

A Functional Semantics of Attribute Grammars

Kevin Backhouse

Oxford University Computing Laboratory

Abstract. A definition of the semantics of attribute grammars is given, using the lambda calculus. We show how this semantics allows us to prove results about attribute grammars in a calculational style. In particular, we give a new proof of Chirica and Martin's result [6], that the attribute values can be computed by a structural recursion over the tree. We also derive a new *definedness test*, which encompasses the traditional closure and circularity tests. The test is derived by abstract interpretation.

1 Introduction

This paper introduces a new definition of the semantics of attribute grammars (AGs), which we believe is easier to manipulate mathematically. The semantics is based on the lambda calculus, so a calculational proof style is possible. We illustrate the power of the semantics by proving Chirica and Martin's result [6]: that AGs can be evaluated by structural recursion. More importantly though, we show that the semantics is an excellent basis from which to derive tools and algorithms. We illustrate this by deriving a new *definedness test* for AGs. Our original goal was to derive Knuth's circularity test [12,13], but we were surprised to discover a new test that is actually slightly more powerful than Knuth's.

Notation and Semantic Framework. Our notation is based on the notation of the functional language Haskell [4]. Although the definitions given in Section 2 are valid Haskell, we do not restrict ourselves to Haskell's lazy semantics. Instead we merely assume that the computations are over a semantic domain in which the least fixed points exist. This is so as not to rule out work such as Farrow's [9], in which non-flat domains are used to great effect. For example, Farrow uses the set domain to do data flow analysis. Throughout the paper, we use a number of standard functions such as *fst* and *map*. These are defined in the appendix. Our definition of sequences, which is a little unusual, is also given in the appendix.

2 The Semantics of Attribute Grammars

In this section, we define the semantics of AGs as a least fixed point computation. Our approach is to represent abstract syntax trees as *rose trees* and then define the semantics using functions over rose trees. Throughout, we illustrate our definitions with Bird's *repmim* problem [3]. As noted by Kuiper and Swierstra [14], *repmim* can be easily written as an attribute grammar. The input to *repmim* is a binary tree of integers. Its output is a tree with identical structure, but with every leaf value replaced by the minimum leaf value of the input tree.

2.1 Abstract Syntax Trees as Rose Trees

We use rose trees (see appendix) to represent abstract syntax trees.

```
type AST = Rose ProdLabel
```

The type *ProdLabel* is a sum type representing the different productions of the grammar. In repmin, The abstract syntax tree is a binary tree, so *ProdLabel* is:

```
data ProdLabel = Root | Fork | Leaf Int
```

A single *Root* production appears at the top of the tree. It has one child, which is the binary tree. *Fork* nodes have two children and *Leaf* nodes have zero. Note that the type *ProdLabel* can contain terminal information such as integer constants and strings, but it does not contain non-terminal information. The non-terminals of a production are represented as children in the rose tree.

2.2 Attributes

A drawback of our use of rose trees is that every node in the tree must have the same set of attributes. That is, if leaf nodes have a *code* attribute then root nodes have a *code* attribute too. However, attributes can be left undefined and our definedness test ensures that undefined attributes are not used. Repmin has one inherited attribute for the global minimum leaf value and two synthesised attributes: the minimum leaf value for the current subtree and the output tree for the current subtree. So we define the following datatypes:¹

```
type Inh = Int⊥
type Syn = (Int⊥, AST⊥)
```

In an AG with more attributes, we might use records or lookup tables, rather than tuples to store the attributes. Our definition of AGs below uses *polymorphism* to leave this choice open.

2.3 Attribute Grammars as Functions

In the following definition of an attribute grammar, the inherited and synthesised attributes are represented by the type parameters α and β , respectively. As explained above, this means that we are not bound to a particular implementation of attributes. (*Seq* is defined in the appendix.)

```
type AG  $\alpha$   $\beta$  = ProdLabel  $\rightarrow$  SemRule  $\alpha$   $\beta$ 
type SemRule  $\alpha$   $\beta$  = ( $\alpha$ , Seq  $\beta$ )  $\rightarrow$  ( $\beta$ , Seq  $\alpha$ )
```

¹ The \perp subscript on the types indicates that we have lifted them to a flat domain by adding \perp and \top elements. We use a similar informal notation for lifting functions and values: if x has type A , then x_{\perp} has type A_{\perp} and if f has type $A \rightarrow B$, then f_{\perp} has type $A_{\perp} \rightarrow B_{\perp}$.

