Course Script

INF 5110: Compiler construction

INF5110, spring 2021

Martin Steffen

Contents

5 Semantic analysis **1**
5.1 Introduction 1
5.2 Attribute grammars 5

6 References **33**

Chapter 5
Semantic analysis

What is it about?

Learning Targets of this Chapter

1. “attributes”
2. attribute grammars
3. synthesized and inherited attributes
4. various applications of attribute grammars

Contents

5.1 Introduction 1
5.2 Attribute grammars 5

5.1 Introduction

5.1.1 Semantic analysis in general

Semantic analysis or *static analysis* is a very broad and diverse topic. The lecture concentrates on a few, but crucial aspects. This particular chapter here is concerned with *attribute grammars*. It’s a generic or general “framework” to “semantic analysis”. Later chapters also deal with semantic analysis, namely the ones about *symbol tables* and about *type checking*. In the context of the lecture, those chapters work basically on (abstract syntax) trees (except that for the symbol tables and for the type system, it’s not so visible). The fact that it’s a mechanism to “analyze trees” is most visible for attribute grammars: context-free grammars describe trees and the semantic rules (see later) added to the grammar specify how to analyze resulting trees.

Wrt. the general placement of semantic analysis in a compiler: Not all semantic analyses are “tree analyses”. Data flow analysis (on which we touch upon later) often works on *graphs* (typically control flow graphs). Furthermore, it’s not the case, that semantic analysis is restricted to be done directly after parsing. There are many semantic analyses done at later stages (and on other representations). In particular, it could be that a later intermediate representation uses a different form of syntax, closer to machine code (often call *intermediate code*). That syntax could also be given by a grammar, meaning that a program in that syntax corresponds to a tree of that syntax. As a result, one can apply techniques like attribute grammars also at that level (maybe thereby using in on the AST, and later differently on some intermediate code).

5.1.2 Overview

On a very high level, the attribute grammar format does the following: it enhances a given grammar by additional, so called *semantic rules*, which *specify* how trees conforming to the grammar should be analysed.

Two points might be noted here. First, the AG formalism adds rules on top of context-free *grammars*, but the intention is to specify analyses on *trees* formed *according to the given grammar*. Secondly, it's a *specification* of such tree analyses. The AG format is quite general, meaning that it allows to express all kinds of ways attributes should be evaluated. If not constrained in some way, the AG formalism can be seen as too expressive in that it leads to specifications contradict themselves or does not lead to proper implementation.

5.1.3 Side remark: XML

As some side remark, and not part of the technical content of the lecture: XML is some "exchange format" or markup language built around "trees". "markup" is kind of like the opposite of "mark-down" (tongue-in-cheek): mark-down allows easy textual representation, optimized for "human consumption". Mark-up offers easy consumption for "machines" (easy, unique parsing, easy exchange of "texts"). That's why XML reads so horrible to the naked eye.[1] Anyway, since pieces of XML-data are *trees*, there is also the notion of *grammars* according to which such trees are considered well-formed. In the XML terminology, that corresponds basically to *schemas*.[2] That being so, there are tools that check whether a tree adheres to a given schema, a problem that in that form does not present itself in parsing: the parser process generates *only* trees in the AST format. Since XML processing is concerned with "tree processing" (checking, transformation etc), there are some similarities with attribute grammars and some XML related technologies. We don't go deeper than that here.

5.1.4 Overview over the chapter resp. SA in general

- semantic analysis in general
- attribute grammars (AGs)
- symbol tables (not today)
- data types and type checking (not today)

5.1.5 What do we get from the parser?

- output of the parser: (abstract) syntax tree
- often: in anticipation: nodes in the tree contain "space" to be filled out by SA
- examples:
  - for expression nodes: *types*

[1] The `build.xml` from the oblig is some example of some xml-kind of file, used for "building" a project with *ant*.

[2] In UML context, the role of a grammar is taken by something with the slightly confusing title "meta-model".

– for identifier/name nodes: reference or pointer to the *declaration*

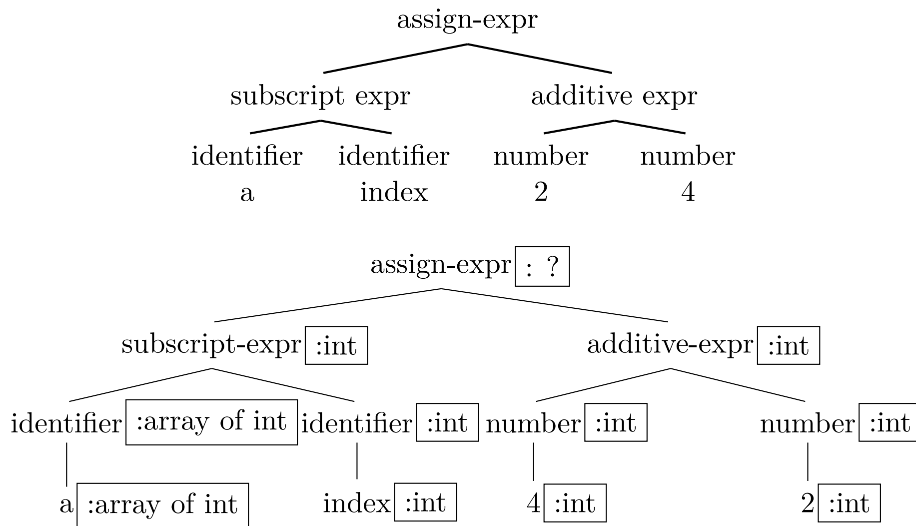

By "space", one might think of *fields* or *instance variables* in an object-oriented setting. Fields can be seen as one way to implent "attributes". When introducing attribute grammars, the notion of *attribute* will be a specific concept, namely the specific form of attributes in an attribute grammar. But very generally, an "attribute" means just a "property attached to some element". Typically here, attached to syntactic representations of the language, in particular to nodes in the abstract syntax tree. Since the notion of attribute is so general, it can take very different forms (like types, data flow information, all kind of extra information). Also, attributes in that sense, need not to be "attached" to abstract syntax trees, only. For instance, data flow information is extra information (calculated by data flow analysis) not to a syntax tree, but to something called a *control-flow graph*. So, since such graphs are *not* described by context-free grammars, and therefore, data flow analyses will *not* be described by attribute grammars.[3]

5.1.6 General: semantic (or static) analysis

5.1.6.1 Rule of thumb

Check everything which is possible *before* executing (run-time vs. compile-time), but cannot already done during lexing/parsing (syntactical vs. semantical analysis)

- Goal: fill out "semantic" info (typically in the AST)
- typically:
  – all *names declared*? (somewhere/uniquely/before use)
  – *typing*:

[3]Besides the reason mentioned —data-flow analyses typically operate on graphs, not trees— there is a second (but closely related) reason why DFA will in general not be done with AGs; the evaluation of AGs on a concrete tree explicitly disallows *cycles* in the dependency graph (see later). DFA in the general form *definitely* will have to handle cyclic situations.

    * is the declared type consistent with use
    * types of (sub)-expression consistent with used operations
- *border* between sematical vs. syntactic checking not always 100% clear
  - `if a then ...`: checked for syntax (and semantics)
  - `if a + b then ...`: semantical aspects as well?

5.1.7 SA is nessessarily approximative

- note: not all can (precisely) be checked at compile-time
  - division by zero?
  - "array out of bounds"
  - "null pointer deref" (like `r.a`, if `r` is null)
- but note also: *exact* type cannot be determined statically either

5.1.7.1 `if x then 1 else "abc"`

- statically: ill-typed[4]
- dynamically ("run-time type"): `string` or `int`, or run-time type error, if `x` turns out not to be a boolean, or if it's null

The fact that one cannot *precisely* check everything at compile-time is due to *fundamental* reasons. It's fundamentally impossible to predict the behavior of a program (provided, the programming language is expressive enough, i.e. Turing complete, which can be taken as granted for all general programming languages). The "fundamental reasons" mentioned above basically is a result of the famous *halting problem*. The particular version here is a consequence of that halting problem and is know as **Rice's theorem**. Actually it's more pessimisic than the sentence on the slide: Rice stipulates: **all** non-trivial semantic problems of a programming language are undecidable. If it were otherwise, the halting problem would be decidable as well (which it isn't, end-of-proof). Note that *approximative* checking is doable, resp. that's what the SA is doing anyhow.

