

Caveat lector! This is the first edition of this lecture note. Some topics are incomplete, and there are almost certainly a few serious errors. Please send bug reports and suggestions to jeffe@illinois.edu.

*Nothing is better than simplicity . . .
nothing can make up for excess or for the lack of definiteness.*
— Walt Whitman, Preface to *Leaves of Grass* (1855)


*Freedom of choice
Is what you got.
Freedom from choice
Is what you want.*
— Devo, "Freedom of Choice", *Freedom of Choice* (1980)

Nondeterminism means never having to say you are wrong.
— BSD 4.3 fortune(6) file (c.1985)

4 Nondeterminism

4.1 Nondeterministic State Machines

The following diagram shows something that looks like a finite-state machine over the alphabet $\{0, 1\}$, but on closer inspection, it is not consistent with our earlier definitions. On one hand, there are two transitions out of s for each input symbol. On the other hand, states a and b are each missing an outgoing transition.



A nondeterministic finite-state automaton

Nevertheless, there is a sense in which this machine “accepts” the set of all strings that contain either 00 or 11 as a substring. Imagine that when the machine reads a symbol in state s , it makes a choice about which transition to follow. If the input string contains the substring 00 , then it is *possible* for the machine to end in the accepting state c , by *choosing* to move into state a when it reads a 0 immediately before another 0 . Similarly, if the input string contains the substring 11 , it is *possible* for the machine to end in the accepting state c . On the other hand, if the input string does not contain either 00 or 11 —or in other words, if the input alternates between 0 and 1 —there are no choices that lead the machine to the accepting state. If the machine incorrectly chooses to transition to state a and then reads a 1 , or transitions to b and then reads 0 , it explodes; the only way to avoid an explosion is to stay in state s .

This object is an example of a *nondeterministic finite-state automaton*, or *NFA*, so named because its behavior is not uniquely *determined* by the input string. Formally, every NFA has 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 $\delta : Q \times \Sigma \rightarrow 2^Q$.
- A **start state** $s \in Q$.
- A subset $A \subseteq Q$ of **accepting states**.

The only difference from the formal definition of *deterministic* finite-state automata is the domain of the transition function. In a DFA, the transition function always returns a single state; in an NFA, the transition function returns a *set* of states, which could be empty, or all of Q , or anything in between.

Just like DFAs, the behavior of an NFA is governed by an **input string** $w \in \Sigma^*$, which the machine reads one symbol at a time, from left to right. Unlike DFAs, however, an NFA does not maintain a single current state, but rather a *set* of current states. Whenever the NFA reads a symbol a , its set of current states changes from C to $\bigcup_{q \in C} \delta(q, a)$. After all symbols have been read, the NFA **accepts** w if its current state set contains *at least one* accepting state and **rejects** w otherwise. In particular, if the set of current states ever becomes empty, it will stay empty forever, and the NFA will reject.

More formally, we define the function $\delta^* : Q \times \Sigma^* \rightarrow 2^Q$ that transitions on *strings* as follows:

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

The NFA $(Q, \Sigma, \delta, s, A)$ **accepts** $w \in \Sigma^*$ if and only if $\delta^*(s, w) \cap A \neq \emptyset$.

We can equivalently define an NFA as a directed graph whose vertices are the states Q , whose edges are labeled with symbols from Σ . We no longer require that every vertex has exactly one outgoing edge with each label; it may have several such edges or none. An NFA accepts a string w if the graph contains *at least one* walk from the start state to an accepting state whose label is w .

4.2 Intuition

There are at least three useful ways to think about non-determinism.

Clairvoyance. Whenever an NFA reads symbol a in state q , it *chooses* the next state from the set $\delta(q, a)$, always *magically* choosing a state that leads to the NFA accepting the input string, unless no such choice is possible. As the BSD fortune file put it, “Nondeterminism means never having to say you’re wrong.”¹ Of course real machines can’t actually look into the future; that’s why I used the word “magic”.

Parallel threads. An arguably more “realistic” view is that when an NFA reads symbol a in state q , it spawns an independent execution thread for each state in $\delta(q, a)$. In particular, if $\delta(q, a)$ is empty, the current thread simply dies. The NFA accepts if *at least one* thread is in an accepting state after it reads the last input symbol.

Equivalently, we can imagine that when an NFA reads symbol a in state q , it branches into several parallel universes, one for each state in $\delta(q, a)$. If $\delta(q, a)$ is empty, the NFA destroys the

¹This sentence is a riff on a horrible aphorism that was (sadly) popular in the US in the 70s and 80s. Fortunately, everyone seems to have forgotten the original saying, except for that one time it was parodied on the Simpsons.

universe (including itself). Similarly, if the NFA finds itself in a non-accepting state when the input ends, the NFA destroys the universe. Thus, when the input is gone, only universes in which the NFA somehow chose a path to an accept state still exist. One slight disadvantage of this metaphor is that if an NFA reads a string that is not in its language, it destroys *all* universes.