This definition states that an AG is a set of semantic rules, indexed by the productions of the grammar. Our concept of a semantic rule deviates slightly from tradition. Traditionally, there is one semantic rule per attribute. In our model, a single semantic rule defines all the attribute values. It is a function from the “input” attributes to the “output” attributes. The input attributes are the inherited attributes of the parent node and the synthesised attributes of the children. The output attributes are the synthesised attributes of the parent and the inherited attributes of the children. The fact that we are using a single function does not mean that the attribute values have to be computed simultaneously. For example, in Haskell [4] attribute values would be evaluated on demand, due to the lazy semantics. In Farrow’s model [9], attribute values are computed iteratively, so the function would be called several times.

RepmIn. Repmin is encoded as a value of type *AG* as follows:

$$\begin{aligned}
\text{repAG} & :: \text{AG Inh Syn} \\
\text{repAG Root } (-, \text{syns}) & = ((\perp, \text{Node}_{\perp} \text{Root } [\text{snd } \text{syns}_1]), [\text{fst } \text{syns}_1]) \\
\text{repAG Fork } (\text{gmin}, \text{syns}) & = ((\text{min}_{\perp} (\text{fst } \text{syns}_1) (\text{fst } \text{syns}_2), \\
& \quad \text{Node}_{\perp} \text{Fork } [\text{snd } \text{syns}_1, \text{snd } \text{syns}_2]), \\
& \quad [\text{gmin}, \text{gmin}]) \\
\text{repAG (Leaf } k) (\text{gmin}, \text{syns}) & = ((k_{\perp}, \text{Node}_{\perp} (\text{Leaf}_{\perp} \text{gmin}) []), [])
\end{aligned}$$

This should be read as follows. The local minimum of the root production is undefined, because it will remain unused. The output tree of the root production is equal to the output tree of its subtree. In the root production, the local minimum of the subtree becomes the inherited global minimum attribute. (The *gmin* attribute of the root node is ignored.) In the fork production, the local minimum is the minimum of the two local minima of the subtrees. The output tree is built from the output trees of the two subtrees and the global minimum is passed unchanged to the subtrees. The local minimum in the leaf production equals the value at the leaf and the new tree is a leaf node with the global minimum as its value.

Semantic Trees. Given an AST and an AG, we can compute a *semantic tree*, which is a tree containing semantic functions. This is done by simply mapping the AG over the tree:

$$\begin{aligned}
\text{type SemTree } \alpha \beta & = \text{Rose } (\text{SemRule } \alpha \beta) \\
\text{mkSemTree} & :: \text{AG } \alpha \beta \rightarrow \text{AST} \rightarrow \text{SemTree } \alpha \beta \\
\text{mkSemTree} & = \text{mapRose}
\end{aligned}$$

We will often work with semantic trees rather than abstract syntax trees, because it makes some of our definitions and proofs more concise.

2.4 Shifting Attributes Up and Down the Tree

The goal of an attribute grammar evaluator is to produce a tree containing the attribute values. In other words we wish to produce a value of type $Rose(\alpha, \beta)$, where α and β are the types of the inherited and synthesised attributes. However, this format is not easily compatible with the semantic rules defined earlier. Instead we use the following two formats to represent annotated trees:

type $InputTree\ \alpha\ \beta = Rose(\alpha, Seq\ \beta)$
type $OutputTree\ \alpha\ \beta = Rose(\beta, Seq\ \alpha)$

In the *input* format, the synthesised attributes of a node are stored by the node's parent. The parent uses a sequence to store the synthesised attributes of all its children. Similarly, in the *output* format, the inherited attributes of a node are stored by the node's parent. The output tree is the result of applying the semantic tree to the input tree. Note that in the input format, the synthesised attributes of the root node are not stored. Similarly, in the output format the inherited attributes are not stored. Therefore, this data needs to be stored separately.

The *shift* function is used to convert from the output to the input format. It encapsulates the idea that inherited attributes are passed down the tree and synthesised attributes are passed up. Note that *shift* merely moves attributes around; the calculation of the attribute values is done by the semantic tree.

$shift :: \alpha \rightarrow OutputTree\ \alpha\ \beta \rightarrow InputTree\ \alpha\ \beta$
 $shift\ a\ (Node\ (b, as)\ ts) = Node\ (a, bs)\ (zipWith\ shift\ as\ ts)$
where $bs = map\ (fst \circ root)\ ts$

Note that the inherited attributes of the root node are the first parameter of *shift*. The inherited attributes of the root node are always an external parameter to an attribute grammar.