As for type checking: the footnote refers to something which is a form of *polymorphism*, which is a form of "laxness" or "liberality" of the type system, which allows that some element of the language can have more than one type. In the particular example, it would be a specific form of polymorphism, namely (operator) overloading, in that + is used for addition as well as string concatenation. Additionally, in this particular situation, `1` is not just an integer, but *also a string*. The type checker may allow that, but if so, the later phases of the compiler must arrange it so that `1` is *actually converted* to a string (integers and strings are not represented uniformely in that `"42"` and `42` typically have not the same bit-level representation, so the compiler has to arrange somethere here.).

[4]Unless some fancy behind-the-scence type conversions are done by the language (the compiler). Perhaps `print(if x then 1 else "abc")` is accepted, and the integer `1` is implicitly converted to `"1"`.

5.1.8 An unrealistic dream

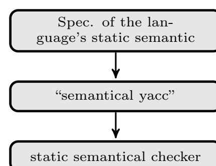

- no standard description language
- no standard "theory"
  - part of SA may seem ad-hoc, more "art" than "engineering", complex
- *but*: well-established/well-founded (and non-ad-hoc) fields do exist
  - *type systems*, type checking
  - *data-flow* analysis ....

- in general
  - semantic "rules" must be individually specified and implemented per language
  - rules: defined based on trees (for AST): often straightforward to implement
  - clean language design includes *clean semantic rules*

When saying that there is no general standard theory, well, of course there is the notion of *context-sensitive grammars*, a class of grammars more expressive than context free grammars, while not yet as expressive as Turing machines (= full compuation power). The notion of context-sensitive languages is sure well-defined, but as a formalism, it's too general, too unstructured to give much guiding light when it comes to concrete problems being analysed. They sure have been studied extensively, but as far a compilers are concerned, one deals more with specialized problems: a type systems specification is a case of a context-sensitive definition, though it's normally not formulated using the terminology and notation from the area of context-sensitive grammars. Context-sensitive grammars are not on the pensum.

5.2 Attribute grammars

5.2.1 Attributes

5.2.1.1 Attribute

- a "property" or characteristic feature of something
- here: of language "constructs". More specific in this chapter:
- of syntactic elements, i.e., for non-terminal and terminal nodes in syntax trees

5.2.1.2 Static vs. dynamic

- distinction between **static** and *dynamic attributes*
- association attribute ↔ element: *binding*
- *static* attributes: possible to determine at/determined at compile time
- dynamic attributes: the others ...

With the concept of *attribute* so general, very many things can be subsumed under being an attribute of "something". After having a look at how attribute grammars are used for "attribution" (or "binding" of values of some attribute to a syntactic element), we will normally be concerned with more concrete attributes, like the *type* of something, or the *value* (and there are many other examples). In the very general use of the word "attribute" and "attribution" (the act of attributing something to something) is almost synonymous with "analysis" (here semantic analysis). The analysis is concerned with figuring out the value of some attribute one is interested in, for instance, the *type* of a syntactic construct. After having done so, the result of the analysis is typically *remembered* (as opposed to being calculated over and over again), but that's for efficiency reasons. One way of remembering attributes is in a specific data structure, for attributes of "symbols", that kind of data structure is known as the *symbol table*.

5.2.2 Examples in our context

- data *type* of a variable : static/dynamic
- *value* of an expression: dynamic (but in seldom cases static as well)
- *location* of a variable in memory: typically dynamic (but in old FORTRAN: static)
- *object-code*: static (but also: dynamic loading possible)

The *value* of an expression, as stated, is typically *not* a static "attribute" (for reasons which I hope are clear). Later in this chapter, we will actually *use* values of expressions as attributes. That can be done, for instance, if there are *no* variables mentioned in the expressions. The values of those values typically are not known at compile-time and would not allow to calculate the value at compile time. However, having no variables is exactly the situation we will see later.

As a side remark: values of a variable are typically dynamic "attributes". In some cases the are not, for instance if they are declared as immutable. Even with standard, i.e., mutable variables, *sometimes* the compiler *can* figure out, that, in some situations, the value of a variable *is* at some point is known in advance. In that case, an *optimization* could be to *precompute* the value and use that instead. To figure out whether or not that is the case is typically done via *data-flow analysis* which operates on *control-flow graphs* (not trees). These form of analyses are therefore not done via attribute grammars in general.

5.2.3 Attribute grammar in a nutshell

- AG: general formalism to bind "attributes to trees" (where trees are given by a CFG)[5]
- two potential ways to calculate "properties" of nodes in a tree:

5.2.3.1 "Synthesize" properties

define/calculate prop's *bottom-up*

[5]Attributes in AG's: *static*, obviously.

5.2.3.2 "Inherit" properties

define/calculate prop's *top-down*

- allows both *at the same time*

5.2.3.3 Attribute grammar

CFG + **attributes** one grammar symbols + **rules** specifing for each production, how to determine attributes

- *evaluation* of attributes: requires some thought, more complex if mixing bottom-up + top-down dependencies

5.2.4 Example: evaluation of numerical expressions

5.2.4.1 Expression grammar (similar as seen before)

$$
\begin{array}{rcl}
exp & \rightarrow & exp + term \mid exp - term \mid term \\
term & \rightarrow & term * factor \mid factor \\
factor & \rightarrow & (\, exp\,) \mid \textbf{number}
\end{array}
$$

- goal now: **evaluate** a given expression, i.e., the syntax tree of an expression, resp:

5.2.4.2 more concrete goal

Specify, in terms of the grammar, how expressions are evaluated

- grammar: describes the "format" or "shape" of (syntax) trees
- syntax-directedness
- value of (sub-)expressions: *attribute* here

As stated earlier: *values* of syntactic entities are generally *dynamic* attributes and cannot therefore be treated by an AG. However, in this simplistic example of expressions, it's statically doable, since there are no variables and no state-change. We have seen (static) expression evaluation already before, but not in the context of the attribute grammar chapter. For parsing, for instance when talking about that a parser has to calculate an AST, we used at some point for illustration a more simple task for the parser: not to give back a tree for a parsed expression, but just evaluating an expression, giving back an integer. It's basically the same problem. Indeed, also for the oblig, in the provided starting point, one example in the CUP-parser is an expression evaluator.

Having the action part of a grammar for expression calculate the value for expressions can be seen as "attribuation". Of course, also the more standard case, where the action part of the grammar gives back an AST can be understood under the angle of attribute grammars, with the nodes of the AST being the attribute values.

One can even go further: not do an AST first, but directely produce code when parsing, probable intermediate code, not directly machine code. This is known as *syntax-directed*

translation, syntax-directed, because the code is generated directly following the syntax tree. Doing that while parsing, in the action part of the grammar used for parsing restricts what can be achieved. The parser does its thing, perhaps working bottom-up, and the code generation has to follow that bottom-up strategy and that limits what can be done; that also means, it's seldomly done nowadays. It's better to generate an AST first as intermediate representation, then doing type checking, then doing perhaps further intermediate represention s, doing further analyses and optimizations etc. That frees the analyses (type checking, code generation) from following the parse-strategy (like being one-pass, left-to-right and bottom-up in the syntax tree).

We will later in this chapter have a short look at that situation from the attribute grammar angle: what kind of attribute and attribute dependencies can be done if the attribute grammar evaluation is coupled to an LR-style parsing strategy.

Coming back to the expression evaluation example: evaluating expressions is a pure bottom-up thing: the value of a compound expression is determined by the values of sub-expressions. In the terminology of attribute-grammars, that will correspond to *synthesized* attributes (see later).

5.2.5 Expression evaluation: how to do if on one's own?

- simple problem, easy solvable without having heard of AGs
- given an expression, in the form of a syntax tree
- evaluation:
  - simple *bottom-up* calculation of values
  - the value of a compound expression (parent node) **determined by the value of its subnodes**
  - realizable, for example, by a simple recursive procedure

5.2.5.1 Connection to AG's

- AGs: basically a formalism to specify things like that
- *however*: general AGs will allow *more complex* calculations:
  - not just **bottom up** calculations like here but also
  - **top-down**, including both at the same time

When talking about recursive procedures, we mean not just direct recursion. Often a number of mutually recursive procedures is needed, for example, one for factors, one for terms, etc. See the next slide. The use of such recursive arrangement may remind us to the sections about top-down parsing.

