

Ox:
An Attribute-Grammar Compiling System
based on Yacc, Lex, and C:
User Reference Manual

by Kurt M. Bischoff

November 14, 1993
©1992, 1993 Kurt M. Bischoff
Revised: May 25, 2022

Contents

1	Overview of Use	3
2	Preliminary	4
3	Attribute declarations	4
3.1	Semantics of attribute declarations	6
4	Rules and attribute occurrences	6
5	Attribute definitions	7
5.1	Inherited vs. synthesized attributes	7
5.2	Attribute reference sections in the Y-file	8
5.2.1	Explicit mode	8
5.2.2	Implicit mode	9
5.2.3	Mixed mode	9
5.3	Attribute reference sections in the L-file(s)	10
5.3.1	Generality of Ox	10
5.3.2	Ox adaptation to Lex's line-oriented syntax	10
5.3.3	Resolution of ambiguity regarding token returned	12
5.4	Cycles	13
6	Translation into C code	13
7	Temporal behavior of Ox-generated evaluators	14
7.1	Stack operations	14
7.2	Placement of generated code	14
7.3	Decoration and the ready set	15
7.4	Pruning and global variables	15
7.5	Parse tree visualization	16
8	Programming style	16
9	Postdecoration traversals	18
9.1	Example: infix to prefix translation	18
9.2	General description	20
9.2.1	Traversal specifications	20
9.2.2	Traversal action specifications	20
10	Ox macros	22
10.1	Macro definitions	22
10.2	Macro uses	22

<i>CONTENTS</i>	2
10.3 Example	23
11 Automatic generation of copy rules	24
11.1 Example	26
12 File-level organization of Ox evaluators	27
12.1 Conventions of naming Ox output files	27
12.2 Review: combining the outputs of Yacc and Lex	27
12.3 Combined use of Ox, Yacc, and Lex	27
12.4 Typical command sequences	28
13 Command-line options and other points	28
13.1 Error recovery	28
13.2 Stripping Ox constructs	29
13.3 Preventing execution of attribute definition code	29
13.4 Parse tree statistics	29
14 Example: an integer calculator	30
15 Example: a binary number translator	32
16 Example: translation to postfix and prefix	34
A Using Ox with non-Lex lexical analyzers	36
A.1 Default context-sensitivity of L-file compilation	36
A.2 Ox compilation of C-coded lexical analyzers	36
A.2.1 Example	37
B Traversal semantics	40
C List of reserved words and reserved file names	42
D Summary of command-line options	43
E Lex- and Yacc-compatible tool interoperation	46
F Using <i>Graphviz</i> for parse tree visualization	50
F.1 The default parse tree configuration	51
F.2 Specifying <i>cgraph</i> node structure	52
F.3 Example: the integer calculator, revisited	53
References	57
Index	59

1 Overview of Use

Lex and Yacc are powerful and widely-used tools for the automatic generation of language recognizers. Lex accepts a set of user-written regular expressions and writes a C program that performs lexical analysis according to those expressions. Yacc translates user-written grammar rules into C source code for a syntax analyzer. While they afford “hooks” for execution of hand-coded C-language semantic actions, Lex and Yacc provide little other facility for automatic implementation of language semantics.

Attributed parse trees are often used as data structures in evaluators for languages. Often the language implementer hand-crafts code for building, traversing, and evaluating attributes of parse trees, and for parse tree memory management. A Yacc specification defines a context-free language and a mapping from the set of legal sentences to the set of parse trees, but code for parse tree management is not generated automatically by Yacc.

The Ox¹ user can specify a language using the familiar languages of Lex and Yacc, or take an existing Lex/Yacc specification, and add semantics to the language by augmenting the specification files with declarations and definitions of typed attributes of parse tree nodes.

That specification constitutes an *attribute grammar*, and from it Ox can automatically generate an *evaluator* written in Yacc, Lex, and C. For a given input, the evaluator builds a parse tree, determines an order of evaluation for attributes of the tree, and performs, for each attribute, the semantic action required to evaluate it. This parse tree is managed independently of any trees managed by hand-written C code, but information may be moved between the evaluator-managed tree and any global data structure.

Additionally, the Ox user can easily specify parse tree traversals that are performed after evaluation of the tree’s attributes and that refer to those attributes. Such traversals greatly simplify tasks such as code generation and the gathering of compilation statistics.

The language designer is freed from the tedious and error-prone details of writing code for parse tree management. Ox-generated evaluators use memory-management techniques that bring large time-efficiency gains over hand-built evaluators that use the common technique of calling `malloc` once for each parse tree node. Also, Ox provides security by testing the definition for consistency and completeness, and the evaluator performs tests to ensure that a circular definition has not prevented evaluation of attributes.

Ox is a compiler that accepts two or more files, and translates these into files

¹The name “Ox” originated as a homophone for an acronym for “An Attribute Grammar Compiling System”. It was noticed later that every yak is an ox and that Ox generalizes the function of Yacc.

suitable for input to Lex² and Yacc. With few exceptions, all Lex-input/Yacc-input pairs of files that constitute recognizers or translators are legal inputs to Ox. Thus much existing software is amenable to modification using Ox, and implementations that use Ox can be converted stepwise by hand to “pure” Lex/Yacc implementations. This makes Ox well-suited to language designers, experimenters, and implementers already familiar with Lex, Yacc, and C.

2 Preliminary

It is assumed that the reader is familiar with the use of Yacc [Johnson 1975], Lex [Lesk 1975], and C [KR 1988]; Ox syntactic constructs are described mainly as augmentations of the languages accepted by those tools.³ Prior acquaintance with the basic ideas of attribute grammars (for instance, as found in [Waite 1984] or [Aho 1986]) is helpful.

An Ox input specification consists of at least two files: a syntactic specification (which resembles a Yacc input specification and is called a *Y-file*) that Ox translates into a Yacc input specification, and one or more lexical specifications (which resemble Lex input specifications and are called *L-files*) that Ox translates into Lex input specifications. Usually there is exactly one *L-file*, but an evaluator that uses more than one lexical analyzer [Lesk 1975] may be constructed by submitting to Ox more than one *L-file*. This manual presents descriptions of the Ox-specific constructs that may appear in these files, as well as pertinent underlying concepts. These constructs are illustrated mainly by using fragments of three examples of Ox input specifications, the complete texts of which appear in sections 14, 15, and 16.

Within Ox-specific constructs, C-style and C++-style comments may appear anywhere whitespace may appear. The identifiers visible to user-written code within the Ox-generated skeleton source are prefixed by `yyy`, so the user can avoid name conflicts within the generated evaluator by abstaining from the use of identifiers that begin with `yyy`.

3 Attribute declarations

As described in [Johnson 1975], the definition section of a Yacc input specification is the part that precedes the first `%%` mark, and in it the user may declare the start

²The general descriptions in this manual assume the use of a Lex-based lexical analyzer. It is possible, however, to use Ox with lexical analyzers hand-written in C: details are given separately (in appendix A).

³“Yacc”, “Lex”, and “C” can, in this manual, be taken to mean “Yacc, BYacc [BYacc], BtYacc [BtYacc], Bison [Bison] or Msta [Msta]”; “Lex, Flex [Flex] or RE/flex [REflex]”; and “C or C++” respectively. See appendix E for Ox interoperability issues with Lex- and Yacc-compatible tools.

symbol, tokens, associativities, unions, C code sections, etc. The *Y-file* contains such a definition section, and in it are permitted all of the constructs of a Yacc definition section, as well as Ox *attribute declarations*⁴. An attribute declaration consists of the reserved word `@attributes` followed by `{`, an attribute declaration list, `}`, and a list of grammar symbols.

Suppose that a grammar has a symbol `bitlist` and the following attribute declaration:

```
@attributes {float value; int scale,length;} bitlist
```

Then the Ox evaluator, when building a parse tree node labeled `bitlist`, allocates storage for a float named `value` and integers named `scale` and `length`.

An attribute declaration list (in the previous example, the part between curly braces) resembles a C structure declaration list. Digit strings and C-style identifiers, as well as the following characters and reserved words, arranged according to C/C++ type specifier syntax, are legal in attribute declaration lists⁵:

```
( ) * : ; , [ ] ... :: & < > struct union enum char short
int long float double signed unsigned void const volatile
register static restrict _Bool _Complex _Imaginary class
typename wchar_t bool @void
```

Any fundamental or derived type permitted in a C/C++ program may be used as an attribute type specifier⁶. In addition, the attribute type specifier `@void` defines an attribute with no value, and is intended for those situations where the only information of importance is a dependent / dependee relationship between attributes.

Yacc input specifications often contain C code sections between `%{` and `%}`, and these are also permitted in Ox input specifications. Any name given meaning as a type by a `struct`, `class`, `union`, or `typedef` declaration, or by a `#define` in a previous C code section may be used as an attribute type specifier.

The list of grammar symbols following `}` is a possibly empty list of Yacc tokens (including character constants) and nonterminals, members of the list being separated by whitespace.

⁴Yacc declarations (e.g., uses of the reserved words `%token`, `%left`, `%right`, `%nonassoc`, etc.) must precede all Ox declarations.

⁵Curly braces may not appear, so structures, unions, and classes may not be defined within an attribute declaration list.

⁶Currently, only default constructors for a C++ class type (`struct`, `class`, or `union`) will be called from within Ox-generated skeleton code. If a parameterized constructor is required, it must be called within the evaluation part of each definition of an attribute with that type through use of a *placement new operator*.

3.1 Semantics of attribute declarations

An attribute declaration informs Ox that each symbol in the grammar symbol list has attributes of the names and types appearing in the attribute declaration list. If a appears in the attribute declaration list and s appears in the grammar symbol list, then a is said to *belong* to s or to be an attribute of s . Each grammar symbol has its own attribute name space. When the evaluator creates a node labeled by one of the listed symbols, it allocates storage of the specified type for each of the named attributes. A storage location so allocated is called an *attribute instance* (concisely, an *instance*) in the parse tree. Instances may be said to *belong* to nodes.

4 Rules and attribute occurrences

Yacc grammar rules (productions), and the objects of `return` statements in Lex actions (each such object being a token), are here referred to generically as *rules*. Since Ox accepts the constructs of Yacc and Lex, and passes these unchanged, the corresponding constructs of Ox input specifications are also called *rules*. Each rule is viewed as a sequence of grammar symbols, the object of each `return` statement in a Lex action being a sequence consisting of a single grammar symbol. The leftmost symbol of a rule is called the *left-hand side (LHS)*. The *right-hand side (RHS)* comprises the rule's other symbols.

A symbol's position in a rule together with an attribute of that symbol constitute an *attribute occurrence* (concisely, an *occurrence*) in that rule. If the attribute in question is a , the occurrence is said to be an *occurrence of a* . Supposing the `@attributes` declaration of section 3 and the rule:

```
num      :      bitlist DOT      bitlist
```

the attribute occurrence **scale** of the leftmost appearance of `bitlist` is denoted in Ox code as `bitlist.0.scale`, while the attribute occurrence **scale** of the rightmost appearance of `bitlist` is denoted `bitlist.1.scale`.

In general, attribute occurrences are named by a grammar symbol, followed by a period, followed optionally by a non-negative decimal integer and another period, followed by the name of an attribute of that symbol. The integer and the second period are needed only when a given grammar symbol appears more than once in the rule, in which case those distinct appearances are numbered from left to right with consecutive increasing integers starting with 0. For a symbol `X` with an attribute `a`, `X.a` is a synonym for `X.0.a`.

A given rule and an attribute occurrence in that rule constitute an *attribute occurrence* in the grammar.

5 Attribute definitions

For each rule, the Ox user may provide an *attribute reference section*, delimited by @{ and @}, and optionally containing definitions of attribute occurrences of the given rule. Attribute occurrences may be defined therein in terms of the rule's other attribute occurrences and C code such as global variables, constants, macros, and function calls.

5.1 Inherited vs. synthesized attributes

An attribute occurrence o in a rule R is *synthesized* if and only if

1. o is on the LHS of R and the attribute reference section of R contains a definition of o , or
2. o is on the RHS of R and the attribute reference section of R contains no definition of o .

An attribute occurrence o in a rule R is *inherited* if and only if

1. o is on the LHS of R and the attribute reference section of R contains no definition of o , or
2. o is on the RHS of R and the attribute reference section of R contains a definition of o .

An error message is issued if an attribute is found to have both synthesized and inherited occurrences in the grammar. An attribute is *synthesized* if and only if it has at least one occurrence, and its every occurrence is synthesized. An attribute is *inherited* if and only if it has at least one occurrence, and its every occurrence is inherited. It follows from the above that the grammar's start symbol may have only synthesized attributes. Referring to returned tokens as rules emphasizes the equal status of tokens and nonterminals, inasmuch as each kind of symbol (except the start symbol) may have both synthesized and inherited attributes. Since each symbol has a distinct name space (section 3.1), same-named attributes of different symbols are different attributes, and may differ as to whether they are inherited or synthesized.

