


## 6.4 Pattern Matching



- ▶ regular expressions
- ▶ REs and NFAs
- ▶ NFA simulation
- ▶ NFA construction
- ▶ applications

Algorithms in Java, 4th Edition · Robert Sedgwick and Kevin Wayne · Copyright © 2008 · April 14, 2009 9:29:44 AM

- ▶ regular expressions
- ▶ NFAs
- ▶ NFA simulation
- ▶ NFA construction
- ▶ applications

2

### Pattern matching

**Substring search.** Find occurrences of a single string in text.

**Pattern matching.** Find occurrences of one of a **specified set** of strings in text.

Ex. [genomics]

- Fragile X syndrome is a common cause of mental retardation.
- Human genome contains triplet repeats of `CGG` or `AGG`, bracketed by `GCG` at the beginning and `CTG` at the end.
- Number of repeats is variable, and correlated with syndrome.

pattern `GCG (CGG | AGG) *CTG`

text `GCGCGTGTGTGCGAGAGAGTGGGTTTAAAGCTGCGCGGAGGCGGCTGCGCGGAGGCTG`

3

### Pattern matching: applications

**Test if a string matches some pattern.**

- Process natural language.
- Scan for virus signatures.
- Access information in digital libraries.
- Filter text (spam, NetNanny, Carnivore, malware).
- Validate data-entry fields (dates, email, URL, credit card).
- Search for markers in human genome using PROSITE patterns.

**Parse text files.**

- Compile a Java program.
- Crawl and index the Web.
- Read in data stored in ad hoc input file format.
- Automatically create Java documentation from Javadoc comments.

4

## Regular expressions

A **regular expression** is a notation to specify a (possibly infinite) set of strings.

↑  
a "language"

operation	example RE	in language	not in language
concatenation	<b>AABAAB</b>	AABAAB	every other string
or	<b>AA   BAAB</b>	AA BAAB	every other string
closure	<b>AB*A</b>	AA ABBBBBBBA	AB ABABA
parentheses	<b>A(A B)AAB</b>	AAAAB ABAAB	every other string
	<b>(AB)*A</b>	A ABABABABABA	AA ABBA

5

## Regular expression shortcuts

Additional operations are often added for convenience.

Ex. **[A-E]+** is shorthand for **(A|B|C|D|E)(A|B|C|D|E)\***

operation	example RE	in language	not in language
wildcard	<b>.U.U.U.</b>	CUMULUS JUGULUM	SUCCUBUS TUMULTUOUS
at least 1	<b>A(BC)+DE</b>	ABCDE ABCBCDE	ADE BCDE
character classes	<b>[A-Za-z][a-z]*</b>	word Capitalized	camelCase 4illegal
exactly k	<b>[0-9]{5}-[0-9]{4}</b>	08540-1321 19072-5541	111111111 166-54-111
complement	<b>[^AEIOU]{6}</b>	RHYTHM	DECADE

6

## Regular expression examples

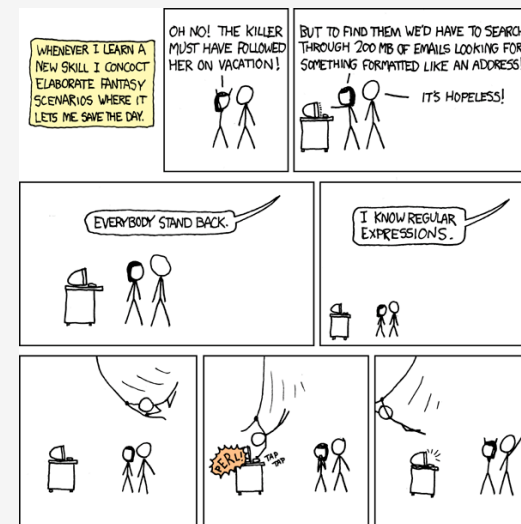
Notation is surprisingly expressive

regular expression	in set	not in set
<b>. *SPB. *</b> <i>(contains the trigraph spb)</i>	RASPBERRY CRISPBREAD	SUBSPACE SUBSPECIES
<b>[0-9]{3}-[0-9]{2}-[0-9]{4}</b> <i>(Social Security numbers)</i>	166-11-4433 166-45-1111	11-55555555 8675309
<b>[a-z]+@[a-z]+\.[a-z]+(edu com)</b> <i>(valid email addresses)</i>	wayne@princeton.edu rs@princeton.edu	spam@nowhere
<b>[\$_A-Za-z][\$_A-Za-z0-9]*</b> <i>(valid Java identifiers)</i>	ident3 PatternMatcher	3a ident#3

and plays a well-understood role in the theory of computation.