As mentioned, AGs make use of more complex "strategies", not just pure bottom-up or pure top-down even mixed ones exists. To evaluate the *simple* expressions here, as pure bottom-up evaluation strategy works well.

5.2.6 Pseudo code for evaluation

```
eval_exp(e) =
  case
  :: e matches PLUSnode ->
       return eval_exp(e.left) + eval_term(e.right)
  :: e matches MINUSnode ->
       return eval_exp(e.left) - eval_term(e.right)
  ...
  end case
```

5.2.7 AG for expression evaluation

	productions/grammar rules			semantic rules
1	exp_1	$\to$	$exp_2 + term$	exp_1 .val = exp_2 .val + $term$ .val
2	exp_1	$\to$	$exp_2 - term$	exp_1 .val = exp_2 .val − $term$ .val
3	exp	$\to$	$term$	exp .val = $term$ .val
4	$term_1$	$\to$	$term_2 * factor$	$term_1$ .val = $term_2$ .val ∗ $factor$ .val
5	$term$	$\to$	$factor$	$term$ .val = $factor$ .val
6	$factor$	$\to$	(exp)	$factor$ .val = exp .val
7	$factor$	$\to$	**number**	$factor$ .val = **number**.val

- *specific* for this example is:
  - only *one* attribute (for all nodes), in general: different ones possible
  - (related to that): only one semantic rule per production
  - as mentioned: rules here define values of attributes “bottom-up” only
- note: subscripts on the symbols for disambiguation (where needed)

5.2.8 Attributed parse tree

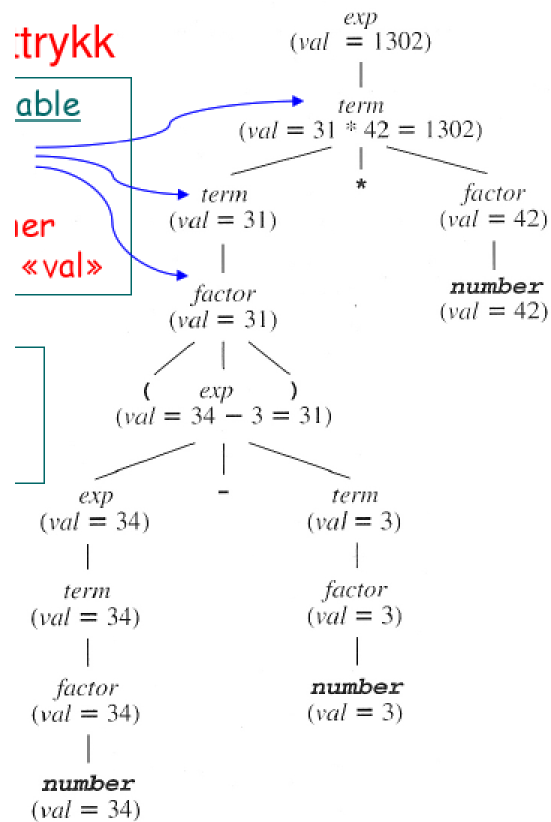

The attribute grammar (being purely synthesized = bottom-up) is very simple and hence, the values in the attribute `val` should be self-explanatory.

5.2.9 Possible dependencies (perhaps move)

In general, an attribute grammar allows multiple attributes (of different types) with multiple dependencies. Though not everything is allowed (or makes sense). Multiple dependencies are possible per grammar production leads to the fact that one grammar production can be accompagnied by multiple semantic rules. Basically, the only thing that is impossible for now is dependencies between attributes across different rules. Since a semantical rule is associated with (the attributes of) **one** grammar production, one simply cannot express such dependencies. Since ultimately the dependencies will be *evaluated* not in the grammar, but in the (parse) trees formed according to the grammar, there will for instance not be dependencies spanning "more than two generations" (like from grandparent to grandchild), also not from "niece to nephew". It's only between parents and children or the other way around or between siblings. It's a simple consequence of the format of semantical rules, attached to the productions of the grammar.

Of course, we are talking here about **direct** dependencies. The value of atribute on a node can be **indirectly** depending on the value of a far-away node. That was seen in the

simple expression evaluation example: indirectly, the value of the root node dependes on the value of all nodes, the direct and the indirect children, with the values propagating up the tree.

When talking about possible resp. impossible (direct) dependencies. It's not the last word, the description what's possible or not is just a simple restriction as a consequence of the way semantic rules are written.

As we will see, not all AG make sense. That is meant as follows: AG are intended to be *evaluated* and that means, evaluated in *trees* (not grammars), like in the simple bottom-up evaluation strategy for expression evaluation. It's easy two write AGs that cannot be evaluated. It's easy when that happens: Evaluation is impossible, if the dependencies contain (in one syntactically correct tree) a *cycle*, or if it contains *contradicting dependencies*, like an attribute depending on *two or more* others, or some attribute is uncovered by a dependency, and this remains undefined.

That's wish for AG is quite easy to state, unfortunately harder to check if a given grammar is not defective in the given sense. From the mentioned criteria, the last two (are all attributes defined, is no attribute defined twice) are actually easy. It's he **acyclicity** that gives headache.

Therefore, one typically restricts interest on subclasses of attribute grammars, where the acyclicity is clear.

5.2.9.1 Possible dependencies (>1 rule per production possible)

- parent attribute on *childen* attributes
- attribute in a node dependent on other attribute of the *same* node
- child attribute on *parent* attribute
- sibling attribute on *sibling* attribute
- *mixture* of all of the above at the same time
- but: **no** immediate dependence **across generations**

The way that the attribute grammars specify dependencies, namely taking a grammar and add semantic rules on top of the given productions puts those restrictions on the (direct) dependencies. It's quite natural. One cannot specify a dependency between attributes making use of *two* or more productions. Of course, there can by an *indirect* dependency.

5.2.10 Attribute dependence graph

- dependencies ultimately between attributes in a syntax *tree* (instances) not between grammar symbols as such

$\Rightarrow$ attribute dependence graph (per syntax tree)

- complex dependencies possible:
  - evaluation complex
  - invalid dependencies possible, if not careful (especially **cyclic**)

5.2.11 Sample dependence graph (for later example)

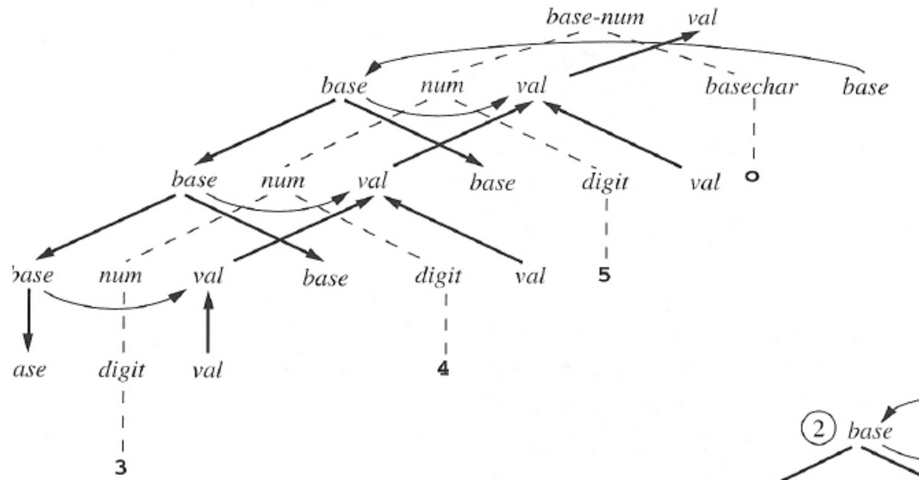

The graph belongs to an example we will revisit later. The dashed lines represent the tree (parse tree or AST, it does not matter). The bold arrows the dependence graph. Later, we will classify the attributes in that `base` (at least for the non-terminals *num*) is inherited ("top-down"), whereas `val` is synthesized ("bottom-up").

We will later have a closer look at what synthesized and inherited means. As we see in the example already here, being synthesized is (in its more general form) not as simplistic as "dependence only from attributes of children". In the example the synthesized attribute `val` depends on its inherited "sister attribute" `base` in most nodes. So, synthesized is not only "strictly bottom-up", it also goes "sideways" (from `base` to `val`). Now, this "sideways" dependence goes from inherited to synthesized only but never the other way around. That's fortunate, because in this way it's immediately clear that there are no *cycles* in the dependence graph. An evaluation (see later) following this form of dependence is **"down-up"**, i.e., first top-down, and afterwards bottom-up (but not then down again etc., the evaluation does not go into cycles).