2.5 The Semantics as a Least Fixed Point Computation

We explained above how we can produce a *semantic tree* by mapping an AG over an AST. Given a semantic tree *semtree* and the inherited attributes *a* of the root node, we expect the following equations to hold:

$$inputTree = shift\ a\ outputTree \quad (1)$$

$$outputTree = appRose\ semtree\ inputTree \quad (2)$$

As we discussed above, the attributes of the tree can be stored either in the input or the output format. Equation (1) states that the input tree can be computed from the output tree with the *shift* function. Equation (2) says that the output tree can be computed by applying the semantic tree to the input tree. Equations (1) and (2) form a pair of simultaneous equations that can be used to compute *inputTree* and *outputTree*. Following Chirica and Martin [6], we define the semantics of attribute grammars to be the *least* solution of the equations. This choice has no effect if the attributes are not circularly defined, because the

solution to the equations is unique. If our domains are flat, then circularly defined attributes will evaluate to \perp . Farrow [9] discusses some interesting applications of working with non-flat domains. The least solution can be computed with the following function:

$$\begin{aligned} eval &:: SemTree \alpha \beta \rightarrow \alpha \rightarrow OutputTree \alpha \beta \\ eval \text{ semtree } a &= \mu (appRose \text{ semtree } \circ \text{ shift } a) \end{aligned}$$

This function is our definition of the semantics of attribute grammars. Please note that we do not propose it as an efficient implementation though! In the following section, we define a more convenient version called *io* and demonstrate it on the repmin problem (Section 3.4).

3 AG Semantics as a Structural Recursion

The *eval* function given above computes attribute values for the entire tree. In practice, we often only want to know the values of the synthesised attributes of the root node of the tree. Therefore, it is convenient to define the function *io*:

$$\begin{aligned} io &:: SemTree \alpha \beta \rightarrow \alpha \rightarrow \beta \\ io \text{ semtree } a &= fst (\text{root } (eval \text{ semtree } a)) \end{aligned}$$

The name *io* stands for input-output. It converts the semantic tree into a function from inherited attributes to synthesised attributes.

Johnsson [11] showed that AGs can be viewed as an idiom of lazy functional programming. Using the idea that a semantic tree is simply a function from inherited to synthesised attributes, he recursively builds the input-output function for the entire AST. The idea that input-output functions can be computed by structural recursion was also stated much earlier by Chirica and Martin [6]. We state the idea formally with the following theorem:

Theorem 1. *The following definition of *io* is equivalent to the one given above:*

$$io = foldRose \text{ knit}$$

where:

$$\begin{aligned} knit &:: (SemRule \alpha \beta, Seq (\alpha \rightarrow \beta)) \rightarrow (\alpha \rightarrow \beta) \\ knit (\text{semrule}, fs) a &= fst (\mu \langle b, as \bullet \text{semrule } (a, appSeq fs as) \rangle) \end{aligned}$$

The definition of *knit* is rather unintuitive, but it has appeared in different forms in many papers on AGs, including one of our own [8]. Other examples are Gondow & Katayama [10] and Rosendahl [18]. Very briefly, *as* and *appSeq fs as* are respectively the inherited and synthesised attributes of the subtrees. Their values are mutually dependent (using *semrule*), so they are computed by a least fixed point computation. For a better understanding of *knit*, see Johnsson [11].

Our proof is essentially the same as the proof given by Chirica and Martin [6], although our notation is rather different. The most important step is to use the mutual recursion rule and we will start by proving a lemma, that adapts the mutual recursion rule to rose trees. Another important step is to use the abstraction rule though. The use of the abstraction rule is the motivation for our use of sequences rather than lists throughout this document.