Proofs/oracles. Finally, we can treat NFAs not as a mechanism for *computing* something, but only as a mechanism for *verifying proofs*. If we want to *prove* that a string w contains one of the suffixes **00** or **11**, it suffices to demonstrate a single walk in our example NFA that starts at s and ends at c , and whose edges are labeled with the symbols in w . Equivalently, whenever the NFA faces a nontrivial choice, the prover can simply tell the NFA which state to move to next.

This intuition can be formalized as follows. Consider a *deterministic* finite state machine whose input alphabet is the product $\Sigma \times \Omega$ of an *input* alphabet Σ and an *oracle* alphabet Ω . Equivalently, we can imagine that this DFA reads simultaneously from two strings of the same length: the *input* string w and the *oracle* string ω . In either formulation, the transition function has the form $\delta: Q \times \Sigma \times \Omega \rightarrow Q$. As usual, this DFA accepts the pair $(w, \omega) \in (\Sigma \times \Gamma)^*$ if and only if $\delta^*(s, w, \omega) \in A$. Finally, M **nondeterministically accepts** the string $w \in \Sigma^*$ if there is an oracle string $\omega \in \Omega^*$ with $|\omega| = |w|$ such that $(w, \omega) \in L(M)$.

4.3 ε -Transitions

It is fairly common for NFAs to include so-called **ε -transitions**, which allow the machine to change state without reading an input symbol. An NFA with ε -transitions accepts a string w if and only if there is a sequence of transitions $s \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \xrightarrow{a_3} \dots \xrightarrow{a_\ell} q_\ell$ where the final state q_ℓ is accepting, each a_i is either ε or a symbol in Σ , and $a_1 a_2 \dots a_\ell = w$.

More formally, the transition function in an NFA with ε -transitions has a slightly larger domain $\delta: Q \times (\Sigma \cup \{\varepsilon\}) \rightarrow 2^Q$. The **ε -reach** of a state $q \in Q$ consists of all states r that satisfy one of the following conditions:

- $r = q$
- $r \in \delta(q', \varepsilon)$ for some state q' in the ε -reach of q .

In other words, r is in the ε -reach of q if there is a (possibly empty) sequence of ε -transitions leading from q to r . Now we redefine the extended transition function $\delta^*: Q \times \Sigma^* \rightarrow 2^Q$, which transitions on arbitrary strings, as follows:

$$\delta^*(q, w) := \begin{cases} \{q\} & \text{if } w = \varepsilon, \\ \bigcup_{r \in \varepsilon\text{-reach}(q)} \bigcup_{r' \in \delta(r, a)} \delta^*(r', x) & \text{if } w = ax. \end{cases}$$

As usual, the modified NFA accepts a string w if and only if $\delta^*(s, w) \cap A \neq \emptyset$.

Given an NFA $M = (\Sigma, Q, s, A, \delta)$ with ε -transitions, we can easily construct an equivalent NFA $M' = (\Sigma, Q', s', A', \delta')$ without ε -transitions as follows:

$$\begin{aligned} Q' &:= Q \\ s' &= s \\ A' &= \{q \in Q \mid \varepsilon\text{-reach}(q) \cap A \neq \emptyset\} \\ \delta'(q, a) &= \bigcup_{r \in \varepsilon\text{-reach}(q)} \delta(r, a) \end{aligned}$$

Straightforward definition-chasing implies that M and M' accept exactly the same language. Thus, whenever we reason about or design NFAs, we are free to either allow or forbid ε -transitions, whichever is more convenient for the task at hand.

4.4 Kleene's Theorem

We are now finally in a position to prove the following fundamental fact, first observed by Steven Kleene:

Theorem 4.1. *A language L can be described by a regular expression if and only if L is the language accepted by a DFA.*

We will prove Kleene's fundamental theorem in four stages:

- Every DFA can be transformed into an equivalent NFA.
- Every NFA can be transformed into an equivalent DFA.
- Every regular expression can be transformed into an NFA.
- Every NFA can be transformed into an equivalent regular expression.

The first of these four transformations is completely trivial; a DFA is just a special type of NFA where the transition function always returns a single state. Unfortunately, the other three transformations require a bit more work.

4.5 DFA from NFAs: The Subset Construction

In the parallel-thread model of NFA execution, an NFA does not have a single current state, but rather a *set* of current states. The evolution of this set of states is *determined* by a modified transition function $\delta': 2^Q \times \Sigma \rightarrow 2^Q$, defined by setting $\delta'(P, a) := \bigcup_{p \in P} \delta(p, a)$ for any set of states $P \subseteq Q$ and any symbol $a \in \Sigma$. When the NFA finishes reading its input string, it accepts if and only if the current set of states intersects the set A of accepting states.