5.2.11.1 Two-phase evaluation

Perhaps a too fine point concerning evaluation in the example. The above explanation highlighted that the evaluation is "phased" in first a top-down evaluation and afterwards a bottom-up phase. Conceptually, that is correct and gives a good intuition about the design of the dependencies of the attribute. Two "refinements" of that picture may be in order, though. First, as explained later, a dependence graph does not represent **one** possible evaluation (so it makes no real sense in speaking of "the" evaluation of the given graph, if we think of the edges as individual steps). The graph denotes which values need to be present *before* another value can be determined. Secondly, and relatd to that: If we take that view seriously, it's **not** strictly true that *all inherited depenencies are evaluated before all synthesized*. "Conceptually" they are, in a way, but there is an amount of "independence" or "parallelism" possible. Looking at the following picture, which depicts one of many possible evaluation orders, shows that, for example, step 8 is

filling an inherited attribute, and that comes *after* 6 which deals with an synthesized one. But both steps are independent, so they could as well be done the other way around.

So, the picture "first top-down, then bottom-up" is *conceptually correct* and a good intuition, it needs some fine-tuning when talking about when an indivdual step-by-step evaluation is done.

5.2.12 Possible evaluation order

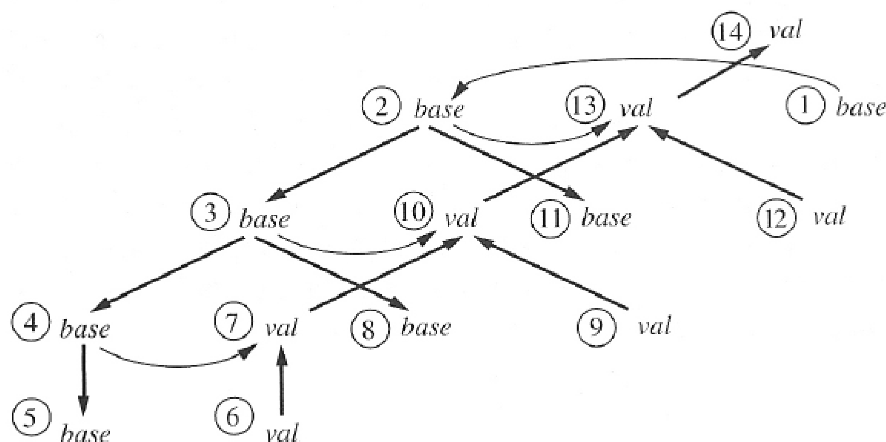

The numbers in the picture give *one possible* evaluation order. As mentioned earlier, there is in general more than one possible way to evaluate dependency graph, in particular, when dealing with a syntax *tree*, and not with the generate case of a " syntax list" (considering lists as a degenerated form of trees). Generally, the rules that say when an AG is properly done assure that all possible evaluations give a unique value for all attributes, and the order of evaluation does not matter. Those conditions assure that each attribute instance gets a value *exactly once* (which also implies there are no cycles in the dependence graph).

5.2.13 Restricting dependencies

- general GAs allow bascially any kind of dependencies[6]
- complex/impossible to meaningfully evaluate (or understand)
- typically: restrictions, disallowing "mixtures" of dependencies
  - fine-grained: per attribute
  - or coarse-grained: for the whole attribute grammar

5.2.13.1 Synthesized attributes

bottom-up dependencies only (same-node dependency allowed).

[6]Apart from immediate cross-generation dependencies.

5.2.13.2 Inherited attributes

top-down dependencies only (same-node and sibling dependencies allowed)

The classification in inherited = top-down and synthesized = bottom-up is a general guiding light. The discussion about the previous figures showed that there might be some fine-points like that "sideways" dependencies are acceptable, not only strictly bottom-up dependencies.

5.2.14 Synthesized and inherited attributes

- terminals and non-terminals carry attributes
- attributes can be typed
- it's *"either-or"* (per symbol)

Later we will say something about attributes on terminals whether they should be considered synthesized or inherited (both views exists). Attributes are typed, and one can have different types in the same grammar (like one attribute is for integer values and the other one for booleans). That's not a big deal, Later, we don't bother mentioning types when defining attribute grammers, since, as said, it is really no big deal. In the examples, we may or may not use attributes of different types. Of one does so, it goes without saying that the semantic rules must honor the types.

All of that is a bit of a side show (and straightforward). *Not* a side show is the distinction between synthesized and inherited attributes.

Attributes of an AG are cleanly split into the synthesized ones and the inherited ones. The split is "per symbol". One may stumble upon presentations that divide 'the attributes globally into those two categories. Actually, it makes no difference. It's more like how one "thinks" of attributes. For example, the illustration earlier made use of "the" attribute *val*, in the nodes or non-terminals *basenum* and *num*. So it's a question of whether one thinks *basenum.val* and *num.val* are the "same" attribute (called *val*) or not. We consider them as different attributes, it's about attributes at a particular symbol. What matters is that would be principally ok, if *num.val* were synthesized and *basenum.val* inherited or the other way around. But per symbol it must be "either-or".

To split the attributes so that an attribute is not both inherited and synthesized (at a symbol) is *necessary*. An attribute both inherited and synthesized would be covered by more than one rule, one specifying its value with information coming from "above", treating the attribute as inherited, and another rule treating it as synthesized with information coming from "below". But each attribute occurrence has to be covered by exactly one rule, not two, which would be contradictory nor zero, which would leave the attribute undefined.

We later will have later a second look at the the latter remark that ever attribute occurrence has to be covered by at least one rule. That will be done in the context for attributes for *terminals*, providing some fine points there.

5.2.15 Semantic rules

Adding attributes to a context-free grammar, split in synthesized and inherited ones, is one thing. The core of an attribute grammar are the rules, specifying how the attributes obtain their values depending of the values of other attributes.

The general for of the rules is shown in equation (5.1), just saying that an attribute obtains it values depending on other variables, and the functional dependence is expressed by f in the equation. Let's call a the target attribute or target variable of the constraint (5.1), and the $\vec{a}$ the source variables or source attributes.

$$a = f(\vec{a}) \tag{5.1}$$

That's of course very non-descript, not even mentioning a context-free grammar. More concretely, the semantic rules are "attached" to the productions of a context-free grammar and the variables a and a_i in equation (5.1) refer to *occurrences* of attributes belonging to the symbols in the production.

Ultimately, the attribute evaluation takes place on parse trees. The productions of the grammar are used to generate the tree (or parse it) and, when using a particular production, the parent node in the tree corresponds to the non-terminal symbol on the left-hand side of the production, and the children correspond to the symbols on the right-hand side. So far, so clear from the chapters about grammars and parsing.

Since the rules are connected to the attributes of individual rules, the specified direct dependencies are between parents and children, also between "siblings", but not further across "generations", like a direct dependency of a grandparent node's attribute from those of grandchildren. Indirect dependencies of course are of course a different story. For instance, in the expression evaluation examples from before, the value of the root node depends on all the nodes below.

So, as explained, the attributes mentioned in a semantic rule need to all refer to attribute occurences in the production to which the semantic rules belongs to. Another requirement, the restricts the rule format in connection with the split of the attributes in synthesized and inherited ones and in connection with the **target** attribute a in equation (5.1).

This requirement is illustrated in Figures 5.2 and 5.3. The picture makes use of a *convention* that draws the inherited attributes positioned to the left of the node, and the synthesized ones to the right (see Figure 5.1).

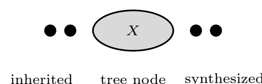

Figure 5.1: Pictorial covention

Note that in the previous example discussing the dependence graph with attributes `base` and `val` was of this format and followed this convention, showing the inherited attribute `base` on the left, the synthesized `val` on the right.

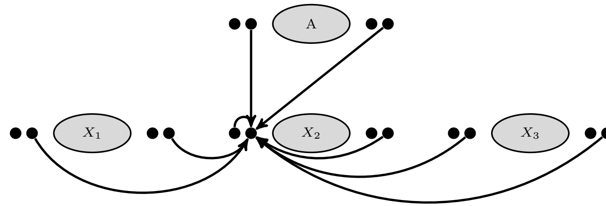

Figure 5.2: Inherited

The inherited resp. synthesized attribute in question is in Figures 5.2 and 5.3 the one the dependency arrows point to, the "target".