7

## Regular expressions to the rescue



<http://xkcd.com/208/>

8



- › regular expressions
- › **NFAs**
- › NFA simulation
- › NFA construction
- › applications

### Pattern matching implementation: basic plan (first attempt)

Overview is the same as for KMP !

- No backup in text stream.
- Linear-time guarantee.



Ken Thompson

Underlying abstraction:

Deterministic finite state automata (DFA)

Basic plan.

- Build DFA from RE.
- Simulate DFA with text as input.



Bad news: Basic plan is infeasible (DFA may have exponential number of states)

### Pattern matching implementation: basic plan (revised)

Overview is similar to KMP

- No backup in text stream.
- Quadratic-time guarantee (linear-time typical).



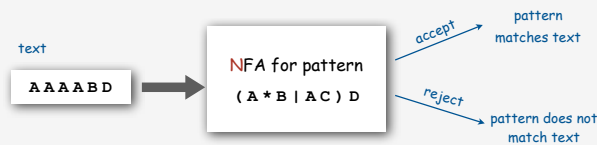
Ken Thompson

Underlying abstraction:

Nondeterministic finite state automata (NFA)

Basic plan.

- Build NFA from RE.
- Simulate NFA with text as input.




### Nondeterministic finite-state automata

Pattern matching NFA.

- Pattern enclosed in parentheses.
- One state per pattern character (start = 0, accept = M).
- Red  $\epsilon$ -transition (change state, but don't scan input).
- Black match transition (change state and scan to next char).
- Accept if any sequence of transitions ends in accept state.



Nondeterminism.

- one view: machine can guess the proper sequence of state transitions
- another view: sequence is a proof that the machine accepts the text




### Nondeterministic finite-state automata

Ex. Is **AAAABD** recognized by NFA?




### Nondeterministic finite-state automata

Ex. Is **AAAABD** recognized by NFA?




YES!

Note: any sequence of legal transitions that ends in 11 is a proof.




### Nondeterministic finite-state automata

Ex. Is **AAAAC** recognized by NFA?



NO

Note: this is not a complete proof  
(need to mention the infinite number of sequences involving  $\epsilon$ -transitions between 2 and 3)





### Nondeterminism

Q. How to determine whether a string is recognized by an automaton?

DFA. Deterministic  $\Rightarrow$  exactly one applicable transition.

NFA. Nondeterministic  $\Rightarrow$  can be several applicable transitions;  
need to select the right one!



Q. How to simulate NFA?

A. Systematically consider all possible transition sequences.

## Pattern matching implementation: basic plan (revised)

Overview is similar to KMP

- No backup in text stream.
- Quadratic-time guarantee (linear-time typical).



Ken Thompson

Underlying abstraction:

Nondeterministic finite state automata (NFA)

Basic plan.

- Build NFA from RE.
- Simulate NFA with text as input.



- ▶ regular expressions
- ▶ NFAs
- ▶ **NFA simulation**
- ▶ NFA construction
- ▶ applications


## NFA representation

State names. Integers from 0 to  $m$ .

Match-transitions. Keep regular expression in array  $\text{re}[]$ .

$\epsilon$ -transitions. Store in a digraph  $g$ .