This formulation makes the NFA completely deterministic! We have just shown that any NFA $M = (\Sigma, Q, s, A, \delta)$ is equivalent to a DFA $M' = (\Sigma, Q', s', A', \delta')$ defined as follows:

$$\begin{aligned} Q' &:= 2^Q \\ s' &:= \{s\} \\ A' &:= \{S \subseteq Q \mid S \cap A \neq \emptyset\} \\ \delta'(q', a) &:= \bigcup_{p \in q'} \delta(p, a) \quad \text{for all } q' \subseteq Q \text{ and } a \in \Sigma. \end{aligned}$$

Similarly, any NFA with ε -transitions is equivalent to a DFA with the transition function

$$\delta'(q', a) = \bigcup_{p \in q'} \bigcup_{r \in \varepsilon\text{-reach}(p)} \delta(r, a)$$

for all $q' \subseteq Q$ and $a \in \Sigma$. This conversion from NFA to DFA is often called the **subset construction**, but that name is somewhat misleading; it's not a "construction" so much as a change in perspective.

One disadvantage of this "construction" is that it usually leads to DFAs with far more states than necessary, in part because most of those states are unreachable. These unreachable states can be avoided by constructing the DFA incrementally, essentially by performing a breadth-first search of the DFA graph, starting at its start state.


To execute this algorithm by hand, we prepare a table with $|\Sigma| + 3$ columns, with one row for each DFA state we discover. In order, these columns record the following information:

- The DFA state (as a subset of NFA states)
- The ϵ -reach of the corresponding subset of NFA states
- Whether the DFA state is accepting (that is, whether the ϵ -reach intersects A)
- The output of the transition function for each symbol in Σ .

We start with DFA-state s in the first row and first column. Whenever we discover an unexplored state in one of the last $|\Sigma|$ columns, we copy it to the left column in a new row.

For example, given the NFA on the first page of this note, this incremental algorithm produces the following table, yielding a five-state DFA. For this example, the second column is redundant, because the NFA has no ϵ -transitions, but we will see another example with ϵ -transitions in the next subsection. To simplify notation, we write each set of states as a simple string, omitting braces and commas.

q'	ϵ -reach(q')	$q' \in A'$?	$\delta'(q', 0)$	$\delta'(q', 1)$
s	s		as	bs
as	as		acs	bs
bs	bs		as	bcs
acs	acs	✓	acs	bcs
bcs	bcs	✓	acs	bcs



Our example NFA, and the output of the incremental subset construction for that NFA.

4.6 NFAs from Regular Expressions: Thompson’s Algorithm

Lemma 4.2. *Every regular language is accepted by a non-deterministic finite automaton.*

Proof: In fact, we will prove the following stronger claim: Every regular language is accepted by an NFA with exactly one accepting state, which is different from its start state. The following construction was first described by Ken Thompson in 1968. Thompson’s algorithm actually proves a stronger statement: For any regular language L , there is an NFA that accepts L that has exactly one accepting state t , which is distinct from the starting state s .


Let R be an arbitrary regular expression over an arbitrary finite alphabet Σ . Assume that for any sub-expression S of R , the language described by S is accepted by an NFA with one accepting state distinct from its start state, which we denote pictorially by $\rightsquigarrow \boxed{s} \odot$. There are six cases to consider—three base cases and three recursive cases—mirroring the recursive definition of a regular expression.

- If $R = \emptyset$, then $L(R) = \emptyset$ is accepted by the empty NFA: $\rightsquigarrow \quad \odot$.