For each parse tree node except the root node, two rules of the Ox input specification are of particular interest. The *home rule* is the rule applied at the node, i.e., the rule whose LHS is the label of the given node, and whose RHS symbols are the labels of the children of the node. The *parent rule* is the rule applied at the node's parent. The attribute definition of a synthesized attribute instance of a given node is associated with the node's home rule (i.e., it appears in the attribute reference section for that rule), and definitions of inherited attribute instances are similarly associated with the parent rule.

In a legal input specification, each attribute of a symbol appearing in a rule is either synthesized or inherited, but not both, so the definitions of all attributes “fit together” completely and without contradiction.

5.2 Attribute reference sections in the Y-file

The *rules section* of a Yacc file follows the first %% mark [Johnson 1975], and contains the productions (rules) of the grammar. As mentioned above, the Ox user may augment each rule by an attribute reference section, each of which is delimited by @{ and @}, and which contains zero or more *attribute definitions*. When present, the attribute reference section is the last item (other than a terminating semicolon) in a rule.⁷ Conceptually, an attribute definition has a *dependency part* and an *evaluation part*, but syntactically, the parts may be combined or separate. There are three modes of expression of attribute definitions, and different modes may be used within a single attribute reference section. Each attribute definition begins with a *definition mode annunciator* (@e, @i, or @m).

5.2.1 Explicit mode

In this, the most powerful and most verbose attribute definition mode, an attribute definition takes the form of @e (mnemonic for *explicit*) followed by a *dependency expression* (which expresses the dependency part of the definition) followed by an evaluation part. In the following example, the attribute reference section contains three attribute definitions, each expressed in the explicit mode:

```

num      :      bitlist DOT      bitlist
@{ @e num.value : bitlist.0.value bitlist.1.value;
    @num.value@ = @bitlist.0.value@ + @bitlist.1.value@ ;
    @e bitlist.0.scale : ;
    @bitlist.0.scale@ = 0 ;
    @e bitlist.1.scale : bitlist.1.length ;
    @bitlist.1.scale@ = -@bitlist.1.length@ ;
@}
;

```

A dependency expression makes explicit the constraints on the order of execution of evaluation parts and is a non-empty list of attribute occurrences of the rule, followed by a colon, followed by a possibly empty list of attribute occurrences and a terminating semicolon. The occurrences to the left of the colon are said to *depend upon* (hence are called *dependents* of) those to the right, and are the occurrences *defined* in the given attribute definition. The occurrences to the right are called *dependees* of those on the left. An evaluation part immediately follows the semicolon

⁷Thus it does not precede any Yacc action or the Yacc reserved word %prec in the rule, and any following identifier must be the LHS of the next rule.

of the dependency expression, and is basically a C statement⁸ that may contain *attribute references*, each of which is an attribute occurrence enclosed within @ symbols. Attribute references behave as C variables, and all of the usual C operators, such as those for arithmetic, logical, and pointer operations, may be applied to them, as in a C program⁹.

The Ox evaluator chooses an evaluation order such that the evaluation parts for all of the dependees in the definition are executed before those of the dependents. Usually there is a single dependent in a given attribute definition, but in some cases, code may be made more compact by placing more than one attribute occurrence in a dependent list, thereby combining the definitions of those in the list. The evaluation part is executed on behalf of the dependents *taken as a set*, rather than once for each dependent. This is known as *solving* the attribute instances corresponding to the occurrences in that set.

5.2.2 Implicit mode

The *implicit mode*, which is the usual mode of expressing attribute definitions, syntactically combines the dependency part with the evaluation part. The following Ox code is equivalent to that of the preceding example.

```

num      :      bitlist DOT      bitlist
          @{ @i @num.value@ = @bitlist.0.value@ + @bitlist.1.value@;
            @i @bitlist.0.scale@ = 0;
            @i @bitlist.1.scale@ = -@bitlist.1.length@;
          @}
          ;

```

In this mode, an attribute definition takes the form of @i (mnemonic for *implicit*) followed by an evaluation part. The mode annunciator @i informs Ox that the definition has a single dependent, namely the first attribute occurrence referenced in the evaluation part. The dependees in the definition consist of all *other* attribute occurrences referenced in the evaluation part.

5.2.3 Mixed mode

Mixed mode attribute definitions are announced by @m (mnemonic for *mixed*) followed by a dependency part, consisting of one or more attribute occurrences (the dependents in the definition) with a terminating semicolon, followed by an evaluation part. The attribute occurrences referenced in the evaluation part, except those that also appear between @m and the semicolon, are taken to be the dependees in the

⁸An evaluation part may consist of either a single C statement or a C compound statement.

⁹Ensuring that a dependent attribute is assigned a value in an evaluation part is the responsibility of the Ox user—it is not checked by Ox. Particular attention should be given to the loops and conditionals within an evaluation part.

definition. Thus the dependents are given explicitly and the dependees implicitly. The code in the following example has the same meaning as that in the previous two.

```

num      :      bitlist DOT      bitlist
          @ { @m num.value ;
              @num.value@ = @bitlist.0.value@ + @bitlist.1.value@;
              @i @bitlist.1.scale@ = - @bitlist.1.length@;
              @m bitlist.0.scale ; @bitlist.0.scale@ = 0;
          @ }
          ;

```

5.3 Attribute reference sections in the L-file(s)

Definitions of inherited attributes of tokens are associated with rules appearing in the *Y-file*, while their synthesized attributes are defined in the *L-file(s)*. Ox compiles the *Y-file* before compiling the *L-file(s)*. If a given attribute occurrence of a token is not defined in the *Y-file*, then the attribute is taken to be synthesized.

Lexical rules are associated with `return` statements in Lex actions. After the terminating semicolon of each such statement, there may appear a possibly empty *attribute reference section*, delimited by `@{` and `@}`, in which are defined all of the synthesized attributes of the returned token.

Note that each point of `return` from `yylex()` must be *explicit* in the sense that the text must bear the C reserved word `return`. In particular, `returns` must not be done within C macros, unless the *L-file* is passed through the C preprocessor prior to compilation by Ox. Guaranteeing this property of `yylex()` is the responsibility of the Ox user—it is not checked by Ox.

5.3.1 Generality of Ox

The class of attribute grammars accepted by Ox is restricted only as follows: synthesized attributes of tokens do not have dependees. Attribute definitions in the *L-file(s)* can thus be written more simply than in the *Y-file*: each attribute occurrence is defined by referring to it in C code, exactly once in the attribute reference section associated with the `return` statement, as in the following example (wherein `CONST`'s only synthesized attribute is `val`):

```
[0-9]+  return(CONST); @ { sscanf(yytext,"%ld",&@CONST.val@); @ }
```

Thus mode declarations and dependency expressions are unnecessary in the *L-file(s)*.

5.3.2 Ox adaptation to Lex's line-oriented syntax

When Ox is compiling the *L-file* and has recognized a rule (i.e., the object of a `return` statement), if the returned token has synthesized attributes, Ox looks for an

attribute reference section following the `return` statement. Ox's rules for recognizing attribute reference sections in the *L-file* are adapted from the way Lex actions are terminated: Ox gives up looking for an attribute reference section when it pairs the rightmost right curly brace in the action with the leftmost left curly brace, or when it encounters a newline unprotected by curly braces. Newlines are insignificant inside attribute reference sections.

Examples of correct and incorrect syntax are shown below. All of the correct forms shown are semantically equivalent to one another.

- incorrect (attribute reference section appears to the right of the rightmost curly brace):

```
[a-zA-Z]+ { count(); return ID; } @{ @ID.name@ = id(); @}
```

- incorrect (attribute reference section not part of rule, since the Lex action is terminated by an unprotected newline):

```
[a-zA-Z]+ count(); return ID;
      @{ @ID.name@ = id(); @}
```

- correct:

```
[a-zA-Z]+ { count(); return ID; @{ @ID.name@ = id(); @} }
```

- correct:

```
[a-zA-Z]+ return ID; @{ count(); @ID.name@ = id(); @}
```

- correct:

```
[a-zA-Z]+ { count(); return ID;
      @{ @ID.name@ = id(); @}
}
```

- correct:

```
[a-zA-Z]+ count(); return ID;  @{
                                @ID.name@ = id();
                                @}
```

5.3.3 Resolution of ambiguity regarding token returned

A slight difficulty arises in rules like

```
return(yytext[0]);
```

and

```
return(cond ? TOKEN1 : TOKEN2);
```

for which Ox cannot determine at evaluator-generation time which token will be returned.

In the first case, wherein no declared token or character constant is recognized in the `returned` expression, Ox assumes that the token `returned` has no attributes, and issues a warning like:

```
ox: scan.1: warning: line 8: ambiguous form of return of token.
   unknown node type--assuming no attributes.
```

In the second case, wherein more than one declared token or character constant is recognized, the node appended to the tree during evaluation is of the type of the declared token or character constant appearing *leftmost* in the expression. Ox issues a warning like:

```
ox: scan.1: warning: line 8: ambiguous form of return of token.
   multiple tokens in object of return statement.
```

The above warnings should be taken seriously, because the conditions of which they warn can result in the generated evaluator attempting to access attribute instances that are nonexistent or of the wrong type. These kinds of warnings are most often seen when first converting an existing Yacc/Lex translator to Ox.

A condition causing one of the above-described warnings may be tolerated if the Ox user verifies that for the rule (i.e., object of the `return` statement) in question:

- all of the tokens that can be `returned` for the rule are contained in the grammar-symbol list (section 3) of a single attribute declaration.
- no token that can be `returned` for the rule appears in a grammar-symbol list of an attribute declaration.

5.4 Cycles

It is easy to write an attribute grammar such that some attribute instance of some parse tree has a chain of dependencies that leads back to itself. Such a grammar is called *circular*, and such a chain of dependencies is called a *cycle*. For such a tree, there is an attribute instance that the evaluator cannot begin to solve until that instance has already been solved. A cycle also makes it impossible to solve any attribute instance that has a chain of dependencies leading to an instance involved in the cycle. Circularity is usually not intended by the evaluator designer. A general circularity test performed at evaluator-generation time would require exponential running time for some inputs [Jazayeri 1975]. Polynomial-time tests for special kinds of non-circularity are known, but the present version of Ox deals with the problem by checking for cycles at evaluation time.

6 Translation into C code

Ox translates attribute declarations into C structure declarations. If the *Y-file* contains a Yacc `%union` semantic value type declaration, Ox augments that declaration with an Ox semantic value type member consisting of the union of the structure declarations translated from all of the attribute declarations. Ox also supports the Bison `”%define api.value.type union”` directive as the mechanism for defining the semantic value type. The semantic value type definition mechanisms supported by Yacc:

```
#define YYSTYPE <C_type_specifier> (default: int)
or
typedef <C_type_specifier> YYSTYPE;
```

are not supported, since Ox cannot, in general, detect these definitions. Ox does not support the Bison semantic value type definition mechanisms:

```
%define api.value.type variant
%define api.value.type {<C_type_specifier>}
```

If no explicit semantic value type is defined, Ox will generate a `%union` semantic value type declaration consisting of the single Ox semantic value type member.

Ox generates an error message if a reference to a Yacc positional semantic value pseudo variable (`$$`, `$1`, `$2`, etc.) is detected in a Yacc action or an Ox attribute reference section without an explicit semantic type declaration. This error is most often seen when first converting an existing Yacc/Lex translator to Ox.

The evaluation expression of each attribute definition is copied verbatim into Ox’s output, except that attribute references are translated into parenthesized references to C variables.

7 Temporal behavior of Ox-generated evaluators

7.1 Stack operations

Inasmuch as an ordinary Yacc/Lex recognizer employs an LR parsing algorithm [Aho 1986], each input entails a sequence of lookaheads, shifts, and reductions, and a stack of parser states is maintained. Ox generates an evaluator whose `yyparse` function goes through the same sequence of lookaheads, shifts, and reductions as does `yyparse()` of the ordinary Yacc/Lex recognizer.

The Ox evaluator, in building a parse tree, maintains a stack of subtrees that are created within Ox-generated Yacc semantic actions, utilizing the `yyparse()` semantic stack management mechanism. The operations on the stack of Ox subtrees are thus automatically synchronized with the operations `yyparse()` performs on its stack of parser states. Parsing operations involving the “marker nonterminals” (see [Johnson 1975]) inserted into the grammar by Yacc are ignored by Ox.

