

ON DEFINING SEMANTICS BY MEANS OF EXTENDED ATTRIBUTE GRAMMARS*

Ole Lehrmann Madsen

Computer Science Department, Aarhus University,
Ny Munkegade, DK-8000 Aarhus C, Denmark.

1. Introduction.

Knuth has introduced attribute grammars (AGs) as a tool to define the semantics of context-free languages. The use of AGs in connection with programming language definitions has mostly been to define the context-sensitive syntax of the language and to define a translation into code for a hypothetical machine. The semantics of a program is then defined by the interpreter for this machine ([Knuth 68], [Wilner 72], [Marcotty et al. 76]). This is a rather compiler oriented approach to semantics but it has among others resulted in a number of translator writing systems based upon AGs.

Defining semantics this way is useful for an implementer of a language but is less useful for a user or designer of a language. It is often assumed that this is the only way that AGs can be used to define semantics and for this reason AGs are not really considered as an acceptable way of defining semantics. AGs are viewed as a compiler writing tool. One of the reasons for this may be that AGs were not introduced as a complete formalism in the sense that it is not part of the formalism how to define the domains of the attributes.

We would like to point out that we find AGs to be a very useable tool for defining the context-sensitive syntax of a programming language. See e.g. the definition of Pascal in [Watt 78].

In the original paper by Knuth it was stated that any semantics for a language which can be defined as a function of the set of parse trees can be defined by an AG. The purpose of this paper is to demonstrate different ways and techniques for using AGs to define different kinds of semantics.

The motivation for this work comes from an interest in practical translator writing systems (TWSs). Having a TWS intended for implementing (parts of) practical compilers it would be desirable if the same TWS could also be used to make an experimental implementation based on a formal semantics of the language. This will ease experiments with definitions of new languages.

If several different kinds of complementary formal semantics can be used in the same TWS then one may start with a rather human oriented semantics as the basis for an

* This paper is an abbreviated version of [Madsen 79b]

initial (and inefficient) implementation. One may then successively develop more implementation oriented semantics that give more efficient implementations, and if still too inefficient it may be used as an 'implementation guide' for a hand written implementation.

In this paper we shall investigate the possibilities of defining predicate transformers, denotational semantics, and operational semantics. The approach to operational semantics is to specify an AG that defines the possible transformations upon a representation of the program. The operational approach is used (1) to specify abstract data types, (2) to specify semantics of programming languages, and (3) as a model for defining nondeterministic and concurrent computations.

As mentioned we are interested in exploring the possible uses of an AG based TWS, so we are concerned with what is possible to define by means of AGs, i.e. expressibility. We are however just as well interested in the ways things are done with respect to intuition, readability, efficiency, etc. This is often a matter of personal opinion. In connection with this there is no single way of using AGs. The attributes may be used in different ways just as it may be natural to use a translation grammar instead of having the translation as a synthesised attribute.

The power (or expressibility) of AGs is dependent upon the actual domains available. We propose that the domains shall be defined by other AGs. In this way AGs become multi-level instead of two-level. At the bottom we define pure AGs which have a 'built in' set of domain types. We have chosen tree languages as this basic domain. Pure AGs are then quite similar to vW-grammars ([van Wijngaarden et al. 75]), and extended affix grammars [Watt 74a], where the basic domains are context-free (string) languages.

We use a version of AGs called extended attribute grammars (EAGs) ([Watt & Madsen 77]). EAGs are generative in the same sense as affix-grammars ([Koster 71]) and vW-grammars, whilst retaining the essential advantages of AGs. In our opinion EAGs are better suited for analysis and lead to more readable and natural descriptions.

The notions of pure AG and multi-level AG are further refinements of EAGs.

In [Watt & Madsen 77] it is also mentioned that the idea of EAGs can be carried over to translation grammars. We shall also make use of these extended attributed translation grammars (EATGs). EATGs may be a useful tool to define programming languages where the (context-sensitive) syntax and one or more semantics are defined by an integrated formalism. We imagine that the input grammar defines the syntax and a particular semantics is defined by a particular output grammar.

It is still an open problem how to make a general and efficient implementation of AGs without enforcing strange requirements upon the dependencies allowed between attributes. Such requirements are often introduced in order to have a well defined (and

efficient) order of evaluation of the attributes. We give an evaluator for AGs where the order of evaluation of attributes is no problem. The evaluator is general as it accepts all AGs and it is 'very fast'. However it still needs space proportional to the size of the parse tree.

We hope to achieve the following:

- demonstrate that AGs are a very powerful meta language for defining different kinds of semantics,
- demonstrate that a TWS based upon AGs can be used for many purposes,
- contribute to a better understanding of AGs and to the theory of AGs and show how AGs may be turned into a complete formalism,
- contribute to a unification of different formal semantics. We do not claim that we add anything new to these methods,
- present a general and fast evaluator for AGs.

The rest of the paper is organised as follows:

Section 2 reviews the basic terminology being used. The sections 3-7 fall into three parts: Part I (section 3-5) is about semantic definitions. Verification generators are treated in section 3, denotational semantics in section 4 and operational semantics in section 5. Part II (section 6) is a concluding section about the AG formalisms. Part III (section 7) describes a general evaluator for AGs called the DAG-evaluator. The paper is concluded in section 8.

Acknowledgement. I am grateful to Brian Mayoh and Mogens Nielsen for their advice during the preparation of this paper. The following people have also made useful contributions by participating in discussions or by giving comments: Poul Jespersen, Kurt Jensen, Niel Jones, Michael Madsen, Robin Milner, Peter Mosses, Hanne Riis, Erik Meineche Schmidt, Sven Skyum, and David Watt.

2. Basic Terminology.

We use a generative version of AGs called extended attribute grammars (EAGs) ([Watt & Madsen 77]) with the modification that we allow the start symbol to have synthesised attributes. The definition of EAG is repeated below. For a more expository exposition, the reader is referred to [Watt & Madsen 77] or [Madsen 79a].

Defintion 2.1.

An extended attribute grammar is a 5-tuple

$$G = (D, V, Z, B, R)$$

whose elements are defined in the following paragraphs.

$D = (D_1, D_2, \dots, f_1, f_2, \dots)$ is an algebraic structure with domains D_1, D_2, \dots and (partial) functions f_1, f_2, \dots operating on Cartesian products of these domains. Each object in one of these domains is called an attribute.

V is the vocabulary of G , a finite set of symbols which is partitioned into the non-terminal vocabulary V_N and the terminal vocabulary V_T . Associated with each symbol in

V is a fixed number of attribute-positions. Each attribute-position has a fixed domain chosen from D , and is classified as either inherited or synthesised.

Z , a member of V_N , is the start-symbol of G .

The start-symbol Z and the terminal symbols have only synthesised attribute-positions.

B is a finite collection of attribute variables (or simply variables). Each variable has a fixed domain chosen from D .

An attribute-expression is one of the following:

- (a) a constant attribute, or
- (b) an attribute variable, or
- (c) a function application $f(e_1, \dots, e_m)$, where e_1, \dots, e_m are attribute expressions and f is an appropriate (partial) function chosen from D .

In the examples, we shall make use of infix operators where convenient.

Let $v \in V$, and let v have p attribute-positions whose domains are D_1, \dots, D_p , respectively. If a_1, \dots, a_p are attributes in the domains D_1, \dots, D_p , respectively, then

$$\langle v \uparrow a_1 \dots \uparrow a_p \rangle$$

is an attributed symbol corresponding to v . In particular, it is an attributed non-terminal (terminal) if v is a nonterminal (terminal). Each \uparrow stands for either \downarrow or \uparrow , prefixing an inherited or synthesised attribute-position as the case may be.

AV_N (AV_T) stands for the set of attributed nonterminals (terminals), $AV = AV_N \cup AV_T$, and AZ is the set of attributed nonterminals corresponding to the start-symbol Z .

If e_1, \dots, e_p are attribute expressions whose ranges are included in D_1, \dots, D_p , respectively, then

$$\langle v \uparrow e_1 \dots \uparrow e_p \rangle$$

is an attributed symbol form.

R is a finite set of production rule forms (or simply rules), each of the form

$$F ::= F_1 \dots F_m$$

where $m \geq 0$, and F, F_1, \dots, F_m are attributed symbol forms, F being a nonterminal.

A production rule form defines a set of production rules in the following way:

Let $F ::= F_1 \dots F_m$ be a rule. Take a variable x which occurs in this rule, select any attribute a in the domain of x , and systematically substitute a for x throughout the rule. Repeat such substitutions until no variables remain, then evaluate all the attribute expressions. Provided all the attribute expressions have defined values, this yields a production rule, which will be of the form

$$A ::= A_1 \dots A_m$$

where $m \geq 0$, and A, A_1, \dots, A_m are attributed symbols.

The relation \Rightarrow is defined as follows:

Let $\underline{a}, \underline{g} \in AV$, $A \in AV_N$, and let $A ::= \underline{b}$ be a production rule, then $\underline{a} A \underline{g} \Rightarrow \underline{a} \underline{b} \underline{g}$.

\Rightarrow^* and \Rightarrow^+ are defined in the usual way.

The language generated by G , $L(G)$ a subset of AV_T^* , is defined as

$$L(G) = \{ w \mid S \Rightarrow^* w \text{ and } S \in AZ \}$$

Let D_1, \dots, D_p be the attribute domains of Z . The translation defined by G , $T(G)$ a subset of $AV_T^*(X_1 D_1 \dots X_p D_p)$, is defined as

$$T(G) = \{ (w, m) \mid m = (a_1, \dots, a_p), \langle Z \uparrow a_1 \dots \uparrow a_p \rangle \Rightarrow^* w \}$$

If $(w, m) \in T(G)$ then m is a meaning of w .

The relation \Rightarrow defines in the usual way an attributed parse tree. An attributed parse tree defines in a unique way a corresponding parse tree from its underlying CFG.

One may distinguish between the following three kinds of ambiguity/unambiguity of an EAG:

- (1) G is semantically ambiguous if there are meanings m_1, m_2 , $m_1 \neq m_2$ and $(w, m_1) \in T(G)$ and $(w, m_2) \in T(G)$.
- (2) G is structurally ambiguous if there is a $w \in L(G)$ and w is the frontier of two or more distinct attributed parse trees.
- (3) G is syntactically ambiguous if the underlying CFG of G is ambiguous.

Observe that the distinction between inherited and synthesised attribute makes no difference to the language and translation defined by the EAG. The distinction is traditional, may improve the readability and is important when considering implementations of EAGs. This is also the case for the following definitions.

Inherited attribute-positions on the left-side and synthesised attribute-positions on the right-side of a rule are called defining positions. Synthesised attribute-positions on the left-side and inherited attribute-positions on the right-side are called applied positions.

An EAG is well-formed iff

- (a) every variable occurs in at least one defining position in each rule in which it is used; and
- (b) every function used in the composition of an attribute expression in a defining position is injective.

//

We shall also use the EAG meta syntax for ordinary Knuth-like AGs, which we define in the following way

[2.2] A Knuth-like AG (or just an AG) is an EAG that satisfies: (1) it is well-formed, (2) only (attribute-) variables appear in defining positions, and (3) the same variable appears in only one defining position.

This definition of an AG differs from the one in [Knuth 68] in the following ways:

[2.3a] In Knuth's definition terminals cannot have synthesised attributes.