- If $R = \epsilon$, then $L(R) = \{\epsilon\}$ is accepted by the NFA $\text{Start} \xrightarrow{\epsilon} \text{Accept}$.
- If $R = a$ for some symbol $a \in \Sigma$, then $L(R) = \{a\}$ is accepted by the NFA $\text{Start} \xrightarrow{a} \text{Accept}$. (The case where R is a single *string* with length greater than 1 reduces to the single-symbol case by concatenation, as described in the next case.)
- Suppose $R = ST$ for some regular expressions S and T . The inductive hypothesis implies that the languages $L(S)$ and $L(T)$ are accepted by NFAs $\text{Start} \xrightarrow{S} \text{Accept}$ and $\text{Start} \xrightarrow{T} \text{Accept}$, respectively. Then $L(R) = L(ST) = L(S) \cdot L(T)$ is accepted by the NFA $\text{Start} \xrightarrow{S} \text{Intermediate} \xrightarrow{T} \text{Accept}$, built by connecting the two component NFAs in series.
- Suppose $R = S + T$ for some regular expressions S and T . The inductive hypothesis implies that the language $L(S)$ and $L(T)$ are accepted by NFAs $\text{Start} \xrightarrow{S} \text{Accept}$ and $\text{Start} \xrightarrow{T} \text{Accept}$, respectively. Then $L(R) = L(S + T) = L(S) \cup L(T)$ is accepted by the NFA $\text{Start} \xrightarrow{\epsilon} \text{S_Start} \xrightarrow{S} \text{S_Accept} \xrightarrow{\epsilon} \text{Accept}$ and $\text{Start} \xrightarrow{\epsilon} \text{T_Start} \xrightarrow{T} \text{T_Accept} \xrightarrow{\epsilon} \text{Accept}$, built by connecting the two component NFAs in parallel with new start and accept states.
- Finally, suppose $R = S^*$ for some regular expression S . The inductive hypothesis implies that the language $L(S)$ is accepted by an NFA $\text{Start} \xrightarrow{S} \text{Accept}$. Then the language $L(R) = L(S^*) = L(S)^*$ is accepted by the NFA $\text{Start} \xrightarrow{\epsilon} \text{Start} \xrightarrow{S} \text{Intermediate} \xrightarrow{\epsilon} \text{Accept}$.

In all cases, the language $L(R)$ is accepted by an NFA with one accepting state, which is different from its start state, as claimed. \square


As an example, given the regular expression $(0 + 10^*1)^*$ of strings containing an even number of **1**s, Thompson's algorithm produces a 14-state NFA shown on the next page. As this example shows, Thompson's algorithm tends to produce NFAs with many redundant states. Fortunately, just as there are for DFAs, there are algorithms that can reduce any NFA to an equivalent NFA with the smallest possible number of states.



The NFA constructed by Thompson's algorithm for the regular expression $(0 + 10^*1)^*$. The four non- ϵ -transitions are drawn with with bold red arrows for emphasis.

Interestingly, applying the incremental subset algorithm to Thompson's NFA tends to yield a DFA with relatively few states, in part because the states in Thompson's NFA tend to have large ϵ -reach, and in part because relatively few of those states are the targets of non- ϵ -transitions. Starting with the NFA shown above, for example, the incremental subset construction yields a DFA for the language $(0 + 10^*1)^*$ with just five states:

q'	ϵ -reach(q')	$q' \in A'$?	$\delta'(q', 0)$	$\delta'(q', 1)$
s	$sabjm$	✓	k	c
k	$sabjklm$	✓	k	c
c	$cdegh$		f	i
f	$degh$		f	i
i	$sabjilm$	✓	k	c



The DFA computed by the incremental subset algorithm from Thompson's NFA for $(0 + 10^*1)^*$.

This DFA can be further simplified to just two states, by observing that all three accepting states are equivalent, and that both non-accepting states are equivalent. But still, five states is pretty good, especially compared with the $2^{13} = 8096$ states that the naïve subset construction would yield!

***4.7 NFAs from Regular Expressions: Glushkov's Algorithm**

Thompson's algorithm is actually a modification of an earlier algorithm, which was independently discovered by Robert McNaughton and Hisao Yamada in 1960 and by V. I. Glushkov in 1961. Given a regular expression containing n symbols (not counting the parentheses and pluses and stars), Glushkov's algorithm produces an NFA with exactly $n + 1$ states.

Glushkov's algorithm combines six functions on regular expressions:

- $index(R)$ is the regular expression obtained by replacing the symbols in R with the integers 1 through n , in order from left to right. For example, $index((0 + 10^*1)^*) = (1 + 23^*4)^*$.
- $symbols(R)$ denotes the string obtained by removing all non-symbols from R . For example, $symbols((0 + 10^*1)^*) = 0101$.
- $has-\epsilon(R)$ is TRUE if $\epsilon \in L(R)$ and FALSE otherwise.
- $first(R)$ is the set of all initial symbols of strings in $L(R)$.
- $last(R)$ is the set of all final symbols of strings in $L(R)$.
- $middle(R)$ is the set of all pairs (a, b) such that ab is a substring of some string in $L(R)$.