The evaluator maintains its stack of subtrees as follows. Lookaheads coincide with calls to `yylex()`. Just before a `return` is executed in a Lex action, a leaf node is created in `yyval` by Ox-generated code, and any synthesized attribute instances are solved. At a shift operation, the subtree corresponding to that leaf node is pushed onto the stack by the standard `yyparse()` shift handling of the `yyval` semantic value. Ox-generated code in a Yacc action at the end of each rule’s RHS uses the zero or more subtrees on the semantic stack corresponding to the RHS symbols as the children of a newly-created node, yielding a new subtree in `yyval` corresponding to the rule’s LHS symbol. The root of the new subtree is given a label to indicate the production being applied at the node. At each reduction, the RHS subtrees are popped from the stack and the new LHS subtree is pushed onto the stack by the standard `yyparse()` reduce handling of the `yyval` semantic value. The parse tree is completed upon end of input together with reduction to the start symbol.

7.2 Placement of generated code

Code for parse tree management and attribute evaluation is placed in Yacc and Lex actions in Ox’s output. If a given rule in the *Y-file* has an ordinary Yacc action, the Ox-generated code is placed *after* any programmer-supplied C code contained in the action. If a given rule in the *Y-file* lacks a Yacc action, an action is created, and the Ox-generated code is placed there. The actions so created are introduced only at the *ends* of rules, so Yacc does not create a marker nonterminal for the action, and the LALR(1) property of the grammar is unaffected.

When an attribute reference section in an *L-file* contains definitions for more than one attribute occurrence, code for implementing those definitions is executed in the same order in which the definitions appear in that section.

For the attribute occurrences defined in the *Y-file*, Ox and the Ox-generated evaluator perform analyses to determine when to execute the code segment that evaluates a given attribute. The order of execution of the code segments associated with the definitions in a given attribute reference section is determined by the dependencies of the definitions, and is not necessarily related to the order of appearance of the definitions.

Some attribute occurrences, for example those that have no dependees, are evaluated as part of the generated Yacc action executed just prior to reduction by the associated production. Definitions of such occurrences are allowed to reference the Yacc positional semantic value pseudo variables (`$$`, `$1`, `$2`, etc.). Ox issues an error message if an attribute occurrence referencing a pseudo variable cannot be evaluated at production reduction time.

7.3 Decoration and the ready set

The Ox evaluator maintains a set of attribute instances that are ready to be solved, i.e., those whose every dependee has been solved, but which have not themselves been solved. During parsing of the input, it is possible to remove an attribute instance from this *ready set*, solve it, and then check whether the solving of that instance has caused any of its dependents to be ready to be solved. Instances that are thus made ready are then placed in the ready set. Repeating this process until the ready set is empty is known as *decoration*. Following a decoration, further parsing of the input may result in creation of parse tree nodes and insertion of attribute instances into the ready set. Scheduling of decorations is performed automatically by the evaluator. Evaluation of a given syntactically-correct input involves at least one decoration, performed after the final reduction to the start symbol.

In specialized cases, e.g., if the evaluator is intended for interactive operation, it is likely that automatic scheduling of decorations will not be adequate. The evaluator designer may place the `@decorate` reserved word immediately following the `@{` at the beginning of an attribute reference section to force decoration to occur at the end of the Yacc action executed just prior to reduction by the associated production. For instance, in a line-oriented interactive language application, this could be done in a grammar rule that handles the top level processing for a each line of input, e.g., the rule recognizing the completed input line, including the ending *newline* character.

7.4 Pruning and global variables

An Ox evaluator, by default, implements a run-time memory optimization that depends on the incremental decoration implementation described in section 7.3 above. The optimization forces decorations to occur more frequently than would otherwise be scheduled by the evaluator: after every reduction by the parser of a production. When *all* attributes associated with the nodes of a sub-tree of the parse have been

completely evaluated, the evaluator assumes that there is no further need for those nodes and they are *pruned* (removed from allocated memory).

As noted previously, the execution order of individual attribute value computations is determined by the dependencies specified in the attribute definitions. In cases where one or more global variables are involved in the computations of non-root attributes, Ox is unaware of any ordering dependencies that may exist, and the generated attribute evaluation schedule may very well violate those hidden dependencies. Pruning can mask problems associated with hidden ordering dependencies between non-root attributes and global variables. For example, an attribute value representing the table index from a lookup in a global table will depend on the existence of the desired table entry. If that table entry is populated by some other attribute computation, there is an attribute evaluation order dependency hidden from Ox. Pruning may accidentally force the attribute value computations to occur in the expected order (creating the table entry prior to looking up and finding the table entry).

Ox provides the `-d` command-line option to suppress pruning, although the automatic scheduling of incremental decorations may still mask such hidden dependency problems. It is strongly recommended that global variables *not* be used as intermediate values in non-root attribute computations for this reason.

7.5 Parse tree visualization

If Ox is invoked with the `-t` command-line option, the generated evaluator¹⁰ includes code to build an internal visual representation of a parse tree, using the *cgraph* library [Gansner 2014] [North 2014] from the *Graphviz* [Graphviz] package. The *cgraph* data structures must be initialized prior to execution of any of the generated parse tree management code. Ox provides the function `yyinit` for this purpose. For a generic (non-Lex) hand-written scanner (see appendix A.2), the user must ensure that `yyinit()` is executed prior to any of the parse tree management code added by Ox. An Ox-generated Lex scanner takes care of this automatically.

8 Programming style

Definitions of *attribute grammar*, (for instance those in [Lorho 1988] and [Waite 1984]) employ no notion of execution sequence. The usual Ox programming style involves defining synthesized attribute occurrences of tokens in terms of `yytext` and `yylen` and other such data structures of the lexical analyzer. Then the attribute definitions of each production are written only in terms of constants and other attribute occurrences of that production. For a given sentence, the synthesized attribute instances of the tokens then completely determine the values of all

¹⁰The C preprocessor macro name `YYOX_TREE` is `#defined`'d in the generated evaluator source so that user-written code can test for use of the `-t` command-line option.

attribute instances of the parse tree. The attribute instances of the root node are often of particular interest, and their definitions often contain code that copies their values to global C variables, so that they may be used in code executed after the return from `yyparse()`.

Since attribute definitions in Ox code may contain any C code, the Ox programmer may deviate from the safe approach described above by using non-root attribute definitions that read and write global variables. Before attempting the use of side effects, the programmer should be familiar with the material of section 7. Preference should be given to the use of inherited attributes to pass values down from root node *global* attribute instances.

The order of evaluation of attributes is, by design, not explicit in the Ox input specifications. It is not recommended, and usually it is not convenient, to use attribute definitions for order-sensitive side effects such as code generation. A common general approach to translation is to build and decorate a parse tree (meanwhile performing some of the checks for semantic errors), and to then make one or more determinate-order tree traversals for final error checks, gathering of compilation statistics, code generation, etc. Ox has a facility for specification of such traversals, and this is the topic of section 9.

9 Postdecoration traversals

The idea of decoration was described in section 7.3. *Postdecoration* refers to any time after the *final* decoration of the parse tree, which follows parsing of a correct input. This section shows how the Ox user can cause *postdecoration traversals*, each of which permits access (in a user-specified order) to the tree's attribute instances¹¹.

9.1 Example: infix to prefix translation

The problem of parsing infix arithmetic expressions, and their translation to prefix form serves to introduce Ox's postdecoration traversal facility.

The tokens of the example language are determined by the following *L-file*:

```
%{
#include "y.tab.h"
#include "oxout.h"
%}

%%
[ \n\t\f]*      ;
[0-9]+          return(CONST); @{ sscanf(yytext,"%d",&@CONST.val@); @}
\(              return('(');
\)              return(')');
\+              return('+');
\*              return('*');
.               fprintf(stderr,"illegal character\n");
%%
```

¹¹Pruning (see section 7.4) is suppressed if a postdecoration traversal is defined, since the traversal may need to visit the entire attributed tree.

The following *Y-file* completes the specification of the evaluator.

```
%token CONST
%left '+'
%left '*'

@attributes {int val;} CONST
@traversal @lefttoright @preorder LRpre

%{
#include "oxout.h"
#include <stdio.h>
%}

%%
expr      :      expr      '*'      expr      /* rule 1 */
           @ { @LRpre printf(" * "); @ }

          |      expr      '+'      expr      /* rule 2 */
           @ { @LRpre printf(" + "); @ }

          |      '('      expr      ')'      /* rule 3 */

          |      CONST      /* rule 4 */
           @ { @LRpre printf(" %d ",@CONST.val@); @ }

          ;

%%

int main()
{
  yyparse();
  printf("\n");
}
```

The sequence: `@traversal @lefttoright @preorder LRpre` specifies that a left-to-right preorder traversal of the parse tree be performed by the evaluator after the final decoration, and that the traversal be identified as `LRpre`. Note that `LRpre` is programmer-defined, and is *not* an Ox reserved word.

Each attribute reference section in the above *Y-file* contains a *traversal action specifier* starting with the *traversal mode annunciator* `@LRpre`, which is defined in the above-mentioned `@traversal` specification.

When the `LRpre` traversal reaches a node at which rule 1 is applied, an asterisk is printed, then each subtree rooted at a child of the node is traversed, the leftmost subtree first. The behavior of the traversal at a node at which rule 2 is applied is the same, except that a plus sign is printed instead of an asterisk. When `LRpre` reaches a node for rule 3, no traversal action is performed, but the children of the node are traversed recursively as described above for nodes for rules 1 and 2. The

`val` attribute of the `CONST` child is printed when a node for rule 4 is reached. No action is performed during a traversal of a subtree that consists of a terminal node.

9.2 General description

9.2.1 Traversal specifications

The Ox programmer may place one or more *traversal specifications* in the *Y-file* definition section. Such a specification consists of the reserved word `@traversal`, followed by a *traversal specifier sequence* and a non-empty sequence of identifiers, the identifiers being separated by whitespace. A traversal specifier sequence may contain the following *traversal specifiers* (in any order):

- at most one of: `@postorder`, `@preorder`
- at most one of: `@lefttoright`, `@righttoleft`
- optionally: `@disable`

If neither `@postorder` nor `@preorder` appears in the sequence, the traversal is postorder by default. A left-to-right traversal is specified by default when neither `@lefttoright` nor `@righttoleft` appears.

Following the final decoration, the parse tree is traversed once for each traversal specification. The order of performing the traversals corresponds to the order of appearance of the traversal specifications. The `@disable` reserved word causes the generated evaluator to skip any traversal in whose specification it appears, which may be useful for debugging.

The code fragment:

```
@traversal @preorder LRpre
@traversal LRpost
```

appearing in the *Y-file* definition section specifies that, after the final decoration, the generated evaluator is to perform a left-to-right preorder traversal named `LRpre`, followed by a left-to-right postorder traversal named `LRpost`.

9.2.2 Traversal action specifications

In addition to attribute definitions (section 5.2), the attribute reference sections of the *Y-file* may contain *traversal action specifications*. Each of these consists of a *traversal mode annunciator*, followed by a sequence of *dynamic traversal modifiers* and a *traversal action*. A traversal mode annunciator is `@` followed immediately by the name of a previously-declared traversal.

Suppose traversal specifications of `LRpre` and `LRpost` as above. Then in the code fragment:

```

s      :      expr
        @ { @LRpost printf("\n");           /* 1 */
            @LRpost @revorder (1) printf("postfix: "); /* 2 */
            @LRpre  @revorder (1) printf("\n"); /* 3 */
            @LRpre  printf("prefix:  ");      /* 4 */
        @ }
      ;

```

the attribute reference section has four traversal action specifications and no attribute definitions. Each specification is announced by either `@LRpre` or `@LRpost`. Each of the `printf` statements constitutes a traversal action. The form of a traversal action is that of a C code fragment¹², except that it may contain references to the attribute occurrences of the associated rule.

The second and third specifications each have `@revorder (1)` as a dynamic traversal modifier. A dynamic traversal modifier is either `@revorder` or `@revdirection`, followed by a parenthesized expression that conforms to C syntax, except that it may refer to the rule’s attribute occurrences. `@revorder` and `@revdirection` may each occur at most once in a given traversal action specification. If `@revdirection` appears in two traversal action specifications within a given attribute reference section, the two specifications must have different annunciators. Dynamic traversal modifiers are used to override the traversal specifications of a given traversal when it reaches a given kind of node. The modifier `@revorder expr` means roughly “reverse order if *expr*”. When the LRpre traversal reaches a node at which the rule “`s : expr`” is applied, the expression “(1)” is evaluated, and because it is nonzero, the third traversal action, which prints a line feed, is executed as if LRpre were a postorder traversal, i.e., *after* the recursive traversal of the subtree rooted at the node’s sole child. The execution of the fourth traversal action, “`printf("prefix: ");`” is not affected by any dynamic traversal modifier, and occurs according to LRpre’s (static) specification, i.e., *before* the traversal of the child subtree.

