## Lexical Analysis

## Announcements

- Programming Assignment 1 Out
- Due Monday, July 9 at 11:59 PM.
- Four handouts (all available online):
- Decaf Specification
- Lexical Analysis
- Intro to flex
- Programming Assignment 1


## Where We Are

 Code

$$
\begin{gathered}
\text { while (ip < z) } \\
\text { ++ip; }
\end{gathered}
$$



$$
\begin{gathered}
\text { while (ip }<~ z) \\
++i p ;
\end{gathered}
$$



while (ip $<z$ )
++ip;

$$
\text { do }[\text { for }]=\text { new } 0 ;
$$

| $\mathbf{d}$ | $\mathbf{o}$ | $[$ | $\mathbf{f}$ | $\mathbf{o}$ | $\mathbf{r}$ | $]$ |  | $=$ |  | $\mathbf{n}$ | $\mathbf{e}$ | $\mathbf{w}$ |  | 0 | i |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

$$
\text { do }[\text { for }]=\text { new } 0 ;
$$


do[for] = new 0;


## Scanning a Source File



## Scanning a Source File

|  | w h | i | 1 | e | ( |  | 1 | 3 | 7 | $<$ | i | , | + i |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

## Scanning a Source File



## Scanning a Source File



## Scanning a Source File



## Scanning a Source File



## Scanning a Source File



## Scanning a Source File



T_While

## Scanning a Source File



The piece of the original program from which we made the token is called a lexeme.

T_While

This is called a token. You can think of it as an enumerated type representing what logical entity we read out of the source code.

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While
Sometimes we will discard a lexeme rather than storing it for later use. Here, we ignore whitespace, since it has no bearing on the meaning of the program.

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File

| w | w h | i | 1 | e | ( |  |  | 3 | 7 | $<$ | i |  |  | + |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

T_While

## Scanning a Source File

|  | w h | i | 1 | e | ( |  | 1 | 3 | 7 | $<$ | i |  |  | + |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

T_While

## Scanning a Source File



T_While

## Scanning a Source File



T_While

## Scanning a Source File



| T_While | T_IntConst |
| :---: | :---: |
| 137 |  |

## Scanning a Source File



T_While (T_IntConst \begin{tabular}{|c}

| Some tokens can have |
| :---: |
| attributes that store |
| extra information about |
| the token. Here we |
| store which integer is |
| represented. | <br>

\hline
\end{tabular}

## Goals of Lexical Analysis

- Convert from physical description of a program into sequence of of tokens.
- Each token represents one logical piece of the source file - a keyword, the name of a variable, etc.
- Each token is associated with a lexeme.
- The actual text of the token: "137," "int," etc.
- Each token may have optional attributes.
- Extra information derived from the text - perhaps a numeric value.
- The token sequence will be used in the parser to recover the program structure.


## Choosing Tokens

## What Tokens are Useful Here?

for (int $k=0 ; k<m y A r r a y[5] ;++k)\{$ cout << k << endl;
\}

## What Tokens are Useful Here?

for (int $k=0 ; k<m y A r r a y[5] ;++k)\{$ cout << k << endl;
\}


## What Tokens are Useful Here?

for (int $k=0 ; k<m y A r r a y[5] ;++k)\{$ cout << k << endl;
\}


Identifier
IntegerConstant

## Choosing Good Tokens

- Very much dependent on the language.
- Typically:
- Give keywords their own tokens.
- Give different punctuation symbols their own tokens.
- Group lexemes representing identifiers, numeric constants, strings, etc. into their own groups.
- Discard irrelevant information (whitespace, comments)


## Scanning is Hard

- FORTRAN: Whitespace is irrelevant

DO 5 I = 1,25
DO 5 I = 1.25

## Scanning is Hard

- FORTRAN: Whitespace is irrelevant

$$
\begin{aligned}
\text { DO } 5 \text { I } & =1,25 \\
\text { DO5I } & =1.25
\end{aligned}
$$

## Scanning is Hard

- FORTRAN: Whitespace is irrelevant

$$
\begin{aligned}
\text { DO } 5 \text { I } & =1,25 \\
\text { DO5I } & =1.25
\end{aligned}
$$

- Can be difficult to tell when to partition input.


