

9. The Compiler I – Syntax Analysis¹

"This chapter is work in progress, and thus it includes inconsistencies with the English grammar. As you parse the chapter's text, please restrict error reporting to the Specification and Implementation sections only."

(The Authors, 2002)

In this chapter we start the process of building a compiler for the Jack high level language. The process of compilation is usually partitioned into two conceptual parts: syntactic understanding of the program structure, and semantic generation the compiled code. This chapter deals with the first issue, that of *parsing* a program written in the Jack language as to "understand its structure". The second part, code generation, is the subject of chapter 10.

The concept of "understanding the structure" of a program needs some explanation. When humans reads a program, they immediately see the "structure" of the program: where classes and methods begin and end, what are declarations, what are statements, what are expressions and how they are built, and so on. Notice that this is a complex nested structure: classes contain methods that contain statements that contain expressions, etc. The allowable structure of programs may be formalized, and programming languages today have formal syntax rules, usually given as a "context free language".

Parsing a program that was written according to these rules means determining the exact correspondence between the program and the syntax rules. This correspondence is usually hierarchal, and may be specified by a "derivation tree" for the program. Compilers often keep an explicit data structure that corresponds to this tree and use this data structure for code generation. Alternatively, they may generate this information implicitly and use it on-the-fly for code generation. Since in this chapter we do not generate any code yet, we have chosen to explicitly output the parsed structure in a particular format, specifically in XML. This will demonstrate the correct parsing of the program, in a way that is easily displayed in any web browser. In the next chapter we will simply replace the current parts of the program that output the parsing in XML with parts that do actual code generation.

It is worthwhile to note that the same types of syntax rules used for specifying programming languages are also used for specifying the syntax of many other types of files. While most programmers will never need to write a real compiler, it is very likely that they will often need to parse files of some other type with a complex syntax. This parsing will be done in the same way that the parsing of a programming language is done.

¹ From *The Elements of Computing Systems*, Nisan & Schocken, MIT Press, forthcoming in 2003, www.idc.ac.il/csd

9.1 Background

Lexical Analysis

In its plainest syntactic form, a program is simply a sequence of characters, stored in a text file. The first step in the syntax analysis of the program is to group the characters into *words*, also called *tokens*, while ignoring white space and comments. This step is usually called “lexical analysis,” “scanning,” or “tokenizing”. Once a program has been tokenized, the tokens (rather than the characters) are viewed as its basic atoms. Thus the tokens stream becomes the main input of the compiler. Prog. 9-1 illustrates the tokenizing of a typical code fragment, taken from a Java or C program.

Code fragment

```
while (count<=100) { /** demonstration */
    count++;
    // body of while continues
    ...
```

Tokens

```
while
(
count
<=
100
)
{
count
++
;
...
```

Program 9-1: Lexical Analysis, also called *tokenizing*, converts the input text into a list of tokens. These tokens are then taken to be the elementary atoms from which the program is made.

As Prog. 9-1 illustrates, there are several distinct types of tokens: *while* is a keyword; *count* is an identifier; *<=* is an operator, *100* is a constant, and so on. Also notice that white space (blanks and newline characters) is eliminated in the tokenizing process, and so are comments.

In general, each programming language specifies the types of tokens it allows, as well as the exact syntax rules for combining them into programmatic structures. For example, some languages may specify that “++” is an operator, while other languages may not. In the latter case, an expression containing two consecutive + operators will be considered invalid.

Context Free Languages

Once we have lexically analyzed a program into a stream of tokens, we are now faced with the main challenge of parsing it into its formal structure. We first need to consider how the formal syntax of languages is specified. There is a rich theory called “formal languages” that discusses many types of languages, including the formalisms used to specify them. Almost all programming languages, as well as most other formal languages used for describing the syntax of complex files types, use a formalism known as “context free grammars”.