3.1 The Mutual Recursion Rule on Rose Trees

The most crucial step in our proof of Theorem 1 is to use the *mutual recursion rule* (see the appendix). This theorem allows us to factorise least fixed point computations. In this section, we shall show how it can be applied to rose trees. We do this by proving the following lemma:

Lemma 1. *If f is a monotonic function of type $Rose\ A \rightarrow Rose\ A$, then:*

$$root(\mu f) = \mu \langle x \bullet g_1(x, \mu \langle xs \bullet g_2(x, xs) \rangle) \rangle$$

where:

$$\begin{aligned} g_1 &= fst \circ g \\ g_2 &= snd \circ g \\ g &= split \circ f \circ merge \end{aligned}$$

Proof:

$$\begin{aligned} & root(\mu f) \\ = & \quad \{ merge \circ split = id \} \\ & root(\mu (merge \circ split \circ f)) \\ = & \quad \{ rolling\ rule\ (see\ the\ appendix) \} \\ & root(merge(\mu(split \circ f \circ merge))) \\ = & \quad \{ root \circ merge = fst, Definition\ g \} \\ & fst(\mu g) \\ = & \quad \{ mutual\ recursion\ rule\ (see\ the\ appendix) \} \\ & \mu \langle x \bullet g_1(x, \mu \langle xs \bullet g_2(x, xs) \rangle) \rangle \end{aligned}$$

3.2 Proving Theorem 1

We now prove Theorem 1, deriving the definition of *knit* in the process. First we apply a general result about folds:

$$\begin{aligned} io &= foldRose\ knit \\ \equiv & \quad \{ Universal\ property\ of\ fold, Bird\ \&\ de\ Moor\ [5, page\ 46] \} \\ & \forall \langle s, ss, a \bullet io(Node\ s\ ss)\ a = knit(s, map\ io\ ss)\ a \rangle \end{aligned}$$

We continue by manipulating the expression $io(Node\ s\ ss)\ a$ until we obtain an expression from which it is easy to derive the definition of *knit*:

$$\begin{aligned} & io(Node\ s\ ss)\ a \\ = & \quad \{ Definition\ io \} \\ & fst(root(\mu(appRose(Node\ s\ ss) \circ shift\ a))) \end{aligned}$$

We can apply the mutual recursion rule on rose trees (Lemma 1) to this expression. By simply expanding the definitions of g_1 and g_2 , as specified in Lemma 1, we find that:

$$g_1((b, as), ts) = s(a, \text{map } (fst \circ \text{root}) ts)$$

$$g_2((b, as), ts) = \text{zipWith } \text{appRose } ss (\text{zipWith } \text{shift } as ts)$$

Therefore, Lemma 1 implies that $io (\text{Node } s ss) a$ is equal to:

$$fst (\mu \langle b, as \bullet g_1((b, as), \mu \langle ts \bullet g_2((b, as), ts) \rangle \rangle))$$

(With g_1 and g_2 as defined above.) Consider the sub-expression containing g_2 :

$$\begin{aligned} & \mu \langle ts \bullet g_2((b, as), ts) \rangle \\ = & \quad \{ \text{Definition } g_2 \} \\ & \mu \langle ts \bullet \text{zipWith } \text{appRose } ss (\text{zipWith } \text{shift } as ts) \rangle \\ = & \quad \{ \text{Claim (see the appendix)} \} \\ & \mu \langle ts \bullet \text{appSeq } (\text{zipWith } (\circ) (\text{map } \text{appRose } ss) (\text{map } \text{shift } as)) ts \rangle \\ = & \quad \{ \text{abstraction rule (see the appendix)} \} \\ & \langle k \bullet \mu \langle \text{zipWith } (\circ) (\text{map } \text{appRose } ss) (\text{map } \text{shift } as) k \rangle \rangle \\ = & \quad \{ \text{Definitions } \text{zipWith} \text{ and } \text{map} \} \\ & \langle k \bullet \mu \langle \text{appRose } (ss k) \circ \text{shift } (as k) \rangle \rangle \\ = & \quad \{ \text{Definition } \text{eval} \} \\ & \langle k \bullet \text{eval } (ss k) (as k) \rangle \end{aligned}$$

We continue with the main expression:

$$\begin{aligned} & fst (\mu \langle b, as \bullet g_1((b, as), \mu \langle ts \bullet g_2((b, as), ts) \rangle \rangle)) \\ = & \quad \{ \text{Substitute the new sub-expression} \} \\ & fst (\mu \langle b, as \bullet g_1((b, as), \langle k \bullet \text{eval } (ss k) (as k) \rangle \rangle)) \\ = & \quad \{ \text{Definition } g_1 \} \\ & fst (\mu \langle b, as \bullet s(a, \text{map } (fst \circ \text{root}) \langle k \bullet \text{eval } (ss k) (as k) \rangle)) \rangle \\ = & \quad \{ \text{Definitions } \text{map} \text{ and } \text{io} \} \\ & fst (\mu \langle b, as \bullet s(a, \langle k \bullet \text{io } (ss k) (as k) \rangle)) \rangle \\ = & \quad \{ \text{Definitions } \text{appSeq} \text{ and } \text{map} \} \\ & fst (\mu \langle b, as \bullet s(a, \text{appSeq } (\text{map } \text{io } ss) as)) \rangle \\ = & \quad \{ \text{Definition } \text{knot} \} \\ & \text{knot } (s, \text{map } \text{io } ss) a \end{aligned}$$

In the final step, we have derived the definition of *knot*.

