \notin A$,
- $p \notin A$ and $q \in A$, or
- $\delta(p, a) \not\sim \delta(q, a)$ for some $a \in \Sigma$.

Equivalently, $p \not\sim q$ if and only if there is a string z such that exactly one of the states $\delta^*(p, z)$ and $\delta^*(q, z)$ is accepting. (Sound familiar?) Intuitively, the main algorithm assumes that all states are equivalent until proven otherwise, and then repeatedly looks for state pairs that can be proved distinguishable.

The main algorithm maintains a two-dimensional table, indexed by the states, where $Dist[p, q] = \text{TRUE}$ indicates that we have proved states p and q are distinguishable. Initially, for all states p and q , we set $Dist[p, q] \leftarrow \text{TRUE}$ if $p \in A$ and $q \notin A$ or vice versa, and $Dist[p, q] = \text{FALSE}$ otherwise. Then we repeatedly consider each pair of states and each symbol to find more distinguishable pairs, until we make a complete pass through the table without modifying it. The table-filling algorithm can be summarized as follows:

```

MINDFATABLE( $\Sigma, Q, s, A, \delta$ ):
  for all  $p \in Q$ 
    for all  $q \in Q$ 
      if  $(p \in A \text{ and } q \notin A) \text{ or } (p \notin A \text{ and } q \in A)$ 
         $Dist[p, q] \leftarrow \text{TRUE}$ 
      else
         $Dist[p, q] \leftarrow \text{FALSE}$ 
   $notdone \leftarrow \text{TRUE}$ 
  while  $notdone$ 
     $notdone \leftarrow \text{FALSE}$ 
    for all  $p \in Q$ 
      for all  $q \in Q$ 
        if  $Dist[p, q] = \text{FALSE}$ 
          for all  $a \in \Sigma$ 
            if  $Dist[\delta(p, a), \delta(q, a)]$ 
               $Dist[p, q] \leftarrow \text{TRUE}$ 
               $notdone \leftarrow \text{TRUE}$ 
  return  $Dist$ 

```

The algorithm must eventually halt, because there are only a finite number of entries in the table that can be marked. In fact, the main loop is guaranteed to terminate after at most n iterations, which implies that the entire algorithm runs in $O(n\sigma^3)$ time. Once the table is filled,³ any two states p and q such that $Dist(p, q) = \text{FALSE}$ are equivalent and can be merged into a single state. The remaining details of constructing the minimized DFA are straightforward.




Need to prove that the main loop terminates in at most n iterations.

³More experienced readers should be enraged by the mere suggestion that any algorithm merely *fills in a table*, as opposed to *evaluating a recurrence*. This algorithm is no exception. Consider the boolean function $Dist(p, q, k)$, which equals TRUE if and only if p and q can be distinguished by some string of length at most k . This function obeys the

With more care, Moore's minimization algorithm can be modified to run in $O(\sigma n^2)$ time. A faster DFA minimization algorithm, due to John Hopcroft, runs in $O(\sigma n \log n)$ time.

Example

To get a better idea how this algorithm works, let's visualize its execution on our earlier brute-force DFA for strings containing the substring **11**. This DFA has four unreachable states: (FALSE, 11), (TRUE, ε), (TRUE, 0), and (TRUE, 1). We remove these states, and relabel the remaining states for easier reference. (In an actual implementation, the states would almost certainly be represented by indices into an array anyway, not by mnemonic labels.)



Our brute-force DFA for strings containing the substring **11**, after removing all four unreachable states

The main algorithm initializes (the bottom half of) a 10×10 table as follows. (In the following figures, cells marked \times have value TRUE and blank cells have value FALSE.)

	0	1	2	3	4	5	6	7	8
1									
2									
3									
4									
5									
6	\times	\times	\times	\times	\times	\times			
7	\times	\times	\times	\times	\times	\times			
8	\times	\times	\times	\times	\times	\times			
9	\times	\times	\times	\times	\times	\times			

In the first iteration of the main loop, the algorithm discovers several distinguishable pairs of states. For example, the algorithm sets $Dist[0, 2] \leftarrow TRUE$ because $Dist[\delta(0, 1), \delta(2, 1)] = Dist[2, 9] = TRUE$. After the iteration ends, the table looks like this:

following recurrence:


$$Dist(p, q, k) = \begin{cases} (p \in A) \oplus (q \in A) & \text{if } k = 0, \\ Dist(p, q, k - 1) \vee \bigvee_{a \in \Sigma} Dist(\delta(p, a), \delta(q, a), k - 1) & \text{otherwise.} \end{cases}$$

Moore's "table-filling" algorithm is just a space-efficient dynamic programming algorithm to evaluate this recurrence.

