

Course Project 2

Regular Expressions

CSE 30151 Spring 2025

Version of 2024-10-30

Due date 2025-02-28

In this project, you'll write a regular expression matcher similar to `grep`. This has three major steps: first, parse a regular expression into regular operations; second, execute the regular operations to create a NFA; third, run the NFA on input strings. Our linear-time NFA recognition algorithm is much faster than a naive regular expression matcher, and is in general much faster than Python and Perl's standard regular-expression matchers!

You will need a correct solution for CP1 to complete this project. If your CP1 doesn't work correctly (or you just weren't happy with it), you may use the official solution or another team's solution, as long as you properly cite your source.

Getting started

The project repository should include the following files (among others):


```
bin.{linux,darwin}/
  parse_re
  normalize_tree
  union_nfa
  concat_nfa
  star_nfa
  string_nfa
  re_to_nfa
  agrep
  compare_nfa
tests/
  test-cp2.sh
nonsolutions/
  bgrep.pl
  bgrep.py
cp2/
```

Please place the programs that you write into the `cp2/` subdirectory.

1 Parser

Note: Part 1 (this part) and part 2.1 can be done in parallel. Part 2.2 depends on both, and may require modifications to part 1.

In this first part, we'll write a parser for regular expressions. The goal is to input a regular expression like $(ab|)^*$ and output a tree like



In plain text, we write this tree as

```
star(union(concat(symbol("a"),symbol("b")),epsilon())).
```

1.1 Background

Define the following alphabets:

$$\begin{aligned}
 A &= \text{all printable, non-whitespace ASCII characters} \\
 \Sigma &= A \setminus \{ |, *, (,), \backslash, \& \} \\
 T &= \Sigma \cup \{ |, *, (,), \neg \}
 \end{aligned}$$

Alphabet Σ is the set of “ordinary” characters; our regular expression matcher will read strings over Σ . Alphabet T is the set of characters that regular expressions can contain; we'll explain \neg later.

Figure 1a shows a grammar for regular expressions. The terminal alphabet is T , and the non-terminal alphabet is $V = \{E, E', T, T', F, F', P\}$, with start symbol E . If this grammar looks bigger than the grammar you would have written, it's because this grammar has been designed to avoid *left recursion* (e.g., rules of the form $X \rightarrow X\beta$). This is because, for the style of parser we're going to build, left recursion causes infinite loops.¹

Parsers (for programming languages) are typically based on PDAs. If we convert the grammar to a PDA using the construction in Lemma 2.21, the resulting PDA is shown in Figure 1b. The input alphabet is T , and the stack alphabet is $\Gamma = V \cup T \cup \{\$\}$. The transitions involving a are actually many transitions, one for each $a \in \Sigma$. And \neg is an *endmarker* that must be appended to the input regular expression. Table 3 shows an example run of this PDA. For now, ignore the symbols `union` and `concat`, and the “semantic stack” column.

There are two problems with this PDA. First, it's nondeterministic, and we would rather implement a deterministic PDA (one where there's at most one transition that can be taken in any given configuration). Second, it will accept valid regular expressions and reject invalid ones, but we want it to output a tree (or eventually, an NFA).

For an example of nondeterminism, look at Table 3, row 5. The top stack symbol is P , so there are multiple transitions that the PDA could take: $\varepsilon, P \rightarrow (E)$ or $\varepsilon, P \rightarrow a$. Which transition is the right one? Since the next input symbol is a $($, you know it has to be $\varepsilon, P \rightarrow (E)$. For another example, look at Table 3, row 3. The top stack symbol is T , so there are two transitions that the PDA could take: $\varepsilon, T \rightarrow FT'$ or $\varepsilon, T \rightarrow \varepsilon$. Now, it's less obvious which transition is the right one.

¹For an explanation of how to eliminate left recursion, please see *Compilers and Language Design*, 2nd ed., Section 4.3.1–2.