A *context free grammar* is a specification of allowable syntactic elements. For example, the rules of the Java language allow us to combine the atoms `100`, `count`, and `<=` into the pattern `count<=100`. In a similar fashion, we can observe that according to the Java grammar, the input pattern `count<=100` is valid, i.e. it is consistent with the language's rules. Indeed, every language has a dual perspective. From a constructive standpoint, the language specifies allowable ways to combine *words*, also called *terminals*, into higher-level structures, called *non-terminals*. From an analytic standpoint, the language is also a prescription for doing the reverse: decomposing a given input pattern into non-terminals, lower-level non-terminals, and eventually into terminals that cannot be decomposed further. These terminals correspond to the token types of the lexical analysis phase.

The syntactic structure of the language -- the context free grammar -- is a set of rules that specify how non-terminals can be derived from other non-terminals and terminals. Thus the grammar is inherently recursive. There may be more than one possible rule for deriving any particular non-terminal, and the different alternatives are usually indicated using the “|” notation. Grammar 9-2 gives an example of a small part of the context-free grammar of the C-language.

```

...
statement:  whileStatement
           |  ifStatement
           |  ... // other statement possibilities follow
           |  '{' statementSequence '}'

whileStatement: 'while' '(' expression ')' statement

ifStatement: ... // if definition comes here

statementSequence: " // null, i.e. the empty sequence
                  | statement ';' statementSequence

expression: ... // definition of an expression comes here
...          // more definitions follow

```

GRAMMAR 9-2: A Context Free Grammar is a set of rules that describes the syntactic structure of a language. Here we see part of the C language grammar.

What does Grammar 9-2 mean? First, the grammar implies that *whileStatement*, *ifStatement*, *expression*, and *statementSequence* are non-terminals, whereas “while”, ‘{’, ‘}’, and ‘;’ are terminals. Further, the grammar implies that a *statement* in the C language may be one of several forms, including a *while* statement, an *if* statement, other possibilities not shown here for lack of space, and any sequence of statements enclosed in curly brackets. The resulting grammar is highly recursive, allowing nested structures like the following example:

```

while (some expression) {
    some statement;
    some statement;
    while (some expression) {
        while (some expression) {
            some statement;
        }
        some statement;
    }
}

```

Parsing: The act of checking whether a grammar “accepts” an input text as valid (according to the grammar rules) is called *parsing*. As a side-effect of the parsing process, the entire syntactic structure of the input text is uncovered. Since the grammar rules are hierarchical, the result is a tree-oriented data structure, called *parse tree* or *derivation tree* (whether the tree is stored in memory or recognized on-line is a different issue that will be addressed later). For example, if we apply grammar 9-2 to the tokenized version of program 9-1, we will obtain the parse tree depicted in Fig 9-3.