[2.3b] In Knuth's definition the semantic functions are apparently required to be total whereas [2.2] allows them to be partial.

[2.3c] AGs defined by [2.2] are always in normal form ([Bochmann 76]). By requiring normal form we avoid a number of tedious (and unimportant) complications in the following sections and we exclude only some obscure AGs.

[2.3d] In Knuth's definition, a string is assigned a meaning in the following well known way: (1) A parse tree for the string is constructed. (2) A node in the parse tree and an applied attribute-position of that node is selected. If the attribute-positions referred to (through attribute variables) in the expression of the selected attribute have defined values then the selected attribute-position is assigned the value of the expression. (3) Step (2) is repeated until either all attribute-positions have been assigned a value or no more attribute-position can be assigned a value by this process. (4) The value of a distinguished attribute-position of the root in the parse tree constitutes the meaning of the string corresponding to that parse tree.

[2.3e] Knuth defines an AG to be well-defined if all attributes can always be defined, in any conceivable parse tree using the strategy in [2.3d]. He then shows that an AG is well-defined if and only if it is non-circular.

[2.3f] Knuth's model is intended to define the semantics of context-free languages in the sense that all parse trees may be assigned a meaning in all well-defined AGs. Other AG formalisms, such as EAGs (and [2.2]) may be viewed as a language generating device in the sense that not all parse trees of the underlying CFG may get values assigned to its attributes. Knuth suggests to let an attribute-position in the root of the parse tree decide whether the parse tree (string) is 'malformed' or not.

The use of partial functions in [2.2] implies that not all parse trees may be assigned attribute values, even if the AG is non-circular. Furthermore a parse tree may be assigned attribute values even if the AG is circular. Thus circularity is not an inherent problem in EAGs and AGs as defined in [2.2]. We return to that later in the paper.

One may reformulate Knuth's definition (2.3d) in order to obtain a definition that is equivalent to [2.2]:

[2.4] A string is assigned a meaning in the following way: (1) A parse tree for the string is constructed. (2) A set of equations corresponding to the parse tree is constructed. Each attribute-position is an unknown; each attribute expression determines an equation in the sense that if a is an attribute expression occupied by the expression e , then $a=e$ is an equation; variables in the expressions are also unknowns and may have to be renamed properly. (3) The parse tree may be assigned attribute values if and only if the equations have a solution. (4) The attributes of the root in a solution constitute the meaning of the string corresponding to the parse tree.

In some definitions of AGs ([Marcotty et al 76]) an AG rule has an associated constraint which is a predicate over attribute values. This constraint must be true in

order that the attributes of the rule can be assigned values. We interpret constraints in the following way:

[2.5] Each symbol is given a synthesised attribute with a domain consisting of one value. The constraint is converted to a partial function that assigns the value to this attribute of the left side symbol if the constraint is true.

Consequently if some constraint is false there is an attribute that cannot be assigned a value and this is now captured by the definition of AGs in [2.2] or [2.4].

In [Watt & Madsen 77] it is shown how to convert an EAG into an AG by using constraints.

We also make use of extended attributed translation grammars (EATGs) ([Watt & Madsen 77]). An EATG consists of a translation grammar (like syntax directed translation schemes in [Aho & Ullman 72]) equipped with attributes in the same way as an EAG is a CFG equipped with attributes. An EATG is naturally divided into an input-grammar and an output-grammar. The terminals of the input-grammar (output-grammar) are called input-symbols (output-symbols). As with EAGs input-symbols may have synthesised attributes whereas output-symbols may have only inherited attributes. Each rule in the input-grammar has an associated rule in the output-grammar. The output rule may refer to attribute variables in the input rule but not vice versa. Pairs of input production rules and output production rules are obtained by applying the systematic substitution rule to both the input rule and the corresponding output rule taken together.

When requiring the restrictions in [2.2] to EATGs one obtains attributed translation grammars similar to those in [Lewis et al. 74].

An EATG defines a translation from strings of attributed input-symbols to strings of attributed output-symbols.

There is a choice between defining a translation by using an EATG or by using synthesised attributes of an EAG as in definition 2.2. The actual choice depends upon the kind of semantics (or translation) to be defined. It is often a matter of modularity and by using EATGs one may separate the definition of the (context-sensitive) syntax from the definition of the semantics.

As attribute domain constructors we make use of discriminated unions, cartesian products, sequences, and partial mappings which (among others) may be found in [Watt & Madsen 77] and [Madsen 79a].

3. Defining Verification Generators.

Here we treat the possibilities of expressing predicate transformer semantics by means of EAGs. It is well known that any predicate transformer semantics may be reformulated as a denotational semantics. In section 4 it is shown how any denotational semantics may be defined by an AG. The technique of that section may then be used to define any predicate transformer semantics.

A predicate transformer semantics may be used as a basis for a system which generates verification conditions. In this section we sketch an example on how such a verification generator may be defined by means of an EATG. The example is based upon a forward predicate transformer for partial correctness in the style of [Gerhart 76].

If P is a predicate which is supposed to be true before the execution of a statement S , then the value of the forward predicate transformer $FPT(P, S)$ is a predicate which is true after the execution of S .

Consider the statements :

if B then S_1 else S_2 and

while B assert A do S .

Assert A defines an invariant which must be supplied by the programmer. We may define the following FPTs:

$$FPT(P, \text{if } B \text{ then } S_1 \text{ else } S_2) =$$

$$FPT(P \text{ AND } B, S_1) \text{ OR } FPT(P \text{ AND NON } B, S_2)$$

$$FPT(P, \text{while } B \text{ assert } A \text{ do } S) = A \text{ AND NON } B$$

$$\text{verify: } P \Rightarrow A, FPT(B \text{ AND } A, S) \Rightarrow A.$$

The FPT for the while-statement is only true if the so-called verification condition following verify can be proved to be true.

A verification generator for a language can be defined by an EATG, where the input grammar defines the (context-sensitive) syntax. The output grammar generates a sequence of verification conditions, and the symbols have predicates as attributes. If $\langle \text{stmt} \rangle$ is the nonterminal generating statements, then in the output grammar $\langle \text{stmt} \rangle$ may typically have two attributes:

$$\langle \text{stmt} \rangle \begin{matrix} P & \uparrow P & \\ \text{before} & \text{after} & \end{matrix}$$

where P_{before} is the predicate which is true before the execution of the statement

generated by $\langle \text{stmt} \rangle$ and P a predicate which is true after.
after

An output rule for if-then-else and while-assert-do might look like (in order to ease the reading, the input symbols are included in the rule):

```

<stmtψP↑Q OR R> ::=
  if <exp↑B> then <stmtψP AND B ↑Q>
  else <stmtψP AND NON B ↑R>

<stmtψP ↑A AND NON B> ::=
  while <exp↑B> assert <predicate↑A>
  <stmtψA AND B ↑Q> <verify ψP => A, Q => A>

```

$\langle \text{exp} \rangle$ has a synthesised attribute B which is the predicate corresponding to the expression generated by $\langle \text{exp} \rangle$ and similarly the synthesised attribute of $\langle \text{predicate} \rangle$ is the invariant supplied by the programmer.

$\langle \text{verify} \rangle$ is an output symbol.

If $\langle \text{prg} \rangle$ is the start symbol of the grammar then we may have a rule

```

<prg↑P> ::= <stmtψ true ↑P>,

```

where P then will be a predicate which is true after the execution of the program. Instead of initialising the inherited attribute of $\langle \text{stmt} \rangle$ with true one might as well do as follows:

```

<prg↑P> ::= assert <predicate↑A> <stmtψA ↑P>,

```

where A then is an input assertion.

The idea of generating a verifier from a grammar appears in [Mayoh 76] and is used in the JQNS-system [Nielsen 75]. Most verifiers are designed for a specific language. By means of a TWS based on AGs such verifiers can be automatically constructed from an AG description.

An attribute domain for predicates must be available. In JQNS predicates are basically text strings with an associated set of operations. This is a simple solution. It might be desirable to have a more structural definition of predicates, especially if the verifier is combined with a theorem prover. The domains mentioned in chap. 2 should be sufficient for this.

By using an EATG to define the syntax and a verification generator for a language it should be possible to tie the two definitions together. This is e.g. not the case with the original definition of Pascal where the context-free syntax is defined by BNF, an axiomatic definition appears in [Hoare & Wirth 73], but the context-sensitive

part of the syntax is only informally and very imprecisely defined in [Wirth 71].

The output grammar of the EATG defining semantics could rely upon the input grammar defining the syntax. Type checking would normally appear in the input grammar. In case of Pascal the input grammar could check certain other assumptions made by the semantics. E.g. that aliasing does not appear. The input grammar will normally have attributes corresponding to a symbol table to collect declared identifiers. If the semantics needs a renaming of all identifiers then this may be done using the 'symbol table' of the input grammar.

4. Denotational Semantics and Attribute Grammars.

In [Mayoh 78] it is shown that any AG can be reformulated into an equivalent Denotational Semantics (DS) ([Tennent 76]). An algebraic formulation of AGs is given in [Chirica 76]. In this section we shall discuss the possibilities of reformulating a Denotational Semantics within AGs.

In [Knuth 68] it was proposed that one let a 'meaning' of a string generated by an AG be the synthesised attributes of the start symbol in a parse tree for the string. Furthermore it was shown that this meaning could be any function of the parse tree. According to Knuth an AG defines a function from the set of parse trees into some domain. So AGs are in fact a meta language for defining a kind of mathematical semantics. Reformulation is then a question of using a different meta language.

Another result in [Knuth 68] is that any AG has an equivalent one using only synthesised attributes. The reformulation of AGs as defined by Chirica and Mayoh can be used in order to transform any AG into an equivalent one using only synthesised attributes. This transformation is more natural and constructive compared to the one of Knuth.

Definition 4.1. Attribute grammar notation.

Let A be a symbol of an AG. Define

$$\text{INH}(A) = \text{ID}_1 \times \text{ID}_2 \times \dots \times \text{ID}_k, \text{ and } \text{SYN}(A) = \text{SD}_1 \times \text{SD}_2 \times \dots \times \text{SD}_n,$$

where $\text{ID}_1, \text{ID}_2, \dots, \text{ID}_k, \text{SD}_1, \text{SD}_2, \dots, \text{SD}_n$ are the domains of the inherited and synthesised attributes of A .

Let $A_0 \rightarrow A_1 A_2 \dots A_m$ be the p 'th production in an AG. Define

$$\text{DEF}(p) = \text{INH}(A_0) \times \text{SYN}(A_1) \times \text{SYN}(A_2) \times \dots \times \text{SYN}(A_m), \text{ and}$$

$$APP(p) = SYN(A_0) \times INH(A_1) \times INH(A_2) \times \dots \times INH(A_m).$$

In general the attributes of rule p are defined by a function:

$$F_p : DEF(p) \times APP(p) \rightarrow APP(p)$$

However we assume (as mentioned in chap. 2) that our AGs are in normal form, i.e. F_p is defined by

$$F_p : DEF(p) \rightarrow APP(p).$$

There is a canonical correspondence between sets of functions

$$f_0 : DEF(p) \rightarrow SYN(A_0), f_i : DEF(p) \rightarrow INH(A_i), i=1,2,\dots,m,$$

and F_p .