	0	1	2	3	4	5	6	7	8
1									
2	x	x							
3			x						
4	x	x		x					
5			x		x				
6	x	x	x	x	x	x			
7	x	x	x	x	x	x			
8	x	x	x	x	x	x			
9	x	x	x	x	x	x			

The second iteration of the while loop makes no further changes to the table—We got lucky!—so the algorithm terminates.

The final table implies that the states of our trimmed DFA fall into exactly three equivalence classes: $\{0, 1, 3, 5\}$, $\{2, 4\}$, and $\{6, 7, 8, 9\}$. Replacing each equivalence class with a single state gives us the three-state DFA that we already discovered.



Equivalence classes of states in the trimmed DFA, and the resulting minimal equivalent DFA.

Exercises

1. For each of the following languages in $\{0, 1\}^*$, describe a deterministic finite-state machine that accepts that language. There are infinitely many correct answers for each language. “Describe” does not necessarily mean “draw”.
 - (a) Only the string 0110.
 - (b) Every string except 0110.
 - (c) Strings that contain the substring 0110.
 - (d) Strings that do not contain the substring 0110.
 - * (e) Strings that contain an even number of occurrences of the substring 0110. (For example, this language contains the strings 0110110 and 01011.)
 - (f) Strings that contain the *subsequence* 0110.

- (g) Strings that do not contain the *subsequence* 0110.
- * (h) Strings that contain an even number of occurrences of the *subsequence* 0110.
- (i) Strings that contain an even number of 1s and an odd number of 0s.
- (j) Every string that represents a number divisible by 7 in binary.
- (k) Every string whose reversal represents a number divisible by 7 in binary.
- (l) Strings in which the substrings 01 and 10 appear the same number of times.
- (m) Strings such that in every prefix, the number of 0s and the number of 1s differ by at most 1.
- (n) Strings such that in every prefix, the number of 0s and the number of 1s differ by at most 4.
- (o) Strings that end with $0^{10} = 0000000000$.
- (p) All strings in which the number of 0s is even if and only if the number of 1s is *not* divisible by 3.
- (q) All strings that are both the binary representation of an integer divisible by 3 and the ternary (base-3) representation of an integer divisible by 4.
- (r) Strings in which the number of 1s is even, the number of 0s is divisible by 3, the overall length is divisible by 5, the binary value is divisible by 7, the binary value of the reversal is divisible by 11, and does not contain thirteen 1s in a row. [Hint: This is more tedious than difficult.]
- * (s) Strings w such that $\binom{|w|}{2} \bmod 6 = 4$.
- * (t) Strings w such that $F_{\#(10,w)} \bmod 10 = 4$, where $\#(10,w)$ denotes the number of times 10 appears as a substring of w , and as usual F_n is the n th Fibonacci number:

$$F_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ F_{n-1} + F_{n-2} & \text{otherwise} \end{cases}$$

- * (u) Strings w such that $F_{\#(1\dots 0,w)} \bmod 10 = 4$, where $\#(1\dots 0,w)$ denotes the number of times 10 appears as a *subsequence* of w , and as usual F_n is the n th Fibonacci number:

$$F_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ F_{n-1} + F_{n-2} & \text{otherwise} \end{cases}$$

2. (a) Let $L \subseteq 0^*$ be an arbitrary *unary* language. Prove that L^* is regular.
 (b) Prove that there is a binary language $L \subseteq (0+1)^*$ such that L^* is not regular.

3. Prove that none of the following languages is automatic.

- (a) $\{0^{n^2} \mid n \geq 0\}$
 (b) $\{0^{n^3} \mid n \geq 0\}$
 (c) $\{0^{f(n)} \mid n \geq 0\}$, where $f(n)$ is any fixed polynomial in n with degree at least 2.