When the LRpost traversal reaches a node at which the rule “`s : expr`” is applied, the second traversal action is executed, the traversal proceeds to the child subtree, then the first traversal action is executed.

The preceding description is generally sufficient for understanding postdecoration traversals, but appendix B contains a pseudocode description that describes the behavior somewhat more formally.

Facility for inorder traversal is to be implemented in future versions of Ox.

¹²A traversal action may consist of either a single C statement or a C compound statement.

10 Ox macros

Ox's input specification may be such that the same or similar text appears in more than one place in attribute reference sections. There is a macro substitution feature that can be used to decrease verbosity in such cases.

10.1 Macro definitions

Ox macros are defined in the *Y-file* definition section. Such a definition consists of the `@macro` reserved word, an identifier (the name of the macro), a left parenthesis, a parameter list, a right parenthesis, the body of the macro, and the `@end` reserved word. The parameter list is a possibly empty sequence of identifiers, delimited by commas. Each identifier is a sequence of letters and digits, beginning with a letter.

The body of the macro is a segment of arbitrary text, terminated by the first occurrence of `@end`. When inside a comment or a string, or when preceded immediately by the backslash escape character, an occurrence of `@end` is considered part of the macro body (hence does not terminate the macro). Such a backslash character is deleted from the macro body.

10.2 Macro uses

Ox macros are used only in attribute reference sections and in other Ox macros. Substitution occurs where a macro use is encountered outside of a string, comment, or attribute name.

A macro use consists of the name of a previously-defined macro, and an argument list in parentheses. The argument list is a possibly empty sequence of text fragments, delimited by commas. In expanding a macro use, each text fragment is substituted for each occurrence in the macro body of the corresponding parameter in the macro definition. Paired parentheses, and commas enclosed within paired parentheses, may occur within a macro use text fragment. If commas, parentheses, or backslashes are to appear otherwise in a text fragment, they must be preceded by backslash escape characters, which are removed during substitution.

It is not necessary that the definition of a macro precede that of another macro in which it is used, as no macro substitution occurs until Ox processes the attribute reference sections. The body of a macro may contain nested macro uses.

10.3 Example

The following excerpts from a *Y-file* illustrate the use of Ox macros.

```

:
@macro exprdefs(op)
  @i @expr.1.env@ = @expr.env@;
  @i @expr.2.env@ = @expr.env@;
  @i @expr.type@ = typeResolve(@expr.1.type@,@expr.2.type@);
  @i @expr.value@ = exprEval(op,@expr.type@,@expr.1.type@,@expr.2.type@,
                          @expr.1.value@,@expr.2.value@
                          );
@end

@macro typeResolve(type1,type2)
  ((type1 == type2) ? type1 : FLOATTYPE)
@end

:
%%
:
expr :      expr  '*'  expr
        @{ exprdefs('*') @}
        |      expr  '/'  expr
        @{ exprdefs('/') @}
        |      expr  '+'  expr
        @{ exprdefs('+') @}
        |      expr  '-'  expr
        @{ exprdefs('-') @}
        ;

:

```

The identifier `exprEval` referenced in the definition of the `exprdefs` macro is the name of either a C macro or C function. The Ox macro `typeResolve` above contains no Ox-specific constructs and, as a matter of style, could have been declared instead as a C macro or C function.

11 Automatic generation of copy rules

Often a *Y-file* has attribute definitions that function only to copy an instance belonging to one node to a like-named instance belonging to the node's parent or child. Large attribute grammars tend to have many such definitions, which are sometimes called *copy rules*. The situation is conspicuous when contextual information is moved leafward via inherited attributes.

The Ox user may place the following construct in the *Y-file* definition section:

```
@autoinh <ID_list>
```

where `<ID_list>` is a whitespace-separated list of attribute names. Suppose that `attrbID` is such an attribute name, and the above construct is followed by an `@attributes` declaration whereby `attrbID` is declared as an attribute of the grammar symbol `gSym`. Then Ox knows that `attrbID` is an inherited attribute of `gSym`. Further, for any rule having `gSym` on the RHS, Ox searches that rule's attribute reference section for definitions of the RHS occurrences of `attrbID`. When such a definition is missing, Ox checks whether the LHS has an occurrence of `attrbID`. If so, Ox generates definitions that copy that LHS occurrence to each RHS `attrbID` occurrence that lacks a definition. If there is no such LHS occurrence, Ox issues an error message.

There is an analogous construct for automatic generation of definitions of synthesized occurrences:

```
@autosyn <ID_list>
```

When the `@autosyn` construct is used, Ox tries to supply missing definitions of synthesized occurrences by searching the RHS for same-named occurrences. If exactly one such RHS occurrence is found, Ox generates a definition to copy it to the LHS, otherwise there is an error.

The above-described constructs have a global character in that a single `@autosyn` or `@autoinh` declaration can easily be used to supply missing definitions for all occurrences of attributes of a given name. These reserved words may be used in a more conservative way that generates missing definitions only for occurrences belonging to a selected set of grammar symbols:

Attribute declarations are written as usual, except that `@autoinh` or `@autosyn` may appear before the attribute's type specifier (i.e., after `{` or `;`). Where `<ID_list>` is the usual comma-separated list of attribute names, and `attrbID` is a member of `<ID_list>`:

```

@attributes  {
              :
              @autoinh <typespec> <ID_list> ;
              :
            }
            <grammar_symbol_list>

```

declares `attrbID` as an inherited attribute whenever it occurs in a symbol in `<grammar_symbol_list>`. Further, this instructs Ox to attempt to supply missing definitions of such occurrences by copying from the LHS. The `@autosyn` reserved word may be used locally in an analogous manner.

For safety in the use of `@autosyn` and `@autoinh`, Ox provides the `@warn` reserved word. When it immediately follows `@autosyn` or `@autoinh`, Ox issues a warning for each definition supplied by virtue of the preceding `@autosyn` or `@autoinh`. `@warn` is mainly to be used when modifying the attribute grammar.

11.1 Example

The following code fragment in the *Y-file* definition section:

```

:

@autoinh env

@attributes {struct env *env;
             regNumType maxRegNum;
             }
             execElem statement

@autosyn maxRegNum

@attributes {struct env *env;
             @autoinh regNumType regNum;
             regNumType maxRegNum;
             struct sym *formParamList;
             struct sym *func;
             lineNumType line;
             }
             actParamList

@attributes {struct env *env;
             @autosyn @warn struct sym *varLoc,*funcLoc;
             regNumType regNum;
             regNumType maxRegNum;
             }
             block blockElemList

:

```

causes Ox to attempt to automatically supply missing definitions for occurrences of:

- env for execElem, statement, actParamList, block, and blockElemList
- maxRegNum for actParamList, block, and blockElemList
- regNum for actParamList
- varLoc for block and blockElemList, with warning
- funcLoc for block and blockElemList, with warning

12 File-level organization of Ox evaluators

12.1 Conventions of naming Ox output files

By default, Ox translates the *Y-file* into a file named `oxout.y` destined for processing by a Yacc-compatible parser generator. The *L-file(s)* are translated into files destined for a Lex-compatible lexer generator. If there is exactly one *L-file*, its corresponding output file is named `oxout.1`. If there is more than one *L-file*, the corresponding outputs are sequentially named `oxout1.1`, `oxout2.1`, etc.

12.2 Review: combining the outputs of Yacc and Lex

In developing an ordinary (i.e., non-Ox) Yacc/Lex evaluator, `y.tab.c` and `lex.yy.c` can be compiled immediately into an executable file by placing the line

```
#include "lex.yy.c"
```

in a C-code section of the Yacc input specification [Lesk 1975].

Alternatively, Yacc can be instructed (by using the `-d` command-line option) to produce a separate file `y.tab.h` that contains declarations needed by both `y.tab.c` and `lex.yy.c`. The two files may then be compiled separately if the line

```
#include "y.tab.h"
```

is placed in a C-code sections of the Lex input specification. The two resulting object files can then be linked to produce an executable file.

12.3 Combined use of Ox, Yacc, and Lex

There are certain declarations that must be visible from all of the files produced by Ox. By default, Ox produces files suitable for separate compilation, inasmuch as the Yacc-destined file and the Lex-destined file(s) each contain the common declarations. Ox also supports the one-step development approach described above. By placing `-h` on Ox's command line, the designer calls for generation of a file `oxout.h` containing the common declarations, which are then absent from Ox's other output files. In this case, the line

```
#include "oxout.h"
```

is placed in the *Y-file*.

12.4 Typical command sequences

The following sequence of shell commands is an example of the separate compilation approach described. In this example, Ox translates the *Y-file* `ev.Y` into `oxout.y` and the *L-file* `ev.L` into `oxout.l`. The last command of the sequence links the two object files, yielding the executable file `ev`.

```
ox ev.Y ev.L
yacc -d oxout.y
lex oxout.l
cc -c y.tab.c
cc -c lex.yy.c
cc -o ev y.tab.o lex.yy.o -ll -ly
```

The following command sequence does a one-step compilation.

```
ox -h ev.Y ev.L
yacc oxout.y
lex oxout.l
cc y.tab.c -ll -ly
```

13 Command-line options and other points

This section describes some fine points, mostly related to Ox command-line options. Use of those options is summarized in appendix D.

13.1 Error recovery

Yacc has provisions for building parsers that attempt to recover from syntax errors, and the designer can use the Yacc reserved-words `error`, `yerrorok`, and `yyclearin` to implement such error recovery [Johnson 1975]. When a parser that employs such techniques detects a syntax error, it may attempt to recover by popping items from its stack or by discarding tokens. The Ox evaluator's semantic stack operations are synchronized with those of the Yacc-generated parser (see section 7.1). When the evaluator is built using `error`, `yerrorok`, and `yyclearin`, and a syntax error occurs, this synchronization is not lost, since `yyparse()` manages both stacks. However, especially if the evaluator makes use of global variable side effects (see section 8), it is possible for the evaluator's data structures to become corrupted in such cases.

Ox provides the function `yyyabort` to prevent such chaos. The parser calls `yyperror()` upon any syntax error, and the designer can write `yyperror` such that `yyyabort()` is executed at least once each time `yyperror()` is called. Any syntax error will then cancel further execution of all evaluator code, and `yyparse()` can continue safely. Use of `yyyabort()` is unnecessary but harmless if the *Y-file* makes no use of the reserved-words `error`, `yerrorok`, and `yyclearin`.

13.2 Stripping Ox constructs

Occasionally, the designer may wish copies of the *Y-file* and *L-file(s)* free of Ox-specific constructs. Suppose, for instance, that changes to the underlying grammar are under consideration, and that it is desired to test whether the new grammar has parsing conflicts. To satisfy Ox semantics might require writing attribute definitions for any new rules. Ox's output on `oxout.y` could then be submitted to Yacc to test for parsing conflicts.

To avoid the above-mentioned writing of attribute definitions, the designer can use Ox's `-S` command-line option, which filters all Ox-specific constructs from the inputs and yields files acceptable to Yacc and Lex. The original copies of the *Y-file* and *L-file(s)* are unchanged, but the Ox outputs on `oxout.*` contain neither Ox constructs nor the usual Ox-generated parse tree management code.

13.3 Preventing execution of attribute definition code

Faulty user-written code in attribute reference sections may cause abnormal termination of the evaluator. For instance, dereferencing a stray pointer may corrupt the evaluator's data structures and cause it to falsely report a cycle during attribute evaluation. The `-n` command-line option is a debugging feature that can be used to isolate the effects of anomalous attribute definition code. When Ox is used with this option, the generated evaluator uses the ready set as usual to determine an evaluation order for attribute instances, and still checks for cycles. Each time it is ready to solve an instance, however, it stops short of executing the code for the definition of that instance. When `-n` is used, the designer should take special notice of the effects upon other translation phases of such suppression of semantic analysis.

13.4 Parse tree statistics

Placing `-u` on Ox's command line causes generation of an evaluator that prints, for each input, statistics regarding the parse tree built for the input. These include numbers of:

- terminal nodes and their attribute instances,
- nonterminal nodes and their attribute instances,

and other statistics.

14 Example: an integer calculator

This section has Ox code for an evaluator of simple expressions involving multiplication and addition. Since the grammar has only synthesized attributes, the Ox implementation offers little advantage over one that uses only Yacc and Lex; it is presented as a very easy example of Ox usage.