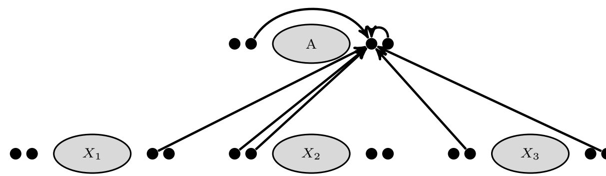

Figure 5.3: Synthesized

Note that in the previous example discussing the dependence graph with attributes `base` and `val` was of this format and followed the convention: show the inherited `base` on the left, the synthesized `val` on the right.

That's the **core** of attribute grammars and of inherited and synthesized attributes: as far as their use as **target** in the semantical rules are concerned, **inherited** attributes belong to symbols on the left-hand side of a production, **synthesized** attributes to the symbols on the right hand side. The same is given more formulaically in equations (5.3) and (5.4), where we assume a production of the following form:

$$A \to X_1 \ldots, X, \ldots X_n \tag{5.2}$$

$$X.\mathtt{i} == f(A.\mathtt{a}, X_1.\mathtt{b}_1, \ldots, X.\mathtt{b}, \ldots X_n.\mathtt{b}_n) \tag{5.3}$$

$$A.\mathtt{s} == f(A.\mathtt{b}, X_1.\mathtt{b}_1, \ldots X_n.\mathtt{b}_k) \tag{5.4}$$

The rule format forbids particular uses inherited and synthesized attributes as *targets* of semantic rules. For the *sources* of semantic rules, on the other hand, no restriction apply, at least not in the general definition of attribute grammars.

Often, one prefers to also impose restructions on the source variables, as well. This is known as *Bochmann's normal form* of attribute grammars. It's actually not a real restriction. An attribute grammar can straightforwardly transformed in that form, if wanted. The general form introduced is quite general, and often one is better off putting more structure on the rules. Also keep in mind that, to make sense, the attribute grammar

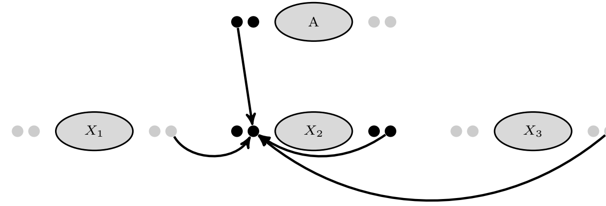

Figure 5.4: Inherited (additional source restrictions)

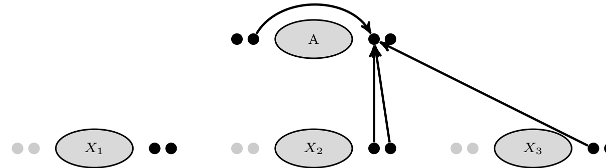

Figure 5.5: Synthesized (additional source restrictions)

must avoid cyclic dependencies in syntax trees. The format so far does not even try to rule out the most obvious cycles. It even allows an attribute occurence in a symbol *directly* depend on itself, in a form of direct recursion. That's not drawn directly in the figures, the pictures would get too crowded with arrows. But, as said, the core definition does not impose any rescrition on the use of source attributes. That's what Bochman's form does additionally. This is shown schemantically in Figures 5.4 and 5.5.

The pictures shows typical ways how one makes use inherited and synthesized attributes in practical situations. For instance, an inherited attribute of a child node depends on an inherited attribute of the parent. Often that's the "same" attribute, like in a situation, where, say X_2 in the picture is A again. That leads to a top-down flow if information. We seen that pattern for instance for the attribute `base` in the non-terminal *num* in the earlier example. Conversely, synthesized attribute depends in particular on the "same" synthesized attribute(s) of children nodes, leading to an bottom-up flow of dependencies.

In both cases, the source attribute is not allowed to depend on an attribute of the same kine "in the same generation", i.e., the same level in the tree. As far as the same generation is concered, inherited attributes can depend only synthesized ones, and vice versa. And accross generation, it's the opposite: inherited one can depend on inherted ones, and synthesized on synthesized ones.

That rules out the most obvious instances of cycles, like direct recursive dependencies. However, it does not exclide indirect cycles, like a synthesized variable depends on an inheriated one in the same node, which in turn depends on said synthesized one. So, the addtional restrictions do not guarantee that the attribute grammar makes sense. That was also not to be expected: we said that any AG can be transformed in an equivalent one on Bochmann form, but that transformation cannot turn a meaningless one (with cycles) to one that makes sense. Still, it's a useful and conventional format

5.2.16 Generally guaranteeing acyclicity

Checking a given grammar for acyclicity is a tricky thing, i.e. computationally complex. That applies also for grammars under the discussed additional source variable restriction. The complexity comes from the fact that acyclicity is required for the dependency graph of *all* trees formed according to the grammar (and there are infinitely many in all reasonable grammars). Checking for acyclicity or a *given* tree, on the other hand, is straightforward. It can be solved by so-called *topological sorting* (and algorithm which is covered by beginners' courses on algorithms and data structures).

5.2.17 Special forms of attribute grammars

An general acyclicity check for attribute grammars is doable, the problem is decidable, but it's mostly no a route taken when working with attribute grammars in a compiler. It's computationally too costly and one is better of to impose further restrictions on the rule format that straightforwardly guarantee acyclicity; no need for some advanced and expensive checking then.

We touch upon 2 such forms. There are more, but those are the most prominent and the simplest ones, especially the first one.

5.2.17.1 S-attributed grammar

Indeed, it cannot get much simpler than the first one. It says: the grammar uses **synthesized attributes only**. That format is known as **S-attributed grammar**. For a grammar in Bochmann form, that immedately and obviously rules out cycles, as all dependency edges go *strictly* upwards, no "horizontal" dependencies. So, that's the simplest, cleanest and prototypical situation for information being synthesized: The information at a parent node depends on the information at the children nodes and on children node information, only. It's so simple that we don't bother to illustrated it with a schematic picture.

5.2.17.2 L-attributed grammar

The second format is a bit more complex; in particular it does not completely rule out inherited attributes. The motivation for that form comes that it can be integrated into parsing. In particular, being used **during** parsing.

The dependencies for an L-attributed grammars are such that they can be evaluated by a *depth-first, left-to-right* traversal of the (parse) tree. The allowed dependencies for inherited attributes are illustrated in Figure 5.6. It should be stressed, that L-attributed grammars *do* allow synthesized attributes. We only show a picture covering the specific restrictions for *inherited* attributes. Note also: For source attributes from the parent node A, **only** inherited attributes are allowed. For the source attributes from X_i, both inherited and synthesized attributes are allowed, but only coming "from the left".

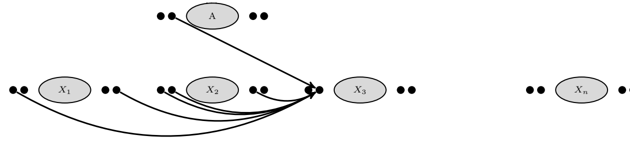

Figure 5.6: L-attributed grammar)

As said, dependencies can be evaluated in a top-down, depth-first (and left-to-right) traversal. This form of traversal can straightforwardly be integrated in a top-own parser, like a recursive descent parser.

Bottom-up parsing builds the parse tree in a bottom-up manner with it's shift- and reduce-steps. That fits well with synthesized attributes, but less well with inherited attributes. Actually, it hardly fits with inherited attibutes with their top-down flow of information. However, the general definition of inherited attributes does allows *horizontal* dependencies where the information flows *sideways*, for instance from sibling to sibling.

However, as shown in Figure 5.6, L-attributed grammars allow clearly at top-down dependence between inherited attributes. That seems to say that those attribute grammars cannot be used for LR-parsers. Actually, it's possible to integrate it into the way a shift-reduce parser works, but it's more complex than for top-down parsers. We don't dig deeper here, and how that can be done is outside the pensum.

5.2.17.3 More formal definitions

We have covered attribute grammars mostly with words, pictures, and some simplified formulas. Later there will also be (more) examples for illustration. Still, for completeness' sake and for reference, let's also nail down the corresponding definitions shortly, without discussing them further.