- (d) $\{0^n \mid n \text{ is composite}\}$
- (e) $\{0^n 1 0^n \mid n \geq 0\}$
- (f) $\{0^m 1^n \mid m \neq n\}$
- (g) $\{0^m 1^n \mid m < 3n\}$
- (h) $\{0^{2n} 1^n \mid n \geq 0\}$
- (i) $\{w \in (0+1)^* \mid \#(0, w) = \#(1, w)\}$
- (j) $\{w \in (0+1)^* \mid \#(0, w) < \#(1, w)\}$
- (k) $\{0^m 1^n \mid m/n \text{ is an integer}\}$
- (l) $\{0^m 1^n \mid m \text{ and } n \text{ are relatively prime}\}$
- (m) $\{0^m 1^n \mid n - m \text{ is a perfect square}\}$
- (n) $\{w \# w \mid w \in (0+1)^*\}$
- (o) $\{ww \mid w \in (0+1)^*\}$
- (p) $\{w \# 0^{|w|} \mid w \in (0+1)^*\}$
- (q) $\{w 0^{|w|} \mid w \in (0+1)^*\}$
- (r) $\{xy \mid x, y \in (0+1)^* \text{ and } |x| = |y| \text{ but } x \neq y\}$
- (s) $\{0^m 1^n 0^{m+n} \mid m, n \geq 0\}$
- (t) $\{0^m 1^n 0^{mn} \mid m, n \geq 0\}$
- (u) Strings in which the substrings 00 and 11 appear the same number of times.
- (v) Strings of the form $w_1 \# w_2 \# \dots \# w_n$ for some $n \geq 2$, where $w_i \in \{0, 1\}^*$ for every index i , and $w_i = w_j$ for some indices $i \neq j$.
- (w) The set of all palindromes in $(0+1)^*$ whose length is divisible by 7.
- (x) $\{w \in (0+1)^* \mid w \text{ is the binary representation of a perfect square}\}$
- ★(y) $\{w \in (0+1)^* \mid w \text{ is the binary representation of a prime number}\}$
4. For each of the following languages over the alphabet $\Sigma = \{0, 1\}$, either prove that the language is regular (by constructing an appropriate DFA or regular expression) or prove that the language is not regular (using fooling sets). Recall that Σ^+ denotes the set of all nonempty strings over Σ . [Hint: Believe it or not, most of these languages **are** actually regular.]
- (a) $\{0^n w 1^n \mid w \in \Sigma^* \text{ and } n \geq 0\}$
- (b) $\{0^n 1^n w \mid w \in \Sigma^* \text{ and } n \geq 0\}$
- (c) $\{w 0^n 1^n x \mid w, x \in \Sigma^* \text{ and } n \geq 0\}$
- (d) $\{0^n w 1^n x \mid w, x \in \Sigma^* \text{ and } n \geq 0\}$
- (e) $\{0^n w 1 x 0^n \mid w, x \in \Sigma^* \text{ and } n \geq 0\}$
- (f) $\{0^n w 0^n \mid w \in \Sigma^+ \text{ and } n > 0\}$
- (g) $\{w 0^n w \mid w \in \Sigma^+ \text{ and } n > 0\}$
- (h) $\{wxw \mid w, x \in \Sigma^*\}$
- (i) $\{wxw \mid w, x \in \Sigma^+\}$

- (j) $\{wxw^R \mid w, x \in \Sigma^+\}$
- (k) $\{wwx \mid w, x \in \Sigma^+\}$
- (l) $\{ww^Rx \mid w, x \in \Sigma^+\}$
- (m) $\{wxwy \mid w, x, y \in \Sigma^+\}$
- (n) $\{wxw^Ry \mid w, x, y \in \Sigma^+\}$
- (o) $\{xwwy \mid w, x, y \in \Sigma^+\}$
- (p) $\{xww^Ry \mid w, x, y \in \Sigma^+\}$
- (q) $\{wxxw \mid w, x \in \Sigma^+\}$
- * (r) $\{wxw^Rx \mid w, x \in \Sigma^+\}$
- (s) All strings w such that no prefix of w is a palindrome.
- (t) All strings w such that no prefix of w with length at least 3 is a palindrome.
- (u) All strings w such that no substring of w with length at least 3 is a palindrome.
- (v) All strings w such that no prefix of w with positive even length is a palindrome.
- (w) All strings w such that no substring of w with positive even length is a palindrome.