- $0 \rightarrow 1, 1 \rightarrow 2, 1 \rightarrow 6, 2 \rightarrow 3, 3 \rightarrow 2, 3 \rightarrow 4, 5 \rightarrow 8, 8 \rightarrow 9, 10 \rightarrow 11$



## NFA simulation

Q. How to efficiently simulate an NFA?

A. Maintain set of **all** possible states that NFA could be in after reading in the first  $i$  text characters.



Q. How to perform reachability?

## Digraph reachability

Find all vertices reachable from a given **set** of vertices.

```

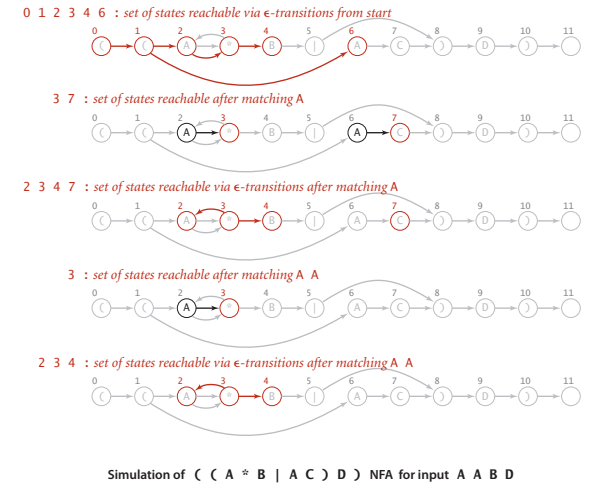
public class DFS
{
    private SET<Integer> marked;
    private Digraph G;

    public DFS(Digraph G)
    { this.G = G; }

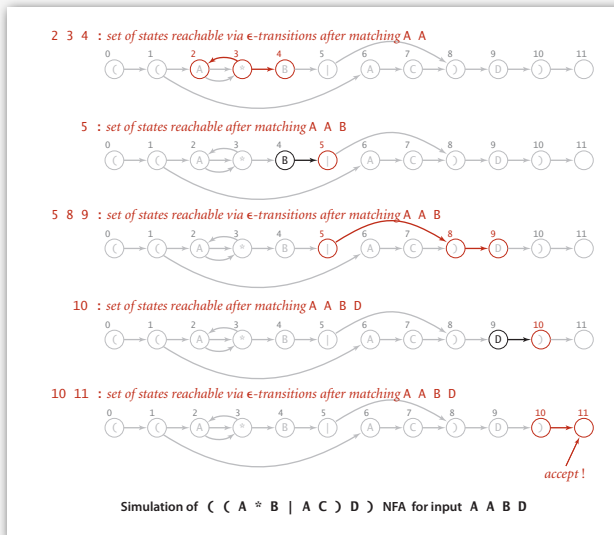
    private void search(int v)
    {
        marked.add(v);
        for (int w : G.adj(v))
            if (!marked.contains(w)) search(w);
    }

    public SET<Integer> reachable(SET<Integer> s)
    {
        marked = new SET<Integer>();
        for (int v : s) search(v);
        return marked;
    }
}
    
```

## NFA simulation example



## NFA simulation example



## NFA simulation: Java implementation

```

public boolean recognizes(char[] txt)
{
    DFS dfs = new DFS(G);

    SET<Integer> pc = new dfs.reachable(0); // states reachable from start by  $\epsilon$ -transitions

    for (int i = 0; i < txt.length; i++)
    {
        SET<Integer> match = new SET<Integer>(); // all possible states after scanning past txt[i]
        for (int v : pc) {
            if (v == M) continue;
            if ((re[v] == txt[i]) || re[v] == '.')
                match.add(v+1);
        }


        pc = dfs.reachable(match); // follow  $\epsilon$ -transitions
    }

    return pc.contains(M); // accept if you can end in state M
}
    
```

### NFA simulation: analysis

**Proposition 1.** Determining whether an N-character text string is recognized by the NFA corresponding to an M-character pattern takes time proportional to NM in the worst case.


**Pf.** For each of the N text characters, we iterate through a set of states of size no more than M and run DFS on the graph of  $\epsilon$ -transitions.  
 (The construction we consider ensures the number of edges is at most M.)