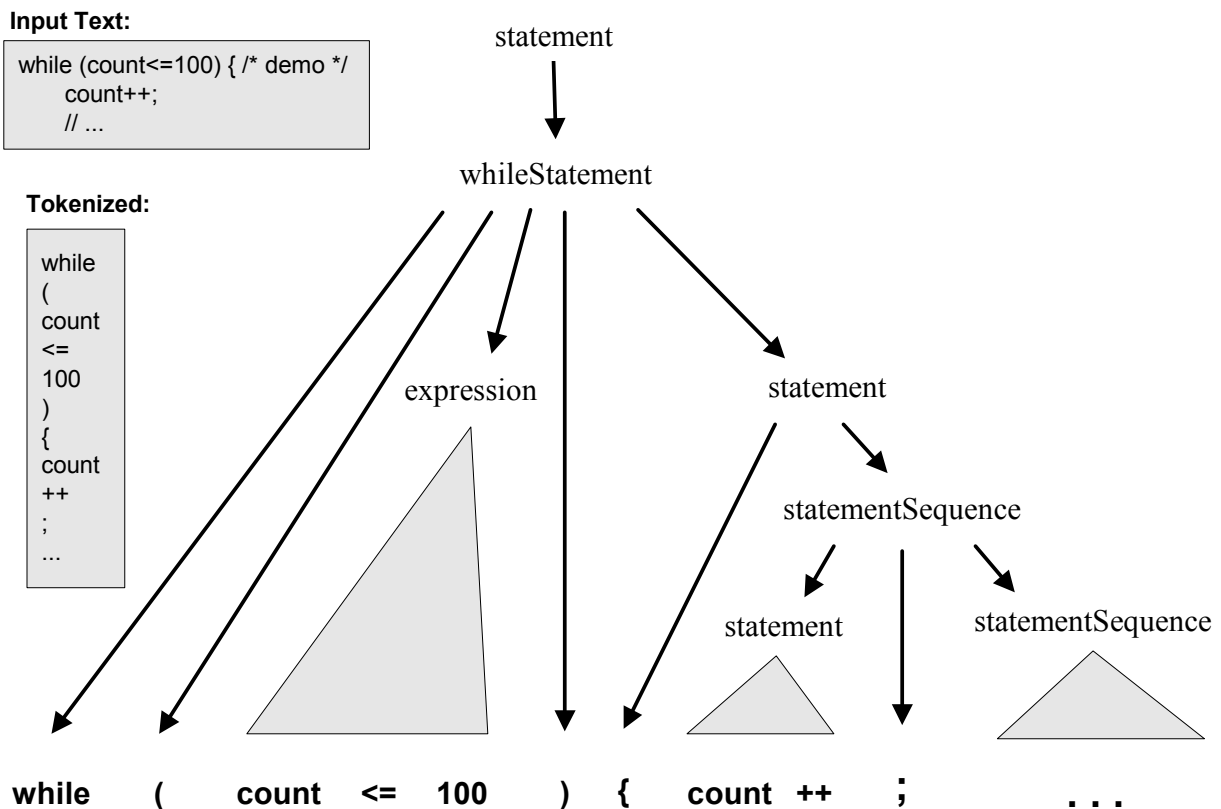


FIGURE 9-3: Parse tree of program 9-1 according to Grammar 9-2. Solid triangles represent lower-level parse trees. The input of the parsing process is the tokenized version of the program.

Recursive Descent Parsing

The last section ended with a description of a *parse tree*. We now turn to describe the algorithms that can be used to construct such trees from given input programs, according to the syntax rules of a given language. There are general algorithms that can do that for any context free language. The general algorithms are not efficient enough for practical use on very long programs, and there are more efficient parsing algorithms that apply to certain restricted classes of context free languages, classes that contain the syntax rules of essentially all program languages. These more efficient algorithms usually run “online” – they parse the input as they read it, and do not have to keep the entire input program in memory. There are essentially two types of strategies for this parsing. The simple strategy works top-down, and this is the one we present here. The more advanced algorithms work bottom-up, and are not described here since they require a non-trivial elaboration of theory.

The top-down approach to parsing, also called *recursive descent parsing*, parses the input stream recursively, using the nested structure prescribed by the language grammar. Let us consider how a *parser* program that implements this strategy can be constructed. For every non-terminal building block of the language, we can equip the parser with a recursive procedure designed to parse that non-terminal. If the non-terminal consists of terminal atoms only, the procedure will simply read them. Otherwise, for every non-terminal building block, the procedure will recursively call the procedure designed to parse the non-terminal. The process will continue recursively, until all the terminal atoms have been reached and read.

For example, suppose we have to write a recursive descent parser that implements Grammar 9-2. Since the grammar has five derivation rules, the parser implementation can consist of five major procedures: `parseStatement()`, `parseWhileStatement()`, `parseIfStatement()`, `parseStatementSequence()`, and `parseExpression()`. The parsing logic of these procedures should follow the syntactic patterns found in the corresponding grammar rules. Thus `parseStatement()` should probably start its processing by reading the first token and determining what it is. Having established the token’s identity, the procedure could determine which statement we are in, and then call the parsing procedure associated with this statement type. For example, if the input stream were Prog. 9-1, the procedure will establish that the first token is `while`, and then call the procedure `parseWhileStatement()`. According to the corresponding grammar rule, this procedure should next attempt to read the terminals “`while`” and “`(`”, and then call `parseExpression()` to parse the non-terminal *expression*. After `parseExpression()` would return (having read and parsed the “`count<=100`” sequence in our example), the grammar dictates that `parseWhileStatement()` should continue parsing the remainder of the while statement. In particular, the grammar states that it should attempt to read the terminal “`)`” and then recursively call `parseStatement()` to parse the non-terminal statement. This call would continue recursively, until at some point only terminal atoms are read.