If $D=(I_0, S_1, \dots, S_m)$, where I_0 is a value of the inherited attributes of A_0 and S_i is a value of the synthesised attributes of A_i , $i=1,2,\dots,m$, then the attributes of an instance of rule p are defined as follows:

$$S_0 = f_0(D), I_i = f_i(D), i=1,2,\dots,m, \text{ and then}$$

$$F_p(D) = \langle f_0(D), f_1(D), \dots, f_m(D) \rangle$$

If f_i ($i=0,1,\dots,m$) defines k attributes then f_i is defined by k functions f_{i1}, \dots, f_{ik} each defining an attribute. (In practice the functions f_{i1}, \dots, f_{ik} will not depend upon the whole of $DEF(p)$.)

Definition 4.2. Reformulated AG with only synthesised attributes.

Let G be an AG. G_s is an AG with only synthesised attributes and defined by the following transformations.

Each symbol A will have one synthesised attribute with domain $[INH(A) \rightarrow SYN(A)]$.

Each production has one function that defines the synthesised attribute of the leftside in terms of the synthesised attributes of the right side.

For rule p we get

$$[4.2*] \quad S'_0 = \lambda I. f_0(I, S'_1(I), \dots, S'_m(I)), \text{ where}$$

$$I'_i = \lambda I_i . f_i (I_i, S'_1(I'_1), \dots, S'_m(I'_m)), \quad i=1,2,\dots,m,$$

and S'_i is the value of the synthesised attribute of A_i , $i=0,1,2,\dots,m$, in rule p of the reformulated AG, G_s .

//

As mentioned in def. 4.1 each f_i ($i \in [0..n]$) defines a number of attributes, thus each of the above equations defines a number of equations corresponding to the attributes defined by each f_i . Thus if rule p has k defining positions then the above $m+1$ equations define k equations.

Below we formulate in what sense G is equivalent to G_s . This is similar to the formulations in [Chirica 76] and [Mayoh 78].

For each parse tree, [4.2*] defines a set of equations. Each instance of a production p defines a set of equations using [4.2*]. These equations have exactly one solution if the AG is non-circular.

Theorem 4.3.

Let G be a non-circular AG with all semantic functions being total and let G_s be the corresponding AG defined by 4.2. Let t be a parse tree of the underlying CFG, and let $A_0 \rightarrow A_1 A_2 \dots A_m$ be production p and let an instance of p appear in t .

A_0 is then a node in t with sons A_1, A_2, \dots, A_m .

Let I'_i, S'_i be the unique values of the inherited and synthesised attributes of A_i , $i=0,1,2,\dots,m$. The equations associated with p have exactly one solution, and $S'_0(I'_0) = S_0$ and $I'_i(I'_0) = I_i$, $i=1,2,\dots,m$.

proof:

We use structural induction on t .

Bottom : Assume that A_1, A_2, \dots, A_m are terminals, i.e. leaves in t . The only unknown in the equations is S'_0 which is well defined as it only depends on I_0 and S_1, \dots, S_m , and clearly $S_0 = S'_0(I_0)$.

Induction step: Assume that S'_i ($i \in [1..m]$) are defined and that $S'_i(I'_i) = S_i$.

As G is non-circular and t is fixed there is a partial ordering of the attributes in p , such that $x < y$ means that the value of y depends on the value of x . This partial ordering can be used to solve the equations by substitution.

Let B_i ($i \in [1..k]$) be the attributes of the defining positions of p , and let $B_i < B_{i+1}$ ($i \in [1..k-1]$) where the partial dependency ordering is extended to some total ordering). Each B_i has an associated equation defining a function B_i' : $[INH(A) \rightarrow DOM(B_i)]$. The equation defining B_i' is independent of B_1', \dots, B_{i-1}' ; if not then G is circular. Consequently the equations may be solved by substitution in the order B_1', B_2', \dots, B_k' .

//

4.1 A redefinition of Attribute Grammars.

In his original definition of AGs Knuth has excluded circular AGs in order to assure that all attributes in all possible derivation trees can be assigned unique values. This may be too strong a requirement. Consider the following examples:

- consider def. 4.2. If we transform a given AG into one with only two attributes for each symbol A , an inherited with domain $INH(A)$ and a synthesised with domain $SYN(A)$, then the new AG is circular if some inherited attribute of A depends on some synthesised attribute of A . However if the original AG is well defined then the new one should not give problems with assigning values to attributes.
- using conditional expressions. Consider the rules:

$\langle A \uparrow y \uparrow z \rangle ::= \langle B \uparrow \text{cond}(c, y, z) \uparrow c \uparrow y \uparrow z \rangle$
 where $\text{cond}(c, y, z) = \text{If } c \text{ then } y \text{ else } z$

$\langle B \uparrow x \uparrow \text{true} \uparrow x \uparrow 2 \rangle ::= 'a'$

$\langle B \uparrow x \uparrow \text{false} \uparrow x \uparrow 1 \uparrow 8 \rangle ::= 'b'$

The dependency graph for the derivation $\langle A \rangle \Rightarrow \langle B \rangle \Rightarrow a$ has cycles in it but there is no problem in assigning unique values to the attributes.

The equations [4.2*] give a basis for discussing circularity. A given parse tree defines a set of equations defined by the productions in the tree and the corresponding equations [4.2*]. These equations may have either (1) Exactly one solution, (2) more than one solution, or (3) no solution.

Case 1 captures all non-circular AGs and some circular ones where unique values may be assigned. Case 2 and 3 capture circular AGs where unique values cannot be assigned.

In the style of denotational semantics we shall now assume that an AG is extended in the following way. The domains are extended to Scott-type domains with a least element and a partial ordering. The equations associated with a parse tree and defined by [4.2*] will then always have a unique least solution.

In this connection it is thus natural to define the meaning of a parse tree based on the minimal solution to the corresponding equations. In the rest of section 4 we shall thus use this definition of AGs.

This in fact gives a difference when compared to EAGs. If there is no solution to the equations then the EAG cannot generate the corresponding input string. If there are more than one solution then the EAG may generate the string in ways corresponding to each solution.

As mentioned in section 2 the definition of EAGs in [2.1] and the corresponding AG definition in [2.4] give also an interpretation to circular AGs. When we extend the domains to Scott-type domains we could also base 'the least solution' approach to AGs on the minimal solution to the equations defined in [2.4].

If we when using Scott-domains use the EAG approach then we must define an ordering upon attributed parse trees in order to get a unique attribute assignment to all parse trees.

With these constructions of 'the least solution' approach we loose the ability to recognise certain kinds of semantic ambiguity. This may be the case if the EAG can generate distinct attributed parse trees with the same corresponding (context-free) parse tree. But this is perhaps reasonable since any AG-evaluator will probably compute the least solution.

4.2 Relation to Denotational Semantics.

By using the reformulated AG of definition 4.2 it is straightforward how to make an equivalent Denotational Semantics. However this reformulation is in general very complicated and unreadable as one must express a general solution to the equations independent of the parse trees. In [Mayoh 78], [4.2*] is viewed as a set of equations defining a function from the set of parse trees into the attribute domains of the start symbol of the AG. Such a solution using the fixpoint operator may be found in [Mayoh 78]. He shows that if the AG is non-circular then the solution may be expressed without using fixpoints. Next Mayoh defines a hierarchy of AGs which in turns simplifies the solution and improves the readability. The hierarchy classifies AGs by the evaluation order of attributes and is also interesting when considering implementations of AGs.

Here we shall consider the opposite direction, namely converting a denotational

semantics into an equivalent AG. This is done by passing denotations as attributes. It should be clear that the two formalisms are equally powerful so the purpose of this is to show how AGs may be used to express a denotational semantics. We note that in general mechanical transformations from one meta language into another may not work very well in the sense that the readability of the transformed definition will usually be bad. It is like transforming PASCAL programs mechanically into FORTRAN programs.

4.2.1 A simple reformulation.

A straightforward transformation of a denotational semantics(DS) is as follows:

A DS is defined by means of functions from syntactic domains (parse trees) into some semantic domains. Let Com be a syntactic domain for the nonterminal <com> (then Com is the set of parse trees derivable from <com>). Consider the semantic function

cc: [Com \rightarrow A], where A is some semantic domain, and the production

$$\langle \text{com} \rangle ::= w_0 \langle A \rangle_1 w_1 \langle A \rangle_2 \dots \langle A \rangle_n w_n$$

$$\text{cc}[w_0 \langle A \rangle_1 w_1 \langle A \rangle_2 \dots \langle A \rangle_n w_n] = f(\text{aa}_1[A_1], \text{aa}_2[A_2], \dots, \text{aa}_n[A_n])$$

where f is some function and aa_i ($i=1,2,\dots,n$) are semantic functions.

In the corresponding AG, <com> will have a synthesised attribute with domain A and we will have the following rule:

$$\langle \text{com} \uparrow f(C_1, C_2, \dots, C_n) \rangle ::= w_0 \langle A \uparrow C_1 \rangle_1 w_1 \langle A \uparrow C_2 \rangle_2 \dots \langle A \uparrow C_n \rangle_n w_n$$

If there are more functions defined on the domain Com, then <com> has a synthesised attribute for each of these.

4.2.2. A reformulation using a circular AG.

If we turn to more specific DSs, then the meaning of a construct like <com> is often defined relatively to an environment (Env) and a (command) continuation (CC), i.e.

$$A = [\text{Env} \rightarrow \text{CC} \rightarrow B]$$

In the AG terminology Env and CC intuitively correspond to inherited attributes and B to a synthesised attribute. This will in fact correspond more to the domain [Env X CC \rightarrow B] which however is isomorphic to [Env \rightarrow CC \rightarrow B].

This is illustrated by the following example:

Example 4.4. Denotational Semantics.

Syntactic Domains.

PROG=Programs; COM=Commands; EXP=Expressions; VAR=Variables; CONST=Constants;

Semantic Domains.

V=Values=Integers; S=States=[VAR \rightarrow V];

CC=Command Continuations = [S \rightarrow S];

EC=Expression Continuations = [V \rightarrow EM];

EM=Expression Meaning = [CC + EC];

Note that for technical reasons we use an expression continuation which is (slightly) different from what is normal been used. We return to that later. An EC will require a sequence of values in order to produce a CC. Consider the following auxillary domains: $EC_0 = CC$, and for each $n > 0$ $EC_n = [V \rightarrow EC_{n-1}]$. For a given $k \in EC$ there exists an $n > 0$ such that $k \in EC_n$.

We do not use an environment in this example as this will not change the principle of the reformulation. In fact environments are straightforward to pass as inherited attributes whereas continuations cause more problems.

Auxillary Functions.

COND : [CC X CC \rightarrow [V \rightarrow CC]]

COND(c_1, c_2)(v) = if $v > 0$ then c_1 else c_2

CONTENT : [VAR X EC \rightarrow EC]

If $k \in EC_n$, $n > 0$ then CONTENT(A, k) $\in EC_{n-1}$ and

CONTENT(A, k) $v_2 \dots v_n$ s = k(s(A)) $v_2 \dots v_n$

UPDATE: [VAR X CC \rightarrow [V \rightarrow CC]]

UPDATE(n, c)(v)s = c(s[v/n])

Syntax

<program> ::= <com>

<com> ::= <com> ; <com> | <var> := <exp>

 | if <exp> then <com> else <com>

 | while <exp> do <com>

<exp> ::= <exp> + <exp> | <var> | <const>

Semantic Functions.

pp: [PROG \rightarrow CC]; cc: [COM \rightarrow [CC \rightarrow CC]];

ee: [EXP \rightarrow [EC \rightarrow EM]]; vv: [CONST \rightarrow V];

For each exp in a program there will exist an $n > 0$ such that $ee[exp] \in [EC_n \rightarrow EC_{n-1}]$. This is easily seen by checking the semantic equations below.