3.3 A More Convenient Version of the *io* Function: *trans*

The *io* function operates on semantic trees. It is usually more convenient to operate on abstract syntax trees, so we define the *trans* function:

$$\begin{aligned} trans &:: AG \alpha \beta \rightarrow AST \rightarrow \alpha \rightarrow \beta \\ trans \ ag &= io \circ mkSemTree \ ag \end{aligned}$$

In the next section, it is important that *trans* can be written as a fold:

$$trans \ ag = foldRose (kmit \circ (ag \times id))$$

The proof is a simple application of fold-map fusion [4, page 131].

3.4 Replemin as a Recursive Function

Let us expand *trans repAG* to demonstrate that it is equivalent to the functional program for *replemin*. Expanding the definitions of *trans* (above), *foldRose* (appendix), *kmit* (page 146) and *repAG* (page 144), we obtain:

$$\begin{aligned} replemin &:: AST \rightarrow Inh \rightarrow Syn \\ replemin (Node \ Root \ [t]) _ &= (\perp, Node_{\perp} \ Root \ [nt]) \\ &\quad \mathbf{where} \ (gm, nt) = replemin \ t \ gm \\ replemin (Node \ Fork \ [t_1, t_2]) \ gm &= (min_{\perp} \ lm_1 \ lm_2, Node_{\perp} \ Fork \ [nt_1, nt_2]) \\ &\quad \mathbf{where} \ (lm_1, nt_1) = replemin \ t_1 \ gm \\ &\quad \quad (lm_2, nt_2) = replemin \ t_2 \ gm \\ replemin (Node \ (Leaf \ k) _) \ gm &= (k_{\perp}, Node_{\perp} \ (Leaf_{\perp} \ gm) \ []) \end{aligned}$$

This is indeed very similar to the program given by Bird [3]. Of course Bird's emphasis is on the recursive definition of *gm* at the root of the tree, which is valid in a language with lazy evaluation.

4 The Definedness Test

When Knuth [12] introduced attribute grammars, he also gave algorithms for testing them for closure and circularity. Closure is the property that there are no 'missing' semantic rules for attributes. An AG is circular if there exists an input tree on which the attributes are circularly defined. We present a new algorithm, which is very similar to the circularity test, but encompasses the roles of both the closure and the circularity tests. It computes equivalent results to the circularity test and improves upon the closure test (which is based on a very shallow analysis of the dependency structure) by reducing the number of false negatives. In other work on extending attribute grammars [19] we have found this improvement essential.

In our model, an undefined attribute is an attribute with the value \perp . Similarly, circularly defined attributes over a flat domain will evaluate to \perp , because

our semantics is a least fixed point semantics.² So in our model, the goal of both the closure and the circularity tests is to detect attributes which will evaluate to \perp . Our *definedness test* accomplishes both of these goals by analysing *strictness*. This concept is very similar to the concept of “dependence” used in the circularity test. A function f is *strict* if $f \perp = \perp$. In our model, an undefined attribute is an attribute with the value \perp , so a strict function ‘depends’ on its parameter. However, the expression $f x$ is said to depend on x even if f is not strict, so the two concepts are not identical. Luckily, semantic rules in a traditional AG are always strict in the attributes that they depend on.

Our presentation of the circularity test consists of two parts. In the first part we explain how the strictness analysis is done as an abstract interpretation. In the second part we explain how the AG can be statically tested for circularity on *all* possible input trees.

4.1 Strictness Analysis and Abstract Interpretation

Strictness analysis is very important for the optimisation of lazy functional programs. Peyton Jones [16, pages 380-395] explains this and gives a good introduction to *abstract interpretation*, which is the technique used to perform the analysis. The use of strictness analysis in AGs has been studied before by Rosendahl [18]. However, Rosendahl’s emphasis is on the optimisation of AGs, rather than definedness.