LL(0) grammars: Recursive parsing algorithms are simple and elegant. If you will think about them, you will realize that the only thing that complicates matters is the existence of several alternatives for parsing non-terminals. For example, when `parseStatement()` attempts to parse a statement, it does not know in advance whether this statement is a while-statement, an if-statement, a curly-bracket enclosed statement list, and so on. The span of possibilities is determined by the underlying grammar, and in some cases it is easy to tell which alternative we are in. For example, consider Grammar 9-2. If the first token is “`while`”, it is clear that we are faced with a while

statement, since this is the only alternative that starts with a “while” token. This observation can be generalized as follows: whenever a non-terminal has several alternative derivation rules, the first token specifies without ambiguity which rule to use. Grammars that have this property are called *LL(0)* grammars, and they can be handled simply and neatly by recursive descent algorithms.

When the first token does not suffice to resolve the element’s type, it is possible that a “look ahead” to the next token will settle the dilemma. Such parsing can obviously be done, but as we need to look ahead at more and more tokens down the stream, things start getting complicated. The Jack language grammar, which we now turn to present, is “almost” *LL(0)*, and thus it can be handled rather simply by a recursive descent algorithm. The only exception is the parsing of expressions, where just a little look ahead is necessary.

9.2 Specification

In this chapter we will write a syntax analyzer for the Jack programming language. In the next chapter we will add the functionality of code generation to the syntax analyzer, and obtain a full compiler. The main purpose of the syntax analyzer is to read a Jack program and “understand” its structure according to the Jack language syntax specification. The meaning of “understanding” is that the program “knows” at each point the meaning of what it is reading: an expression, a command, a variable name, etc. It has to have this knowledge in a complete recursive sense. This is what will be needed for later enabling the addition of code generation.

One way to demonstrate that the analyzer has “understood” the programmatic structure of the input is to have it print the text in a way that provides a visual image of the program structure. Therefore, while syntax analyzers are normally not stand-alone programs, we define our syntax analyzer as having a specific output: an XML description of the program. Thus the syntax analyzer you build here will output an XML file whose structure reflects the structure of the underlying program. In the next chapter you will replace the parts of the program that output the XML code with software that generates executable VM code instead.

Usage: The Jack syntax analyzer accepts a single command line argument that specifies either a file name or a directory name:

```
prompt> JackCompiler source
```

If *source* is a file name of the form *fileName*.jack, the analyzer compiles it into a file named *fileName*.xml, created in the same folder in which the input *fileName*.jack is located. If *source* is a directory name, all the .jack files located in this directory are compiled. For each *fileName*.jack file in the directory, a corresponding *fileName*.xml file is created in the same directory.

Jack Language Syntax

The “programmer oriented” specification of the Jack language was given in chapter 8. We now turn to give a formal specification of the Jack language *syntax*, using a context free grammar. The grammar is based on the following conventions:

- **'xxx'** quoted boldface is used for characters that appear verbatim (“terminals”)
- xxx regular typeface is used for names of language constructs (“non-terminals”)
- () parentheses are used for grouping of language constructs
- x | y means that either x or y can appear
- x? means that x appears 0 or 1 times
- x* means that x appears 0 or more times

Input: The primitive form of a text analyzed by the grammar is a series of tokens, separated by an arbitrary amount of white space (space and newline characters) and comments, which are ignored. Comments are of the standard formats `/* comment until closing */` and `// comment to end of line.`

The grammar is given in Grammar 9-4.