Consider the command $a := e_1 + e_2 + e_3 + e_4$, and assume a left associative parse tree. Then for each $i \in [1..4]$, $ee[e_i] \in [EC_i \rightarrow EC_{i-1}]$.

Semantic equations

$pp[com] = cc[com]c_0$ where c_0 is the initial continuation

$cc[com_1; com_2]c = cc[com_1][cc[com_2]c]$

$cc[var := exp]c = ee[exp]\{UPDATE(var, c)\}$

$cc[\text{if } exp \text{ then } com_1 \text{ else } com_2]c = ee[exp]\{COND(cc[com_1]c, cc[com_2]c)\}$

$cc[\text{while } exp \text{ do } com]c = \text{fix } \lambda c'. ee[exp]\{COND(cc[com]c', c)\}$

$ee[exp_1 + exp_2]k = ee[exp_1][ee[exp_2][\lambda v_1. \lambda v_2. k(v_1 + v_2)]]$

$ee[var]k = CONTENT(var, k)$

$ee[const]k = k(vv[const])$

//

Example 4.5. An equivalent attribute grammar.

In the AG the idea is to let each symbol have two attributes, an inherited defining its continuation and a synthesised defining its meaning in that continuation. We will e.g. have the following rule:

$\langle com \psi C \uparrow C2 \rangle ::= \langle com_1 \psi C1 \uparrow C2 \rangle ; \langle com_2 \psi C \uparrow C1 \rangle$

The meaning of $\langle com \rangle$ in a given continuation C , is $C2$ which is the meaning of $\langle com_1 \rangle$ in the continuation $C1$ which again is the meaning of $\langle com_2 \rangle$ in the continuation C .

We use the same semantic domains as in the DS.

Symbols and their attribute domains

$\langle prog \uparrow CC \rangle; \langle com \psi CC \uparrow CC \rangle; \langle exp \psi EC \uparrow EM \rangle; \langle var \uparrow VAR \rangle; \langle const \uparrow V \rangle;$

AG rules

$\langle prog \uparrow C \rangle ::= \langle com \psi C \uparrow C \rangle$

```

<comψC↑C2> ::= <comψC1↑C2> ; <comψC↑C1>
<comψC↑C1> ::= <var↑N> := <expψUPDATE(N,C)↑C1>
<comψC↑C3> ::= if <expψCOND(C1,C2)↑C3> then <comψC↑C1> else <comψC↑C2>
<comψC↑C2> ::= while <expψCOND(C1,C)↑C2> do <comψC2↑C1>
<expψK↑K2> ::= <expψK1↑K2> + <expψλV2.λV1.k(V1+V2)↑K1>
<expψK↑CONTENT(N,K)> ::= <var↑N>
<expψK↑K(C)> ::= <const↑C>

```

//

If the above AG is transformed into an equivalent one using the method of section 4.2, then we obtain synthesised function attributes which are similar to the functions defined by the semantic equations. AS the while-imperative is the most difficult one (involves circularity) we shall make the transformation.

The corresponding rule is:

$$\langle \text{com}\uparrow C' \rangle ::= \text{while } \langle \text{exp}\uparrow E' \rangle \text{ do } \langle \text{com}\uparrow C'1 \rangle$$

where $C' = \lambda C. C2$, $C2 = E'(K)$, $K = \text{COND}(C1, C)$, and $C1 = C'1(C2)$.

This implies that $C2 = E'(\text{COND}(C'1(C2), C))$

and thus $C' = \lambda C. \text{fix } \lambda C2. E'(\text{COND}(C'1(C2), C))$

Now comparing this with the while-equation and letting $C' = \text{cc}[\text{while exp do com}]$, $E' = \text{ee}[\text{exp}]$, and $C'1 = \text{cc}[\text{com}]$ we see that our reformulated AG defines the same function.

About expression continuations.

As mentioned we use another definition of the meaning of expressions than the usual one. A more standard one is

$$EC = [V \rightarrow CC]$$

$$ee: [EXP \rightarrow [EC \rightarrow CC]]$$

Consider the rule:

$$ee[\text{exp}_1 + \text{exp}_2]k = ee[\text{exp}_1] \lambda V_1. ee[\text{exp}_2] \lambda V_2. k(V_1 + V_2)$$

In our corresponding AG $\langle \text{exp} \rangle$ should have the domains $\langle \text{exp}\psi EC\uparrow CC \rangle$, and the AG rule should be

$$\langle \text{exp}\psi K\uparrow K2 \rangle ::= \langle \text{exp}\psi \lambda V1. K1\uparrow K2 \rangle + \langle \text{exp}\psi \lambda V2. K(V1+V2)\uparrow K1 \rangle$$

Now this does not work as $\lambda V2. K(V1+V2)$ is passed as an EC attribute with $V1$ as a free variable. This works in the DS-rule as $V1$ is bound past the meaning of exp_2 .

The meaning of a construct (command or expression) is interpreted relative to a function (its continuation) which specifies what is to be done after executing the construct. When applying the meaning of a construct to its continuation one gets the meaning of the whole program if executing it beginning with the construct.

If we consider expressions then the meaning of an expression (relative to a continuation) depends on the context of the expression. Consider the expression exp in the following constructs:

(1) If exp then com_1 else com_2

(2) $exp' + exp$

In the above DS the meaning of an expression (relative to an EC) is a CC. In case (1) the EC of exp can be determined by the program text alone (a static continuation). In case (2) the EC of exp depends upon the execution of the program (a dynamic continuation) since it includes the value of exp' .

In an AG one can only express as attributes values which are a static property of the program text unless one turns to include rules for executing the program. One may then discuss whether or not it is reasonable to require a continuation to be static or not. Perhaps it is.

Towards a general rule for using inherited attributes.

There is of course no general rules for transforming a DS into and AG with both inherited and synthesised attributes. It seems likely that a semantic function in a DS can be made more readable by decomposing it into a number of attributes when each attribute is a static property of the program text. One may often benefit by converting a semantic function f with domain $[A \rightarrow B]$ into an inherited attribute with domain A and a synthesised attribute with domain B .

This may be reasonable if A is naturally expressed as a (static) property determined by the context of the constructs defined by f . The following is a rule for this to be natural and possible:

[4.6] Consider a semantic rule

$c[..\] = e$

where e involves one or more applications of f and possibly $c=f$.

(1) all occurrences of f in e must be applied to an expression of type A .

(2) If f is applied to the expression a , then all free variables in a must be convertible to attributes.

The semantic functions cc and ee of example 4.4 satisfy this whereas ee with the 'standard' definition of EC in the previous section does not.

An extension of the AG model which makes it possible to express dynamic properties as

attributes is discussed in [Ganzinger 79a]. In section 5 we shall (among others) demonstrate a technique for doing this within the existing model of AGs.

4.2.3 Other reformulations.

Instead of passing functions around as attributes one might pass lambda-expressions. This will look like the reformulation of section 4.2.1 but be quite different. The meaning of a program will then be a lambda-expression instead of a function. This would then correspond to an AG defining a code generation. In a practical TWS based on AGs this might be a reasonable way of implementing a denotational semantics.

Yet another approach would be to define a denotational semantics by means of syntax directed translation schemes as used in [Aho & Ullman 72]. In their generalised translation schemes they allow nonterminals to have translation elements other than just strings, e.g. integers, booleans. If one allows translation elements to be functions (or lambda-expressions) then a denotational semantics may be defined in a notation which is quite close to the usual notation of denotational semantics. Such a definition will however just be another notation for the one of section 4.2.1 (or the above mentioned).

4.3 Conclusion.

It has been shown that AGs are a suitable tool for defining a TWS in which compilers may be generated based upon a denotational semantics, like e.g. SIS ([Mosses 79]). We also think that the AG notation in many situations gives a more natural and readable definition than the corresponding DS. This is due to the fact that one in the AG may have simpler domains and thus simpler expressions.

A further modularization (and simplification) can be obtained by using a model based on EATGs. Here it is possible to separate the context-sensitive syntax from the semantics.

If one does not like the AG notation then with the right TWS it should be no problem to define one's own notation and just use AGs as an implementation.

In addition to a higher degree of modularity in the semantic definitions one may also benefit when the semantic definition has to be converted into a more implementation oriented semantics. In the AG it is possible to isolate the static properties of the definition; the context-sensitive syntax may (as mentioned) be isolated and the semantic functions may (as mentioned) be split up into attributes describing static properties, like environments and continuations. The possibilities of transforming a DS into an implementation oriented AG have been studied by [Bjørner 78] and [Ganzinger 79b].

5. Defining operational semantics by Attribute Grammars.

In this section a technique to specify operational semantics by means of EAGs will be presented. The approach is to specify a set of EAG rules which defines the possible transformations upon an abstract representation of the program. Consequently it is not a traditional operational semantics where the program is transformed into code for a hypothetical machine which then executes the code. The examples presented in this section are inspired by recent work in the area of specifying abstract data types and make use of techniques which have been used with vW-grammars ([Marcotty et al.76]). Section 5.1 is related to the specification of abstract data types, section 5.2 is about specifying semantics of programming languages, and finally section 5.3 shows how AGs may be viewed as a model for defining nondeterministic and concurrent computations.

5.1 Specifying Abstract Data Types.

An abstract data type is considered to consist of an (abstract) set of values and an (abstract) set of operations. The operations may be combined into expressions denoting abstract values.

An abstract data type is specified by an EAG in the following way: The EAG generates the set of all expressions yielding values of the data type. The synthesised attributes of the start symbol is then the value of the generated expression. The values of the data type are defined by the domains of the EAG and the set of expressions and their values are defined by the production rules of the EAG.

We illustrate the approach by specifying the famous stack:

Example 5.1. Stack specification.

Domain

S: SEQ=(empty | cat(SEQ , ELM)); E: ELM=...;

The values of a stack is a sequence of elements (not specified here).

Rules

<stack↑empty> ::= newstack

<stack↑cat(S,E)> ::= push (<stack↑S> , <element↑E>)

<stack↑S> ::= pop (<stack↑cat(S,E)>)

<element↑E> ::= top (<stack↑cat(S,E)>)

<boolean↑true> ::= empty (<stack↑empty>)

<boolean↑false> ::= empty (<stack↑cat(S,E)>)

The above EAG generates all valid stack expressions and the synthesised attribute of $\langle \text{stack} \rangle$ is the value of the stack expression. An example of a stack expression is

```
push(pop(push(push(newstack ,e1),e2)),e3)
```

with the value $\text{cat}(\text{cat}(\text{empty},e1),e3)$.

Note that this EAG also contains rules for generating expressions of type element and boolean.

One nice property of the EAG is that errors are treated implicitly in the sense that only valid stack expressions can be generated. E.g. it is impossible to generate an expression like: $\text{pop}(\text{newstack})$.