## Scanning is Hard

- C++: Nested template declarations

vector<vector<int>> myVector

## Scanning is Hard

- C++: Nested template declarations
vector < vector < int >> myVector


## Scanning is Hard

- C++: Nested template declarations
(vector < (vector < (int >> myVector)))


## Scanning is Hard

- C++: Nested template declarations
(vector < (vector < (int >> myVector)))
- Again, can be difficult to determine where to split.


## Scanning is Hard

- PL/1: Keywords can be used as identifiers.


## Scanning is Hard

- PL/1: Keywords can be used as identifiers.

IF THEN THEN THEN = ELSE; ELSE ELSE = IF

## Scanning is Hard

- PL/1: Keywords can be used as identifiers.

IF THEN THEN THEN = ELSE; ELSE ELSE = IF

## Scanning is Hard

- PL/1: Keywords can be used as identifiers.

```
IF THEN THEN THEN = ELSE; ELSE ELSE = IF
```

- Can be difficult to determine how to label lexemes.


## Challenges in Scanning

- How do we determine which lexemes are associated with each token?
- When there are multiple ways we could scan the input, how do we know which one to pick?
- How do we address these concerns efficiently?


## Associating Lexemes with Tokens

## Lexemes and Tokens

- Tokens give a way to categorize lexemes by what information they provide.
- Some tokens might be associated with only a single lexeme:
- Tokens for keywords like if and while probably only match those lexemes exactly.
- Some tokens might be associated with lots of different lexemes:
- All variable names, all possible numbers, all possible strings, etc.


## Sets of Lexemes

- Idea: Associate a set of lexemes with each token.
- We might associate the "number" token with the set $\{0,1,2, \ldots, 10,11,12, \ldots\}$
- We might associate the "string" token with the set \{ "", "a", "b", "c", ... \}
- We might associate the token for the keyword while with the set \{ while \}.

How do we describe which (potentially infinite) set of lexemes is associated with each token type?

## Formal Languages

- A formal language is a set of strings.
- Many infinite languages have finite descriptions:
- Define the language using an automaton.
- Define the language using a grammar.
- Define the language using a regular expression.
- We can use these compact descriptions of the language to define sets of strings.
- Over the course of this class, we will use all of these approaches.


## Regular Expressions

- Regular expressions are a family of descriptions that can be used to capture certain languages (the regular languages).
- Often provide a compact and humanreadable description of the language.
- Used as the basis for numerous software systems, including the flex tool we will use in this course.


## Atomic Regular Expressions

- The regular expressions we will use in this course begin with two simple building blocks.
- The symbol $\boldsymbol{\varepsilon}$ is a regular expression matches the empty string.
- For any symbol a, the symbol a is a regular expression that just matches a.


## Compound Regular Expressions

- If $\mathrm{R}_{1}$ and $\mathrm{R}_{2}$ are regular expressions, $\mathbf{R}_{\mathbf{1}} \mathbf{R}_{\mathbf{2}}$ is a regular expression represents the concatenation of the languages of $\mathrm{R}_{1}$ and $\mathrm{R}_{2}$.
- If $R_{1}$ and $R_{2}$ are regular expressions, $\mathbf{R}_{\mathbf{1}} \mid \mathbf{R}_{\mathbf{2}}$ is a regular expression representing the union of $R_{1}$ and $R_{2}$.
- If R is a regular expression, $\mathbf{R}^{*}$ is a regular expression for the Kleene closure of R.
- If R is a regular expression, ( $\mathbf{R}$ ) is a regular expression with the same meaning as R .


## Operator Precedence

- Regular expression operator precedence is

$$
\begin{gathered}
(\mathrm{R}) \\
\mathrm{R}^{*} \\
\mathrm{R}_{1} \mathrm{R}_{2} \\
\mathrm{R}_{1} \mid \mathrm{R}_{2}
\end{gathered}
$$

- So $\mathbf{a b *} \mathbf{c} \mid \mathbf{d}$ is parsed as $\left(\left(\mathbf{a}\left(\mathbf{b}^{*}\right)\right) \mathbf{c}\right) \mid \mathbf{d}$


