

UNIT IV- SYNTAX DIRECTED TRANSLATION & RUN TIME ENVIRONMENT

SEMANTIC ANALYSIS

- ¾ Semantic Analysis computes additional information related to the meaning of the program once the syntactic structure is known.
- ¾ In typed languages as C, semantic analysis involves adding information to the symbol table and performing type checking.
- ¾ The information to be computed is beyond the capabilities of standard parsing techniques, therefore it is not regarded as syntax.
- ¾ As for Lexical and Syntax analysis, also for Semantic Analysis we need both a Representation Formalism and an Implementation Mechanism.
- ¾ As representation formalism this lecture illustrates what are called Syntax Directed Translations.

SYNTAX DIRECTED TRANSLATION

- ¾ The Principle of Syntax Directed Translation states that the meaning of an input sentence is related to its syntactic structure, i.e., to its Parse-Tree.
- ¾ By Syntax Directed Translations we indicate those formalisms for specifying translations for programming language constructs guided by context-free grammars.
 - We associate Attributes to the grammar symbols representing the language constructs.
 - Values for attributes are computed by Semantic Rules associated with grammar productions.
- ¾ Evaluation of Semantic Rules may:
 - Generate Code;
 - Insert information into the Symbol Table;
 - Perform Semantic Check;
 - Issue error messages;
 - etc.

There are two notations for attaching semantic rules:

1. **Syntax Directed Definitions.** High-level specification hiding many implementation details (also called **Attribute Grammars**).
2. **Translation Schemes.** More implementation oriented: Indicate the order in which semantic rules are to be evaluated.

Syntax Directed Definitions

- **Syntax Directed Definitions** are a generalization of context-free grammars in which:
 1. Grammar symbols have an associated set of **Attributes**;
 2. Productions are associated with **Semantic Rules** for computing the values of attributes.
 - f* Such formalism generates **Annotated Parse-Trees** where each node of the tree is a record with a field for each attribute (e.g., X.a indicates the attribute a of the grammar symbol X).
 - f* The value of an attribute of a grammar symbol at a given parse-tree node is defined by a semantic rule associated with the production used at that node.

We distinguish between two kinds of attributes:

1. **Synthesized Attributes.** They are computed from the values of the attributes of the children nodes.
2. **Inherited Attributes.** They are computed from the values of the attributes of both the siblings and the parent nodes

Syntax Directed Definitions: An Example

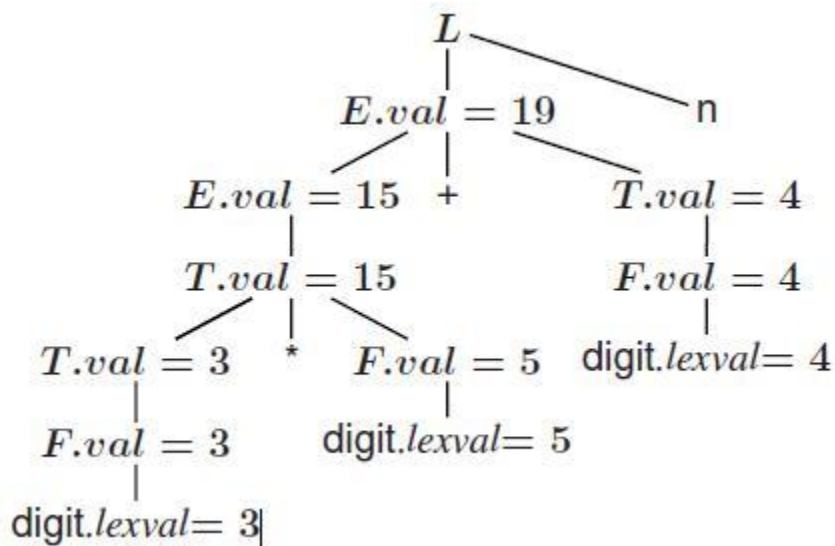
- **Example.** Let us consider the Grammar for arithmetic expressions. The Syntax Directed Definition associates to each non terminal a synthesized attribute called *val*.