The type ELM could be integers in which case we could define the domain

```
E: ELM=INTEGER=(zero | suc(INTEGER) | pred(INTEGER))
```

The stack example makes only use of synthesised attributes. Below we define the data type partial mapping from a set D into a set R ($\{D \rightarrow R\}$). Let $d:D$, $r:R$, $f,g:\{D \rightarrow R\}$, then $\{\}$ is the empty map, $\{d \rightarrow r\}$ is the map defined in one point, $f \cup g$ is the union of f and g (only defined if the domains of f and g are disjoint), $f \setminus g$ is the overriding of f by g (the values of g are used 'before' those of f).

Example 5.2. Specification of partial mapping.

Domain

```
f,g,h: M=(empty | add(M,D,R)); d: D=...; r: R=...
```

Rules

```
<map↑empty> ::= {}
```

```
<map↑add(empty,d,r)> ::= { <dom↑d> -> <range↑r> }
```

```
<map↑h> ::= <map↑f> U <map↑g> <disjoint↓f↓g> <union↓f↓g↑h>  
| <map↑f> \ <map↑g> <union↓f↓g↑h>
```

```
<range↑r> ::= apply ( <map↑f> , <dom↑d> ) <apply↓f↓d↑r>
```

```
<disjoint↓empty↓f> ::= EMPTY
```

```
<disjoint↓add(f,d,r)↓g> ::= <undef↓d↓g> <disjoint↓f↓g>
```

```
<undef↓d↓empty> ::= EMPTY
```

```
<undef↓d↓add(f,d,r)> ::= <not-equal↓d↓d1> <undef↓d↓f>
```

```
<union↓f↓empty↑f> ::= EMPTY
```

```
<union↓f↓add(g,d,r)↑add(h,d,r)> ::= <union↓f↓g↑h>
```

$\langle \text{apply} \downarrow \text{add}(f,d,r) \downarrow d \uparrow r \rangle ::= \text{EMPTY}$

$\langle \text{apply} \downarrow \text{add}(f,d_1,r_1) \downarrow d \uparrow r \rangle ::= \langle \text{not-equal} \downarrow d \downarrow d_1 \rangle \langle \text{apply} \downarrow f \downarrow d \uparrow r \rangle$

The nonterminals $\langle \text{disjoint} \rangle$, $\langle \text{undef} \rangle$, $\langle \text{union} \rangle$ and $\langle \text{apply} \rangle$ are used in a special way. They can only generate the empty string and they do if some relations hold between their actual attributes. The rules for $\langle \text{not-equal} \rangle$ depend on the domain D . This technique is also used in connection with vw -grammars where such nonterminals are called primitive predicate symbols.

5.2 Specifying Semantics of Programs.

The techniques used for defining abstract data types may easily be used to define an operational semantics of a programming language. This is done by using an EAG to define all possible executions of a given program. In [Marcotty et al. 76] a small programming language is defined this way but using vw -grammars.

We define the semantics of the language presented in the example of 4.2.2.

Example 5.3

Domain

$T: \text{TREE} = (\text{seq}(\text{TREE}, \text{TREE}) \mid \text{assign}(\text{NAME}, \text{EXP})$
 $\mid \text{cond}(\text{EXP}, \text{TREE}, \text{TREE}) \mid \text{rep}(\text{EXP}, \text{TREE}));$
 $E: \text{EXP} = (\text{plus}(\text{EXP}, \text{EXP}) \mid \text{v}(\text{NAME}) \mid \text{c}(\text{INTEGER}));$
 $N: \text{NAME}; I: \text{INTEGER}; S: \text{STATE} = \{\text{NAME} \rightarrow \text{INTEGER}\};$

Rules

$\langle \text{program} \uparrow S \rangle ::= \langle \text{stmt} \uparrow T \rangle \langle \text{execute} \downarrow \downarrow \uparrow T \uparrow S \rangle$
 $\langle \text{stmt} \uparrow \text{seq}(T_1, T_2) \rangle ::= \langle \text{stmt} \uparrow T_1 \rangle ; \langle \text{stmt} \uparrow T_2 \rangle$
 $\langle \text{stmt} \uparrow \text{assign}(N, E) \rangle ::= \langle \text{var} \uparrow N \rangle := \langle \text{exp} \uparrow E \rangle$
 $\langle \text{stmt} \uparrow \text{cond}(E, T_1, T_2) \rangle ::= \text{if} \langle \text{exp} \uparrow E \rangle \text{ then} \langle \text{stmt} \uparrow T_1 \rangle \text{ else} \langle \text{stmt} \uparrow T_2 \rangle$
 $\langle \text{stmt} \uparrow \text{rep}(E, T) \rangle ::= \text{while} \langle \text{exp} \uparrow E \rangle \text{ do} \langle \text{stmt} \uparrow T \rangle$
 $\langle \text{var} \uparrow N \rangle ::= \langle \text{name} \uparrow N \rangle$
 $\langle \text{exp} \uparrow \text{plus}(E_1, E_2) \rangle ::= \langle \text{exp} \uparrow E_1 \rangle + \langle \text{exp} \uparrow E_2 \rangle$
 $\langle \text{exp} \uparrow \text{v}(N) \rangle ::= \langle \text{var} \uparrow N \rangle$
 $\langle \text{exp} \uparrow \text{c}(I) \rangle ::= \langle \text{const} \uparrow I \rangle$
 $\langle \text{execute} \downarrow S \downarrow \text{seq}(T_1, T_2) \uparrow S_2 \rangle ::= \langle \text{execute} \downarrow S \downarrow T_1 \uparrow S_1 \rangle \langle \text{execute} \downarrow S_1 \downarrow T_2 \uparrow S_2 \rangle$


```

<execute $\psi$ S $\psi$ assign(N,E) $\uparrow$ S\ $\{N\rightarrow I\}$ > ::= <eval $\psi$ S $\psi$ E $\uparrow$ I>

<execute $\psi$ S $\psi$ cond(E,T1,T2) $\uparrow$ S1> ::=
    <eval $\psi$ S $\psi$ E $\uparrow$ pos(I)> <execute $\psi$ S $\psi$ T1 $\uparrow$ S1>
    | <eval $\psi$ S $\psi$ E $\uparrow$ zero> <execute $\psi$ S $\psi$ T2 $\uparrow$ S1>

<execute $\psi$ S $\psi$ rep(E,T) $\uparrow$ S> ::= <eval $\psi$ S $\psi$ E $\uparrow$ zero>

<execute $\psi$ S $\psi$ rep(E,T) $\uparrow$ S2> ::=
    <eval $\psi$ S $\psi$ E $\uparrow$ pos(I)> <execute $\psi$ S $\psi$ T $\uparrow$ S1> <execute $\psi$ S1 $\psi$ rep(E,T) $\uparrow$ S2>

<eval $\psi$ S $\psi$ plus(E1,E2) $\uparrow$ I1+I2> ::= <eval $\psi$ S $\psi$ E1 $\uparrow$ I1> <eval $\psi$ S $\psi$ E2 $\uparrow$ I2>

<eval $\psi$ S $\psi$ v(N) $\uparrow$ S[N]> ::= EMPTY

<eval $\psi$ S $\psi$ c(I) $\uparrow$ I> ::= EMPTY

```

The definition of the language consists of some rules that define the syntax and collect the given program in a tree structure. The remaining rules define an execution of the given program starting with an empty state and returning a final state. The result of the program is this final state which is the meaning of the program (a synthesised attribute of the start symbol).

The integers are defined in the following way:

```
Integer=(neg(N) | zero | pos(N)); N=(one | suc(N))
```

5.3 Nondeterminism and Concurrency.

We shall now investigate the possibilities for using EAGs as a model for defining nondeterministic and concurrent computations.

Let us for a moment consider CFGs as a model for defining computations consisting of derivations, and let a meaning of a (terminal-) string be a derivation of it. CFGs are by nature nondeterministic because of the alternative operator |.

For a given derivation one may define a partial order between applications of productions in the derivation: Let r and s be applications of productions, then $r < s$ if s cannot be done before r is done. Juxta position in a CFG may then be viewed as a concurrency operator as all nonterminals of a sentential form may be rewritten independently of each other.

If one turns to context-sensitive grammars then juxta position is no longer necessarily a concurrency operator as some dependence may exist between symbols in a sentential form.

We shall not go further with this view upon Chomsky-grammars but instead turn to

EAGs. We shall only consider non-circular EAGs in this connection.

Nondeterminism of an EAG is still described by the alternative operator in the sense that a given string may have several parse trees and thus several meanings. Consider a parse tree t . The dependency graph defines a partial ordering between attributes in t in the sense that $a < b$ if the value of b depends upon the value of a . This partial ordering may be interpreted to define the amount of concurrency that can be applied during the computation of the attribute values by some machine.

A dependency and independency relation may be defined between symbols on the right side of a production:

Let A and B be symbols on the rightside of some production. A depends on B if an attribute in a defining position of B is used in an applied position of A . A and B is independent iff neither A is dependent on B nor vice versa.

Consider example 5.3 and the rules defining $\langle \text{execute} \downarrow S \downarrow \text{seq}(T1, T2) \uparrow S2 \rangle$ and $\langle \text{eval} \downarrow S \downarrow \text{plus}(E1, E2) \uparrow I1 + I2 \rangle$. In the rule defining $\text{seq}(T1, T2)$ the two instances of $\langle \text{execute} \rangle$ on the right side are dependent whereas the two instances of $\langle \text{eval} \rangle$ on the right side of the rule defining $\text{plus}(E1, E2)$ are independent. In the first case sequentiality is imposed. In the second case no evaluation order is imposed and it may in fact go on concurrently.

Below we give four examples of how EAGs may be used to define nondeterminism and concurrency in programming languages.

Example 5.4 Nondeterminism.

Assume that we add the following statement to the grammar of example 5.3:

```
<stmt> ::= one of <stmt> or <stmt> end
```

We may then extend the definition of TREE by $\text{oneof}(\text{TREE}, \text{TREE})$, and add the rules

```
<stmt↑oneof(T1, T2)> ::= oneof <stmt↑T1> or <stmt↑T2> end
```

```
<execute↓S↓oneof(T1, T2)↑S1> ::= <execute↓S↓T1↑S1> | <execute↓S↓T2↑S1>
```

Now either of the two alternatives may be selected in a given execution.

Example 5.5. Interleaving.

Consider the following statement:

```
allof <- S1 -> ; <- S2 -> ; ... ; <- SN -> ;  
and <- S1' -> ; <- S2' -> ; ... ; <- SM' -> ;  
end
```

The two sequences of statements may be executed in 'parallel'. The statements en-

closed by \leftarrow , \rightarrow are considered to be indivisible actions. Parallel execution will in this case then mean interleaving of the two sequences of indivisible actions. This kind of 'parallel' construct is in fact more nondeterministic than parallel.

We extend the definition of TREE by $\text{allof}(\text{ITREE}, \text{ITREE})$, and add the domain

IT: ITREE= \langle TREE \rangle^*

and add the rules

$\langle \text{stmt} \uparrow \text{allof}(\text{IT1}, \text{IT2}) \rangle ::= \underline{\text{allof}} \langle \text{indiv-stmt} \uparrow \text{IT1} \rangle \underline{\text{and}} \langle \text{indiv-stmt} \uparrow \text{IT2} \rangle \underline{\text{end}}$

$\langle \text{indiv-stmt} \uparrow [] \rangle ::= \text{EMPTY}$

$\langle \text{indiv-stmt} \uparrow \text{T.IT} \rangle ::= \leftarrow \langle \text{stmt} \uparrow \text{T} \rightarrow ; \langle \text{indiv-stmt} \uparrow \text{IT} \rangle$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{allof}(\text{T.IT}, \text{IT1}) \uparrow \text{S2} \rangle ::=$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{T} \uparrow \text{S1} \rangle \langle \text{execute} \downarrow \text{S1} \downarrow \text{allof}(\text{IT}, \text{IT1}) \uparrow \text{S2} \rangle$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{allof}(\text{IT1}, \text{T.IT}) \uparrow \text{S2} \rangle ::=$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{T} \uparrow \text{S1} \rangle \langle \text{execute} \downarrow \text{S1} \downarrow \text{allof}(\text{IT1}, \text{IT}) \uparrow \text{S2} \rangle$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{allof}([], []) \uparrow \text{S} \rangle ::= \text{EMPTY}$

Example 5.6. Concurrency.