## Simple Regular Expressions

- Suppose the only characters are 0 and 1 .
- Here is a regular expression for strings containing 00 as a substring:

$$
(0 \text { | 1)* } 00(0 \text { | 1)* }
$$

## Simple Regular Expressions

- Suppose the only characters are 0 and 1 .
- Here is a regular expression for strings containing 00 as a substring:

$$
(0 \mid 1)^{*} 00(0 \mid 1)^{*}
$$

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings containing 00 as a substring:

$$
(0 \mid 1)^{*} 00(0 \mid 1)^{*}
$$

11011100101
0000
11111011110011111

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings containing 00 as a substring:

$$
(0 \mid 1)^{*} 00(0 \mid 1)^{*}
$$

11011100101
0000
11111011110011111

## Simple Regular Expressions

- Suppose the only characters are 0 and 1 .
- Here is a regular expression for strings of length exactly four:


## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings of length exactly four:


## (0|1)(0|1)(0|1)(0|1)

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings of length exactly four:


## (0|1)(0|1)(0|1)(0|1)

## Simple Regular Expressions

- Suppose the only characters are 0 and 1 .
- Here is a regular expression for strings of length exactly four:


## (0|1)(0|1)(0|1)(0|1)

0000
1010
1111
1000

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings of length exactly four:


## (0|1)(0|1)(0|1)(0|1)

0000
1010
1111
1000

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings of length exactly four:


## (0|1)\{4\}

0000
1010
1111
1000

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings of length exactly four:


## (0|1)\{4\}

0000
1010
1111
1000

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:


## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:

$$
\mathbf{1}^{*}(0 \mid \varepsilon) 1^{*}
$$

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:

$$
1^{*}(0 \mid \varepsilon) 1^{*}
$$

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:

$$
1^{*}(0 \mid \varepsilon) 1^{*}
$$

11110111
111111
0111
0

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:

$$
1^{*}(0 \mid \varepsilon) 1^{*}
$$

11110111
111111
0111
0

## Simple Regular Expressions

- Suppose the only characters are 0 and 1.
- Here is a regular expression for strings that contain at most one zero:
1*0?1*

11110111
111111
0111
0

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
\text { aa* }^{*}\left(. a a^{*}\right)^{*} @ \text { aa*.aa* (.aa*)* }
$$

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is
aa* (.aa*)* @ aa*.aa* (.aa*)*
cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov


## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is
aa* (.aa*)* @ aa*.aa* (.aa*)*
cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov


## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is
aa* (.aa*)* @ aa*.aa* (.aa*)*


## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is
aa* (.aa*)* @ aa*.aa* (.aa*)*


## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
\mathrm{a}^{+} \text {(.aa*)* @ aa*.aa* (.aa*)* }
$$

## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
\mathrm{a}^{+} \quad\left(. \mathrm{a}^{+}\right)^{*} @ \mathrm{a}^{+} . \mathrm{a}^{+} \quad\left(. \mathrm{a}^{+}\right)^{*}
$$

## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
\mathrm{a}^{+} \quad\left(. \mathrm{a}^{+}\right)^{*} @ \mathrm{a}^{+} . \mathbf{a}^{+} \quad\left(. \mathrm{a}^{+}\right)^{*}
$$

cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
a^{+} \quad\left(. a^{+}\right)^{*} @ \quad a^{+} \quad\left(. a^{+}\right)^{+}
$$

## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose our alphabet is a, @, and ., where a represents "some letter."
- A regular expression for email addresses is

$$
\mathrm{a}^{+}\left(. \mathrm{a}^{+}\right)^{*} @ \mathrm{a}^{+}\left(. \mathrm{a}^{+}\right)^{+}
$$

## cs143@cs.stanford.edu first.middle.last@mail.site.org barack.obama@whitehouse.gov

## Applied Regular Expressions

- Suppose that our alphabet is all ASCII characters.
- A regular expression for even numbers is
(+|-)?(0|1|2|3|4|5|6|7|8|9)*(0|2|4|6|8)


## Applied Regular Expressions

- Suppose that our alphabet is all ASCII characters.
- A regular expression for even numbers is
(+|-)?(0|1|2|3|4|5|6|7|8|9)*(0|2|4|6|8)

> 42
> +1370
> -3248
> -9999912

## Applied Regular Expressions

- Suppose that our alphabet is all ASCII characters.
- A regular expression for even numbers is
(+|-)?(0|1|2|3|4|5|6|7|8|9)*(0|2|4|6|8)

> 42
> +1370
> -3248
> -9999912

## Applied Regular Expressions

- Suppose that our alphabet is all ASCII characters.
- A regular expression for even numbers is