Definition 5.2.1 (Attibute grammar)**.** An *attribute grammar* is a triple $(G, (Attr_i, Attr_s), R)$, where G is a context-free grammar. The functions $Attr_i$ and $Attr_s$ associate to each grammar symbol X a set $Attr_i(X)$ of *inherited* attributes and $Attr_s(X)$ of *synthesized* attributes, with $Attr_i(X) \cap Attr_s(X) = \emptyset$. The set $Attr = \bigcup Attr(X)$ is the overall set of attributes. The form of the semantic rules R will be defined below.

Definition 5.2.2 (Attibute occurence)**.** A production $X_0 \to X_1 \ldots X_n$ has an *attribute occurrence* $X_i.a$ iff $a \in Attr(X_i)$, for some $0 \leq i \leq n$.

Definition 5.2.3 (Rule format)**.** Given a production p of the form $X_0 \to X_1, \ldots X_n$, then a finite set of *semantic rules* R_p is associated with p, with constraints of the form

$$X_i.a = f(x_1, \ldots, x_k) \tag{5.5}$$

where either

1. $i = 0$ and $a \in Attr_s(X_i)$

2. for $i \geq 1$ and $a \in \mathit{Attr}_i(X_i)$,

for each x_j is an attribute occurrence in p. For R_p, there is exactly one such constraint for each synthesized attribute of X_0, and exactly one such constraint for each inherited attribute for all inherited attributes for all X_i (with $1 \leq i \leq n$).

Definition 5.2.4 (Additional restriction (Bochmann normal form))**.** Assume a semantic rule

$$y_0 = f(y_1, \ldots, y_k) \tag{5.6}$$

in R_r where $y_0 = X_i.a$ for a production p of the form

$$X_0 \to X_1 \ \ldots \ Xn \ .$$

Each attribute occurrence y_j with $1 \leq j \leq k$ is of the form $X_l.b$ where either

1. $l = 0$ and $b \in \mathit{Attr}_i(X_i)$, or
2. $1 \leq l \leq k$ and $b \in \mathit{Attr}_s(X_i)$

Definition 5.2.5 (S-attributed grammar)**.** An attribute grammar is *S-attributed* if all attributes are synthesized.

The last definition is often used explicitly or implicitly assuming Bochman's normal form. Only in that case, the S-attribution restriction guarantees acyclicity, which is the purpose.

Definition 5.2.6 (L-attributed grammar)**.** An attribute grammar for attributes $\mathtt{a}_1, \ldots, \mathtt{a}_k$ is *L-attributed*, if for each *inherited* attribute $\mathtt{a}_j$ and each grammar rule

$$X_0 \to X_1 X_2 \ldots X_n \ ,$$

the associated equations for $\mathtt{a}_j$ are all of the form

$$X_i.\mathtt{a}_j = f_{ij}(X_0.\vec{\mathtt{a}}, X_1.\vec{\mathtt{a}} \ldots X_{i-1}.\vec{\mathtt{a}}) \ .$$

where additionally for $X_0.\vec{\mathtt{a}}$, only *inherited* attributes are allowed.

5.2.17.4 A word on terminals

The definitions and rule formats did not make a distinction between terminals and non-terminals. So, if we allow that terminals carry attributes whose value is specified by an attribute grammar, then that implies that the attributes of terminals are necessarily and uniformely *inherited*, since terminals can only occur on the right-hand side of a production.

That's clear enough, and that's also how the *classical* definitions of attribute grammar deals with the issue. There is nothing wrong with that, all what has been said works fine for terminals and non-terminals alike.

However, some more modern presentations deviate from that, counting the attributes of terminals as *synthesized*. Or at least allowing also synthesized attributes. That of course contradices the classical rule format which would insist: terminals can occur only on the right-hand side of a production, they cannot be anything else than inherited.

Looking, however, at which roles terminals of a grammar play in a parser, we see other aspects that factor in. The terminals in a grammar correspond to **tokens**, which often consists of a token class and a token **value**. Take for instance a familar situation of, say numbers, covered by a terminal, say, **num**. Concretely, the lexer hands over to the parser not just the token class, but also a corresponding value, an integer. That integer can also be seen as the value of an attribute of that leaf node of the syntax tree, concrete or otherwise. This value or attribute may play a role in an analysis based on or inspired by attribute grammars. But that does not change the fact that the value of the attribute is *not* defined and calculated by the attribute grammar.

Conceptually, the value comes from "outside" the attribute grammar (injected from the lexer and behaving as if it were a constant in a given tree.) That fits with treating it as synthesized (at least when following the Bochmann format). As leaf node, a terminal has no children. That means, its attributes cannot depend on anything because there are no nodes below, and that means it has to be a constant value in the sense of not being provided by the attribute grammar evaluation mechanism.

Still other presentations may say, only non-terminals can have synthesized and inherited attributes, and attributes of terminals are somehow of a third kind (filled by the lexer). At any rate, it's a corner case of the framwork, and actually an unproblematic one. I just want to point out, that one may find contradicting information about the issue, depending on where one looks, but it's inessential.

5.2.18 Simplistic example (normally done by the scanner)

5.2.18.1 CFG

$$
\begin{array}{rcl}
\textit{number} & \rightarrow & \textit{number}\,\textit{digit} \mid \textit{digit} \\
\textit{digit} & \rightarrow & \mathbf{0} \mid \mathbf{1} \mid \mathbf{2} \mid \mathbf{3} \mid \mathbf{4} \mid \mathbf{5} \mid \mathbf{6} \mid \mathbf{7} \mid \mathbf{8} \mid \mathbf{9} \mid
\end{array}
$$

5.2.18.2 Attributes (just synthesized)

number	`val`
digit	`val`
terminals	[*none*]

We will look at an AG solution. In practice, this conversion is typically done by the scanner already, and the way it's normally done is relying on provided functions of the implementing programming language (all languages will support such conversion functions, either built-in or in some libraries). For instance in Java, one could use the method `valueOf(String s)`, for instance used as static method `Integer.valueOf("900")` of the class of integers. Obviously, not everything done by an AG can be done already by

the scanner. But this particular example used as warm-up is so simple that it could be done by the scanner, and that where's it's done mostly anyway.

5.2.19 Numbers: Attribute grammar and attributed tree

5.2.19.1 A-grammar

Grammar Rule	Semantic Rules
$number_1 \rightarrow number_2\ digit$	$number_1.val = number_2.val * 10 + digit.val$
$number \rightarrow digit$	$number.val = digit.val$
$digit \rightarrow$ **0**	$digit.val = 0$
$digit \rightarrow$ **1**	$digit.val = 1$
$digit \rightarrow$ **2**	$digit.val = 2$
$digit \rightarrow$ **3**	$digit.val = 3$
$digit \rightarrow$ **4**	$digit.val = 4$
$digit \rightarrow$ **5**	$digit.val = 5$
$digit \rightarrow$ **6**	$digit.val = 6$
$digit \rightarrow$ **7**	$digit.val = 7$
$digit \rightarrow$ **8**	$digit.val = 8$
$digit \rightarrow$ **9**	$digit.val = 9$

5.2.19.2 attributed tree

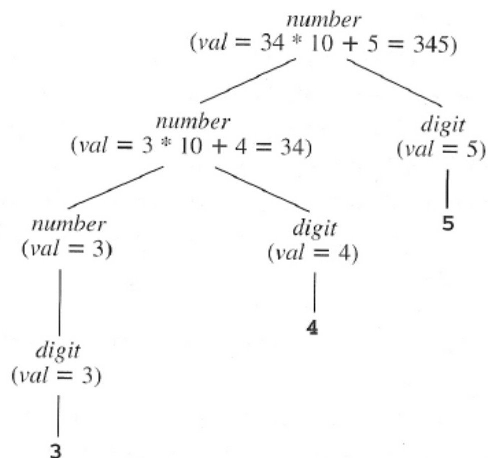

5.2.20 Attribute evaluation: works on trees

i.e.: works equally well for

- *abstract syntax trees*
- *ambiguous* grammars

5.2.20.1 Seriously ambiguous expression grammar

$$exp \rightarrow exp + exp \mid exp - exp \mid exp * exp \mid (\, exp \,) \mid \textbf{number}$$

Alternatively: The grammar is not meant as describing the syntax for the parser, it's meant as grammar describing nice and clean ASTs for an underlying, potentially less nice grammar used for parsing. Remember: grammars describe trees, and one can use EBNF to describe ASTs.