Lexical elements	There are five types of lexical elements in the Jack language:
keyword:	<code>'class' 'constructor' 'function' 'method' 'field' 'static' 'var' 'int' 'char' 'boolean' 'void' 'true' 'false' 'null' 'this' 'let' 'do' 'if' 'else' 'while' 'return'</code>
symbol:	<code>{ ' ' } ' (') ' ' ['] ' ' ; ' ' + ' ' - ' ' * ' ' / ' ' & ' ' ' ' > ' ' < ' ' = ' ' ~ '</code>
integerConstant:	a decimal number in the range 0 .. 32767
stringConstant:	" sequence of ASCII characters not including double quote or newline '"
identifier:	sequence of letters, digits, and underscore (" _ ") not starting with a digit
Program structure:	A program is a collection of classes, each appearing in a separate file. The compilation unit is a class, and is given by the following context free syntax:
class:	<code>'class' className '{' classVarDec* subroutineDec* '}'</code>
classVarDec:	<code>('static' 'field') type varName (',' varName)* ';' </code>
type:	<code>'int' 'char' 'boolean' className</code>
subroutineDec:	<code>('constructor' 'function' 'method') ('void' type) subroutineName '(' parameterList ')' subroutineBody</code>
parameterList:	<code>((type varName) (',' type varName)*)?</code>
subroutineBody:	<code>'{' varDec* statements '}'</code>
varDec:	<code>'var' type varName (',' varName)* ';' </code>
className:	identifier
subroutineName:	identifier
varName:	identifier
Statements	
statements:	statement*
statement:	whileStatement ifStatement letStatement returnStatement doStatement
whileStatement:	<code>'while' '(' expression ')' '{' statements '}'</code>
ifStatement:	<code>'if' '(' expression ')' '{' statements '}' ('else' '{' statements '})?</code>
returnStatement:	<code>'return' expression? ';' </code>
letStatement:	<code>'let' varName '[' expression ']'? '=' expression ';' </code>
doStatement:	<code>'do' subroutineCall ';' </code>
Expressions:	
expression:	term (op term)*
term:	integerConstant stringConstant keywordConstant varName varName '[' expression ']' subroutineCall '(' expression ')' unaryOp term
subroutineCall:	subroutineName '(' expressionList ')' (className varName) '.' subroutineName '(' expressionList ')'
expressionList:	(expression (',' expression)*)?
op:	<code>'+' '-' '*' '/' '&' ' ' '>' '<' '='</code>
unaryOp:	<code>'-' '~'</code>
keywordConstant:	<code>'true' 'false' 'null' 'this'</code>

Grammar 9-4: Complete Grammar of the Jack language

XML Output Format

The output of the syntax analyzer should be an XML description of the program. Figure 9-5 gives a detailed example. Basically, the analyzer has to recognize two major types of language constructs: terminal elements, and non-terminal elements. These constructs are handled as follows.

Non-terminals: Whenever a non-terminal element of type xxx of the language is encountered, the analyzer should generate the output:

```
<xxx>  
    recursive code for the body of the xxx element  
</xxx>
```

Where xxx is one of the following (and only the following) non-terminals of the Jack grammar:

- class, classVarDec, subroutineDec, parameterList, subroutineBody, varDec,
- statements, whileStatement, ifStatement, returnStatement, letStatement, doStatement,
- expression, term, expressionList

Terminals: Whenever a terminal element of type xxx of the grammar is encountered, the following output should be generated:

```
<xxx> terminal </xxx>
```

Where xxx is one of the five terminals specified in the “lexical elements” part of the Jack grammar: keyword, symbol, integerConstant, stringConstant, identifier.

```

class Bar {
    method Fraction foo(int y) {
        var int temp; // a variable
        let temp = (xxx+12)*-6 ;
        ...
    }
}