> (+|-)?[0123456789]*[02468]

42<br>+1370<br>-3248<br>-9999912

## Applied Regular Expressions

- Suppose that our alphabet is all ASCII characters.
- A regular expression for even numbers is

$$
(+\mid-) ?[0-9]^{*}[02468]
$$

42<br>+1370<br>-3248<br>-9999912

Matching Regular Expressions

## Implementing Regular Expressions

- Regular expressions can be implemented using finite automata.
- There are two main kinds of finite automata:
- NFAs (nondeterministic finite automata), which we'll see in a second, and
- DFAs (deterministic finite automata), which we'll see later.
- Automata are best explained by example...


## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## " H E Y A "

## A Simple Automaton



The automaton takes a string as input and decides whether to accept or reject the string.

## A Simple Automaton

$$
\mathrm{A}, \mathrm{~B}, \mathrm{C}, \ldots, \mathrm{Z}
$$



## A Simple Automaton

$$
\mathrm{A}, \mathrm{~B}, \mathrm{C}, \ldots, \mathrm{Z}
$$



## A Simple Automaton

$$
\mathrm{A}, \mathrm{~B}, \mathrm{C}, \ldots, \mathrm{Z}
$$



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton

$$
\mathrm{A}, \mathrm{~B}, \mathrm{C}, \ldots, \mathrm{Z}
$$



## A Simple Automaton



## " H E Y A "

## A Simple Automaton



The double circle indicates that this state is an accepting state. The automaton accepts the string if it ends in an accepting state.

## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## A Simple Automaton



## " A B C

## A Simple Automaton



## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton



## 011101

## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton



## A More Complex Automaton


$\begin{array}{lllllll}0 & 1 & 1 & 1 & 0 & 1\end{array}$

## A More Complex Automaton


$\begin{array}{lllllll}0 & 1 & 1 & 1 & 0 & 1\end{array}$

## A More Complex Automaton



## A More Complex Automaton


$\begin{array}{lllllll}0 & 1 & 1 & 1 & 0 & 1\end{array}$

## A More Complex Automaton



## A More Complex Automaton


$\begin{array}{lllllll}0 & 1 & 1 & 1 & 0 & 1\end{array}$

## A More Complex Automaton



## 011101

## A More Complex Automaton


$\begin{array}{lllllll}0 & 1 & 1 & 1 & 0 & 1\end{array}$

## A More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



These are called $\boldsymbol{\varepsilon}$-transitions. These transitions are followed automatically and without consuming any input.

## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## An Even More Complex Automaton



## Simulating an NFA

- Keep track of a set of states, initially the start state and everything reachable by $\varepsilon$-moves.
- For each character in the input:
- Maintain a set of next states, initially empty.
- For each current state:
- Follow all transitions labeled with the current letter.
- Add these states to the set of new states.
- Add every state reachable by an $\varepsilon$-move to the set of next states.
- Complexity: $\mathrm{O}\left(m n^{2}\right)$ for strings of length $m$ and automata with $n$ states.


## From Regular Expressions to NFAs

- There is a (beautiful!) procedure from converting a regular expression to an NFA.
- Associate each regular expression with an NFA with the following properties:
- There is exactly one accepting state.
- There are no transitions out of the accepting state.
- There are no transitions into the starting state.
- These restrictions are stronger than necessary, but make the construction easier.