We now turn to specify the semantics of concurrent computations. By concurrent we mean that two computations can be carried out independently of each other. Consider the following construct:

cobegin S1 and S2 end

We shall require that S1 and S2 do not refer to the same variables in order to ensure that S1 and S2 can in fact be executed concurrently.

In order to specify the semantics of cobegin we make the following extensions of example 5.3:

- $\langle \text{stmt} \rangle$ is extended with an extra synthesised attribute that is used to collect the set of names used in the statement,
- TREE is extended with $\text{co}(\text{TREE}, \text{TREE})$,
- the following rules are added

$\langle \text{stmt} \uparrow \text{co}(\text{T1}, \text{T2}) \uparrow \text{R1} \cup \text{R2} \rangle ::= \underline{\text{cobegin}} \langle \text{stmt} \uparrow \text{T1} \uparrow \text{R1} \rangle \underline{\text{and}} \langle \text{stmt} \uparrow \text{T2} \uparrow \text{R2} \rangle \underline{\text{end}}$

$\langle \text{execute} \downarrow \text{S} \downarrow \text{co}(\text{T1}, \text{T2}) \uparrow \text{S1} \setminus \text{S2} \rangle ::= \langle \text{execute} \downarrow \text{S} \downarrow \text{T1} \uparrow \text{S1} \rangle \langle \text{execute} \downarrow \text{S} \downarrow \text{T2} \uparrow \text{S2} \rangle$

Now why does this define a concurrent execution of T1 and T2. Consider the last rule:

The two instances of execute on the right side are independent according to our previous definition, i.e. this defines two independent computations. Note that the expression $S1 \setminus S2$ is not symmetric. One might instead pass only the restriction of S corresponding to the names being used by $T1$ to the first $\langle \text{execute} \rangle$ and similarly for the second. The result will then be $S \setminus (S1 \cup S2)$ instead of $S1 \setminus S2$.

5.4 Conclusion.

It has been shown how to use EAGs to define operational semantics for various programming language constructs. In our opinion these definitions are straightforward to make and easy to read. Again this is a matter of personal opinion. The useability of such definitions have to be tested on real programming languages.

Another interesting question is the possibility for generating a compiler automatically from such a description. We shall not discuss methods for doing this but just indicate some of the problems and give some ad hoc rules that may be a basis for further investigations.

Let nonterminals that only generate the empty string be called primitive predicates. (like $\langle \text{execute} \rangle$ and $\langle \text{eval} \rangle$ in example 5.3)

None of the standard AG-evaluators (including the one of section 7) are directly useable for this purpose. The reason is that there are an infinite number of parse trees (due to the heavy use of primitive predicates) to be considered.

It is possible to strip the underlying CFG for primitive predicates and then construct a parse tree in the usual way (this may still cause difficulties if the underlying CFG is ambiguous, but the EAG is structurally unambiguous). Afterwards one may then fill in the primitive predicates by simulating all possible derivations. In order to make this process efficient one must probably restrict the degree of non-determinism in order to recognise deterministic parts of the EAG.

The generation process may be simplified if all rules defining primitive predicates are required to be Left-attributed. This is in fact the case in all our examples.

In some cases it is possible to factorise the EAG. Consider the rules in ex. 5.3 defining the execution of $\text{cond}(E, T1, T2)$. These rules may be factorised into one rule with the following right-side: $\langle \text{eval} \downarrow S \downarrow E \uparrow V \rangle \langle \text{selection} \downarrow S \downarrow T1 \downarrow T2 \downarrow V \uparrow S1 \rangle$ where $\langle \text{selection} \rangle$ selects between $T1$ or $T2$ depending upon the value of V . This definition makes it easier to avoid two evaluations of $\langle \text{eval} \rangle$.

6. Pure and Multi-Level Extended Attribute Grammars.

Here we shall give certain proposals for future work with the AG/EAG formalisms.

When using an EAG to define a language one needs to specify the domains and operations upon these domains. A disadvantage of EAGs (AGs in general) is that the formalism itself does not contain a method for specifying these domains. This disadvantage is not present in vW-grammars, (pure) affix-grammars and extended affix grammars as the domains here are specified as context free (string) languages.

In section 5.1 we have seen that EAGs may be used to define abstract data types. We shall thus propose to use EAGs to define the domains of a particular EAG. This has the effect that EAGs now become a complete formalism. In order to accomplish this we must define a bottom in the hierarchy of EAGs, i.e. a domain constructor which is part of the formalism. Because of the similarities between extended affix grammars and EAGs it seems natural to let extended affix grammars be the bottom. However as argued in [Watt & Madsen 77] the discriminated union seems much more adequate as the basic domain constructor. A discriminated union may be viewed as an abstract context free grammar and is useful to define abstract properties of a language instead of its concrete syntax. A sentential form of a CFG does not contain information about the structure of the form (how it is derived) whereas this is explicitly available in a discriminated union. Altogether we find the latter to be more useful as it leads to more compact and readable descriptions.

A set of discriminated unions can be viewed as a regular tree grammar ([Engelfriet 74]). The theory of tree languages is well founded and gives a good notation. In the formal model we shall thus propose to use tree languages as the basic domain. We may then reformulate definition 2.1 as follows:

[6.1] A Pure (Extended) Attribute Grammar is a 5-tuple:

$$G = (D, V, Z, B, R).$$

The elements of G are similar to the ones of an EAG except that $D = (N, W, P, S)$ is a regular tree grammar and the domains of attribute positions and variables are defined as nonterminals of D .

One may then define a Multi-Level EAG as an EAG where the domains are defined by other Multi-Level EAGs or Pure EAGs.

The practical and theoretical use of such concepts need to be further investigated. In [Engelfriet & Filé 79], [Madsen 79b], and [Riis 80] some results on attribute grammars and tree transducers are given.

7. A General Fast Attribute Grammar Evaluator.

When implementing an evaluator for AGs one is among others concerned with the following problems:

- (a) Decide whether or not to test the AG for circularity,
- (b) deciding the order of evaluation of the attributes,
- (c) the storing of the syntax tree and the attribute values during the evaluation,
- (d) the possibility of evaluating some of the attributes during the (context-free) parsing phase,
- (e) the possibility of letting the attributes influence the (context-free) parsing (so-called affix-(or attribute) directed parsing [Watt 74b]).

In order to overcome (or ease) some of these problems a number of subclasses of AGs have been defined. These include L-attributed grammars ([Lewis et al. 74]), multipass AGs ([Bochmann 76]), and many more ([Watt 77],[Mayoh 78]). Affix grammars ([Koster 71]) may also be considered as such a subclass.

A class of general evaluators build the syntax tree and then evaluate the attributes in some order. The efficiency of these evaluators depends on the method used to determine the order of evaluation of the attributes. These evaluators may roughly be divided into two classes:

- (1) Static evaluation order: Those where the order of evaluation can be determined from the AG; i.e. the attributes of a symbol X are evaluated in an order which is independent of the subtree below X ([Kennedy & Warren 76], [Saarinen 78]). These AGs are called benign in [Mayoh 78b].
- (2) Dynamic evaluation order: Those where the order of evaluation is determined by the parse tree built during the parsing phase ([Lorho 77], [Cohen & Harry 79], [Kennedy & Ramanathan 79]).

In case (1) the order of evaluation may be defined once and for all by some analysis done by the TWS. In case (2) the analysis for determining the order of evaluation must be done for each parse tree. It is obvious that evaluators in class (1) accept a smaller class of AGs than those in class (2).

Less general methods avoid building the parse tree but use a linear representation in form of e.g. a right-parse or left-parse ([Aho & Ullman 72]). Attributes are then assigned values during one or more scans (forward or backwards) of the linear tree. These methods include L-attributed AGs, multipass AGs, The Alternating Semantic Evaluator ([Jazayeri & Walter 75]).

It is well known that the underlying CFG plays an important role here. If the underlying CFG is LL(k) then the construction of a left-parse may coincide with a first left-to-right scan of the linear syntax tree. If the underlying CFG is known to be

LR(k) then one can in general only evaluate synthesised attributes during the parsing phase. Affix-free LR-grammars ([Watt 77]) are a large and useful subclass of AGs with LR(k) underlying CFGs. They include all LL(k) grammars and have the same advantages as LL(k)-grammars and so have SD-grammars ([Lewis et al. 74]).

The evaluator presented below (from now on called the DAG-evaluator) is based on dynamic evaluation order. The parse tree has to be represented in the form of a right-parse. A DAG is constructed during one left-to-right scan of the right-parse. This DAG represents the dependency-graph of the parse tree and the necessary information for evaluating all attributes. The value of the attributes may be evaluated by a recursive scan of the DAG.

The advantages of the DAG-evaluator are: (1) it works for all AGs, (2) the syntax tree is represented as a right-parse, (3) the DAG is constructed through a single scan of the right-parse, (4) if the underlying CFG is LR(k) then the DAG may be constructed during the parsing phase.

Furthermore the method will also work for circular AGs. The DAG will then contain cycles. The recursive scan may detect these cycles. As mentioned in section 4 it may be meaningful to continue the evaluation. In many cases the DAG itself may be considered as the result of the evaluation. E.g. the AG in ex. 4.5 defines a translation from parse trees to functions (command continuations). In this case the DAG itself may be a suitable representation of the function and such a DAG may have cycles in it.

In many cases it may be desirable to let the TWS check the AG for circularity and not defer it to the evaluator.

The DAG-evaluator is based upon the fact that any AG has an equivalent one using only synthesised attributes. It may in fact be viewed as an implementation based upon the transformation of def. 4.2 The DAGs being constructed by the evaluator may thus be viewed as a representation of the function valued attributes.

The DAG-evaluator is used in NEATS which is a TWS based on EATGs ([Jespersen et al. 78]).

7.1 The DAG-evaluator.

Below we describe the DAG evaluator. It is assumed that a right-parse of the input string has been constructed. During a left-to-right scan of the right-parse a DAG is constructed. In the following description this left-to-right scan is described as a simulation of the parsing phase (as mentioned it may in fact be performed during the parsing if the grammar is e.g. LR(k)).

The evaluator uses three stacks, the ordinary parse stack (PARSE), a stack for

storing information about synthesised attributes (SYN), and one for storing information about inherited attributes (INH).