The last four functions obey the following recurrences:

$$\text{has-}\varepsilon(\emptyset) = \emptyset$$

$$\text{has-}\varepsilon(w) = \begin{cases} \text{TRUE} & \text{if } w = \varepsilon \\ \text{FALSE} & \text{otherwise} \end{cases}$$

$$\text{has-}\varepsilon(S + T) = \text{has-}\varepsilon(S) \vee \text{has-}\varepsilon(T)$$

$$\text{has-}\varepsilon(ST) = \text{has-}\varepsilon(S) \wedge \text{has-}\varepsilon(T)$$

$$\text{has-}\varepsilon(S^*) = \text{TRUE}$$

$$\text{first}(\emptyset) = \emptyset$$

$$\text{first}(w) = \begin{cases} \emptyset & \text{if } w = \varepsilon \\ \{a\} & \text{if } w = ax \end{cases}$$

$$\text{first}(S + T) = \text{first}(S) \cup \text{first}(T)$$

$$\text{first}(ST) = \begin{cases} \text{first}(S) \cup \text{first}(T) & \text{if } \text{has-}\varepsilon(S) \\ \text{first}(T) & \text{otherwise} \end{cases}$$

$$\text{first}(S^*) = \text{first}(S)$$

$$\text{last}(\emptyset) = \emptyset$$

$$\text{last}(w) = \begin{cases} \emptyset & \text{if } w = \varepsilon \\ \{a\} & \text{if } w = xa \end{cases}$$

$$\text{last}(S + T) = \text{last}(S) \cup \text{last}(T)$$

$$\text{last}(ST) = \begin{cases} \text{last}(S) \cup \text{last}(T) & \text{if } \text{has-}\varepsilon(T) \\ \text{last}(T) & \text{otherwise} \end{cases}$$

$$\text{last}(S^*) = \text{last}(S)$$

$$\text{middle}(\emptyset) = \emptyset$$

$$\text{middle}(w) = \begin{cases} \emptyset & \text{if } |w| \leq 1 \\ \{(a, b)\} \cup \text{middle}(bx) & \text{if } w = abx \end{cases}$$

$$\text{middle}(S + T) = \text{middle}(S) \cup \text{middle}(T)$$

$$\text{middle}(ST) = \text{middle}(S) \cup (\text{last}(S) \times \text{first}(T)) \cup \text{middle}(T)$$

$$\text{middle}(S^*) = \text{middle}(S) \cup (\text{last}(S) \times \text{first}(S))$$

For example, the set $\text{middle}((1 + 23^*4)^*)$ can be computed recursively as follows. If we're doing this by hand, we can skip many of the steps in this derivation, because we know what the functions first , middle , last , and $\text{has-}\varepsilon$ actually mean, but a mechanical recursive evaluation would necessarily evaluate every step.

$$\text{middle}((1 + 23^*4)^*)$$

$$\begin{aligned} &= \text{middle}(1 + 23^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= \text{middle}(1) \cup \text{middle}(23^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= \emptyset \cup \text{middle}(23^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= \text{middle}(2) \cup (\text{last}(2) \times \text{first}(3^*4)) \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= \emptyset \cup (\{2\} \times \text{first}(3^*4)) \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= (\{2\} \times (\text{first}(3^*) \cup \text{first}(4))) \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= (\{2\} \times (\text{first}(3) \cup \text{first}(4))) \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= (\{2\} \times \{3, 4\}) \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &= \{(2, 3), (3, 4)\} \cup \text{middle}(3^*4) \cup (\text{last}(1 + 23^*4) \times \text{first}(1 + 23^*4)) \\ &\vdots \\ &= \{(1, 1), (1, 2), (2, 3), (2, 4), (3, 3), (3, 4), (4, 1), (4, 2)\} \end{aligned}$$

Finally, given any regular expression R , Glushkov's algorithm constructs the NFA $M_R = (\Sigma, Q, s, A, \delta)$ that accepts exactly the language $L(R)$ as follows:

$$Q = \{0, 1, \dots, |\text{symbols}(R)|\}$$

$$s = 0$$

$$A = \begin{cases} \{0\} \cup \text{last}(\text{index}(R)) & \text{if } \text{has-}\epsilon(R) \\ \text{last}(\text{index}(R)) & \text{otherwise} \end{cases}$$

$$\delta(0, a) = \{j \in \text{first}(\text{index}(R)) \mid a = \text{symbols}(R)[j]\}$$

$$\delta(i, a) = \{j \mid (i, j) \in \text{middle}(\text{index}(R)) \text{ and } a = \text{symbols}(R)[j]\}$$

There are a few natural ways to think about Glushkov's algorithm that are somewhat less impenetrable than the previous wall of definitions. One viewpoint is that Glushkov's algorithm first computes a DFA for the indexed regular expression $\text{index}(R)$ —in fact, a DFA with the fewest possible states, except for an extra start state—and then replaces each index with the corresponding symbol in $\text{symbols}(R)$ to get an NFA for the original expression R . Another useful observation is that Glushkov's NFA is identical to the result of removing all ϵ -transitions from Thompson's NFA for the same regular expression.