## Base Cases



Automaton for $\varepsilon$


Automaton for single character a

## Construction for $\mathrm{R}_{1} \mathrm{R}_{2}$

## Construction for $\mathrm{R}_{1} \mathrm{R}_{2}$



## Construction for $\mathrm{R}_{1} \mathrm{R}_{2}$



## Construction for $\mathrm{R}_{1} \mathrm{R}_{2}$



## Construction for $\mathrm{R}_{1} \mathrm{R}_{2}$



## Construction for $\mathrm{R}_{1} \mid \mathrm{R}_{2}$

## Construction for $\mathrm{R}_{1} \mid \mathrm{R}_{2}$


start


## Construction for $R_{1} \mid R_{2}$



## Construction for $R_{1} \mid R_{2}$



## Construction for $R_{1} \mid R_{2}$



## Construction for $R_{1} \mid R_{2}$



## Construction for $R_{1} \mid R_{2}$



## Construction for $\mathrm{R}^{*}$

## Construction for $\mathrm{R}^{*}$



## Construction for $\mathrm{R}^{*}$



## Construction for $\mathrm{R}^{*}$



## Construction for $\mathrm{R}^{*}$



## Construction for $\mathrm{R}^{*}$



## Construction for $\mathrm{R}^{*}$



## Overall Result

- Any regular expression of length $n$ can be converted into an NFA with $\mathrm{O}(n)$ states.
- Can determine whether a string of length $m$ matches a regular expression of length $n$ in time $\mathrm{O}\left(m n^{2}\right)$.
- We'll see how to make this $\mathrm{O}(\mathrm{m})$ later (this is independent of the complexity of the regular expression!)


## A Quick Diversion...

I am having some difficulty compiling a C++ program that I've written.

This program is very simple and, to the best of my knowledge, conforms to all the rules set forth in the C++ Standard. [...]

The program is as follows:

I am having some difficulty compiling a C++ program that I've written.

This program is very simple and, to the best of my knowledge, conforms to all the rules set forth in the C++ Standard. [...]

The program is as follows:

```
#include <iostream>
int main(int argc, char** argv)
{
    std::cout << "Hello World!"<<std:: endl;
    return 㫙
}
```

I am having some difficulty compiling a C++ program that I've written.

This program is very simple and, to the best of my knowledge, conforms to all the rules set forth in the C++ Standard. [...]

The program is as follows:
\#include $\langle i o s t r e a m\rangle$

$$
\begin{aligned}
& \text { int main(int argo, char**argv) } \\
& \text { std:: out <<"Hello world!" } \ll \text { std:: endl; } \\
& \} \text { return } \varnothing_{i}
\end{aligned}
$$

> g++ helloworld.png
helloworld.png: file not recognized: File format not recognized collect 2: ld returned 1 exit status

## Challenges in Scanning

- How do we determine which lexemes are associated with each token?
- When there are multiple ways we could scan the input, how do we know which one to pick?
- How do we address these concerns efficiently?


## Challenges in Scanning

- How do we determine which lexemes are associated with each token?
- When there are multiple ways we could scan the input, how do we know which one to pick?
- How do we address these concerns efficiently?


## Lexing Ambiguities

T For
T_Identifier
for
[A-Za-z_][A-Za-z0-9_]*

## Lexing Ambiguities



## Lexing Ambiguities

| T_For | for |
| :--- | :--- |
| $\mathrm{T}_{-}$Identifier | $\left[\mathrm{A}-\mathrm{Za}-z_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |

$$
\begin{array}{l|l|l|l}
\mathrm{f} & \mathrm{o} & \mathrm{r} & \mathrm{t}
\end{array}
$$



## Conflict Resolution

- Assume all tokens are specified as regular expressions.
- Algorithm: Left-to-right scan.
- Tiebreaking rule one: Maximal munch.
- Always match the longest possible prefix of the remaining text.


## Lexing Ambiguities

| T_For | for |
| :--- | :--- |
| $\mathrm{T}_{-}$Identifier | $\left[\mathrm{A}-\mathrm{Za}-z_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |

$$
\begin{array}{l|l|l|l}
\mathrm{f} & \mathrm{o} & \mathrm{r} & \mathrm{t}
\end{array}
$$



## Lexing Ambiguities

T For
for
T_Identifier
[A-Za-z_][A-Za-z0-9_]*


## Implementing Maximal Munch

- Given a set of regular expressions, how can we use them to implement maximum munch?
- Idea:
- Convert expressions to NFAs.
- Run all NFAs in parallel, keeping track of the last match.
- When all automata get stuck, report the last match and restart the search at that point.


## Implementing Maximal Munch

T Do<br>T Double<br>T_Mystery

do
double
[A-Za-z]

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



D $\mathrm{O} \left\lvert\,$|  | U | B | D | O | U | B |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | L\right. E

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



D $\mathrm{O} \left\lvert\,$|  | U | B | D | O | U | B |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | L\right. E

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



T

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |




## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



$$
\begin{aligned}
& \text { U B D D O U B L E } \\
& \\
&
\end{aligned}
$$

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



$$
\begin{aligned}
& \text { U B B D O U B L E } \\
& \\
&
\end{aligned}
$$

## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[A-Z a-z]$ |



## Implementing Maximal Munch

| T_Do | do |
| :--- | :--- |
| T_Double | double |
| T_Mystery | $[\mathrm{A}-\mathrm{Za}-\mathrm{z}]$ |



## A Minor Simplification

## A Minor Simplification




## A Minor Simplification



## A Minor Simplification



## A Minor Simplification



Build a single automaton
that runs all the matching automata in parallel.