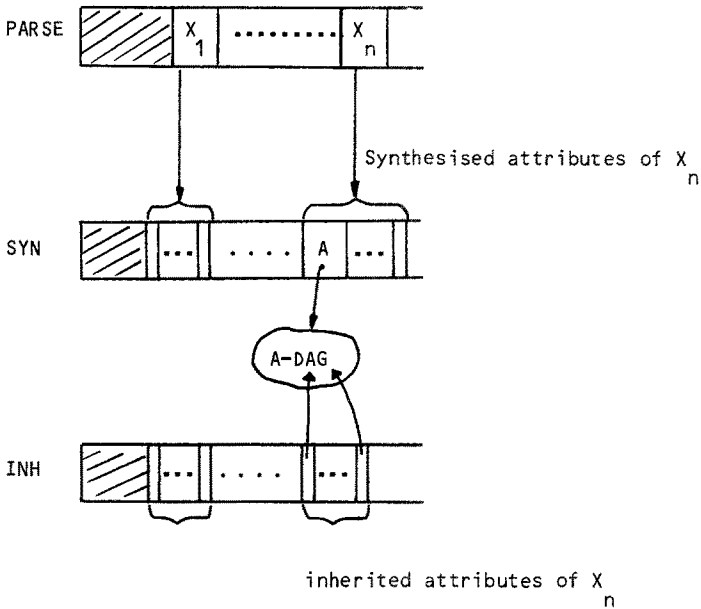
Assume that the parser reduces by a production

$$p: X_0 \rightarrow X_1 X_2 \dots X_n$$

and that

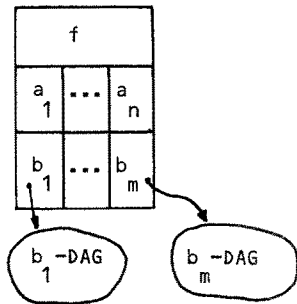
- SYN contains a value for each right side synthesised attribute of p (called $RS[p]$). Each value is a pointer to an expression-DAG which when evaluated will yield the value of the corresponding synthesised attribute. Let SA be a synthesised attribute of X_i . Some of the leaves in the DAG for SA will correspond to inherited attributes of X_i , and a value has not yet been filled in for these leaves.
- INH contains a value for each right side inherited attribute of p ($RI[p]$). Let IA be an inherited attribute of X_i . The value on INH corresponding to IA is a pointer to a list of leaves in the DAGs corresponding to the synthesised attributes of X_i . A value for IA has to be inserted at each such leaf in order to define the synthesised attributes.

The following picture illustrates the situation:



When the reduction is applied the following steps are performed:

- (1) For each leftside inherited attribute (LI[p]) an empty List is constructed.
- (2) For each rightside inherited attribute A (RI[p]) a DAG-node is constructed: The value of A is a function: $f(a_1, \dots, a_n, b_1, \dots, b_m)$ where each $a_i \in LI[p]$ and $b_i \in RS[p]$. The DAG-node has the form:



The DAG-node has a pointer corresponding to each a_i and b_i . The pointer for b_i points to the DAG for b_i . The pointer field for a_i is chained in the list constructed for a_i in step (1).

The list stored in INH corresponding to A is scanned and each leaf in the list is replaced by a pointer to the newly constructed DAG-node for A.

- (3) For each leftside synthesised attribute A in $LS[p]$, a DAG is constructed in the same way as in step (2).
- (4) All elements on INH and SYN corresponding to rightside attributes of p are popped off INH and SYN. The values corresponding to leftside attributes of p (the ones constructed in steps (1) and (3)) are pushed on INH and SYN.

Remark. If in steps (2) and (3) $n=0$ and $m=1$ and f is the identity function, then a new DAG-node is not constructed. Pointers to the DAG-node for A will instead point to the DAG for b_1 .

Traversing the DAG.

When the scan of the right parse is finished the SYN-stack will contain a pointer (a root of the DAG) for each synthesised attribute of the start-symbol.

The values of these attributes may be evaluated by a recursive scan of the DAG starting at each root. Assume that the DAG contains no cycles. By using a so-called 'depth-first-search' of the DAG each node need only be visited once. After the visit of a node a value may be assigned to the node. The next visit to the node may then use this value.

Cycles in the DAG may easily be detected by the depth-first-search algorithm. It will then depend upon the actual AG whether or not a further evaluation is meaningful.

Some domains, like partial maps, may contain 'large' values. In such cases it may not be practical to store a value at each node without using a shared representation. The DAG itself is such a shared representation. By using a sub-DAG as the representation of a value, this sub-DAG may of course be traversed several times during the scan of the DAG.

For common used domains (partial maps) it may be possible to transform the DAG into a more suitable representation.

7.2 Examples.

Example 7.1

Consider the following EAG.

Domain Env: [NAME -> INTEGER]; V: INTEGER; N: NAME;

Rules

$\langle \text{evaluation} \uparrow V \rangle ::= \langle \text{exp} \downarrow C \rangle \uparrow V$

$\langle \text{exp} \downarrow \text{Env} \uparrow V \rangle ::= (\langle \text{id} \uparrow N \rangle = \langle \text{exp} \downarrow \text{Env} \uparrow V1 \rangle , \langle \text{exp} \downarrow \text{Env} \setminus \{N \rightarrow V1\} \uparrow V \rangle)$

$\langle \text{exp} \downarrow \text{Env} \uparrow V1 + V2 \rangle ::= \langle \text{exp} \downarrow \text{Env} \uparrow V1 \rangle + \langle \text{exp} \downarrow \text{Env} \uparrow V2 \rangle$

$\langle \text{exp} \downarrow \text{Env} \uparrow V \rangle ::= \langle \text{term} \downarrow \text{Env} \uparrow V \rangle$

$\langle \text{term} \downarrow \text{Env} \uparrow \text{Env}(N) \rangle ::= \langle \text{id} \uparrow N \rangle$

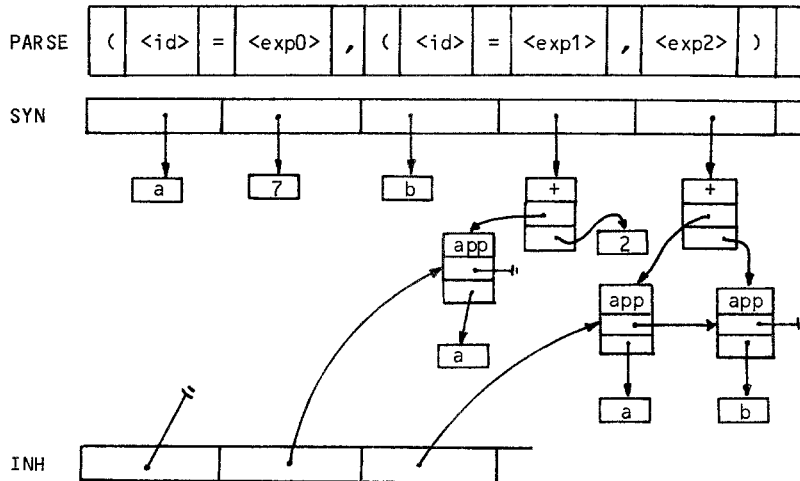
$\langle \text{term} \downarrow \text{Env} \uparrow V \rangle ::= \langle \text{const} \uparrow V \rangle$

The construct $(a=e1, e2)$ is an abbreviation of LET $a=e1$ IN $e2$. Consider a parse of the following string:

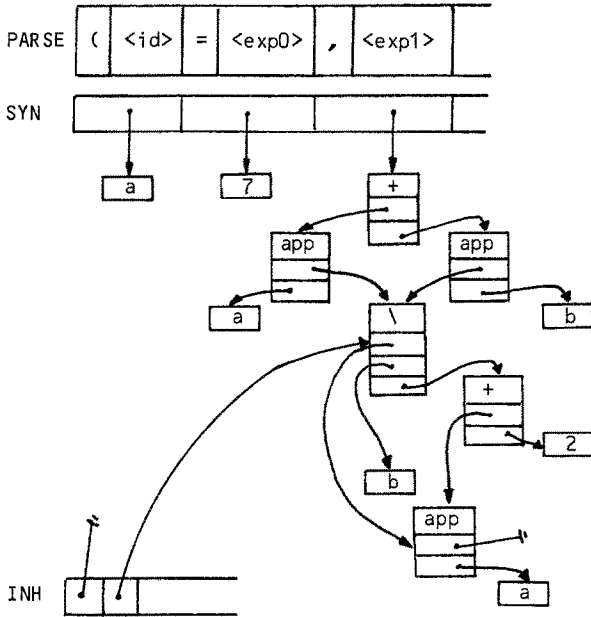
$(a=7, (b=a+2, a+b))$

↑

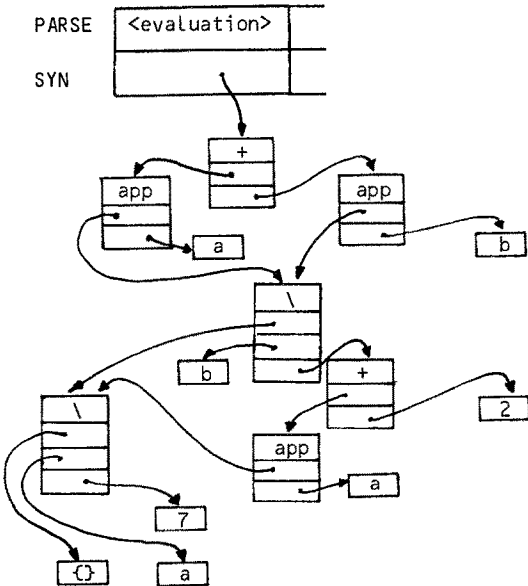
Consider the following snapshot of the stacks when the input is placed at ↑:



After the reduction $\langle \text{exp} \rangle ::= (\langle \text{id} \rangle = \langle \text{exp0} \rangle , \langle \text{exp1} \rangle)$ we have



The next two reductions yield



A reduction of the DAG will yield the value 16.

Example 7.2. Circular AG.

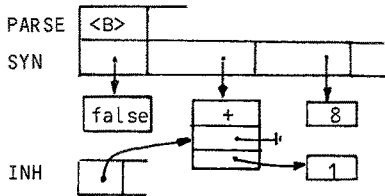
Consider the following EAG

$\langle A \uparrow y + z \rangle ::= \langle B \uparrow \text{cond}(c, y, z) \uparrow c \uparrow y \uparrow z \rangle$

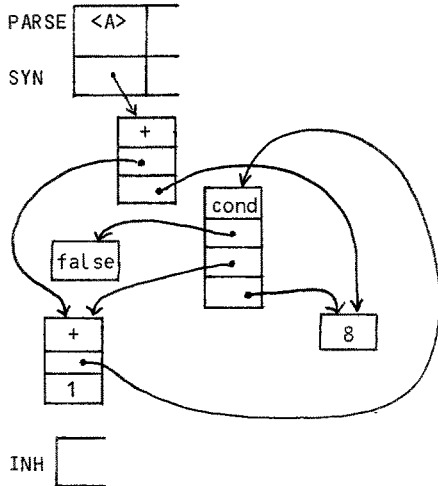
$\langle B \uparrow x \uparrow \text{true} \uparrow 7 \uparrow x + 2 \rangle ::= a$

$\langle B \uparrow x \uparrow \text{false} \uparrow x + 1 \uparrow 8 \rangle ::= b$

Assume that the input string b has been reduced to $\langle B \rangle$. We then have the following snapshot:



Next when $\langle B \rangle$ is reduced to $\langle A \rangle$ we have:



The constructed DAG is circular, but an evaluation may reduce it to the value 17.