```

```

<class>
  <keyword> class </keyword>
  <identifier> Bar </identifier>
  <symbol> { </symbol>
  <subroutineDec>
    <keyword> method </keyword>
    <identifier> Fraction </identifier>
    <identifier> foo </identifier>
    <symbol> ( </symbol>
    <parameterList>
      <keyword> int </keyword>
      <identifier> y </identifier>
    </parameterList>
    <symbol> ) </symbol>
    <subroutineBody>
      <symbol> { </symbol>
      <varDec>
        <keyword> var </keyword>
        <keyword> int </keyword>
        <identifier> temp </identifier>
        <symbol> ; </symbol>
      </varDec>
      <statements>
        <letStatement>
          <keyword> let </keyword>
          <identifier> temp </identifier>
          <symbol> = </symbol>
          <expression>
            <term>
              <symbol> ( </symbol>
              <expression>
                <term>
                  <identifier> xxx </identifier>
                </term>
              <symbol> + </symbol>
              <term>
                <integerConstant> 12 </integerConstant>
              </term>
            </expression>
            <symbol> ) </symbol>
          </term>
          <symbol> * </symbol>
          <term>
            <symbol> - </symbol>
            <term>
              <integerConstant> 6 </integerConstant>
            </term>
          </term>
        </expression>
        <symbol> ; </symbol>
      </letStatement>
      ...
    </subroutineBody>
  </subroutineDec>
  </symbol> } </symbol>

```

Figure 9-5: Analyzer input (top) and output (bottom)

9.3 Implementation

We suggest to arrange the implementation of the syntax analyzer in three modules:

- `JackAnalyzer`: a main driver that organizes and invokes everything;
- `JackTokenizer`: a tokenizer;
- `CompilationEngine`: a recursive top-down syntax analyzer.

These modules handle the syntax of the language. In the next chapter we will extend this implementation with two additional modules that handle the language's semantics: a *symbol table* and a *VM-code writer*. This will complete the construction of a full scale compiler for the Jack language.

Class `JackAnalyzer`

The program receives a name of a directory, and analyzes all the Jack files in this directory. For each `xxx.jack` file, it creates an `xxx.xml` file in the same directory. The logic is as follows:

For each `xxx.jack` file in the directory:

1. Create a tokenizer from the `xxx.jack` file
2. Create a output stream into the `xxx.xml` file
3. Compile(INPUT: tokenizer, OUTPUT: output stream)

Class `JackTokenizer`

The tokenizer removes all comments and white-space from the input stream and breaks it into Jack-language tokens, as specified in the Jack grammar. If your implementation is Java-based, such a tokenizer can be easily implemented using Java's `StreamTokenizer` class. Here is the API:

- **`JackTokenizer(Reader inp)`**: create a tokenizer from this input stream (`Reader`).
- **`boolean hasMoreTokens()`**: do we have more tokens in the input?
- **`void advance()`**: gets the next token from the input and makes it the current token. This method should only be called if *hasMoreTokens()* is true. Initially there is no current token.
- **`int tokenType()`**: returns the type of the current token, which is one of the following class constants: `KEYWORD`, `SYMBOL`, `IDENTIFIER`, `INT_CONST`, `STRING_CONST`.
- **`int keyWord()`**: returns the type of the keyword. Should be called only when `tokenType()` is `KEYWORD`. Returns one of the following class constants: `KW_CLASS`, `KW_METHOD`, `KW_FUNCTION`, `KW_CONSTRUCTOR`, `KW_INT`, `KW_BOOLEAN`, `KW_CHAR`, `KW_VOID`, `KW_VAR`, `KW_STATIC`, `KW_FIELD`, `KW_LET`, `KW_DO`, `KW_IF`, `KW_ELSE`, `KW_WHILE`, `KW_RETURN`, `KW_TRUE`, `KW_FALSE`, `KW_NULL`, `KW_THIS`.
- **`char symbol()`**: returns the character which is the current token. Should be called only when *tokenType()* is `SYMBOL`.
- **`String getIdentifier()`**: returns the string which is the current token. Should be called only when *tokenType()* is `IDENTIFIER`.