- › regular expressions
- › NFAs
- › NFA simulation
- › **NFA construction**
- › applications

### Building an NFA corresponding to an RE

**States.** Include a state for each symbol in the RE, plus an accept state.




### Building an NFA corresponding to an RE

**Concatenation.** Add match-transition edge from state corresponding to letters in the alphabet to next state.

**Alphabet.** A B C D


**Metacharacters.** ( ) . \* |





### Building an NFA corresponding to an RE



**Parentheses.** Add  $\epsilon$ -transition edge from parentheses to next state.



33

### Building an NFA corresponding to an RE



**Closure.** Add three  $\epsilon$ -transition edges for each \* operator.



34

### Building an NFA corresponding to an RE

**Or.** Add two  $\epsilon$ -transition edges for each | operator.



35


### NFA construction: implementation

**Goal.** Write a program to build the  $\epsilon$ -transition digraph.

**Challenge.** Need to remember left parentheses to implement closure and or; need to remember | to implement or.


**Solution.** Maintain a stack.

- Left parenthesis: push onto stack.
- | symbol: push onto stack.
- Right parenthesis: add edges for closure and or.




36

### NFA construction: example



### NFA construction: example



### NFA construction: Java implementation

```

public NFA(int R, char[] re)
{
    Stack<Integer> ops = new Stack<Integer>();
    this.re = re;
    M = re.length;
    G = new Digraph(M+1);
    for (int i = 0; i < M; i++) {
        int lp = i;

        if (re[i] == '(' || re[i] == '|') ops.push(i);

        else if (re[i] == ')') {
            int or = ops.pop();
            if (re[or] == '|') {
                lp = ops.pop();
                G.addEdge(lp, or+1);
                G.addEdge(or, i);
            }
        }

        if (i < M-1 && re[i+1] == '*') {
            G.addEdge(lp, i+1);
            G.addEdge(i+1, lp);
        }

        if (re[i] == '(' || re[i] == '*' || re[i] == ')')
            G.addEdge(i, i+1);
    }
}
    
```

left parentheses and |

or

closure (needs lookahead)


metasymbols

39

### NFA construction: analysis

**Proposition 2.** Building the NFA corresponding to an  $M$ -character pattern takes time and space proportional to  $M$  in the worst case.

**Pf.** For each of the  $M$  characters in the pattern, we add one or two  $\epsilon$ -transitions and perhaps execute one or two stack operations.



- › regular expressions
- › NFAs
- › NFA simulation
- › NFA construction
- › applications

### Generalized regular expression print

**Grep.** Takes a pattern as a command-line argument and prints the lines from standard input having some substring that is matched by the pattern.


```
public class GREP
{
    public static void main(String[] args)
    {
        String regexp = "(.*" + args[0] + ".*)";
        char[] re = regexp.toCharArray();
        while (!StdIn.isEmpty())
        {
            String line = StdIn.readLine();
            char[] txt = line.toCharArray();
            NFA nfa = new NFA(256, re);
            if (nfa.recognizes(txt))
                StdOut.println(line);
        }
    }
}
```

find lines containing RE as a substring

**Bottom line.** Worst-case for grep (proportional to MN) is the same as for elementary exact substring match.

### Typical grep application

#### Crossword puzzle



dictionary (standard in UNIX) also on booksite

```
% more words.txt
a
aback
abacus
abalone
abandon
...
% grep s..ict.. words.txt
constrictor
stricter
stricture
```

### Industrial-strength grep implementation

To complete the implementation:

- Add character classes.
- Handling metacharacters.
- Add capturing capabilities.
- Extend the closure operator.
- Error checking and recovery.
- Greedy vs. reluctant matching.

Ex. Which substring(s) should be matched by the RE `<blink>.*</blink> ?`



## Regular expressions in other languages

Broadly applicable programmer's tool.

- Originated in Unix in the 1970s
- Many languages support extended regular expressions.
- Built into grep, awk, emacs, Perl, PHP, Python, JavaScript.

```
% grep NEWLINE */*.java ← print all lines containing NEWLINE which occurs
                           in any file with a .java extension

% egrep '^[qwertyuiop]*[zxcvbnm]*$' dict.txt | egrep '.....'
```

PERL. Practical Extraction and Report Language.

```
% perl -p -i -e 's|from|to|g' input.txt ← replace all occurrences of from
                                         with to in the file input.txt

% perl -n -e 'print if /^[A-Za-z][a-z]*$/' dict.txt ← print all uppercase words
  ↑
do for each line
```

45

## Regular expressions in Java

Validity checking. Is input in the set described by the re?