For example, given the regular expression $R = (0 + 10^*1)^*$, Glushkov's algorithm computes

$$\text{index}(R) = (1 + 23^*4)^*$$

$$\text{symbols}(R) = 0101$$


$$\text{has-}\epsilon(R) = \text{TRUE}$$

$$\text{first}(\text{index}(R)) = \{1, 2\}$$

$$\text{last}(\text{index}(R)) = \{1, 4\}$$

$$\text{middle}(\text{index}(R)) = \{(1, 1), (1, 2), (2, 3), (2, 4), (3, 3), (3, 4), (4, 1), (4, 2)\}$$

and then constructs the following five-state NFA.




Glushkov's DFA for the index expression $(1 + 23^*4)^*$ and Glushkov's NFA for the regular expression $(0 + 10^*1)^*$.

Hey, look, Glushkov's algorithm actually gave us a DFA! In fact, it gave us *precisely* the same DFA that we constructed earlier by sending Thompson's NFA through the incremental subset algorithm! Unfortunately, that's just a coincidence; in general the output of Glushkov's algorithm is *not* deterministic. We'll see a more typical example in the next section.

4.8 Another Example


Here is another example of all the algorithms we've seen so far, starting with the regular expression $(0 + 1)^*(00 + 11)(0 + 1)^*$, which describes the language accepted by our very first example NFA. Thompson's algorithm constructs the following 26-state monster:



Thompson's NFA for the regular expression $(0 + 1)^*(00 + 11)(0 + 1)^*$

Given this NFA as input, the incremental subset construction computes the following table, leading to a DFA with just nine states. Yeah, the ϵ -reaches get a bit ridiculous; unfortunately, this is typical for Thompson's NFA.

q'	ϵ -reach(q')	$q' \in A'$?	$\delta'(q', 0)$	$\delta'(q', 1)$
s	$sabdghim$		cj	en
cj	$sabdfghijklm$		cjl	en
en	$sabdfghmno$		cj	enp
cjl	$sabdfghijklmqrtuwz$	✓	$cjlv$	enx
enp	$sabdfghmnopqrtuwz$	✓	cjv	$enpx$
$cjlv$	$sabdfghijklmqrtuvwxyz$	✓	$cjlv$	enx
enx	$sabdfghmnopqrtuvwxyz$	✓	cjv	$enpx$
cjv	$sabdfghijkmr tuvwyz$	✓	$cjlv$	enx
$enpx$	$sabdfghmnopqrtuvwxyz$	✓	cjv	$enpx$



The DFA computed by the incremental subset algorithm from Thompson's NFA for $(0 + 1)^*(00 + 11)(0 + 1)^*$.

This DFA has far more states than necessary, intuitively because it keeps looking for 00 and 11 substrings even after it's already found one. After all, when Thompson's NFA finds a 00 and 11 substring, it doesn't kill all the other parallel threads, because it *can't*. NFAs often have significantly fewer states than equivalent DFAs, but that efficiency also makes them kind of stupid.

Glushkov's algorithm recursively computes the following values for the same regular expression

$$R = (0 + 1)^*(00 + 11)(0 + 1)^*:$$

$$\text{index}(R) = (1 + 2)^*(34 + 56)(7 + 8)^*$$

$$\text{symbols}(R) = 01001101$$


$$\text{has-}\varepsilon(R) = \text{FALSE}$$

$$\text{first}(\text{index}(R)) = \{1, 2, 3, 5\}$$

$$\text{last}(\text{index}(R)) = \{4, 6, 7, 8\}$$

$$\text{middle}(\text{index}(R)) = \{(1, 1), (1, 2), (2, 1), (2, 2), (1, 3), (1, 5), (2, 3), (2, 5), (3, 4), \\ (5, 6), (4, 7), (4, 8), (6, 7), (6, 8), (7, 7), (7, 8), (8, 7), (8, 8)\}$$

These values imply the nine-state NFA shown below. Careful readers should confirm that running the incremental subset construction on this NFA yields exactly the same DFA (with different state names) as it did for Thompson's NFA.



Glushkov's NFA for $(0 + 1)^*(00 + 11)(0 + 1)^*$

*4.9 Regular Expressions from NFAs: Han and Wood's Algorithm

The only component of Kleene's theorem we still have to prove is that every language accepted by a DFA or NFA is regular. As usual, it is actually easier to prove a stronger result. We consider a more general class of finite-state machines called *expression automata*, introduced by Yo-Sub Han and Derick Wood in 2005.² Formally, an expression automaton consists of the following components:

- A finite set Σ called the *input alphabet*
- Another finite set Q whose elements are called *states*
- A *start state* $s \in Q$
- A single *terminal state* $t \in Q \setminus \{s\}$
- A *transition function* $R: (Q \setminus \{t\}) \times (Q \setminus \{s\}) \rightarrow \text{Reg}(\Sigma)$, where $\text{Reg}(\Sigma)$ is the set of regular expressions over Σ

Less formally, an expression automaton is a directed graph that includes a directed edge $p \rightarrow q$ labeled with a regular expression $R(p \rightarrow q)$, from every vertex p to every vertex q (including $q = p$), where by convention, we require that $R(q \rightarrow s) = R(t \rightarrow q) = \emptyset$ for every vertex q .