5.2.21 Evaluation: Attribute grammar and attributed tree

5.2.21.1 A-grammar

Grammar Rule	Semantic Rules
$exp_1 \rightarrow exp_2 + exp_3$	$exp_1 .val = exp_2 .val + exp_3 .val$
$exp_1 \rightarrow exp_2 - exp_3$	$exp_1 .val = exp_2 .val - exp_3 .val$
$exp_1 \rightarrow exp_2 * exp_3$	$exp_1 .val = exp_2 .val * exp_3 .val$
$exp_1 \rightarrow (exp_2)$	$exp_1 .val = exp_2 .val$
$exp \rightarrow$ **number**	$exp.val =$ **number** $.val$

5.2.21.2 Attributed tree

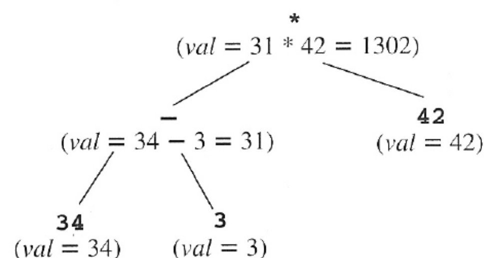

5.2.22 Expressions: generating ASTs

5.2.22.1 Expression grammar with precedences & assoc.

$$
\begin{array}{rcl}
exp & \to & exp + term \mid exp - term \mid term \\
term & \to & term * factor \mid factor \\
factor & \to & (\, exp\,) \mid \mathbf{number}
\end{array}
$$

5.2.22.2 Attributes (just synthesized)

$exp, term, factor$	`tree`
number	`lexval`

5.2.23 Expressions: Attribute grammar and attributed tree

5.2.23.1 A-grammar

Grammar Rule	Semantic Rules
$exp_1 \to exp_2$ **+** $term$	$exp_1.tree =$ $mkOpNode(+, exp_2.tree, term.tree)$
$exp_1 \to exp_2$ **-** $term$	$exp_1.tree =$ $mkOpNode(-, exp_2.tree, term.tree)$
$exp \to term$	$exp.tree = term.tree$
$term_1 \to term_2$ ***** $factor$	$term_1.tree =$ $mkOpNode(*, term_2.tree, factor.tree)$
$term \to factor$	$term.tree = factor.tree$
$factor \to$ **(** exp **)**	$factor.tree = exp.tree$
$factor \to$ **number**	$factor.tree =$ $mkNumNode($**number**$.lexval)$

5.2.23.2 A-tree

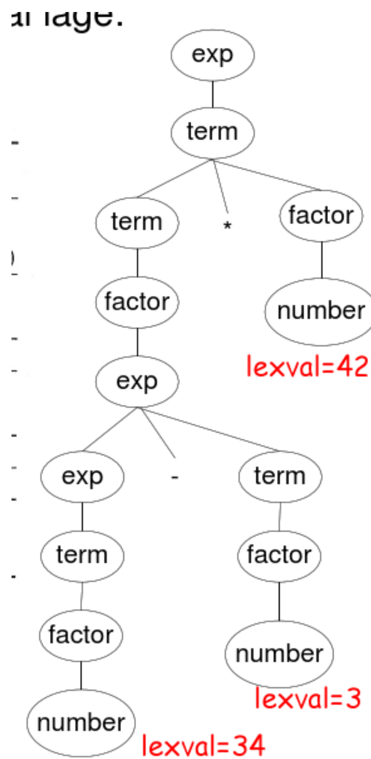

The AST looks a bit bloated. That's because the grammar was massaged in such a way that precedences and associativities during *parsing* are dealt with properly. The the grammar is describing more a parse tree rather than an AST, which often would be less verbose. But the AG formalisms itself does not care about what the grammar describes (a grammar used for parsing or a grammar describing the abstract syntax), it does especially not care if the grammar is ambiguous.

5.2.24 Example: type declarations for variable lists

5.2.24.1 CFG

$$
\begin{array}{rcl}
\mathit{decl} & \rightarrow & \mathit{type}\ \mathit{var\text{-}list} \\
\mathit{type} & \rightarrow & \mathbf{int} \\
\mathit{type} & \rightarrow & \mathbf{float} \\
\mathit{var\text{-}list}_1 & \rightarrow & \mathbf{id,}\ \mathit{var\text{-}list}_2 \\
\mathit{var\text{-}list} & \rightarrow & \mathbf{id}
\end{array}
$$

- Goal: attribute type information to the syntax tree
- *attribute*: `dtype` (with values *integer* and *real*)
- complication: "top-down" information flow: type declared for a list of vars ⇒ **inherited** to the elements of the list

concerning `dtype`: There are thus 2 different attribute values. We don't mean "the attribute `dtype` has integer values", like $0, 1, 2, \ldots$

5.2.25 Types and variable lists: inherited attributes

grammar productions			semantic rules		
decl	$\rightarrow$	*type var-list*	*var-list* .`dtype`	=	*type* .`dtype`
type	$\rightarrow$	**int**	*type* .`dtype`	=	*integer*
type	$\rightarrow$	**float**	*type* .`dtype`	=	*real*
$var\text{-}list_1$	$\rightarrow$	**id**, $var\text{-}list_2$	**id**.`dtype`	=	$var\text{-}list_1$.`dtype`
			$var\text{-}list_2$.`dtype`	=	$var\text{-}list_1$.`dtype`
var-list	$\rightarrow$	**id**	**id**.`dtype`	=	*var-list* .`dtype`

- **inherited**: attribute for **id** and *var-list*
- but also *synthesized* use of attribute `dtype`: for *type* .`dtype`[7]

The dependencies are (especially for the variable lists) in such a way that the attribute of a later element depends on an ealier; in other words, the type information propagates from left to right through the "list". Seen as a tree, that means, the information propagates top-down in the tree. That can be seen in the next (quite small) example: the type information (there **float**) propagates down the right-branch of the tree, which corresponds to the list of two variables x and y.

5.2.26 Types & var lists: after evaluating the semantic rules

float id(x),**id**(y)

5.2.26.1 Attributed parse tree

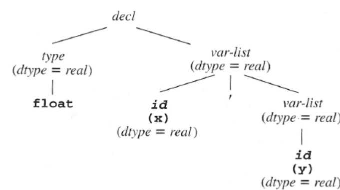

5.2.26.2 Dependence graph

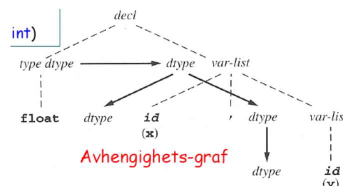

[7] Actually, it's conceptually better not to think of it as "the attribute `dtype`", it's better as "the attribute `dtype` of non-terminal *type*" (written *type* .`dtype`) etc. Note further: *type* .`dtype` is *not* yet what we called *instance* of an attribute.

5.2.27 Example: Based numbers (octal & decimal)

The *based numbers* are a rather well-known example for illustrating synthesized and inherited attributed. Well-known insofar that they are covered in many text-books talking about AGs. The fact that it's inherited and synthesized can easily be seen intuitively: if one wants to evaluate such a number, one would do that left-to-right (which corresponds to top-down), however, the evaluation does not yet know how to calculate until it seems the last piece of information, the specification of what number system to use (decimal or octal). The piece of information has to be calculated resp. carried along "in the opposite direction".

In a way, the notation is designed silly in a way: it's like having a compressed or encrypted file, and then putting the kind of meta-information how to interpret the data not into the header, where it would belong, but at the end...

- remember: grammar for numbers (in decimal notation)
- evaluation: synthesized attributes
- now: *generalization* to numbers with decimal and octal notation

5.2.27.1 Context-free grammar

$$
\begin{array}{rcl}
\mathit{based\text{-}num} & \rightarrow & \mathit{num}\ \mathit{base\text{-}char} \\
\mathit{base\text{-}char} & \rightarrow & \mathbf{o} \\
\mathit{base\text{-}char} & \rightarrow & \mathbf{d} \\
\mathit{num} & \rightarrow & \mathit{num}\ \mathit{digit} \\
\mathit{num} & \rightarrow & \mathit{digit} \\
\mathit{digit} & \rightarrow & \mathbf{0} \\
\mathit{digit} & \rightarrow & \mathbf{1} \\
 & \dots & \\
\mathit{digit} & \rightarrow & \mathbf{7} \\
\mathit{digit} & \rightarrow & \mathbf{8} \\
\mathit{digit} & \rightarrow & \mathbf{9}
\end{array}
$$

5.2.28 Based numbers: attributes

5.2.28.1 Attributes

- *based-num* .`val`: synthesized
- *base-char* .`base`: synthesized
- for *num*:
  - *num* .`val`: synthesized
  - *num* .`base`: **inherited**
- *digit* .`val`: synthesized

- **9** is not an octal character

⇒ attribute `val` may get value "*error*"!