The *L-file* specifies that the tokens are digit strings, parentheses, '*', and '+':

```
%{
/* expr.L: L-file for a simple expression language */
#include "y.tab.h"
#include "oxout.h"
%}

%%
[ \n\t\f]*      ;
[0-9]+          return(CONST); @{
                 sscanf(yytext,"%ld",&@CONST.val@); @}
\(              return('(');
\)              return(')');
\+              return('+');
\*              return('*');
%%
```

The grammar is disambiguated by use of Yacc's `%left` reserved word. Each parse tree node labeled by `s`, `e`, or `CONST` has an integer attribute instance named `val`. Use of the global variable `sVal` obviates postdecoration traversal.

```

/* expr.Y: Y-file for a simple expression language */
%left '+'
%left '*'
%token CONST

@attributes {long val;} s e CONST

%{
#include "oxout.h"
long sVal;
%}

%%
s
    :      e
      @i sVal = @s.val@ = @e.val@;      @}
    ;
e
    :      e      '+'      e
      @i @e.0.val@ = @e.1.val@ + @e.2.val@;  @}
    ;
e
    :      e      '*'      e
      @i @e.0.val@ = @e.1.val@ * @e.2.val@;  @}
    ;
e
    :      '('      e      ')'
      @i @e.val@ = @e.1.val@;      @}
    ;
e
    :      CONST
      @i @e.val@ = @CONST.val@;      @}
    ;
%%

int main()
{yyparse();
 printf("%ld\n",sVal);
}

```

The following command sequence is used to build an executable file `calc` from the above specifications:

```

bash-3.2$ ox -h expr.Y expr.L
bash-3.2$ yacc -d oxout.y
bash-3.2$ lex oxout.l
bash-3.2$ cc -c y.tab.c
bash-3.2$ cc -c lex.yy.c
bash-3.2$ cc -o calc y.tab.o lex.yy.o -ly -ll

```


15 Example: a binary number translator

This illustrates the use of Ox to build an evaluator based on an example attribute grammar that appears in the seminal paper on the subject [Knuth 1968]. The input (after removal of whitespace) is either a nonempty string of binary digits or two such strings separated by a period. This input is interpreted as a binary representation of a floating point number, which is then printed on the standard output in its base-ten form. Removing the Ox-specific constructs and the `printf` statement from the source below yields a pair of files that constitute a semantics-free recognizer of binary numbers.

Construction of this evaluator follows the separate compilation approach described in section 12.

Following is the text of the *L-file*:

```
%{
#include "y.tab.h"
}%

%%

[0]          return ZERO;
[1]          return ONE;
\.          return DOT;
[\\n\\t\\v ] ;
.           {fprintf(stderr,"illegal character\\n");
             exit(-1);
            }
```

Here is the text of the *Y-file*:

```
%token ZERO ONE DOT

@attributes {float value; int scale;}          bit
@attributes {float value; int scale,length;}  bitlist
@attributes {float value;}                   num

%start num

%{
#include <stdio.h>
float numValue;
}%
```

```

%%
bit      :      ZERO
          @i @bit.value@ = 0;
          /* value is synthesized for bit. */
          /* scale is inherited for bit. */
          @}
          ;

bit      :      ONE
          @i @bit.value@ = twoToThe(@bit.scale@);
          @}
          ;

bitlist  :      bit
          @i @bitlist.value@ = @bit.value@;
          @i @bitlist.scale@ = @bitlist.scale@;
          @i @bitlist.length@ = 1;
          /* value and length are synthesized for bitlist. */
          /* scale is inherited for bitlist. */
          @}

          |      bitlist bit
          @i @bitlist.0.value@ = @bitlist.1.value@ + @bit.value@;
          @i @bitlist.scale@ = @bitlist.0.scale@;
          @i @bitlist.1.scale@ = @bitlist.0.scale@ + 1;
          @i @bitlist.0.length@ = @bitlist.1.length@ + 1;
          @}
          ;

num      :      bitlist
          @i numValue = @num.value@ = @bitlist.0.value@;
          @i @bitlist.scale@ = 0;
          /* value is synthesized for num. */
          @}

          |      bitlist DOT      bitlist
          @i numValue = @num.value@ =
          @bitlist.0.value@ + @bitlist.1.value@;
          @i @bitlist.0.scale@ = 0;
          @i @bitlist.1.scale@ = - @bitlist.1.length@;
          @}
          ;

%%
int main()
{if (!yyvsparse())
    printf("%30.15f\n",numValue);
}

float twoToThe(in)          /* returns 2 raised to the power in */
int in;
{if (in < 0) return (1.0 / twoToThe(-in));
    if (in == 0) return 1.0;
    else return (2.0 * twoToThe(in - 1));
}

```

16 Example: translation to postfix and prefix

In this example, the generated evaluator is to perform two postdecoration traversals, one for printing the prefix form of a given infix expression, and one for printing the postfix form. The tokens of the language are specified as follows:

```
%{
/* L-file for translation of infix expressions */
#include "y.tab.h"
#include "oxout.h"

char *lexeme()
{char *dum;
  dum = (char *)malloc(yyleng+1);
  strcpy(dum,yytext);
  return dum;
}
%}

%%
[ \n\t\f]*          ;
[0-9]+\.[0-9]*      return(CONST); @{ @CONST.lexeme@ = lexeme(); @}
[A-Za-z_][A-Za-z_0-9]* return(ID);    @{ @ID.lexeme@ = lexeme();    @}
\(                  return('(');
\)                  return(')');
\+                  return('+');
\*                  return('*');
\/                  return('/');
\-                  return('-');
%%
```

The first traversal performed is named `LRpre`, and the second is named `LRpost`. By default, both are left-to-right traversals. `LRpost` is a postorder traversal by default. `LRpre` is specified as a preorder traversal.

```
/* Y-file for translation of infix expressions to prefix and postfix */
%token ID CONST
%start s
%left '+' '-'
%left '*' '/'

@attributes {char *lexeme;} ID CONST
@traversal @preorder LRpre
@traversal LRpost
```

```

%{
#include "oxout.h"
#include <stdio.h>
%}

%%

s      :      expr
        @{ @LRpost printf("\n");
           @LRpost @revorder (1) printf("postfix:  ");
           @LRpre  @revorder (1) printf("\n");
           @LRpre  printf("prefix:  ");
        @}

expr   :      expr '*' expr
        @{ @LRpost printf(" * ");
           @LRpre  printf(" * ");
        @}
        |      expr '+' expr
        @{ @LRpre  printf(" + ");
           @LRpost printf(" + ");
        @}
        |      expr '/' expr
        @{ @LRpost printf(" / ");
           @LRpre  printf(" / ");
        @}
        |      expr '-' expr
        @{ @LRpost printf(" - ");
           @LRpre  printf(" - ");
        @}
        |      '(' expr ')'
        |      ID
        @{ @LRpost printf(" %s ",@ID.lexeme@);
           @LRpre  printf(" %s ",@ID.lexeme@);
        @}
        |      CONST
        @{ @LRpost printf(" %s ",@CONST.lexeme@);
           @LRpre  printf(" %s ",@CONST.lexeme@);
        @}
        ;

%%

int main()
{yyparse();
}

```

A Using Ox with non-Lex lexical analyzers

A.1 Default context-sensitivity of L-file compilation

Unless instructed otherwise, Ox searches each *L-file* for `return` statements in the context of C-coded Lex actions. Since the string `return` must be ignored outside of that context (for instance, in a Lex regular expression), the default behavior of Ox compilation is to assume that the *L-file* conforms to the syntax of a (possibly Ox-augmented) Lex file. Thus Ox recognizes the three `return` statements in the following (unaugmented) fragment of an *L-file*:

```

:
renames   return(TK_RENAMES);
return    return(TK_RETURN);
reverse   return(TK_REVERSE);
:

```

as points of `return` of tokens by `yylex()`. Ox's sensitivity to the context of the string `return` in the second Lex regular expression above prevents its erroneous recognition as a point of `return` from the lexical analyzer.

A.2 Ox compilation of C-coded lexical analyzers

Ox always ignores the string `return` in the context of C/C++ comments and string constants. When given the `-G` command-line option preceding the name of an *L-file*, it ignores `return` *only* in those contexts. Thus a file containing C code may be augmented with attribute reference sections and input to Ox as an *L-file*. The occurrences of the string `return` must coincide exactly with `returns` of tokens.

A.2.1 Example

Suppose it is desired to convert to Ox a translator that uses the following C code for its lexical analyzer:

```
#include <stdio.h>
#include <string.h>
#include "y.tab.h"

#define bufsize 80
char buf[bufsize];
char *lexBuf;

char *lexeme(inString)
    char inString[];
    {return strcpy((char *)malloc(1+strlen(inString)),inString);}

int yylex()
    {char *bufp = buf;

    while ((*bufp = getchar()) != EOF)
        {if (bufp == (buf + bufsize - 1))
            {fprintf(stderr,"exceeded buffer\n"); exit(-1);}
          if ((*bufp == ' ') || (*bufp == '\n') ||
              (*bufp == '\t') || (*bufp == '\f')
              )
            {if (bufp == buf) continue; else break;}
          if (!isalnum(*bufp)) {fprintf(stderr,"illegal character\n"); exit(-1);}
          bufp++;
        }
    if (bufp != buf)
        {***bufp = '\0';
          lexBuf = lexeme(buf);
          bufp = lexBuf;
          if (isalpha(lexBuf[0]))
              {while (*bufp != '\0')
                  if (isdigit(*bufp++))
                      {fprintf(stderr,"illegal string\n"); exit(-1);}
                  return (IDENT);
                }
          if (isdigit(lexBuf[0]))
              {while (*bufp != '\0')
                  if (isalpha(*bufp++))
                      {fprintf(stderr,"illegal string\n"); exit(-1);}
                  return (ICONST);
                }
        }
    return 0;
}
```

The C reserved word `return` occurs in the file exactly four times. Only two of these occurrences correspond to returns of tokens by the lexical analyzer. The approach is to excise from the file a section of code containing those two occurrences (and no others), place that code section in a separate file, and submit the new file to Ox as a non-Lex *L-file*, by using the `-G` option. Ox translates the new file and places it on `oxout.1`. The excised code is replaced in the original file by the line:

```
#include "oxout.1"
```

Here is the *L-file*, which has been augmented by two attribute reference sections:

```
if (isalpha(lexBuf[0]))
  {while (*bufp != '\0')
    if (isdigit(*bufp++))
      {fprintf(stderr,"illegal string\n"); exit(-1);}
    return (IDENT); @{ @IDENT.string@ = lexBuf; @}
  }
if (isdigit(lexBuf[0]))
  {while (*bufp != '\0')
    if (isalpha(*bufp++))
      {fprintf(stderr,"illegal string\n"); exit(-1);}
    return (ICONST); @{ @ICONST.string@ = lexBuf; @}
  }
```

Here is the skeleton of the lexical analyzer, which now `#includes` the files `oxout.h` and `oxout.l`:

```
#include <stdio.h>
#include <string.h>
#include "y.tab.h"
#include "oxout.h"

#define bufsize 80
char buf[bufsize];
char *lexBuf;

char *lexeme(inString)
    char inString[];
    {return strcpy((char *)malloc(1+strlen(inString)),inString);}

int yylex()
    {char *bufp = buf;

    while ((*bufp = getchar()) != EOF)
        {if (bufp == (buf + bufsize - 1))
            {fprintf(stderr,"exceeded buffer\n"); exit(-1);}
          if ((*bufp == ' ') || (*bufp == '\n') ||
              (*bufp == '\t') || (*bufp == '\f')
              )
              {if (bufp == buf) continue; else break;}
          if (!isalnum(*bufp)) {fprintf(stderr,"illegal character\n"); exit(-1);}
          bufp++;
        }
    if (bufp != buf)
        {***bufp = '\0';
          lexBuf = lexeme(buf);
          bufp = lexBuf;
        }

#include "oxout.l"

    }
    return 0;
}
```