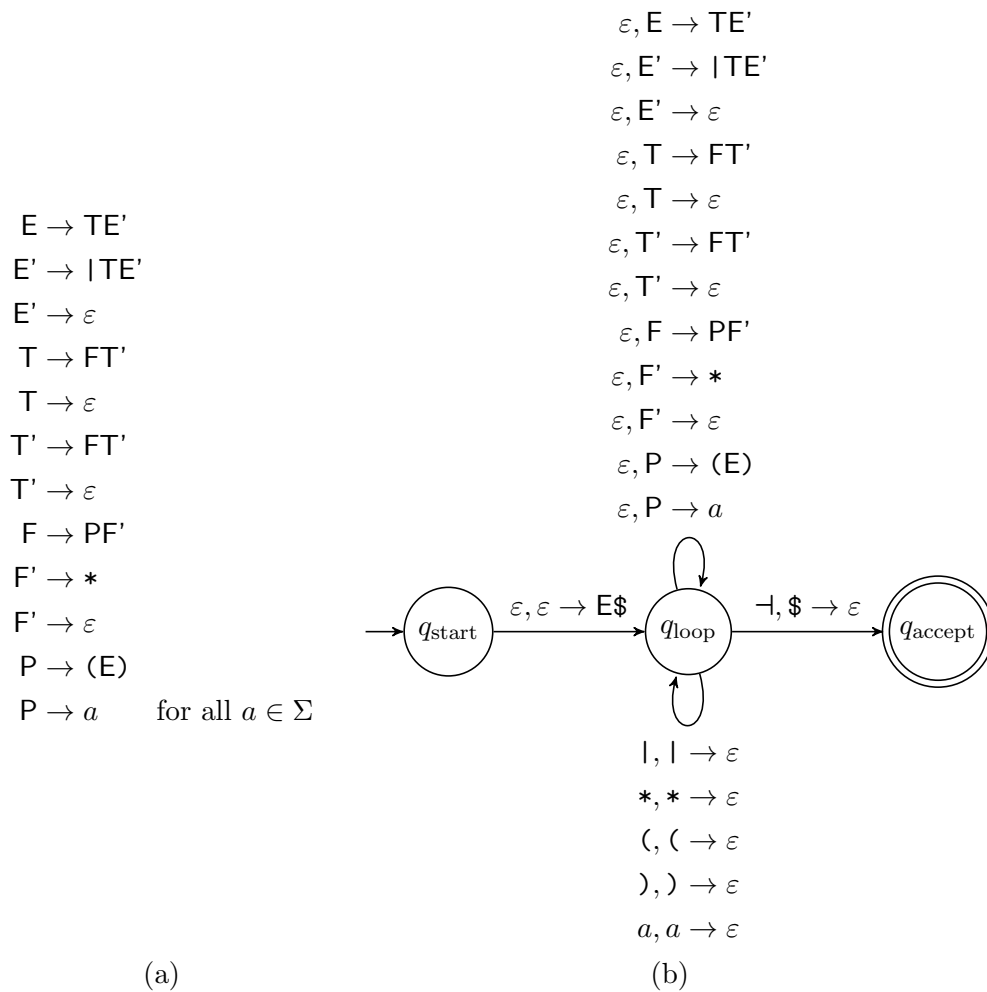


Figure 1: (a) A CFG for regular expressions, with start symbol E . Note that $|$ is a terminal symbol; it is not being used for writing two rules on one line. (b) The equivalent PDA, by the construction in Lemma 2.21.

read	input peek	parse stack		semantic stack	
		pop	push	pop	push
ε	T	E	TE'	ε	ε
ε	{ }	E'	T union E'	ε	ε
ε	$T \setminus \{ $	E'	ε	ε	ε
ε	$T \setminus \{ , \rangle, \neg\}$	T	FT'	ε	ε
ε	{ , \rangle, \neg}	T	ε	ε	epsilon()
ε	$T \setminus \{ , \rangle, \neg\}$	T'	F concat T'	ε	ε
ε	{ , \rangle, \neg}	T'	ε	ε	ε
ε	T	F	PF'	ε	ε
ε	{*}	F'	*	ε	ε
ε	$T \setminus \{*\}$	F'	ε	ε	ε
ε	{(}	P	(E)	ε	ε
ε	{a}	P	a	ε	ε
			ε	ε	ε
*		*	ε	α	star(α)
((ε	ε	ε
))	ε	ε	ε
a		a	ε	ε	symbol("a")
ε	T	union	ε	$\beta\alpha$	union(α, β)
ε	T	concat	ε	$\beta\alpha$	concat(α, β)

Table 2: Parse table. Every row is a transition from q_{loop} to itself. A row mentioning a stands for many transitions, one for each $a \in \Sigma$.