5.2.29 Based numbers: a-grammar

Grammar Rule	Semantic Rules
based-num → *num basechar*	*based-num.val* = *num.val* *num.base* = *basechar.base*
basechar → **o**	*basechar.base* = 8
basechar → **d**	*basechar.base* = 10
$num_1 \rightarrow num_2\ digit$	$num_1.val$ = **if** *digit.val* = *error* **or** $num_2.val$ = *error* **then** *error* **else** $num_2.val * num_1.base + digit.val$ $num_2.base = num_1.base$ $digit.base = num_1.base$
num → *digit*	*num.val* = *digit.val* *digit.base* = *num.base*
digit → **0**	*digit.val* = 0
digit → **1**	*digit.val* = 1
. . .	. . .
digit → **7**	*digit.val* = 7
digit → **8**	*digit.val* = **if** *digit.base* = 8 **then** *error* **else** 8
digit → **9**	*digit.val* = **if** *digit.base* = 8 **then** *error* **else** 9

3/12/2015

The attribute grammar should rather be straightforward and the next slides will shed light on the dependencies and the evaluation. That illustrates the synthesized vs. the inherited parts perhaps more clearly than the equations of the semantic rules. As mentioned in the slides: the evaluation can lead to *errors* insofar that for base-8 numbers, the characters `8` and `9` are not allowed. Technically, to be a proper attribute grammar, a value need to be attached to each attribute instance for each tree. If we would take that serious, it required that we had to give back an "error" value, as can be seen in the code of the semantic rules. If we take that even more seriously, it would mean that the "type" of the `val` attribute is not just integers, but integers *or* an error value.

In a practical implementation, one would probably rather operate with *exceptions*, to achieve the same. Technically, an exception is not a ordinary *value* which is given back, but interrupts the standard control-flow as well. That kind of programming convenience is outside the (purely functional/equational) framework of AGs, and therefore, the given semantic rules deal the extra error value explicitly and evaluation propagate errors explicitly; since the errors occur during the "calculation phase", i.e., when dealing with the synthesized attribute, an error is propagated upwards the tree.

5.2.30 Based numbers: after eval of the semantic rules

5.2.30.1 Attributed syntax tree

5.2.31 Based nums: Dependence graph & possible evaluation order

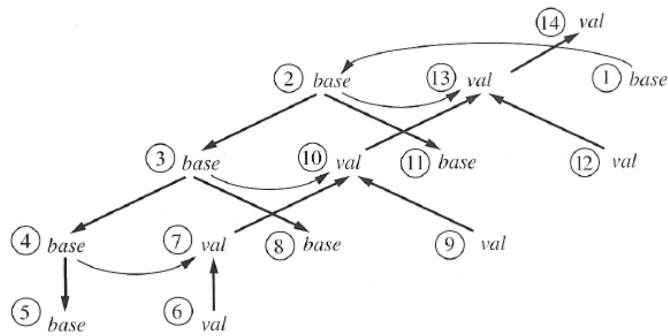

5.2.32 Dependence graph & evaluation

- **evaluation order** must respect the edges in the *dependence graph*
- *cycles* must be avoided!
- directed acyclic graph (DAG)
- dependence graph $\sim$ partial order
- *topological sorting*: turning a partial order to a total/linear order (which is consistent with the PO)
- *roots* in the dependence graph (**not** *the* root of the syntax tree): their values must come "from outside" (or constant)
- often (and sometimes required): terminals in the syntax tree:
  - terminals *synthesized* / *not inherited*
  - ⇒ terminals: *roots* of dependence graph
  - ⇒ get their value from the parser (token value)

A DAG is not a tree, but a generalization thereof. It may have more than one "root" (like a forest). Also: "shared descendents" are allowed. But no cycles.

As for the treatment of terminals, resp. restrictions some books require: An alternative view is that terminals get token values "from outside", the lexer. They are as if they were synthesized, except that it comes "from outside" the grammar.

5.2.33 Evaluation: parse tree method

For acyclic dependence graphs: possible "naive" approach

5.2.33.1 Parse tree method

Linearize the given partial order into a total order (topological sorting), and then simply evaluate the equations following that.

- works only if *all* dependence graphs of the AG are acyclic
- acyclicity of the dependence graphs?
  - decidable for given AG, but computationally expensive[8]

[8] On the other hand: the check needs to be done only once.

  – don't use general AGs but: restrict yourself to subclasses

- disadvantage of parse tree method: also not very efficient check per parse tree

5.2.34 Observation on the example: Is evalution (uniquely) possible?

- all attributes: *either* inherited *or* synthesized[9]
- all attributes: must actually be *defined* (by some rule)
- guaranteed in that for every production:
  – all *synthesized* attributes (on the left) are defined
  – all *inherited* attributes (on the right) are defined
  – local loops forbidden
- since all attributes are either inherited or synthesized: each attribute in any parse tree: defined, and defined only *one* time (i.e., **uniquely defined**)

5.2.35 Loops

- loops intolerable for *evaluation*
- difficult to check (exponential complexity).

Acyclicity checking for a *given* dependence graph: not so hard (e.g., using topological sorting). Here the question is: for *all* syntax trees.

5.2.36 Variable lists (repeated)

5.2.36.1 Attributed parse tree

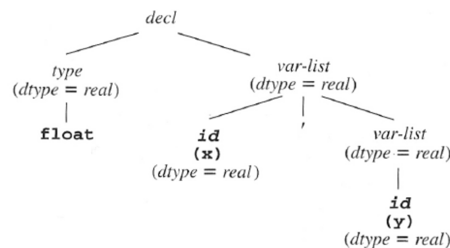

5.2.36.2 Dependence graph

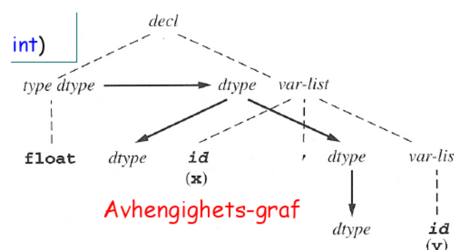

[9] *base-char* .base (synthesized) considered different from *num* .base (inherited)

5.2.37 Typing for variable lists

- code assume: tree given

var-list → **id**

```
procedure EvalType ( T: treenode );
begin
   case nodekind of T of
   decl:
         EvalType ( type child of T );
         Assign dtype of type child of T to var-list child of T;
         EvalType ( var-list child of T );
   type:
         if child of T = int then T.dtype := integer
         else T.dtype := real;
   var-list:
         assign T.dtype to first child of T;
         if third child of T is not nil then
            assign T.dtype to third child;
            EvalType ( third child of T );
   end case;
end EvalType;
```

Dette er også skrevet ut som et program i boka!

The assumption that the tree is *given* is reasonable, if dealing with ASTs. For parse-tree, the attribution of types must deal with the fact that the parse tree is being built during parsing. It also means: it “blurs” typically the border between context-free and context-sensitive analysis.

Index

abstract syntax tree, 3
acyclic graph, 30
approximation, 4
attribute, 5, 6
attribute grammar, 1, 2
attribute grammars, 5

binding, 3, 6

context-sensitive grammar, 5

DAG, 30
data flow analysis, 1
data-flow analysis, 6
dependence graph, 30
directed acyclic graph, 30

evaluation order, 30

grammar
    context-sensitive, 5
graph
    cycle, 30

halting problem, 4

intermediate code, 1
intermediate representation, 1

linear order, 30

overloading, 4

partial order, 30
polymorphism, 4
precomputation, 6

Rice's theorem, 4

semantic rule, 2
semantics analysis, 1
static analysis, 1

topological sorting, 30
total order, 30
trees, 1
Turing completeness, 4
type, 3

variable
    mutable, 6

XML, 2