PRODUCTION	SEMANTIC RULE
$L \rightarrow En$	$print(E.val)$
$E \rightarrow E_1 + T$	$E.val := E_1.val + T.val$
$E \rightarrow T$	$E.val := T.val$
$T \rightarrow T_1 * F$	$T.val := T_1.val * F.val$
$T \rightarrow F$	$T.val := F.val$
$F \rightarrow (E)$	$F.val := E.val$
$F \rightarrow digit$	$F.val := digit.lexval$

S-ATTRIBUTED DEFINITIONS

Definition. An **S-Attributed Definition** is a Syntax Directed Definition that uses only synthesized attributes.

- **Evaluation Order.** Semantic rules in a S-Attributed Definition can be evaluated by a bottom-up, or PostOrder, traversal of the parse-tree.
- **Example.** The above arithmetic grammar is an example of an S-Attributed Definition. The annotated parse-tree for the input $3*5+4n$ is:



L-attributed definition

Definition: A SDD is *L-attributed* if each inherited attribute of X_i in the RHS of $A \rightarrow X_1 \dots X_n$ depends only on

1. attributes of $X_1; X_2; \dots; X_{i-1}$ (symbols to the left of X_i in the RHS)
2. inherited attributes of A .

Restrictions for translation schemes:

1. Inherited attribute of X_i must be computed by an action before X_i .
2. An action must not refer to synthesized attribute of any symbol to the right of that action.
3. Synthesized attribute for A can only be computed after all attributes it references have been completed (usually at end of RHS).

SYMBOL TABLES

A symbol table is a major data structure used in a compiler. Associates attributes with identifiers used in a program. For instance, a type attribute is usually associated with each identifier. A symbol table is a necessary component. Definition (declaration) of identifiers appears once in a program. Use of identifiers may appear in many places of the program text. Identifiers and attributes are entered by the analysis phases. When processing a definition (declaration) of an identifier. In simple languages with only global variables and implicit declarations. The scanner can enter an identifier into a symbol table if it is not already there. In block-structured languages with scopes and explicit declarations:

- The parser and/or semantic analyzer enter identifiers and corresponding attributes
- Symbol table information is used by the analysis and synthesis phases
- To verify that used identifiers have been defined (declared)
- To verify that expressions and assignments are semantically correct – type checking
- To generate intermediate or target code

9 Symbol Table Interface

The basic operations defined on a symbol table include:

- ¾ allocate – to allocate a new empty symbol table
- ¾ free – to remove all entries and free the storage of a symbol table
- ¾ insert – to insert a name in a symbol table and return a pointer to its entry

¾ lookup – to search for a name and return a pointer to its entry

¾ set_attribute – to associate an attribute with a given entry

¾ get_attribute – to get an attribute associated with a given entry

Other operations can be added depending on requirement For example, a delete operation removes a name previously inserted Some identifiers become invisible (out of scope) after exiting a block

- This interface provides an abstract view of a symbol table
- Supports the simultaneous existence of multiple tables •

Implementation can vary without modifying the interface

Basic Implementation Techniques

- First consideration is how to insert and lookup names
- Variety of implementation techniques
- Unordered List
- Simplest to implement
- Implemented as an array or a linked list
- Linked list can grow dynamically – alleviates problem of a fixed size array
- Insertion is fast $O(1)$, but lookup is slow for large tables – $O(n)$ on average
- Ordered List
- If an array is sorted, it can be searched using binary search – $O(\log_2 n)$
- Insertion into a sorted array is expensive – $O(n)$ on average
- Useful when set of names is known in advance – table of reserved words
- Binary Search Tree
- Can grow dynamically
- Insertion and lookup are $O(\log_2 n)$ on average