- **int intValue()**: returns the integer value of the current token. Should be called only when *tokenType()* is `INT_CONST`.
- **String stringValue()**: returns the string value of the current token. Should be called only when *tokenType()* is `STRING_CONST`.

Class CompilationEngine

This class effects the actual compilation into XML form. It reads its input from a `JackTokenizer` and writes its parsed XML structure into the output stream. This is done by a series of `compilexxx()` methods, where `xxx` is a corresponding syntactic element of the Jack grammar. The contract between these methods is that each `compilexxx()` method should read the syntactic construct `xxx` from the input, `advance()` the tokenizer exactly beyond `xxx`, and output the XML parsing of `xxx`. Thus, `compilexxx` may only be called if indeed `xxx` is the next syntactic element of the input. The API is divided into two functional parts, but this division is only for documentation purposes.

High level Program Structure compilation:

- **CompilationEngine(JackTokenizer input, Writer output)**: creates a new compilation engine with the given input and output. The next method called must be `compileClass()`.
- **Void compileClass()**: compiles a complete class.
- **Void compileClassVarDec()**: compiles a static declaration or a field declaration.
- **Void compileSubroutine()**: compiles a complete method, function, or constructor.
- **Void compileParamList()**: compiles a (possibly empty) parameter list, not including the enclosing “()”.
- **Void compileSubroutineBody()**: compiles the body of a subroutine, including the enclosing “{}”.
- **Void compileVarDeclaration()**: compiles a `var` declaration.
- (Statement Compilation)
- **Void compileStatements()**: compiles a sequence of statements, not including the enclosing “{}”.
- **Void compileDo()**: compiles a `do` statement.
- **Void compileLet()**: compiles a `let` statement (where the LHS may also be an array element).
- **Void compileWhile()**: compiles a `while` statement.
- **Void compileReturn()**: compiles a `return` statement.
- **Void compileIf()**: compiles an `if` statement, possibly with a trailing `else` clause.

Expression Compilation:

- **CompileExpression()**: compiles an expression.
- **CompileTerm()**: Compiles a *term*. This method is faced with a slight difficulty when trying to decide between some of the alternative rules. Specifically, if the current token is an identifier, it must still distinguish between a variable, an array entry, and a subroutine call. The distinction can be made by looking ahead one extra token. The look-ahead token, which may be one of “[“, “(“, “.”, suffices to distinguish between the three possibilities. Any other token is not part of this term and should not be advanced over.
- **CompileExpressionList()**: compiles a (possibly empty) comma-separated list of expressions.

9.4 Perspective

In this chapter we have side-stepped almost all of the formal language theory studied in a typical compilation course. We were able to do this by choosing a very simple syntax for Jack that could be easily compiled using recursive descent techniques. This avoided the need for bottom-up parsing of “LR” languages, usually used in other programming languages.

In reality, programmers rarely write syntax analyzers by hand. Instead, they use so called “compiler-compilers” (such as “yacc”) utilities. These programs receive as input a context free grammar, and produce as output a syntax analysis code for this grammar. Following the “show me” spirit of this book, we have chosen not to use such black boxes in the implementation of our compiler.

9.5 Build it

This section describes how to build a syntactic analyzer for the Jack language. In the next chapter, we will extend this analyzer into a full-scale Jack compiler. We begin by describing the programs that will be used to test your compiler implementation.

Objective: Develop a syntactic analyzer that parses Jack programs according to the Jack grammar. The output of the analyzer should be written in XML format, following the example given in Fig. 9-6.

Contract: Write the Jack Analyzer program. Use it to parse the three `.jack` files that we supply. Your Analyzer should generate three corresponding `.xml` files. These files should be identical to the compare-files supplied by us.