It turns out (for this grammar, not for all grammars) that you can always decide which transition to use by looking at the next input symbol. Have a look at Table 2. The “read” column and the “pop” and “push” columns under “parse stack” correspond to the three parts of a PDA transition that you’re familiar with. The “peek” column tells you how to use the next input symbol to decide which transition to use. In our example, the next input symbol is (, so we look in Table 2 for the transition that has (in the “peek” column, which is $\varepsilon, T \rightarrow FT'$. In general, each entry in the “peek” column is a set of symbols; a transition can only be used if the next unread input symbol is in the “peek” set.²

To compute semantics, we’re going to add a second stack, called the *semantic stack*; we’ll call the original stack the *parse stack*. The semantic stack (for now) stores trees. In Table 3 at step 12, the PDA reads an **a** and pops an **a** from the parse stack. At the same time, it pushes the semantics of **a**, which is `symbol("a")`, onto the semantic stack. This is because in Table 2, the row for $a, a \rightarrow \varepsilon$ says to pop ε and push `symbol("a")` onto the semantic stack. Later, at step 19, the PDA pops `symbol("b")`, then `symbol("a")` from the semantic stack, and pushes `concat(symbol("a"), symbol("b"))` onto the semantic stack. This is because in Table 2, the row for $\varepsilon, \text{concat} \rightarrow \varepsilon$ says to pop $\beta\alpha$ from the semantic stack (where α and β stand for any trees) and to push `concat(α, β)` onto the semantic stack. At the end of an accepting run, the semantic stack has exactly one tree; this is the semantics of the input regular expression.

1.2 Implementation

Write code to run the parser described above on a regular expression and construct a tree from it. You will want to write it in a modular way so that in Part 2.2, you can modify the parser to construct an NFA instead. To test your parser, write a program with the following command-line usage:

```
parse_re regexp
```

- *regexp*: a regular expression
- Output: string representation of the syntax tree for *regexp*

For example:

```
$ cp2/parse_re 'a'
symbol("a")
$ cp2/parse_re ''
epsilon()
$ cp2/parse_re '(a)'
symbol("a")
$ cp2/parse_re '()'
epsilon()
$ cp2/parse_re 'a*'
star(symbol("a"))
$ cp2/parse_re 'abc'
concat(concat(symbol("a"), symbol("b")), symbol("c"))
$ cp2/parse_re 'a|b|c'
union(union(symbol("a"), symbol("b")), symbol("c"))
```

²For an explanation of how to derive the “peek” sets, please see *Compilers and Language Design*, 2nd ed., Section 4.3.3.

	state	input	parse stack	semantic stack
1	q_{start}	(ab)* \rightarrow	ϵ	ϵ
2	q_{loop}	(ab)* \rightarrow	E\$	ϵ
3	q_{loop}	(ab)* \rightarrow	TE'\$	ϵ
4	q_{loop}	(ab)* \rightarrow	FT'E'\$	ϵ
5	q_{loop}	(ab)* \rightarrow	PF'T'E'\$	ϵ
6	q_{loop}	(ab)* \rightarrow	(E)F'T'E'\$	ϵ
7	q_{loop}	ab)* \rightarrow	E)F'T'E'\$	ϵ
8	q_{loop}	ab)* \rightarrow	TE')F'T'E'\$	ϵ
9	q_{loop}	ab)* \rightarrow	FT'E')F'T'E'\$	ϵ
10	q_{loop}	ab)* \rightarrow	PF'T'E')F'T'E'\$	ϵ
11	q_{loop}	ab)* \rightarrow	aF'T'E')F'T'E'\$	ϵ
12	q_{loop}	b)* \rightarrow	F'T'E')F'T'E'\$	symbol("a")
13	q_{loop}	b)* \rightarrow	T'E')F'T'E'\$	symbol("a")
14	q_{loop}	b)* \rightarrow	F concat T'E')F'T'E'\$	symbol("a")
15	q_{loop}	b)* \rightarrow	PF' concat T'E')F'T'E'\$	symbol("a")
16	q_{loop}	b)* \rightarrow	bF' concat T'E')F'T'E'\$	symbol("a")
17	q_{loop})* \rightarrow	F' concat T'E')F'T'E'\$	symbol("b") symbol("a")
18	q_{loop})* \rightarrow	concat T'E')F'T'E'\$	symbol("b") symbol("a")
19	q_{loop})* \rightarrow	T'E')F'T'E'\$	concat(symbol("a"),symbol("b"))
20	q_{loop})* \rightarrow	E')F'T'E'\$	concat(symbol("a"),symbol("b"))
21	q_{loop})* \rightarrow	T union E')F'T'E'\$	concat(symbol("a"),symbol("b"))
22	q_{loop})* \rightarrow	T union E')F'T'E'\$	concat(symbol("a"),symbol("b"))
23	q_{loop})* \rightarrow	union E')F'T'E'\$	epsilon() concat(symbol("a"),symbol("b"))
24	q_{loop})* \rightarrow	E')F'T'E'\$	union(concat(symbol("a"),symbol("b")),epsilon())
25	q_{loop})* \rightarrow)F'T'E'\$	union(concat(symbol("a"),symbol("b")),epsilon())
26	q_{loop}	* \rightarrow	F'T'E'\$	union(concat(symbol("a"),symbol("b")),epsilon())
27	q_{loop}	* \rightarrow	*T'E'\$	union(concat(symbol("a"),symbol("b")),epsilon())
28	q_{loop}	\rightarrow	T'E'\$	star(union(concat(symbol("a"),symbol("b")),epsilon()))
29	q_{loop}	\rightarrow	E'\$	star(union(concat(symbol("a"),symbol("b")),epsilon()))
30	q_{loop}	\rightarrow	\$	star(union(concat(symbol("a"),symbol("b")),epsilon()))
31	q_{accept}	ϵ	ϵ	star(union(concat(symbol("a"),symbol("b")),epsilon()))