RUNTIME ENVIRONMENT

- ¾ Runtime organization of different storage locations
- ¾ Representation of scopes and extents during program execution.
- ¾ Components of executing program reside in blocks of memory (supplied by OS).
- ¾ Three kinds of entities that need to be managed at runtime:
 - Generated code for various procedures and programs.
 - forms text or code segment of your program: size known at compile time.
 - Data objects:
 - Global variables/constants: size known at compile time
 - Variables declared within procedures/blocks: size known
 - Variables created dynamically: size unknown.
 - Stack to keep track of procedure
- activations. Subdivide memory conceptually into code and data areas:

f Code:


Program ● instructions

f Stack: Manage activation of procedures at runtime.

f Heap: holds variables created dynamically

STORAGE ORGANIZATION

1. *Fixed-size objects can be placed in predefined locations.*



2. Run-time stack and heap The STACK is used to store:

- Procedure activations.
 - The status of the machine just before calling a procedure, so that the status can be restored when the called procedure returns.
 - The HEAP stores data allocated under program control (e.g. by malloc() in C).
- Activation records

Any information needed for a single activation of a procedure is stored in the ACTIVATION RECORD (sometimes called the STACK FRAME). Today, we'll assume the stack grows DOWNWARD, as on, e.g., the Intel architecture. The activation record gets pushed for each procedure call and popped for each procedure return.

STATIC ALLOCATION

Statically allocated names are bound to storage at compile time. Storage bindings of statically allocated names never change, so even if a name is local to a procedure, its name is always bound to the same storage. The compiler uses the type of a name (retrieved from the symbol table) to determine storage size required. The required number of bytes (possibly aligned) is set aside for the name. The address of the storage is fixed at compile time.

Limitations:

- The size required must be known at compile time.
- Recursive procedures cannot be implemented as all locals are statically allocated.
- No data structure can be created dynamically as all data's static.

```
float f(int k)
{
float c[10],b;
b = c[k]*3.14;
return b;
}
```

Return value	offset = 0
Parameter k	offset = 4
Local c[10]	offset = 8
Local b	offset = 48

TM Stack-dynamic allocation

- ☺ Storage is organized as a stack.
- ☺ Activation records are pushed and popped.
- ☺ Locals and parameters are contained in the activation records for the call.
- ☺ This means locals are bound to fresh storage on every call.
- ☺ If we have a stack growing downwards, we just need a stack_top pointer.
- ☺ To allocate a new activation record, we just increase stack_top.
- ☺ To deallocate an existing activation record, we just decrease stack_top.

TM Address generation in stack allocation

The position of the activation record on the stack cannot be determined statically. Therefore the compiler must generate addresses RELATIVE to the activation record. If we have a downward-growing stack and a stack_top pointer, we generate addresses of the form $\text{stack_top} + \text{offset}$

HEAP ALLOCATION

Some languages do not have tree-structured allocations. In these cases, activations have to be allocated on the heap. This allows strange situations, like callee activations that live longer than their callers' activations. This is not common. Heap is used for allocating space for objects created at run time. For example: nodes of dynamic data structures such as linked lists and trees.

Dynamic memory allocation and deallocation based on the requirements of the program *malloc()* and *free()* in C programs

new() and *delete()* in C++ programs

new() and garbage collection in Java programs

Allocation and deallocation may be *completely manual* (C/C++), *semi-automatic* (Java), or *fully automatic* (Lisp)