B Traversal semantics

The behavior of postdecoration traversals was illustrated in the examples of section 9.2. In view of those examples, the C-like pseudocode in this appendix holds no surprises, but describes such behavior somewhat more formally. The traversals are carried out by a single call of `doTraversals` (below) after the final decoration.

```
enum orderType {PREORDER,POSTORDER};
enum directionType {LEFTTORIGHT,RIGHTTOLEFT};

enum orderType staticOrder(traversal T)
  {if (@preorder appears in the traversal specification of T)
    return PREORDER;
   return POSTORDER;
 }

enum directionType staticDirection(traversal T)
  {if (@righttoleft appears in the traversal specification of T)
    return RIGHTTOLEFT;
   return LEFTTORIGHT;
 }

int isDisabled(traversal T)
  {if (@disable appears in the traversal specification of T)
    return 1;
   return 0;
 }
```

```

void pdTrav(parse_tree_node N, traversal T)
{grammar_rule R;          /* the rule applied at N */
  enum orderType order[Z]; /* Z >= # of traversal action specs
                           for T in R */

  enum directionType direction;
  int i,j,k;

  R = the grammar rule applied at N;
  let the traversal actions for T in the attribute reference section
    of R be numbered from 0 to k-1;
  for (i=0; i<k; i++)
    {if (the ith traversal action specifier has no @revorder)
      order[i] = staticOrder(T);
      else if ((the expression associated with @revorder) == 0)
        order[i] = staticOrder(T);
      else if (staticOrder(T) == POSTORDER)
        order[i] = PREORDER;
      else
        order[i] = POSTORDER;
      if (the ith traversal action specifier has no @revdirection)
        direction = staticDirection(T);
      else if ((the expression associated with @revdirection) == 0)
        direction = staticDirection(T);
      else if (staticDirection(T) == LEFTTORIGHT)
        direction = RIGHTTOLEFT;
      else
        direction = LEFTTORIGHT;
    }
  for (i=0; i<k; i++)
    if (order[i] = PREORDER)
      execute the ith traversal action;
  number the children of N from left to right
    with integers from 0 to j-1;
  if (direction == LEFTTORIGHT)
    for (i=0; i<j; i++) pdTrav(the ith child of N,T);
  else
    for (i=j-1; i>=0; i--) pdTrav(the ith child of N,T);
  for (i=0; i<k; i++)
    if (order[i] = POSTORDER)
      execute the ith traversal action;
}

void doTraversals()
{int i,k;
  parse_tree_node r;

  r = the root of the parse tree;
  k = the number of traversals;
  number the traversals from 0 to k-1, according to
    the order of appearance of their specifications;
  for (i=0; i<k; i++)
    if (!isDisabled(the ith traversal))
      pdTrav(r,the ith traversal);
}

```

C List of reserved words and reserved file names

The Ox reserved words are as follows:

```
@{
@}
@attributes
@autoinh
@autosyn
@decorate
@disable
@e
@end
@i
@lefttoright
@m
@macro
@postorder
@preorder
@revdirection
@revorder
@righttoleft
@traversal
@void
@warn
```

The following file names in the current directory are reserved for use by Ox:

```
oxout.h                (when using -h option)
oxout.y
oxout.l                (when using exactly one L-file)
oxout1.1, oxout2.1, oxout3.1, ... (when using more than one L-file)
```

D Summary of command-line options

The Ox command line takes the form:

```
ox { option } Y-file [-G] L-file { [-G] L-file }
```

By default, Ox generates an output file for each file mentioned on the command-line. The *Y-file* generates a parser specification file `oxout.y` suitable for processing by a Yacc-compatible parser generator. A single *L-file* generates a lexer specification file `oxout.l` suitable for processing by a Lex-compatible lexer generator. More than a single *L-file* generates lexer specification files sequentially named `oxout1.1`, `oxout2.1`, etc.

The `-G` annunciator and the *options* are described as follows:

- `--help` Show the Ox command line usage summary, and exit.
- `-b, --prefix=basename`
Use *basename* instead of `oxout` for constructing Ox source output filenames.
- `-c` Generated evaluator does not produce cycle reports.
- `-d` Generated evaluator does not do automatic pruning.
- `-G L-file`
L-file contains a generic (i.e., non-Lex) scanner hand-written in C. Except for attribute reference sections, the *L-file* must conform to Lex & C syntax. The `return` reserved word is recognized in any context other than comments and string literals. See appendix A.
- `-h` Produce an Ox header file `oxout.h` to be `#included` in a code section (between `%{` and `%}`) in the *Y-file* or *L-file(s)*. This permits one-step compilation of the parser and scanner(s). When this option is not used, Ox places the header information in each output file rather than in a separate header file. See section 12.3.
- `-j` Save bad output files. Ox normally skips writing output files if there is an error.
- `-L, -l, --noline`
Suppress generation of `#line` directives; `#line` directives specified by the user will be retained.
- `-n` Generate an evaluator that determines an evaluation order and checks for cycles, but does not execute the code that evaluates attribute instances. See section 13.3.
- `-p, --name-prefix=prefix`
Use *prefix* instead of `yyy` for constructing Ox-generated external symbols.

- S, --strip
Strip Ox-specific constructs from the *Y-file* and *L-file(s)* and place the pure Yacc and pure Lex results on *oxout.y* and *oxout*.1*, respectively. See section 13.2.
- t
Generate an evaluator that builds an internal visual representation of a parse tree; the evaluator must be linked with *libcgraph*. See section F.
- u
Generate an evaluator that prints parse tree memory usage statistics for each input. See section 13.4.
- V, --version
Show the Ox version number, and exit.
- header-prefix=*basename*
Use *basename* instead of *oxout* for constructing the Ox header source output filename (overrides --prefix=*basename*, if present).
- lex-prefix=*basename*
Use *basename* instead of *oxout* for constructing the Ox lexer source output filename (overrides --prefix=*basename*, if present).
- parse-prefix=*basename*
Use *basename* instead of *oxout* for constructing the Ox parser source output filename (overrides --prefix=*basename*, if present).
- reentrant-parser, --pure-parser
Compile *Y-file* assuming it is a reentrant (pure) parser; compile each *L-file* accordingly: the lexer communication variable *yylval* is passed as a `YYSTYPE *` value to `yylex()` from a reentrant parser. If Flex is used to generate the scanner, it must be invoked with either the --bison-bridge or the --c++ command-line option.
- language=[C|C++]
Compile *Y-file* and each Lex *L-file* with C or C++ as the target language (default: --language=C). This option overrides any prior target language choice on the command line. This option can be overridden for *Y-file* by a subsequent occurrence of one of the target parser generator options below. This option can be overridden for each Lex *L-file* by a subsequent occurrence of one of the target lexer generator options below.
- bison[=C|C++]
Compile *Y-file* for Bison in C (default, unless --language=C++ precedes on the command line) or C++; if C++ then `bison --language=c++` (or equivalent) must be used for *oxout.y* (implies --reentrant-parser).
- btyacc[=C|C++]
Compile *Y-file* for BtYacc in C (default, unless --language=C++ precedes on the command line) or C++.

- `--byacc [=C|C++]`
 Compile *Y-file* for BYacc in C (default, unless `--language=C++` precedes on the command line) or C++.
- `--msta [=C|C++]`
 Compile *Y-file* for Msta in C (default, unless `--language=C++` precedes on the command line) or C++; if C++ then `msta -c++` must be used for `oxout.y`.
- `--yacc [=C|C++]`
 Compile *Y-file* assuming only Yacc functionality in C (this is the default for a *Y-file*, unless `--language=C++` precedes on the command line) or C++.
- `--flex [=C|C++]`
 Compile each Lex *L-file* for Flex in C (default, unless `--language=C++` precedes on the command line) or C++; if C++ then `flex --c++` (or equivalent) must be used for each `oxout.l`.
- `--lex [=C|C++]`
 Compile each Lex *L-file* assuming only Lex functionality in C (this is the default for a Lex *L-file*, unless `--language=C++` precedes on the command line) or C++.
- `--reflex [=C++]`
 Compile each Lex *L-file* for RE/flex in C++.

The filename extension used for each type of output file (parser specification, header, and lexer specification) is based on the corresponding input filename extension, as follows:

- For a 1-character (or omitted) input filename extension, the output filename extensions `.y`, `.h`, or `.l` are used, respectively.
- For an input filename extension with exactly 2 identical characters, the output filename extensions `.yy`, `.hh`, or `.ll` are used, respectively.
- Otherwise, the output filename extension is constructed using the input filename extension, converted to lower case, and with the first character replaced by `y`, `h`, or `l`, respectively.

If a target parser generator is not specified on the command-line, the target will be inferred from the first occurrence of a directive that is unique to a supported target parser generator during processing of the *Y-file* definition section. If no such directive is seen in the definition section, the default target (Yacc) is assumed. An occurrence of a directive that is inconsistent with the target is a syntax error.

E Lex- and Yacc-compatible tool interoperation

In general, Ox is interoperable with the Lex, Flex and RE/flex lexer generators, and with the Yacc, BYacc, BtYacc, Bison and Msta parser generators.

Ox makes use of the following Lex- and Yacc-compatible declarations and directives in analysis and generation of the evaluator code:

`%token, %term, %left, %right, %nonassoc, %precedence`

Ox identifies the terminal symbols of the grammar from these declarations. Character literal symbols are always terminal symbols, whether or not they are declared.

`%token <name> <"literal string token">`

This Bison extension defines `<"literal string token">` as an alias for the token `<name>`, and both can be used interchangeably in further declarations or in the grammar rules. Ox recognizes this extension, and likewise treats `<"literal string token">` as an interchangeable alias for the token `<name>` in attribute declarations and occurrences.

`%union { ... }`

Ox augments a Yacc `%union` semantic value type declaration with an Ox-generated semantic value type member that is the union of the Ox-generated attribute implementations. This added member does not interfere with the original Yacc `%union` semantic value type members.

`%define api.namespace {<namespace>}`

This Bison directive specifies the namespace for the parser class, replacing the default value `yy`; implies `%language "C++"`.

`%define api.parser.class {<name>}`

This Bison directive specifies the name of the parser class, replacing the default value `parser`; implies `%language "C++"`.

`%define api.pure <purity>`

This Bison directive requests generation of a reentrant (pure) parser, if `<purity>` is `true`, `full`, or omitted. Each *L-file* is compiled accordingly: the lexer communication variable `yylval` is passed as `YYSTYPE *` value to `yylex()` from a reentrant parser.

`%define api.token.prefix {<prefix>}`

This Bison directive adds `<prefix>` to named symbols when generating their definition in C. If the `"%define api.value.type union"` directive is

given, the member name of the generated semantic value type union corresponding to a named symbol with a declared type is the prefixed symbol name. Ox recognizes this directive and treats prefixed symbol names as aliases for the named symbols in attribute occurrences.

`%define api.value.type union`

This Bison directive interprets the tags of grammar symbol type declarations as types, rather than as member names of the `%union` semantic value type, and generates the semantic value type union (YYSTYPE) from those types. The Ox-generated semantic value type is used as the tag for the Ox-generated grammar start symbol.

`%define namespace {<namespace>}`

Same as `%define api.namespace {<namespace>}`.

`%define parser_class_name {<name>}`

Same as `%define api.parser.class {<name>}`.

`%language "<language>"`

This Bison directive specifies the programming language for the generated parser. Ox supports "C" and "C++" (case insensitive); the `%language "C++"` directive implies `%define api.pure`.

`%option bison-bridge`

This Flex and RE/flex *L-file* option directs Ox to reference `yylval` as a YYSTYPE * value in the generated `oxout*.1` file, for use with a Bison or BYacc reentrant parser.

`%parse-param {<parameter_declaration>}, %param {<parameter_declaration>}`

These Bison and BYacc directives specify `<parameter_declaration>` be added as a parameter to the signature of `yyparse()` (and `yylex()`, etc.). The syntax for `<parameter_declaration>` is the same as for the formal parameter of a C function prototype. Subsequent directives accumulate, in order of appearance. Ox recognizes this directive and ensures that the parameter can be referenced within any attribute evaluation part.

`%pure-parser`

This Bison and BYacc directive requests generation of a reentrant parser. Each *L-file* is compiled accordingly: the lexer communication variable `yylval` is passed as YYSTYPE * value to `yylex()` from a reentrant parser.

`%skeleton "<file>"`

This Bison directive specifies the skeleton file to use. Ox supports filename extensions `.c` (implies `%language "C"`) and `.cc` (implies `%language "C++"`).