Test Programs

We developed a simple interactive game, called *Square Dance*, which will be used to test your compiler in both chapters 9 and 10. Although the function of this application is irrelevant to the compilation process and to these chapters, we describe it briefly. *Square Dance* is a trivial “game” that enables you to move a black square around the screen using the keyboard’s four arrow keys. While moving, the size of the square can be increased and decreased by pressing the ‘z’ and ‘x’ keys,

respectively. To quit the game, press the 'q' key. The game implementation is organized in three classes:

- ❑ **Class Main:** Initializes a new game and starts it.
- ❑ **Class Square:** Implements an animated square. A square object has a screen location and size properties, and methods for drawing, erasing, moving, and size changing.
- ❑ **Class SquareGame:** Runs the game according to the game rules.

These classes are written in Jack. It may be a good idea to take a look at the source code and get a general feeling of how the application is working. If you want, you can also compile the three classes using the supplied Jack compiler, and then play the game on the VM Emulator. These activities are completely irrelevant to the analyzer implementation, but they serve to highlight the fact that the test programs are not just plain text (although this is perhaps the best way to think about them in the context of this project).

Test files: We provide three sets of files:

Source code: `Main.jack`, `Square.jack`, `SquareGame.jack`

Tokenizer output (compare files): `Main.txt`, `Square.txt`, `SquareGame.txt`

Analyzer output (compare files): `Main.xml`, `Square.xml`, `SquareGame.xml`

Since the output files that your tokenizer and analyzer will create will have the same names and extensions as the compare files, we suggest putting the latter files in a separate directory.

Stage 1: Tokenizer

First, implement the Jack tokenizer. When applied to a text file that contains Jack code, the tokenizer should produce a list of tokens, stored in a text file. Each token should be printed in a separate line along with its classification: *symbol*, *keyword*, *identifier*, *integer constant*, and *string constant*. The classification should be recorded using XML tags. For example, the input “`let x=5+yy; let city="Paris";`” should generate the following output:

```
<keyword> let </keyword>
<identifier> x </identifier>
<symbol> = </symbol>
<integerConstant> 5 </integerConstant>
<symbol> + </symbol>
<identifier> yy </identifier>
<symbol> ; </symbol>
<keyword> let </keyword>
<identifier> city </identifier>
<symbol> = </symbol>
<stringConstant> Paris </stringConstant>
<symbol> ; </symbol>
```

Note that the tokenizer throws away the double quote characters. That's OK.

A slight difficulty, and a solution: Four of the symbols used in the Jack language (<, >, ", &) are also used for XML markup, and thus they cannot appear as data in XML files. To solve the problem, your tokenizer should output these tokens as `<`, `>`, `"`, and `&`, respectively. For example, in order for the text “`<symbol> & </symbol>`” to be displayed in an XML viewer, the source XML should be written as “`<symbol> & </symbol>`”.

Testing: As a temporary convention, give the output files of your tokenizer the extension “.txt”. Apply your tokenizer to the three test files: `Main.jack`, `Square.jack`, and `SquareGame.jack`. This should generate the three output files `Main.txt`, `Square.txt`, and `SquareGame.txt`. Next, use the `TextComparer` utility to compare your output files to the three supplied compare files.

Stage 2: Parser

Next, implement the Compilation Engine. Write each method of the engine, as specified in the API, and make sure that it emits the correct XML output.

Implementation tips:

- The indentation of the XML is only for readability. The XML viewer ignores white space.
- Note that conceptually speaking, the output of the tokenizer is embedded within the XML output. In other words, the parser builds the language “super structure” on top of the terminal tokens.
- You may want to start by writing a parser that can handle everything except expressions. For example, assume that the only expressions that the input source code can contain may be single identifiers, and handle everything else. Next, extend the parser to handle expressions as well.

Testing: Give the output files of your analyzer the extension “.xml”. Apply your analyzer to the three test files: `Main.jack`, `Square.jack`, and `SquareGame.jack`. This should generate the three output files `Main.xml`, `Square.xml`, and `SquareGame.xml`. Next, use the `TextComparer` utility to compare your output files to the three supplied compare files.