ASSIGNMENT 1

CS 473: Compiler Design / Fall 2022

Description

In this assignment, you will implement a lexical parser for the language Albatross. The language is described below. Your lexical parser should reject programs with invalid tokens and print a list of tokens for programs with only valid tokens.

The Albatross Language

The Albatross language is defined in a separate document.

Format of the output

Your lexical parser should output each token recognized in a separate line, with the following format:

```
<col> <line> <token> <info>
```

Where:

- **col** is a number denoting the column in the input file where the token starts
- **line** is a number denoting the line in the input file where the token starts
- **token** is a string denoting which type of token was recognized
- **info** is further information about some tokens

For instance, consider the following valid program in Albatross:

```
return 0;
```

The expected output for this program is:

```
1 1 RETURN
8 1 INT 0
9 1 SEMICOLON
```

Tokens

Your assignment should recognize the tokens as per the table below.

- Parenthesis, brackets, and curly brackets; with the token itself as the name.
- All operators for expressions, the name of the token is as follows:

Operator	Name
+	PLUS
-	MINUS
*	MUL
/	DIV
%	REM
==	EQ
<>	NE
<	LT
<=	LE
>	GT
>=	GE
!	NOT
&	BAND
|	BOR
^	XOR
&&	AND
||	OR

- All operators for statement keywords, the name of the token is the capitalized keyword (e.g., RETURN, IF, etc.)
- Types "int", "string", "char", and "void". The name of the token is TYPE and the info is a null-terminated string representing the type (e.g., TYPE "int")
- Variable names: alphanumeric sequences of characters that start with a letter. The name of the token is NAME and the info is the name of the variable. If the variable is preceded by a type, then the info should show the type token as well. Examples:
    - "variable"
        - 1 1 NAME variable
    - "int variable"
        - 1 1 TYPE int
        - 4 1 NAME variable TYPE int
- Integer constants: sequences of numbers. The name of the token is INT and the info is the content of the constant (e.g., INT 473)
- String constants: sequences of characters between double-quotes. The name of the token is STRING and the info is the length of the string, followed by the contents of the string
    - (e.g., STRING 6 CS 473)

- Octal constants: sequences of numbers in base 8 (i.e., between 0 and 7) that start with one leading zero. The name of the token is INT and the info is the content of the constant in base 10 (e.g., the octal 010 should result in the token INT 9)
- Hexadecimal constants: sequences of numbers in base 16 (i.e., numbers between 0 and 9 and capital characters between A and F) starting with 0x. The name of the token is INT and the info is the content of the constant in base 10 (e.g., the hex 0xA should result in the token INT 10)

Escape characters

Strings may have escape characters, which should be replaced by the character they represent. Albatross has the following espace sequences:

- \n -> new line character
- \t -> tab character
- \" -> double-quote character
- \\ -> the escape character itself \

Separators in decimal constants

Albatross supports using the character "_" in decimal constants (i.e., base 10). For instance, 1_000_000 should result in the token INT 1000000. All characters "_" should be ignored when getting the value of a decimal constant.

Return value and lexing errors

When lexing a file successfully, your submission should print the output described above and return with code zero. When there is a lexing error, the output is left unspecified but your submission should exit with code non-zero.

Useful library functions

You can use any functions which are part of the standard C library. In particular, functions in `string.h` are useful for dealing with strings in C as arrays of characters, and string conversion functions in `stdlib.h` are useful to turn strings that represent numbers into native integers. See the links below:

https://pubs.opengroup.org/onlinepubs/7908799/xsh/string.h.html

https://cplusplus.com/reference/cstdlib/

Entry Point

You will modify file `lexer.lex` with all the rules for the lexer to implement the behavior described above. As we will see in class, you will have to modify several sections of the .lex file.

You will also modify file `driver.c`, adding tokens to the large switch-case statement and printing the relevant information for each token as described above.

You will also modify the file `tokens.h` that will define all the tokens used by the lexer, and their format.

Due Date and Resubmission Policy

This assignment is due on **September 10 2022** (Saturday) at **5pm CST**. There is no late policy.

The code and date used for your submission is defined by the last commit to your Git repository.

To resubmit this assignment, your **original grade** (as defined by the autograder) should be **equal to or higher than 30%** for undergraduate students. You can resubmit your assignment until **September 17 2022** (following Saturday) at **5pm CST**. Together with your resubmission, you will have to submit a written description of what you changed from the original submission (on Gradescope).

Bonus Points

This assignment has a total of **10% bonus points**, which you can earn by using Piazza as described in the syllabus. Your posts should be public, tagged with the `assignment1` label, and non-anonymous to the instructors to count towards the bonus. You can claim bonus points through **a Gradescope quiz**.

Submission and Grading

This assignment is submitted through Github, and has an automatic grade component of 100%. You can check your current grade at any point by submitting your code and checking the autograder. The automatic grade is determined by 10 tests, that will check if your submission outputs the expected result. Each test is worth 10%.

Graduate students

On top of the behavior described above, graduate students have to implement multi-line comments or multi-line strings. Graduate students should pick a lexical construction to denote one of these and implement them.

Each regular test is only worth 8% for graduate students. The remaining 20% will be graded as described below.

This language extension will be graded via one video screen-cast (**through Gradescope**) that answers the questions below by explaining how your code. You can record such a video without installing any software by using the following website: https://screenapp.io/#/

1. What lexical form do multi-line comments or multi-line strings take in your language?
2. Which lex rules did you use?
3. Show one example of a successful program that uses multi-line comments or multi-line strings
4. Show one example of a program that contains lexing errors on the multi-line comment or multi-line string

The maximum length for the video is 5 minutes, instructors will stop watching at the 5 minute mark (nothing past that point in the video will be graded). This video should be a screencast of your IDE open on the code submitted, and you should highlight the code. Note that longer videos are not better videos, and you should record a video as short as needed to show all the expressions and answer the questions above.

The final grade for the assignment will be the grade of the original submission, for assignments without a resubmission; or the average between the original grade and the resubmission grade, for assignments with a resubmission. The grade of the original submission includes any bonus points.

Errors and Omissions

If you find an error or an omission, please post it on Piazza as soon as you find it.

Hardcoding and Academic Integrity

Any hardcoding will result in a 0% grade. Hardcoding is when you submit code that detects which test is being run, and simply outputs the expected result. For instance, detecting that test 22 is running, and replacing the usual execution of your submission with `printf("expected result")`.

The academic integrity policy described in the syllabus applies to this assignment. You are responsible for writing all the code that you submit. We will use an automatic tool that detects plagiarism on all submitted code, and we will investigate all instances where plagiarism is more than likely.

Please refer to the syllabus for the full academic integrity policy.