Table 3: Example run of the parser.

```
$ cp2/parse_re '||'
union(union(epsilon(),epsilon()),epsilon())
```

Test your program by running `tests/test-cp2.sh`.

2 Converter

2.1 Regular operations

Write a function that creates an NFA that accepts exactly one string. To test it, write a program with the following command-line usage:

```
string_nfa w
```

- w : a string (possibly empty)
- Output: an NFA recognizing the language $\{w\}$

Write functions that perform the three regular operations, using the constructions given in the book, and programs to test them:

```
union_nfa M1 M2
```

- M_1, M_2 : NFAs
- Output: NFA recognizing language $L(M_1) \cup L(M_2)$

```
concat_nfa M1 M2
```

- M_1, M_2 : NFAs
- Output: NFA recognizing language $L(M_1) \circ L(M_2)$

```
star_nfa M
```

- M : an NFA
- Output: NFA recognizing language $L(M)^*$

Test all of these programs by running `tests/test-cp2.sh`.

2.2 Building the NFA

Modify your parser from Part 1 so that, instead of constructing a tree, it constructs an NFA. That is, whenever it used to construct a node `epsilon()`, it constructs an NFA that accepts only ε ; whenever it used to construct a node `symbol("a")`, it constructs an NFA that accepts only a ; whenever it used to construct a `union` node, it constructs the union of the two child NFAs, and so on. Beware that if you choose to construct a tree and then convert the tree to an NFA, you may get an error for exceeding the maximum recursion depth. Test your modified parser by writing a program with the following command-line usage:

```
re_to_nfa regexp
```

- `regexp`: Regular expression
- Output: NFA M equivalent to `regexp`

Test your program using `tests/test-cp2.sh`.

3 Putting it together

Finally, combine your regular expression converter with your NFA simulator from CP1 to write a `grep` replacement, called `agrep` (for “automaton-based `grep`”), that has the following command-line usage:

`agrep regexp`

- `regexp`: regular expression
- Input: strings (one per line)
- Output: the input strings that match `regexp`

Note that unlike `grep`, the regular expression must match the entire line, not just part of the line. Test your program by running `tests/test-cp2.sh`.

The test script also tests the time complexity of `agrep`. This test is the same as in CP1, but now we can say a bit more about it. For various values of n , it creates the regular expression $(|a)^n$ and tries to match it against the string $a^n b$, using your `agrep`. You can also try timing `nonsolutions/bgrep.pl` and `nonsolutions/bgrep.py`, use Perl and Python’s standard regular expression engines, by making `cp2/agrep` a symlink to either of them. How do they compare to yours?

Submission instructions

Your code should build and run on `studentnn.cse.nd.edu`. The automatic tester will clone your repository, change to its root directory, run `make -C cp2`, and then run `tests/test-cp2.sh`. You’re advised to try all of the above steps and ensure that all tests pass.

To submit your work:

1. Push your repository to GitHub.
2. In GitHub, create a new release by clicking on “Releases,” then “Draft a new release.”
3. Fill in “Release title” with `cp2` if you’re submitting the whole assignment, `cp2-1` if you’re submitting part 1, `cp2-2` if you’re submitting part 2, etc.
4. Click on “Choose a tag,” then type the same name you used for the release title, then “Create new tag: cp2... on publish.”
5. Finally, click “Publish Release.”

Rubric

Part 1	
parsing	3
building syntax tree	3
parse_re	3
Part 2	
string_nfa	3
union_nfa	3
concat_nfa	3
star_nfa	3
re_to_nfa	3
Part 3 (agrep)	
correctness	3
time complexity	3
Total	<hr/> 30