The prerequisite for abstract interpretation (first introduced by Cousot & Cousot [7]) is a *Galois Connection* between the “concrete” domain (A, \leq) and the “abstract” domain $(A^\#, \preceq)$. For an introduction to Galois Connections, see Aarts [1]. Briefly, there should be an *abstraction* function abs of type $A \rightarrow A^\#$ and a *concretisation* function con of type $A^\# \rightarrow A$. The “connection” is:

$$\forall \langle x, y \bullet abs\ x \preceq y \equiv x \leq con\ y \rangle$$

The Galois Connection used in strictness analysis is between any domain (A, \leq) and the boolean domain $(Bool, \Rightarrow)$:

$$\begin{array}{ll} abs_A \quad :: A \rightarrow Bool & con_A \quad :: Bool \rightarrow A \\ abs_A\ x = (x \neq \perp) & con_A\ b = \mathbf{if\ } b \mathbf{\ then\ } \top \mathbf{\ else\ } \perp \end{array}$$

The idea is that in the abstract domain, \perp is represented by *False* and other values by *True*. Given a function f , we derive an abstract version $f^\#$ that models the strictness of f using booleans. Suppose for example that f has type $A \rightarrow B$. Then $f^\#$ will have type $Bool \rightarrow Bool$. The result of $f^\#$ is *False* if and only if the result of f is \perp . Mathematically, this relationship between f and $f^\#$ is:³

$$\forall \langle x \bullet abs_B (f\ x) = f^\# (abs_A\ x) \rangle \quad (3)$$

² If the domain is non-flat, then circular attributes may not evaluate to \perp , but in this situation circularity is usually intended rather than erroneous. Farrow [9] suggests some useful applications of circular attributes over non-flat domains.

³ Abstract interpretations are often only able to give *conservative* results, so the equation is an inequality: $abs (f\ x) \leq f^\# (abs\ x)$. However, in most AGs such as *repmim*, an exact result can be given.

Often, if this equation has a solution, it is the function $f^\# = \text{abs}_B \circ f \circ \text{con}_A$. We will use this technique when we discuss `repmIn` below.

RepmIn. Let us apply strictness analysis to `repmIn`. First we must define abstract versions of the types `Inh` and `Syn`, defined on page 143:

```

type Inh# = Bool
type Syn# = (Bool, Bool)

```

`Inh` and `Inh`[#] are Galois Connected, as are `Syn` and `Syn`[#]:

$$\begin{aligned} \text{absInh} &= \text{abs}_{\text{Int}} & \text{absSyn} &= \text{abs}_{\text{Int}} \times \text{abs}_{\text{AST}} \\ \text{conInh} &= \text{con}_{\text{Int}} & \text{conSyn} &= \text{con}_{\text{Int}} \times \text{con}_{\text{AST}} \end{aligned}$$

Now we need to derive an abstract version of `repAG`, the attribute grammar for `repmIn`, defined on page 144. The type of `repAG` is `AG Inh Syn`, so `repAG`[#] should have type `AG Inh# Syn#`. Following the suggestion above, we define:

$$\text{repAG}^\# p = (\text{absSyn} \times \text{map absInh}) \circ \text{repAG } p \circ (\text{conInh} \times \text{map conSyn})$$

A simple calculation reveals that `repAG` and `repAG`[#] satisfy an equation similar to (3), which is required by Theorem 2. An equivalent definition of `repAG`[#] is:

$$\begin{aligned} \text{repAG}^\# &:: \text{AG Inh}^\# \text{Syn}^\# \\ \text{repAG}^\# \text{Root } (-, \text{syms}) &= ((\text{False}, \text{snd syms}_1), [\text{fst syms}_1]) \\ \text{repAG}^\# \text{Fork } (gmin, \text{syms}) &= ((\text{fst syms}_1 \wedge \text{fst syms}_2, \\ &\quad \text{snd syms}_1 \wedge \text{snd syms}_2), [gmin, gmin]) \\ \text{repAG}^\# \text{Leaf } k (gmin, -) &= ((\text{True}, gmin), []) \end{aligned}$$

The reader should compare this definition with the definition of `repAG`, given on page 144. Where `repAG` computes the value \perp , this function computes `False`.