7.3 Evaluation.

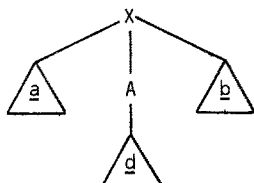
It should be obvious that the DAG-evaluator works for any well formed AG in normal form. It is also reasonably fast as it does not need to detect an evaluation order for the attributes in the tree of the parsed string. Furthermore the evaluation of the DAG may be done during one scan of the DAG.

The main problem is that the DAG-evaluator may take up too much space to be practical. We shall compare the space requirements with an evaluator that builds the par-

se tree and then assigns values to attributes in some order. We assume that at each node there will be space to store the values of the associated attributes (or a pointer to values).

In the DAG-evaluator the situation is as follows: (1) The parse tree is not stored, (2) only a subset of the attributes in the parse tree have a corresponding node in the DAG.

Cf. (2): Consider a symbol appearing in the parse tree as part of rules

$$X \rightarrow a A b \text{ and } A \rightarrow d$$


A synthesised attribute of A will have a DAG-node if the function in the rule $A \rightarrow d$ defining the attribute is nontrivial (not the identity function). Similarly for an inherited attribute of A . Thus attributes that are simple copies of other attributes will not occupy space in the DAG. Notice that the lists for collecting the use of inherited attributes take no additional space as they use the pointer fields that have not yet been filled in.

The DAG-node corresponding to a nontrivial function f may take up more space than the value of the attribute.

The size of the DAG is a linear function of the parse tree. It may be smaller or larger depending on the amount of copying and the format of the nontrivial functions.

It should be noticed that some attribute domains such as partial mappings are of a form where it is unreasonable to store the values at the nodes. These domains must be implemented as data structures with a shared representation. The DAG-representation of such domains may then turn out to be smaller.

In some situations it is possible to reduce the space requirements:

- If A is a synthesised attribute of X and the values of all the inherited attributes of X are known in the DAG for A then the value of A may be evaluated.
- In many situations during parsing, the values of some inherited attributes are in fact known. This is the case if the grammar is L-attributed and the underlying CFG is LL(k). Using a bottom-up parser this information is not directly available. In [Watt 77] an algorithm is given which is able to parse L-attributed AGs having an underlying CFG in a subset of the LR(k)-grammars. This subset includes LL(k) grammars and most practical LR-grammars. Watt's algorithm makes use

of the fact that for many practical L-attributed AGs most rules simply copy the inherited attributes of the left side. In the rules where the inherited attributes are given 'new' values the parser is often in an LL(k) situation (the LR-table has one item in its closure). Extra nonterminals are inserted in the rules before nonterminals where inherited attributes are given 'new' values. If the grammar is still LR, then the original AG is called affix-free. It is possible to detect automatically where such e-productions may be inserted without destroying LR-properties (introducing conflicts).

The algorithm of section 7.1 may be improved by using Watt's method and computing inherited attributes wherever possible.

Unfortunately this only improves the one pass case and does not help the multipass situation. In general one has to use secondary storage to store the DAGs if not enough primary storage is available.

8. Concluding remarks.

We think that it has been demonstrated that EAGs (and AGs and EATGs) may be used for many different purposes, and that a TWS based upon EATGs may serve many different purposes.

EAGs may also be used as a powerful tool for defining computations not necessarily connected with semantics of programming languages.

Various variants of attribute grammars have been discussed, EAGs, pure EAGs, multi-level EAGs, and AGs with Scott-type domains. Furthermore pure EAGs have been formulated within the theory of tree languages.

The three main approaches are the original AGs of Knuth, AGs with Scott-type domains and EAGs. We tend to prefer the generative approaches or the equational approaches for the Knuth AGs.

It has been demonstrated elsewhere ([Madsen et al.76], [Watt & Madsen 77], and [Watt 79]) that EAGs are suitable for defining the analysis phase (context-sensitive syntax and a translation into an intermediate form) for realistic programming languages.

The purpose of this paper has been to demonstrate that EAGs have a much wider range of applications. The main subject has been to show how three dominating methods for specifying semantics can be expressed within the framework of EAGs.

This makes it possible within the same TWS to experiment with several definitions and implementations. Consistency between different semantic definitions can be proved within the framework of the same formalism. This may turn out to be an advantage although it has not been treated here.

Given a TWS based on EAGs it should be possible to define

- (1) The context-sensitive syntax of a programming language, and
- (2) the following kinds of translators/compiler: - a verification generator, - a compiler based on denotational semantics corresponding to the ones generated by SIS ([Mosses 79]), - a compiler based on abstract interpretation, - the analysis phase of a (production) compiler, and - a compiler generating code to an abstract machine.

For some of the approaches to be satisfactory/efficient there are (as mentioned) certain problems to be considered.

First of all an efficient AG evaluator must be available. We think that the DAG-evaluator presented in section 7 is a useful contribution to this. It avoids most of the standard problems concerning the order of evaluation of attributes and it constructs the DAG during one left-to-right scan of a right parse of the input string. Furthermore it works for all AGs and does not imply a strange subclass. This is important for the users of a TWS.

The possibilities of letting the attributes influence the parse tree/translation selected for a given string needs a satisfactory solution in order to make full use of EAGs to handle ambiguities in the underlying CFG.

9. References.

- Aho A.V., J.D. Ullman [72]: The theory of parsing, translation, and compiling. Vol 1: Parsing, 1972. Vol 2: Compiling, 1973. Englewood Cliffs (N.J.): Prentice Hall.
- Bjørner D. [78]: The systematic development of a compiling algorithm. In: Le point sur la compilation. (M. Amirchahy, D.Neel, ed.), IRIA, 1978.
- Bochmann G.V. [76]: Semantic evaluation from left to right. Comm. ACM 19, 55-62(1976).
- Chirica L.M. [76]: Contributions to compiler correctness. University of California, Los Angeles, Ph. D. Thesis, October 1976.
- Cohen R., E. Harry [79]: Automatic generation of near-optimal linear-time translators for non-circular attribute grammars. Conference record of the Sixth ACM Symposium on Principles of Programming Languages, San Antonio, January 1979.
- Engelfriet J. [74]: Tree automata and tree grammars. Computer Science Department, Aarhus University, DAIMI FN-10, 1974.
- Engelfriet J., G. Filé [79]: The formal power of one-visit attribute grammars. Twente University of Technology, The Netherlands, mem. no. 286, October 1979.
- Ganzinger H. [79a]: An approach to the derivation of compiler descriptions from the mathematical semantic concept. GI-9. Jahrestagung, Informatik-Fachberichte 19, Berlin-Heidelberg-New York, Springer 1979.
- Ganzinger H. [79b]: Some principles for the development of compiler descriptions from denotational language definitions. Der Technischen Universität München, (Summary), 1979.
- Gerhart S.L. [76]: Proof theory of partial correctness verification systems. SIAM J. Comput., Vol 5., No. 3., September 1976.
- Jazayeri M., K.G. Walter [75]: Alternating semantic evaluation. Proceedings of the ACM Annual Conference Minneapolis, Minn., October 1975.
- Jespersen P., M. Madsen, H. Riis [78]: New extended attribute system (NEATS). Computer Science Department, Aarhus University, 1978.

- Kennedy K., J. Ramanathan [79]: A deterministic Attribute Grammar evaluator based on dynamic sequencing. *ACM Toplas*, Vol 1, no. 1, July 1979 (142-160).
- Kennedy K., S.K. Warren [76]: Automatic generation of efficient evaluators for attribute grammars. Conference record of the Third ACM Symposium on Principles of Programming Languages, January 1976, 32-49.
- Knuth D.E. [68]: Semantics of context free languages. *Math. Sys. Theory*, Vol. 2, No. 2, 1968.
- Koster C.H.A. [71]: Affix grammars. In *ALGOL 68 Implementation* (J.F. Peck, ed.) pp. 95-109. Amsterdam: North-Holland 1971.
- Lewis P.M., D.J. Rosenkrantz, R.E. Stearns [74]: Attributed translations. *J. Computer and System Sciences* 9, 279-307 (1974).
- Lorho B. [77]: Semantic attributes processing in the system DELTA. In *Methods of algorithmic language implementation* (C.H.A. Koster, ed.) pp 21-40, *Lecture Notes in Computer Science*, Vol 47. Berlin-Heidelberg-New York: Springer 1977.
- Madsen O.L., B.B. Kristensen, J. Staunstrup [76]: Use of design criteria for intermediate languages. Computer Science Department, Aarhus University, DAIMI PB-59, August 1976.
- Madsen O.L. [79a]: An introduction to attribute grammars by examples. Computer Science Department, Aarhus University, April 1979.
- Madsen O.L. [79b]: On defining semantics by means of extended attribute grammars. Computer Science Department, Aarhus University, DAIMI IR-14, September 1979. (Revised version DAIMI PB-109, January 1980).
- Marcotty M., H.F. Ledgard, G.V. Bochmann [76]: A sampler of formal definitions. *ACM, Computing Surveys*, Vol. 8, no. 2, 1976.
- Mayoh B.H. [76]: Verification and compilation in MSL. Computer Science Department, Aarhus University, June 1976, unpublished manuscript.
- Mayoh B.H. [78]: Attribute grammars and mathematical semantics. Computer Science Department, Aarhus University, DAIMI PB-90, August 1978.
- Mosses P. [79]: SIS - Semantics implementation system. Reference manual and user guide. Computer Science Department, Aarhus University, DAIMI MD-30, August 1979.
- Nielsen J.E.G. [75]: JQNS systemet. Et system til automatisk generering af verifiere. Computer Science Department, Aarhus University, Master thesis (in Danish), 1975.
- Riis H. [80]: Subclasses of attribute grammars. Computer Science Department, Aarhus University, Master Thesis, DAIMI PB-114, January 1980.
- Saarinen M. [78]: On constructing efficient evaluators for attribute grammars. Automata, Languages and Programming, Fifth Colloquium, Udine, July 1978. *Lecture Notes in Computer Science*, Berlin-Heidelberg-New York, Springer 1978.
- Tennent R.D. [76]: The denotational semantics of programming languages. *Comm. ACM*, Vol. 19, 8, 437-453(1976).
- Watt D.A. [74a]: An analysis oriented two-level grammars. University of Glasgow, Ph. D. Thesis, 1974.
- Watt D.A. [74b]: LR-parsing of affix-grammars. University of Glasgow, Report no. 7, 1974.
- Watt D.A. [77]: The parsing problem for affix-grammars. *ACTA Informatica*, 8, 1-20(1977).
- Watt D.A. [79]: An extended attribute grammar for Pascal. *Sigplan Notices*, Vol. 14, no. 2, 1979.
- Watt D.A., O.L. Madsen [77]: Extended attribute grammars. University of Glasgow, Report no. 10, July 1977. (revised version Computer Science Department, Aarhus University, DAIMI PB-105, November 1979)
- Wijngaarden A. van, B. Mailloux, J.E.L. Peck, C.H.A. Koster, M. Sintzoff, C.H. Lindsey, L.G.L.T. Meertens, R.G. Fisker [75]: Revised report on the algorithmic language ALGOL 68. *ACTA Informatica*, 5, 1-236(1975).
- Wilner W. [72]: Declarative semantic definition. STAN-CS-233-71, Ph. D. Thesis, Stanford, California, 1971.
- Wirth N. [71]: The programming Language PASCAL. *ACTA Informatica* 1, 35-63(1971).