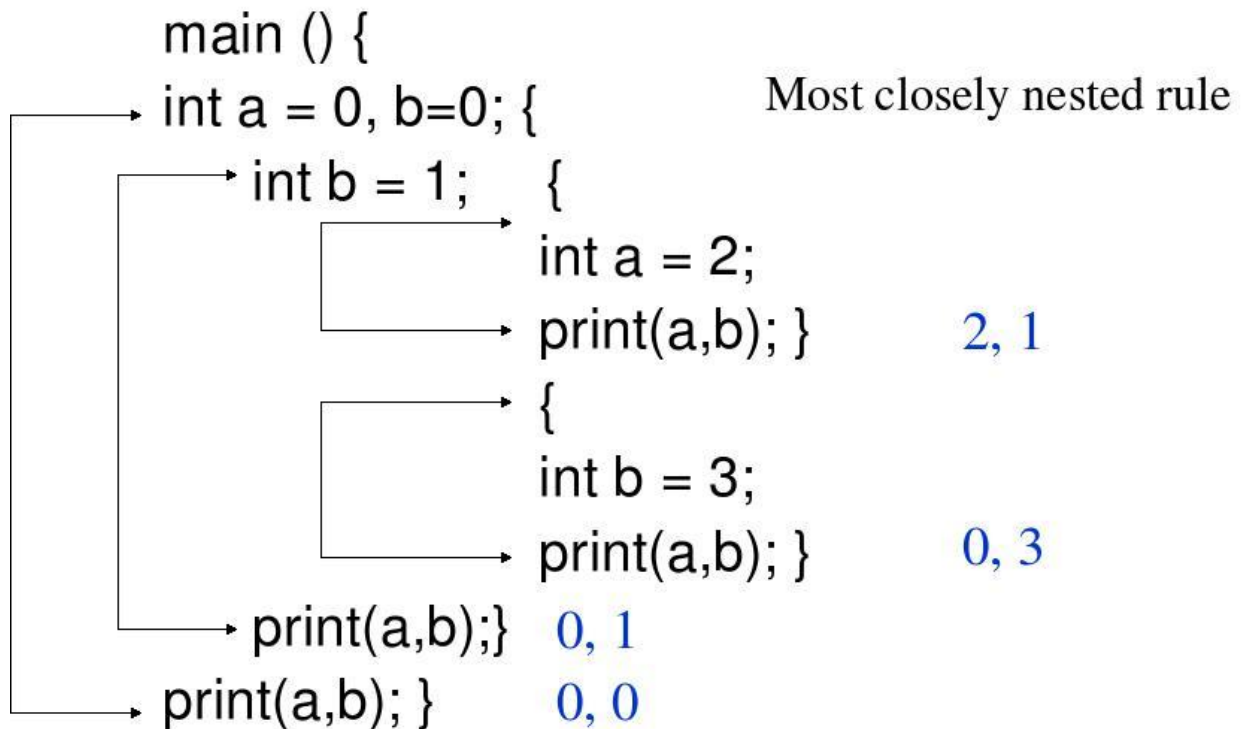
PARAMETERS PASSING

A language has first-class functions if functions can be declared within any scope, passed as arguments to other functions, returned as results of functions. In a language with first-class functions and static scope, a function value is generally represented by a closure, a pair consisting of a pointer to function code and a pointer to an activation record. Passing functions as arguments is very useful in structuring of systems using upcalls.

An example:

```
main()
{
  int
  x =
  4;
  int f
  (int
  y) {
  retur
  n
  x*y;
  }
  int g (int →int h){
  int x = 7;
```

```
return h(3) + x;
}
```



Call-by-Value

The actual parameters are evaluated and their r-values are passed to the called procedure

A procedure called by value can affect its caller either through nonlocal names or through pointers.

Parameters in C are always passed by value. Array is unusual, what is passed by value is a pointer.

Pascal uses pass by value by default, but var parameters are passed by reference.

Call-by-Reference

Also known as call-by-address or call-by-location. The caller passes to the called procedure the l-value of the parameter.

If the parameter is an expression, then the expression is evaluated in a new location, and the address of the new location is passed.

Parameters in Fortran are passed by reference an old implementation bug in Fortran

```
func(a,b) { a = b };  
call func(3,4); print(3);
```

Copy-Restore

A hybrid between call-by-value and call-by reference.

The actual parameters are evaluated and their r-values are passed as in call-by-value. In addition, l values are determined before the call.

When control returns, the current r-values of the formal parameters are copied back into the l-values of the actual parameters.

Call-by-Name

The actual parameters literally substituted for the formals. This is like a macro-expansion or in-line expansion. Call-by-name is not used in practice. However, the conceptually related technique of in-line expansion is commonly used. In-lining may be one of the most effective optimization transformations if they are guided by execution profiles.