Using the Abstract Interpretation. Now that we have `ag`[#], the abstract interpretation of `ag`, what happens when we evaluate an abstract syntax tree using it? Our hope is that the result of evaluating the tree with `ag`[#] can be used to predict something about the result of evaluating the tree with `ag`. The following theorem is what we need:

Theorem 2. *Suppose we are given an attribute grammar `ag` and its abstract interpretation `ag`[#]. By definition, they are related as follows, for all `p`:*

$$(\text{absSyn} \times \text{map absInh}) \circ \text{ag } p = \text{ag}^\# p \circ (\text{absInh} \times \text{map absSyn})$$

Then, for all `t`:

$$\text{absSyn} \circ \text{trans ag } t = \text{trans ag}^\# t \circ \text{absInh}$$

This theorem could be worded as: if `ag`[#] is the abstract interpretation of `ag`, then `trans ag`[#] is the abstract interpretation of `trans ag`. The theorem can be proved manually using fixpoint fusion or fixpoint induction, but we have discovered that it follows immediately from the polymorphic type of `trans`. This is the topic of a forthcoming paper [2], which presents a result based on Reynold’s parametricity theorem [17], similar to Wadler’s “Theorems for free” [20].

4.2 Computing the Strictness for All Possible Trees

In general an attribute grammar will have n inherited attributes and m synthesised attributes. Therefore, the types $Inh^\#$ and $Syn^\#$ will be:

```

type  $Inh^\#$  = ( $Bool$ , ...,  $Bool$ )
type  $Syn^\#$  = ( $Bool$ , ...,  $Bool$ )

```

This means that there are 2^n different values of type $Inh^\#$ and 2^m of type $Syn^\#$. Consequently, there are only a finite number of functions of type $Inh^\# \rightarrow Syn^\#$. So even though there are an infinite number of input trees, $trans\ ag^\#$ can only have a finite number of values. In this section we derive an algorithm, very similar to the traditional circularity test algorithm, for computing those values in a finite amount of time.

The Grammar. Until this point we have ignored the fact that abstract syntax trees are constrained by a *grammar*. We are using rose trees to represent ASTs, which means that we are not prevented from building grammatically incorrect trees. The definedness test depends crucially on the grammar, so let us define it.

A grammar consists of a set of *symbols* and an associated set of *productions*. Just like productions, we represent the symbols with a simple enumeration datatype. In repmin, there are just two symbols:

```

data  $Symbol$  =  $Start$  |  $Tree$ 

```

Every symbol in the grammar is associated with zero or more productions. Every production has a sequence of symbols on its right hand side. We represent this as a *partial function* with the following type:

```

 $grammar :: (ProdLabel, Seq\ Symbol) \leftrightarrow Symbol$ 

```

This function is partial, because it only accepts inputs where the sequence of symbols matches the production. For example, the grammar for repmin is:

```

 $grammar\ (Root, [Tree]) = Start$ 
 $grammar\ (Fork, [Tree, Tree]) = Tree$ 
 $grammar\ (Leaf\ \_, []) = Tree$ 

```

The following function checks whether a tree is grammatically correct:

```

 $foldRose\ grammar :: AST \leftrightarrow Symbol$ 

```

This function is also partial. It is only defined on grammatically correct trees.

Generating Grammatically Correct Trees. We can generate the set of all grammatically correct trees by *inverting* the grammar:

```

 $trees :: Symbol \leftrightarrow AST$ 
 $trees = (foldRose\ grammar)^\circ$ 

```

As indicated by the type signature, *trees* is a relation. Relations can be thought of as set-valued functions, so given a symbol, *trees* produces the set of all grammatically correct trees with that symbol-type.

The above definition is known as an *unfold*. See Bird and de Moor [5] for a proper discussion of folds and unfolds over relations.

Testing all Grammatically Correct Trees. We now have the two necessary ingredients for the definedness test. We can test an individual tree for definedness and we can generate the set of all grammatically correct trees. We can put them together as follows:

$$\begin{aligned} circ &:: Symbol \leftrightarrow (Inh^\# \rightarrow Syn^\#) \\ circ &= trans\ ag^\# \circ trees \end{aligned}$$

Given a symbol, *circ* produces the test results for all trees of that type. As we explained earlier, these results form a set of finite size, despite the fact that the number of trees involved might be infinite. So the question is: how do we compute the results in a finite amount of time? The answer is provided by the hylomorphism theorem (see the appendix). Recalling that *trans* can be written as a fold and *trees* as an unfold, we see that an equivalent definition for *circ* is:

$$circ = \mu \langle c \bullet knit \circ (ag^\# \times map\ c) \circ grammar^\circ \rangle \quad (4)$$

The result of this fixpoint computation can be computed by *iteration*. Iteration is a technique for computing the value of μf in finite sized domains. The technique is to compute the sequence $\perp, f(\perp), f(f(\perp)), \dots$, which is guaranteed to stabilise at μf within a finite number of steps. In the context of computing *circ*, *c* is a relation over a finite domain, so we represent it as a set of pairs. We start with the empty set and iterate until *c* ceases to change.