²Yo-Sub Han* and Derick Wood. The generalization of generalized automata: Expression automata. International Journal of Foundations of Computer Science 16(3):499–510, 2005.

We say that string w *matches* a transition $p \rightarrow q$ if w matches the regular expression $R(p \rightarrow q)$. In particular, if $R(p \rightarrow q) = \emptyset$, then no string matches $p \rightarrow q$. More generally, w matches a sequence of states $q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_k$ if w matches the regular expression $R(q_0 \rightarrow q_1) \cdot R(q_1 \rightarrow q_2) \cdot \dots \cdot R(q_{k-1} \rightarrow q_k)$. Equivalently, w matches the sequence $q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_k$ if either

- $w = \varepsilon$ and the sequence has only one state ($k = 0$), or
- $w = xy$ for some string x that matches the regular expression $R(q_0 \rightarrow q_1)$ and some string y that matches the remaining sequence $q_1 \rightarrow \dots \rightarrow q_k$.

An expression automaton *accepts* any string that matches at least one sequence of states that starts at s and ends at t . The *language* of an expression automaton E is the set of all strings that E accepts.

Expression automata are nondeterministic. A single string could match several (even infinitely many) state sequences that start with s , and it could match each of those state sequences in several different ways. A string is accepted if *at least one* of the state sequences it matches ends at t . Conversely, a string might match *no* state sequences; all such strings are rejected.

Two special cases of expression automata are already familiar. First, every regular language is clearly the language of an expression automaton with exactly two states. Second, with only minor modifications, any DFA or NFA can be converted into an expression automaton with trivial transition expressions. Thompson's algorithm can be used to transform any expression automaton into an NFA, by recursively expanding any nontrivial transition. To complete the proof of Kleene's theorem, we show how to convert any expression automaton into a regular expression by repeatedly deleting vertices.

Lemma 4.3. *Every expression automaton accepts a regular language.*

Proof: Let $E = (Q, \Sigma, R, s, t)$ be an arbitrary expression automaton. Assume that any expression automaton with fewer states than E accepts a regular language. There are two cases to consider, depending on the number of states in Q :

- If $Q = \{s, t\}$, then trivially, E accepts the regular language $R(s \rightarrow t)$.
- On the other hand, suppose there is a state $q \in Q \setminus \{s, t\}$. We can modify the expression automaton so that state q is redundant and can be removed. Define a new transition function $R' : Q \times Q \rightarrow \text{Reg}(\Sigma)$ by setting

$$R'(p \rightarrow r) := R(p \rightarrow r) + R(p \rightarrow q)R(q \rightarrow q)^*R(q \rightarrow r).$$


With this modified transition function in place, any string w that matches the sequence $p \rightarrow q \rightarrow q \rightarrow \dots \rightarrow q \rightarrow r$ with any number of q 's also matches the single transition $p \rightarrow r$. Thus, by induction, if w matches a sequence of states, it also matches the subsequence obtained by removing all q 's. Let E' be the expression automaton with states $Q' = Q \setminus \{q\}$ that uses this modified transition function R' . This new automaton accepts exactly the same strings as the original automaton E . Because E' has fewer states than E , the inductive hypothesis implies E' accepts a regular language.

In both cases, we conclude that E accepts a regular language. □

This proof can be mechanically translated into an algorithm to convert any NFA—in particular, any DFA—into an equivalent regular expression. Given an NFA with n states (including s and

a), the algorithm iteratively removes $n - 2$ states, updating $O(n^2)$ transition expressions in each iteration. If the concatenation and Kleene star operations could be performed in constant time, the resulting algorithm would run in $O(n^3)$ time. However, in each iteration, the transition expressions grows in length by roughly a factor of 4 in the worst case, so the final expression has length $\Theta(4^n)$. If we insist on representing the expressions as explicit strings, the worst-case running time is actually $\Theta(4^n)$.

A figure on the next page shows this conversion algorithm in action for a simple DFA. First we convert the DFA into an expression automaton by adding new start and accept states and merging two transitions, and then we remove each of the three original states, updating the transition expressions between any remaining states at each iteration. For the sake of clarity, edges $p \rightarrow q$ with $R(p \rightarrow q) = \emptyset$ are omitted from the figures.



Converting a DFA into an equivalent regular expression.