Java string library. Use `input.matches(re)` for basic RE matching.

```
public class Validate
{
    public static void main(String[] args)
    {
        String re = args[0];
        String input = args[1];
        StdOut.println(input.matches(re));
    }
}
```

```
% java Validate "[$_A-Za-z][$_A-Za-z0-9]*" ident123 ← legal Java identifier
true

% java Validate "[a-z]+@[a-z]+\.(edu|com)" rs@cs.princeton.edu ← valid email address
true                                                         (simplified)

% java Validate "[0-9]{3}-[0-9]{2}-[0-9]{4}" 166-11-4433 ← Social Security number
true
```

46

## Harvesting information

Goal. Print all substrings of input that match a RE.

```
% java Harvester "gcg(cgg|agg)*ctg" chromosomeX.txt
gcgcggcggcggcggcggcgtg
gcgctg
gcgctg
gcgcggcggcggcggcggcggcgtg
                                     ↑
                                     harvest patterns from DNA

                                     ↓
                                     harvest links from website

% java Harvester "http://(\\w+\\.)*\\w+" http://www.cs.princeton.edu
http://www.princeton.edu
http://www.google.com
http://www.cs.princeton.edu/news
```

47

## Harvesting information

RE pattern matching is implemented in Java's `Pattern` and `Matcher` classes.

```
import java.util.regex.Pattern;
import java.util.regex.Matcher;

public class Harvester
{
    public static void main(String[] args)
    {
        String re = args[0];
        In in = new In(args[1]);
        String input = in.readAll();
        Pattern pattern = Pattern.compile(re);
        Matcher matcher = pattern.matcher(input);
        while (matcher.find())
            StdOut.println(matcher.group());
    }
}
```

`compile()` creates a `Pattern` (NFA) from RE  
`matcher()` creates a `Matcher` (NFA simulator) from NFA and text  
`find()` looks for the next match  
`group()` returns the substring most recently found by `find()`

48

## Algorithmic complexity attacks

Warning. Typical implementations do not guarantee performance!

← grep, Java, Perl

```
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 1.6 seconds
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 3.7 seconds
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 9.7 seconds
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 23.2 seconds
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 62.2 seconds
% java Validate "(a|aa)*b" aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaac 161.6 seconds
```

SpamAssassin regular expression.

```
% java RE "[a-z]+@[a-z]+([a-z\.\.]+\.)+[a-z]+" spammer@x.....
```

- Takes exponential time on pathological email addresses.
- Troublemaker can use such addresses to DOS a mail server.

49

## Not-so-regular expressions

Back-references.

- \1 notation matches sub-expression that was matched earlier.
- Supported by typical RE implementations.

```
% java Harvester "\b(.+)\1\b" dictionary.txt
beriberi
couscous
```

↑ word boundary

Some non-regular languages.

- Set of strings of the form ww for some string w: beriberi.
- Set of bitstrings with an equal number of 0s and 1s: 01110100.
- Set of Watson-Crick complemented palindromes: atttcggaat.

Remark. Pattern matching with back-references is intractable.

50

## Context

Abstract machines, languages, and nondeterminism.

- basis of the theory of computation
- intensively studied since the 1930s
- basis of programming languages

Compiler. A program that translates a program to machine code.

- KMP string  $\Rightarrow$  DFA.
- grep RE  $\Rightarrow$  NFA.
- javac Java language  $\Rightarrow$  Java byte code.

	KMP	grep	Java
pattern	string	RE	program
parser	unnecessary	check if legal	check if legal
compiler output	DFA	NFA	byte code
simulator	DFA simulator	NFA simulator	JVM

51

## Summary of pattern-matching algorithms

Programmer.

- Implement exact pattern matching via DFA simulation.
- Implement RE pattern matching via NFA simulation.

Theoretician.

- RE is a compact description of a set of strings.
- NFA is an abstract machine equivalent in power to RE.
- DFAs and REs have limitations.

You. Practical application of core CS principles.

Example of essential paradigm in computer science.

- Build intermediate abstractions.
- Pick the right ones!
- Solve important practical problems.

52