Specific interoperability issues of note:

- Ox passes through all `%option` directives supported by Flex and RE/flex.
- Ox passes through most `%`-directives supported by Bison, BYacc, BtYacc and Msta, but not all `%`-directives that require Ox to take some action are supported.
- Ox works with the Bison and BYacc reentrant parser skeletons; however, the Ox evaluator skeleton (currently) uses global variables and so the generated evaluator is not reentrant.
- Ox works with the Bison push and GLR parser skeletons, with the BYacc and BtYacc backtracking parser skeletons, and with the Msta LR(k>1) parsing skeleton.
- Ox works with the Flex, RE/flex, Bison, BtYacc and Msta C++ skeletons.
- A reference to a Bison and BYacc positional location pseudo variable (`@$, @1, @2, etc.`), is allowed within an attribute occurrence definition that can be evaluated at production reduction time, as for a Yacc positional semantic value pseudo variable.
- Ox accepts the Bison *named reference* syntax for semantic value and location pseudo variables (`$NAME, @NAME, $[NAME], @[NAME]`), interchangeably with positional semantic value and location pseudo variables.
- BYacc and BtYacc inherited attributes may be used in evaluator specifications, but references to inherited attributes may not be used within Ox attribute occurrence definitions.
- Ox passes through the following literal C/C++ code blocks:
 - `%{ ... %}`
 - `%top { ... }` (Flex, RE/flex)
 - `%class { ... }` (RE/flex)
 - `%printer { ... }` (Bison)
 - `%@ { ... }` (BtYacc)
 - `%location { ... }` (BtYacc)
 - `%position { ... }` (BtYacc)
 - `%local { ... }` (Msta)
 - `%import { ... }` (Msta)
 - `%export { ... }` (Msta)

- %initial-action { ... } (BYacc, BtYacc, Bison)
- %code { ... } (BYacc, Bison)
- %code top { ... } (BYacc, Bison)
- %code requires { ... } (BYacc, Bison)
- %code provides { ... } (BYacc, Bison)
- %destructor { ... } (BYacc, Bison)
- Ox does not support the following RE/flex directives:
 - %option bison-complete
 - %bison-complete
- Ox does not support the following Bison directives:
 - %define api.token.automove
 - %define api.token.constructor
 - %define api.value.type variant
 - %define api.value.type {*type*}
 - %merge <*merge function name*>
- Ox does not recognize the Msta regular right-part grammar rule syntax.

F Using *Graphviz* for parse tree visualization

The *Graphviz* [Graphviz] package enables creation, manipulation, layout, rendering and visualization of attributed graphs. The libraries and programs in the package use the DOT language as a common external textual representation of graphs. The external representation of graph, node and edge objects, and the corresponding internal representation of those objects, can be annotated with DOT language attribute-value pairs to tailor graph layout, rendering and display.

The Ox `-t` command-line option results in the inclusion of code in the evaluator to build an internal visual representation of a parse tree, using only the *cgraph* library[Gansner 2014][North 2014]. It is up to the user to use the capabilities of the *Graphviz* package to layout and render for display the internal visual representation within the evaluator, or to serialize the internal representation into a file for external processing.

Any graph, node or edge object attribute-value pair set by an evaluator can be changed, and other attribute-value pairs can be added, by user-written code. User-defined node and edge objects can also be added to create more elaborate visual decoration external to the default visual representation. Changes or additions can also be accomplished by post-processing the DOT language external representation.

Ox provides the global variable `yyyTreeVizGraph` (type: `Agraph_t *`) for accessing the *cgraph* graph object representing the parse tree. Ox also provides the *Node-Image Object* (NIO) pseudo attribute¹³ (type: `Agnode_t *`) for all grammar symbols as syntactic sugar for accessing the corresponding *cgraph* node object (and any attached edge objects) during evaluation. The NIO pseudo attribute can occur in traversal actions or in the evaluation parts of attribute definitions. It is not included in attribute dependency computations.

The parse tree internal visual representation is constructed during the initial bottom-up parsing phase of the evaluator. Consequently, the terminal node accessed within an *L-file* attribute reference section will not have an *in-edge* connecting its parent nonterminal node, since that parent node will not yet have been created. The nonterminal node corresponding to the LHS symbol accessed within a *Y-file* attribute reference section will have an *out-edge* connecting each child node (corresponding to an RHS symbol). However, it may or may not have an in-edge connecting its parent nonterminal node, depending on the order of evaluation determined by Ox for the attribute definition in which the NIO pseudo attribute occurs. Modifications involving in-edge objects are better left to traversal actions.

Modification of the *cgraph* data structures via an NIO pseudo attribute, other than via *cgraph* functions, is not recommended. It is possible to corrupt the parse tree internal visual representation if the *cgraph* data structures are not fully understood.

¹³In the absence of the Ox `-t` command-line option, a warning diagnostic will be issued if NIO is declared as an attribute. In the presence of the `-t` command-line option, an explicit declaration of an NIO attribute is disallowed.

F.1 The default parse tree configuration

Ox creates a non-strict (multiple edges allowed between two nodes) directed graph to represent a parse tree, with the following graph object attribute-value pairs:

- `ordering = out`
- `outputorder = breadthfirst`
- `label = Ox-generated evaluator parse tree`
- `labeljust = left`
- `rankdir = LR`
- `splines = polyline`

Edge objects are created with names `Edgei`. Node objects representing parse tree terminal nodes are created with names `Leafi`, and those representing nonterminal nodes are created with names `Nodei`¹⁴.

Edge objects are created with the following attribute-value pairs:

- `arrowhead = none`
- `style = dashed,bold`

Node objects are created with the following attribute-value pairs:

- `label = the corresponding Yacc grammar symbol name`
- `style = bold`
- `shape = record`¹⁵ (nonterminal) *or* `Mrecord`¹⁶ (terminal)

A record node is a structured collection of rectangles containing user-defined text, defined by the node object `label` attribute value. Defining a record structure is a simple way to add textual decoration programmatically.

¹⁴In each object category (`Edge`, `Leaf`, `Node`), *i* begins with 1.

¹⁵The `record` shape is a rectangle.

¹⁶The `Mrecord` shape is a rectangle with rounded corners.

F.2 Specifying *cgraph* node structure

A record-based node label value has the following syntax:

```

recordLabel = field ( '|' field )*
field       = fieldId | '{' recordLabel '}'
fieldId    = [ '<' portName '>' ] [ literal ]
portName  = literal

```

Lexical and semantic details:

- The characters '|', '{', '}', '<', '>' and '\' must be escaped with a '\' character to appear as a literal character in a record-based node label value. Spaces are interpreted as separators between tokens, and so also must be escaped to appear as a literal character.
- The optional *portName* in a *fieldId* is used to specify the *field* of a record-based node to which to attach an edge. Internally, the edge object `headport` (`tailport`) attribute value is the *portName* to which the head (tail) of the edge object is attached.
- The *literal* in a *fieldId* is used as the text for the field; it supports the line division escape sequences '\n' (centered), '\l' (left-justified) and '\r' (right-justified).
- Visually, a record is a box, with fields represented by alternating rows of horizontal or vertical subboxes. Flipping between horizontal and vertical layouts is done by nesting fields within braces "{...}" .
- The initial orientation of a record node depends on the graph object `rankdir` attribute value. If this attribute has value TB or BT, corresponding to vertical layouts, the top-level fields in a record are displayed horizontally. If, however, this attribute has value LR (the default) or RL, corresponding to horizontal layouts, the top-level fields are displayed vertically.
- For the default parse tree layout, a record with label "A | B | C | D" will have the fields oriented top to bottom, while "{A | B | C | D}" will have them oriented left to right.

The DOT language also supports *HTML-like label* values, which generalize record-based label values. If a node object has the attribute `shape` set to the value `none` or `plaintext`, then the HTML value of the node's `label` attribute will define the node's shape. If the node object has any other shape (with the exception of `point`), the HTML value of the node's `label` attribute will be embedded within the node as for an ordinary literal label. HTML label values provide a more elaborate mechanism for decorating an individual node, but have a more complex syntax and are not dealt with further here.

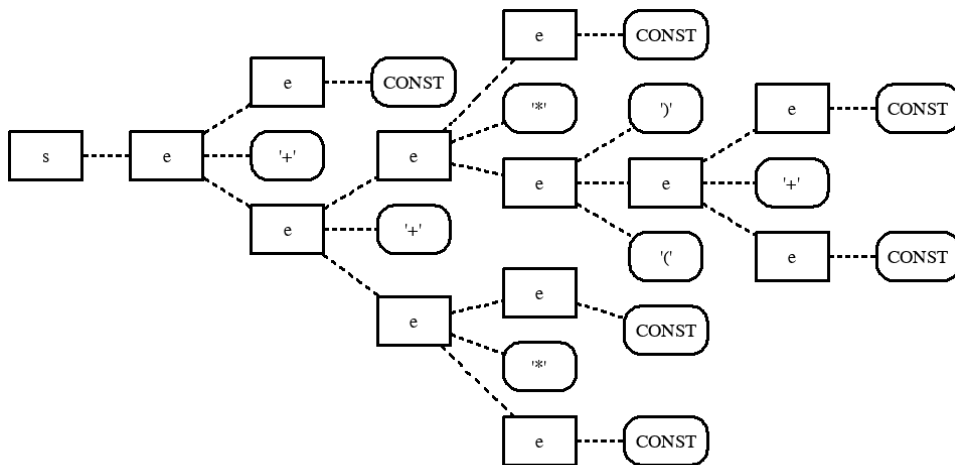
F.3 Example: the integer calculator, revisited

The internal representation is serialized into the DOT language external representation into a file by user-written code using the *cgraph* function `agwrite()`. Revisiting the example from section 14, this modified `main()` program will write the DOT language representation of a parse tree into the file `default-treeout.gv`.

```
#include <graphviz/cgraph.h>
int main()
{FILE *treeout = fopen("default-treeout.gv","w");
  yyparse();
  agwrite(yyyTreeVizGraph,treeout);
  printf("%ld\n",sVal);
}
```

The command to link the executable file `calc` with *libcgraph* is shown below, followed by a sample execution with output, the *Graphviz* `dot` command to layout and render the file `default-treeout.gv` into a Portable Network Graphics image format file (other formats are supported), and then the parse tree image generated:

```
bash-3.2$ cc -o calc y.tab.o lex.yy.o -ly -ll -lcgraph
bash-3.2$ ./calc <<< '4*9+(8+0)*8+2'
102
bash-3.2$ dot -Tpng default-treeout.gv -o default-treeout.png
```



Ox-generated evaluator parse tree

The global `yyyTreeVizGraph` is initialized within `yyyinit()`; once `yyyinit()` has been called, `yyyTreeVizGraph` can be accessed within user-written code during or after the parsing phase (during or after attribute evaluation). Since `yyyTreeVizGraph` is not modified by `agwrite()`, the internal visual representation can be further modified after invoking `agwrite()`.

Note that the internal visual representation can also be laid out and rendered within the evaluator with one of many supported image formats in the *Graphviz* package into a file. The *gvc* library [Gansner 2014] from the *Graphviz* package includes the functions `gvlayout()` and `gvRender()` or `gvRenderFilename()` that can be used for this purpose.

The following change to the attribute reference section in the *L-file* rule for pattern `[0-9]+` of the section 14 example will add the value of each `CONST` token to the parse tree terminal node visual representation.

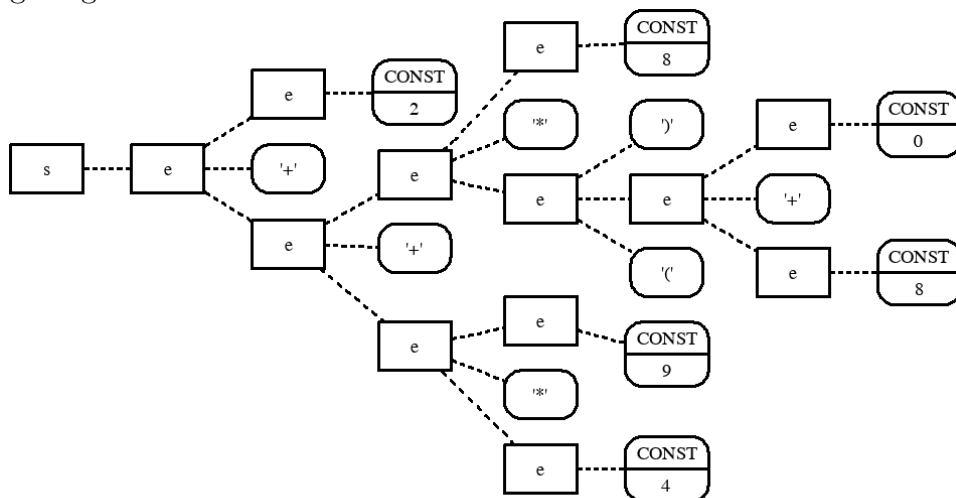
```
[0-9]+      return(CONST); @{
            {long val; char leaf_label[20];
             sscanf(yytext,"%ld",&val); @CONST.val@ = val;
             snprintf(leaf_label,sizeof(leaf_label),"CONST|%ld",val);
             agset(@CONST.NIO@,"label",leaf_label);}
            @}

```

Revising the `main()` program to write the DOT language representation into the file `modleaf-treeout.gv`, and changing the label associated with the parse tree:

```
#include <graphviz/cgraph.h>
int main()
{FILE *treeout = fopen("modleaf-treeout.gv","w");
  yyparse();
  agattr(yyyTreeVizGraph, AGRAPH, "label",
        "AG parse tree w/ modified terminal nodes");
  agwrite(yyyTreeVizGraph,treeout);
  printf("%ld\n",sVal);
}
```

and using the `dot` command to layout and render that file, as above, this parse tree image is generated:



AG parse tree w/ modified terminal nodes

The addition of the traversal actions shown below to the *Y-file* of the section 14 example will add decorations to the parse tree nonterminal node visual representations showing the calculations represented by the corresponding rule.

```

/* expr.Y: Y-file for a simple expression language */
%left '+'
%left '*'
%token CONST

@attributes {long val;} s e CONST
@traversal PV

%{
#include "oxout.h"
long sVal;
%}

%%

s
:
e
@{ @i sVal = @s.val@ = @e.val@;
  @PV {snprintf(node_label,sizeof(node_label),
               "s|result=%ld",@s.val@);
        agset(@e.NI0@,"label",node_label);} @}
;

e
:
e '+' e
@{ @i @e.0.val@ = @e.1.val@ + @e.2.val@;
  @PV {snprintf(node_label,sizeof(node_label),
               "e|%ld+%ld=%ld",@e.1.val@,@e.2.val@,@e.0.val@);
        agset(@e.0.NI0@,"label",node_label);} @}
;

e
:
e '*' e
@{ @i @e.0.val@ = @e.1.val@ * @e.2.val@;
  @PV {snprintf(node_label,sizeof(node_label),
               "e|%ld*%ld=%ld",@e.1.val@,@e.2.val@,@e.0.val@);
        agset(@e.0.NI0@,"label",node_label);} @}
;

e
:
'(' e ')'
@{ @i @e.val@ = @e.1.val@;
  @PV {snprintf(node_label,sizeof(node_label), "e|(%ld)",@e.1.val@);
        agset(@e.0.NI0@,"label",node_label);} @}
;

e
:
CONST
@{ @i @e.val@ = @CONST.val@;
  @PV {snprintf(node_label,sizeof(node_label), "e|%ld",@e.val@);
        agset(@e.NI0@,"label",node_label);} @}
;

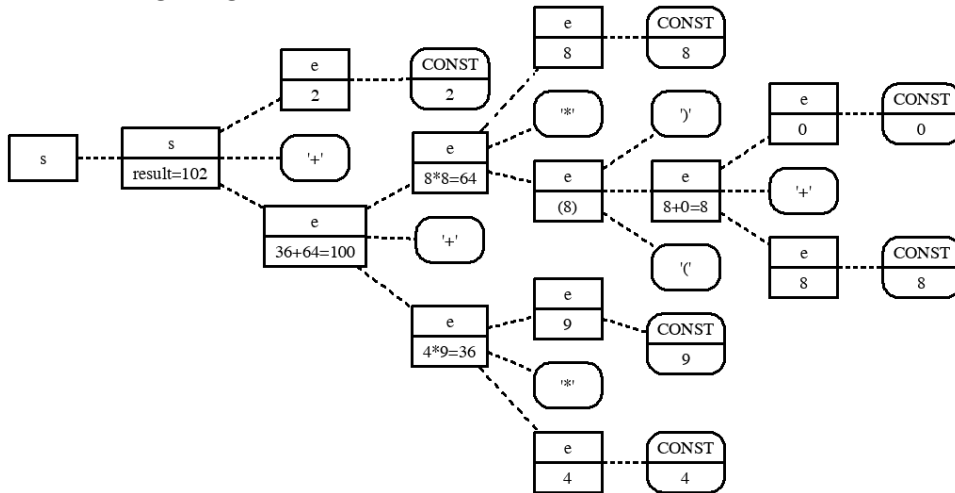
%%

```



```
#include <graphviz/cgraph.h>
int main()
{FILE *treeout = fopen("modnode-treeout.gv","w");
  yyparse()
  agattr(yyyTreeVizGraph, AGRAPH, "label",
    "AG parse tree w/ modified nonterminal & terminal nodes");
  agwrite(yyyTreeVizGraph,treeout);
  printf("%ld\n",sVal);
}
```

Using the dot command to layout and render the DOT language file, as before, this parse tree image is generated:



AG parse tree w/ modified nonterminal & terminal nodes

References

- [Johnson 1975] Stephen C. Johnson, *Yacc: Yet Another Compiler Compiler*, Computing Science Technical Report No. 32, AT&T Bell Laboratories, Murray Hill, New Jersey, 1975. Reprinted as PS1:15 in *UNIX Programmer's Manual*, Usenix Association, 1986.
- [Lesk 1975] M.E. Lesk and E. Schmidt, *Lex-A Lexical Analyzer Generator*, Computing Science Technical Report No. 39, AT&T Bell Laboratories, Murray Hills, New Jersey, October 1975. Reprinted as PS1:16 in *UNIX Programmer's Manual*, Usenix Association, 1986.
- [KR 1988] Brian W. Kernighan and Dennis M. Ritchie, *The C Programming Language, 2nd Ed.* Prentice-Hall, 1988.
- [Waite 1984] William M. Waite and Gerhard Goos, *Compiler Construction*, Springer-Verlag, 1984.
- [Aho 1986] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman, *Compilers: Principles, Techniques, and Tools*, Addison-Wesley, 1986.
- [Jazayeri 1975] M. Jazayeri, W.F. Ogden, and W.C. Rounds, *The Intrinsic Exponential Complexity of the Circularity Problem for Attribute Grammars*, Communications of the ACM, Vol. 18, No. 12, pp. 697-706, December 1975.
- [Lorho 1988] Pierre Deransart, Martin Jourdan, and Bernhard Lorho, *Attribute Grammars: Definitions, Systems, and Bibliography*, Lecture Notes in Computer Science, v. 323, Springer Verlag, 1988.
- [Knuth 1968] Donald E. Knuth, *Semantics of Context-Free Languages* Mathematical Systems Theory, Vol. 2, No. 2, pp. 127-145, 1968.
- [Gansner 2014] Emden R. Gansner, *Using Graphviz as a Library (cgraph version)*, August 21, 2014, available online here: https://graphviz.gitlab.io/_pages/pdf/libguide.pdf.
- [North 2014] Stephen C. North, Emden R. Gansner, *Cgraph Tutorial*, February 7, 2014, available online here: https://graphviz.gitlab.io/_pages/pdf/cgraph.pdf.
- [Graphviz] The Graphviz distribution and documentation can be found on <https://graphviz.gitlab.io>.

- [Flex] The Flex distribution can be retrieved from <https://github.com/westes/flex/>.
- [REflex] The RE/flex distribution can be retrieved from <https://github.com/Genivia/RE-flex/>.
- [Bison] The Bison distribution can be retrieved from <http://www.gnu.org/software/bison/>.
- [BtYacc] The BtYacc distribution can be retrieved from <https://github.com/ChrisDodd/btyacc/>.
- [BYacc] The BYacc distribution can be retrieved from <http://invisible-island.net/byacc/>.
- [Msta] Msta is part of the programming language Dino distribution, and can be retrieved from <https://github.com/dino-lang/dino/>.

Index

- @{, 7, 8, 10, 15
- @}, 7, 8, 10
- @\$, @1, @2, ..., (positional location pseudo variable), **48**
- @attributes, 5, 6
- @autoinh, 24, 25
- @autosyn, 24, 25
- @disable, 20
- @e, 8
- @i, 8
- @lefttoright, 20
- @m, 8
- @postorder, 20
- @preorder, 20
- @revidirection, 21
- @revorder, 21
- @righttoleft, 20
- @warn, 25
- \$(NAME), @(NAME), \$[NAME], @[NAME], ..., (named reference semantic value and location pseudo variable), **48**
- \$\$, \$1, \$2, ..., (Yacc positional semantic value pseudo variable), **13**, 15
- oxout.*, 27
- oxout?.1, 27
- %left, 5
- %nonassoc, 5
- %right, 5
- %token, 5

- agwrite, 53
- agwriteagwrite, 53
- ambiguous form of return of token, 12
- attribute
 - inherited, **7**
 - synthesized, **7**
- attribute (as belonging to a symbol), **6**
- attribute declaration, **5**
- attribute definition, **8**
 - dependency part of, **8**
 - evaluation part of, 5, **8**
 - explicit mode, 8
 - implicit mode, 9
 - mixed mode, 9
- attribute definition modes, 8
- attribute grammar
 - class of AGs accepted by Ox, 10
 - execution sequence not explicit in, 16
- attribute instance, **6**
 - solving an, **9**
- attribute instance (as belonging to a node), **6**
- attribute instances
 - ready set of, **15**
- attribute occurrence, **6**
 - dependees of an, 8
 - dependents of an, 8
 - inherited, **7**
 - synthesized, **7**
- attribute reference, **9**
- attribute reference section, **7**, **10**
- attribute reference section delimiters, 7, 8, 10

- circular grammar, **13**
- code generation, 17
- command-line options, 43
- command-line syntax, 43
- comments, 4
- copy rule, **24**
- cycle (in an attributed parse tree), **13**
- cycle detection, 13

- declaration

- attribute, **5**
- decoration, **15**
- defined, attribute occurrence, **8**
- definition
 - attribute, **8**
- definition mode annunciator, **8**
- dependee, **8**
- dependency expression, **8**
- dependency part
 - of a mixed mode attribute definition, **9**
 - of an attribute definition, **8**
 - of an explicit mode attribute definition, **8**
 - of an implicit mode attribute definition, **9**
- dependent, **8**
- depends upon, **8**
- dynamic traversal modifier, **20**
- error, **28**
- evaluation part
 - of a mixed mode attribute definition, **9**
 - of an attribute definition, **5, 8**
 - of an explicit mode attribute definition, **8**
 - of an implicit mode attribute definition, **9**
- example
 - Knuth's classical, **32**
 - very easy, **30**
- execution sequence not explicit in
 - attribute grammars, **16**
- explicit mode annunciator, **8**
- explicit mode attribute definition, **8**
- file names
 - Ox output, **27**
- final decoration, **15, 18**
- global variables
 - reference to C's, **17**
 - Graphviz*, **16, 50, 53, 54**
- header file `oxout.h`, **27**
- home rule, **7**
- implicit mode annunciator, **9**
- implicit mode attribute definition, **9**
- inherited attribute, **7**
- inherited attribute occurrence, **7**
- instance
 - attribute, **6**
- Knuth's classical example, **32**
- L-file, **4, 10, 11, 14, 18, 27–30, 32, 36, 38, 42–47, 50, 54**
- LALR(1) property preserved by Ox compilation, **14**
- left-hand side, **6**
- `lex.yy.c`, **27**
- LHS (left-hand side), **6**
- libcgraph*, **16, 50, 53**
- libgvc*, **54**
- macros, **22**
- macros forbidden for `return` of `yylex`, **10**
- mixed mode annunciator, **9**
- mixed mode attribute definition, **9**
- mode
 - attribute definition, **8**
- mode annunciator
 - `@e` (explicit), **8**
 - `@i` (implicit), **8, 9**
 - `@m` (mixed), **8, 9**
 - definition, **8**
 - traversal, **19**
- NIO, **50**
- occurrence
 - attribute, **6**

- options
 - command-line, 43
- parent rule, 7
- placement **new** operator, 5
- postdecoration, 18
- postdecoration traversal, 18
- pruning, 15, 16, 18
- pseudo variable
 - named reference semantic value
 - and location, 48
 - positional location, 48
 - Yacc positional semantic value, 13, 15
- ready set (of attribute instances), 15
- reference
 - attribute, 9
- return**
 - from `yylex` must be explicit, 10
- return** statements
 - lexical rules associated with, 10
- RHS (right-hand side), 6
- right-hand side, 6
- rule, 6
 - home, 7
 - parent, 7
 - returned token as a, 6
- rules section of a Yacc file, 8
- side effects, 17
- solving (an attribute instance), 9
- synchronization
 - Ox and Yacc stack, 14, 28
- synthesized attribute, 7
- synthesized attribute occurrence, 7
- token
 - inherited attributes of a, 10
 - synthesized attributes of a, 10
- traversal
 - postdecoration, 18
 - traversal action, 20
 - traversal action specification, 20
 - traversal action specifier, 19
 - traversal mode annunciator, 19, 20
 - traversal modifier
 - dynamic, 20
 - traversal specification, 20
 - traversal specifier, 20
 - traversal specifier sequence, 20
- Y-file, 4, 5, 8, 10, 13–15, 19, 20, 22–24, 26–29, 32, 43–45, 50, 55
- `y.tab.c`, 27
- `y.tab.h`, 27
- `yyclearin`, 28
- `yyerrok`, 28
- `yyerror`, 28
- `yy leng`, 16
- `yy lex`, 10, 14, 36, 44, 47
- `yy lval`, 14, 44, 46, 47
- `yy val`, 14
- `yy parse`, 14, 17, 28, 47
- `yy text`, 16
- `yyy abort`, 28
- `yyy init`, 16, 53
- `yyy TreeVizGraph`, 50, 53