Exercises

1. For each of the following NFAs, describe an equivalent DFA. (“Describe” does not necessarily mean “draw”!)



Half a dozen examples.

2. For each of the following regular expressions, draw an equivalent NFA.



Half a dozen examples.

3. For each of the following regular expressions, describe an equivalent DFA. (“Describe” does not necessarily mean “draw”!)



Half a dozen examples.

4. Let $L \subseteq \Sigma^*$ be an arbitrary regular language. Prove that the following languages are regular.

- (a) $ones(L) := \{w \in \mathbf{1}^* \mid |w| = |x| \text{ for some } x \in L\}$
- (b) $reverse(L) := \{w \in \Sigma^* \mid w^R \in L\}$. (Recall that w^R denotes the reversal of string w .)
- (c) $prefix(L) := \{x \in \Sigma^* \mid xy \in L \text{ for some } y \in \Sigma^*\}$
- (d) $suffix(L) := \{y \in \Sigma^* \mid xy \in L \text{ for some } x \in \Sigma^*\}$
- (e) $substring(L) := \{y \in \Sigma^* \mid xyz \in L \text{ for some } x, z \in \Sigma^*\}$
- (f) $cycle(L) := \{xy \mid x, y \in \Sigma^* \text{ and } yx \in L\}$
- (g) $prefmax(L) := \{x \in L \mid xy \in L \iff y = \varepsilon\}$.
- (h) $sufmin(L) := \{xy \in L \mid y \in L \iff x = \varepsilon\}$.
- (i) $everyother(L) := \{everyother(w) \mid w \in L\}$, where $everyother(w)$ is the subsequence of w containing every other symbol. For example, $everyother(\mathbf{EVERYOTHER}) = \mathbf{VROHR}$.
- (j) $rehtoyreve(L) := \{w \in \Sigma^* \mid everyother(w) \in L\}$.
- (k) $subseq(L) := \{x \in \Sigma^* \mid x \text{ is a subsequence of some } y \in L\}$
- (l) $superseq(L) := \{x \in \Sigma^* \mid \text{some } y \in L \text{ is a subsequence of } x\}$
- (m) $left(L) := \{x \in \Sigma^* \mid xy \in L \text{ for some } y \in \Sigma^* \text{ where } |x| = |y|\}$
- (n) $right(L) := \{y \in \Sigma^* \mid xy \in L \text{ for some } x \in \Sigma^* \text{ where } |x| = |y|\}$
- (o) $middle(L) := \{y \in \Sigma^* \mid xyz \in L \text{ for some } x, z \in \Sigma^* \text{ where } |x| = |y| = |z|\}$
- (p) $half(L) := \{w \in \Sigma^* \mid ww \in L\}$
- (q) $third(L) := \{w \in \Sigma^* \mid www \in L\}$
- (r) $reflect(L) := \{w \in \Sigma^* \mid ww^R \in L\}$
- * (s) $sqrt(L) := \{x \in \Sigma^* \mid xy \in L \text{ for some } y \in \Sigma^* \text{ such that } |y| = |x|^2\}$
- * (t) $log(L) := \{x \in \Sigma^* \mid xy \in L \text{ for some } y \in \Sigma^* \text{ such that } |y| = 2^{|x|}\}$
- * (u) $flog(L) := \{x \in \Sigma^* \mid xy \in L \text{ for some } y \in \Sigma^* \text{ such that } |y| = F_{|x|}\}$, where F_n is the n th Fibonacci number.

- *5. Let $L \subseteq \Sigma^*$ be an arbitrary regular language. Prove that the following languages are regular. [Hint: For each language, there is an accepting NFA with at most q^q states, where q is the number of states in some DFA that accepts L .]

- (a) $repeat(L) := \{w \in \Sigma^* \mid w^n \in L \text{ for some } n \geq 0\}$
- (b) $allreps(L) := \{w \in \Sigma^* \mid w^n \in L \text{ for every } n \geq 0\}$
- (c) $manyreps(L) := \{w \in \Sigma^* \mid w^n \in L \text{ for infinitely many } n \geq 0\}$
- (d) $fewreps(L) := \{w \in \Sigma^* \mid w^n \in L \text{ for finitely many } n \geq 0\}$
- (e) $powers(L) := \{w \in \Sigma^* \mid w^{2^n} \in L \text{ for some } n \geq 0\}$
- ★ (f) $whatthe_N(L) := \{w \in \Sigma^* \mid w^n \in L \text{ for some } n \in N\}$, where N is an **arbitrary** fixed set of non-negative integers. [Hint: You only have to prove that an accepting NFA exists; you don't have to describe how to construct it.]
6. For each of the following expression automata, describe an equivalent DFA *and* an equivalent regular expression.



Half a dozen examples.