## A Minor Simplification



## A Minor Simplification



Annotate each accepting state with which automaton it came from.

## Other Conflicts

| T -Do | do |
| :--- | :--- |
| T -Double | double |
| T _Identifier | $\left[\mathrm{A}-\mathrm{Za}-\mathrm{Z}_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-\right.$ - $\left._{\ldots}\right] *$ |

## Other Conflicts

| T _Do | do |
| :--- | :--- |
| T _Double | double |
| T _Identifier | $\left[\mathrm{A}-\mathrm{Za}-\mathrm{z}_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |

$$
\begin{array}{l|l|l|l|l|l|}
\mathrm{d} & \mathrm{o} & \mathrm{u} & \mathrm{~b} & \mathrm{l} & \mathrm{e} \\
\hline
\end{array}
$$

## Other Conflicts

| T _Do | do |
| :--- | :--- |
| T _Double | double |
| T _Identifier | $\left[\mathrm{A}-\mathrm{Za}-\mathrm{z}_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |


\section*{| d | o | u | $b$ | 1 |
| :--- | :--- | :--- | :--- | :--- |}



## More Tiebreaking

- When two regular expressions apply, choose the one with the greater "priority."
- Simple priority system: pick the rule that was defined first.


## Other Conflicts

| T _Do | do |
| :--- | :--- |
| T _Double | double |
| T _Identifier | $\left[\mathrm{A}-\mathrm{Za}-\mathrm{z}_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |


\section*{| d | o | u | $b$ | 1 |
| :--- | :--- | :--- | :--- | :--- |}



## Other Conflicts

| T _Do | do |
| :--- | :--- |
| T _Double | double |
| T _Identifier | $\left[\mathrm{A}-\mathrm{Za}-\mathrm{z}_{\_}\right]\left[\mathrm{A}-\mathrm{Za}-\mathrm{zO}-9 \_\right] *$ |


\section*{| d | o | u | b | l | e |
| :--- | :--- | :--- | :--- | :--- | :--- |}

## Other Conflicts



\section*{| d | o | u | b | 1 |
| :--- | :--- | :--- | :--- | :--- |
| l |  |  |  |  |}

Why isn't
this a problem?

## One Last Detail...

- We know what to do if multiple rules match.
- What if nothing matches?
- Trick: Add a "catch-all" rule that matches any character and reports an error.


## Summary of Conflict Resolution

- Construct an automaton for each regular expression.
- Merge them into one automaton by adding a new start state.
- Scan the input, keeping track of the last known match.
- Break ties by choosing higherprecedence matches.
- Have a catch-all rule to handle errors.


## Challenges in Scanning

- How do we determine which lexemes are associated with each token?
- When there are multiple ways we could scan the input, how do we know which one to pick?
- How do we address these concerns efficiently?


## Challenges in Scanning

- How do we determine which lexemes are associated with each token?
- When there are multiple ways we could scan the input, how do we know which one to pick?
- How do we address these concerns efficiently?


## DFAs

- The automata we've seen so far have all been NFAs.
- A DFA is like an NFA, but with tighter restrictions:
- Every state must have exactly one transition defined for every letter.
- $\varepsilon$-moves are not allowed.


## A Sample DFA

## A Sample DFA



## A Sample DFA



## A Sample DFA



## Code for DFAs

int kTransitionTable[kNumStates][kNumSymbols] = \{ \{0, 0, 1, 3, 7, 1, ...\},
\};
bool kAcceptTable[kNumStates] = \{ false, true, true,
\};
bool simulateDFA(string input) \{
int state = 0;
for (char ch: input) state = kTransitionTable[state][ch]; return kAcceptTable[state];

## Code for DFAs

int kTransitionTable[kNumStates][kNumSymbols] = \{ \{0, 0, 1, 3, 7, 1, ...\},
\};
dol kAcceptTable[kNumStates] = \{ false, true, true,

```
Runs in time O(m)
    on a string of
    length m.
```


## \};

boot simulateDFA(string input) \{ int state = 0;
for (char ch: input)
state = kTransitionTable[state][ch]; return kAcceptTable[state];

## Speeding up Matching

- In the worst-case, an NFA with $n$ states takes time $\mathrm{O}\left(m n^{2}\right)$ to match a string of length $m$.
- DFAs, on the other hand, take only $\mathrm{O}(m)$.
- There is another (beautiful!) algorithm to convert NFAs to DFAs.



## Subset Construction

- NFAs can be in many states at once, while DFAs can only be in a single state at a time.
- Key idea: Make the DFA simulate the NFA.
- Have the states of the DFA correspond to the sets of states of the NFA.
- Transitions between states of DFA correspond to transitions between sets of states in the NFA.

From NFA to DFA

## From NFA to DFA

## From NFA to DFA

## From NFA to DFA


start $0,1,4,11$

## From NFA to DFA


start $0,1,4,11$

## From NFA to DFA


start $0,1,4,11 \xrightarrow{d} 2,5,12$

## From NFA to DFA


start $0,1,4,11 \xrightarrow{d} 2,5,12$

## From NFA to DFA



$$
\text { start } 0,1,4,11 \xrightarrow{d}-2,12
$$

## From NFA to DFA


start $0,1,4,11 \xrightarrow{d} 2,5,12$

## From NFA to DFA




## From NFA to DFA




## From NFA to DFA




## From NFA to DFA




## From NFA to DFA



## From NFA to DFA




## From NFA to DFA

## From NFA to DFA



## From NFA to DFA



## Modified Subset Construction

- Instead of marking whether a state is accepting, remember which token type it matches.
- Break ties with priorities.
- When using DFA as a scanner, consider the DFA "stuck" if it enters the state corresponding to the empty set.


## Performance Concerns

- The NFA-to-DFA construction can introduce exponentially many states.
- Time/memory tradeoff:
- Low-memory NFA has higher scan time.
- High-memory DFA has lower scan time.
- Could use a hybrid approach by simplifying NFA before generating code.


## Real-World Scanning: Python


while (ip $<z$ )
++ip;

## Python Blocks

- Scoping handled by whitespace:

$$
\begin{aligned}
\text { if } \mathrm{w}= & =\mathrm{z}: \\
\mathrm{a} & =\mathrm{b} \\
\mathrm{c} & =\mathrm{d} \\
\text { else }: & \\
e & =\mathrm{f} \\
\mathrm{~g}=\mathrm{h} &
\end{aligned}
$$

-What does that mean for the scanner?

## Whitespace Tokens

- Special tokens inserted to indicate changes in levels of indentation.
- NEWLINE marks the end of a line.
- INDENT indicates an increase in indentation.
- DEDENT indicates a decrease in indentation.
- Note that INDENT and DEDENT encode change in indentation, not the total amount of indentation.


## Scanning Python

$$
\begin{aligned}
\text { if } \mathrm{w}= & =\mathrm{z}: \\
\mathrm{a} & =\mathrm{b} \\
\mathrm{c} & =\mathrm{d} \\
\text { else }: & \\
\mathrm{e} & =\mathrm{f} \\
\mathrm{~g}=\mathrm{h} &
\end{aligned}
$$

## Scanning Python



## Scanning Python



## Scanning Python



| ident |
| :---: |
| g |$=$| ident |
| :---: |
| h |

## Where to INDENT/DEDENT?

- Scanner maintains a stack of line indentations keeping track of all indented contexts so far.
- Initially, this stack contains 0 , since initially the contents of the file aren't indented.
- On a newline:
- See how much whitespace is at the start of the line.
- If this value exceeds the top of the stack:
- Push the value onto the stack.
- Emit an INDENT token.
- Otherwise, while the value is less than the top of the stack:
- Pop the stack.
- Emit a DEDENT token.


## Interesting Observation

- Normally, more text on a line translates into more tokens.
- With DEDENT, less text on a line often means more tokens:

```
if condl:
    if cond2:
    if cond3:
        if cond4:
        if cond5:
                                statementI
statement2
```


## Summary

- Lexical analysis splits input text into tokens holding a lexeme and an attribute.
- Lexemes are sets of strings often defined with regular expressions.
- Regular expressions can be converted to NFAs and from there to DFAs.
- Maximal-munch using an automaton allows for fast scanning.
- Not all tokens come directly from the source code.


## Next Time