Testing repmin. Let us evaluate *circ* for repmin. We start with *c* equal to the empty relation. After one iteration, *c'* equals:

$$\begin{aligned} c' &= knit \circ (repAG^\# \times map\ c) \circ grammar^\circ \\ &= \langle Tree \mapsto \langle a \bullet (True, a) \rangle \rangle \end{aligned}$$

After the second iteration:

$$c'' = \langle Start \mapsto \langle a \bullet (False, True) \rangle, Tree \mapsto \langle a \bullet (True, a) \rangle \rangle$$

On the third iteration, we find that $c''' = c''$, so the iteration has stabilised and $circ = c''$. What does this value of *circ* tell us about repmin? On the root node, the *lmin* attribute will never be defined, regardless of whether the inherited *gmin* attribute is defined, but the *ntree* attribute is always defined. On internal nodes of the tree, *lmin* is always defined and *ntree* is defined if *gmin* is defined.

5 Conclusion

We have defined a semantics for attribute grammars based on the lambda calculus and given a new proof of Chirica and Martin's important result [6]. We also derived a definedness test, which is very similar to the traditional circularity test. However, our test also encompasses the closure test. The use of the hylomorphism theorem is very important in our derivation of the definedness test and it results in an algorithm based on iteration. We think this may also help to explain why the algorithm for the circularity test is based on iteration.

Our decision to use rose trees comes with one major drawback: loss of type safety. If we had used algebraic datatypes to define the abstract syntax, then they would automatically be grammatically correct. We would not have needed to discuss “The Grammar” in Section 4.2. It is also common for different grammar symbols to have different attributes. For example, in *repmim* the *Start* symbol does not need a *locmin* or a *gmin* attribute. In our model, it is not possible to distinguish between the symbols in this way. It would be interesting to explore the use of category theory to extend our results to an arbitrary recursive datatype.

Acknowledgements. I am supported by a grant from Microsoft Research. I would like to thank Eric Van Wyk, Ganesh Sittampalam and Oege de Moor for all their help with this paper. I would also like to thank my father for his help with the paper and for teaching me about fixpoints.

References

1. C. J. Aarts. Galois connections presented computationally. Graduating Dissertation, Department of Computing Science, Eindhoven University of Technology. Available from: <http://www.cs.nott.ac.uk/~char'176rcb/MPC/galois.ps.gz>, 1992.
2. R. C. Backhouse and K. S. Backhouse. Abstract interpretations for free. Available from: <http://www.cs.nott.ac.uk/~char'176rcb/papers/papers.html>.
3. R. S. Bird. Using circular programs to eliminate multiple traversals of data. *Acta Informatica*, 21:239–250, 1984.
4. R. S. Bird. *Introduction to Functional Programming using Haskell*. Prentice Hall, 2 edition, 1998.
5. R. S. Bird and O. de Moor. *Algebra of Programming*, volume 100 of *International Series in Computer Science*. Prentice Hall, 1997.
6. L. M. Chirica and D. F. Martin. An order-algebraic definition of Knuthian semantics. *Mathematical Systems Theory*, 13(1):1–27, 1979.
7. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.
8. O. de Moor, K. Backhouse, and S. D. Swierstra. First-class attribute grammars. *Informatica*, 24(3):329–341, 2000.

9. R. Farrow. Automatic generation of fixed-point-finding evaluators for circular, but well-defined, attribute grammars. In *SIGPlan '86 Symposium on Compiler Construction*, pages 85–98, Palo Alto, CA, June 1986. Association for Computing Machinery, SIGPlan.
10. K. Gondow and T. Katayama. Attribute grammars as record calculus — a structure-oriented denotational semantics of attribute grammars by using cardelli's record calculus. *Informatica*, 24(3):287–299, 2000.
11. T. Johnsson. Attribute grammars as a functional programming paradigm. In Gilles Kahn, editor, *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*, volume 274 of *LNCS*, pages 154–173, Portland, OR, September 1987. Springer-Verlag.
12. D. E. Knuth. Semantics of context-free languages. *Mathematical Systems Theory*, 2:127–146, 1968.
13. D. E. Knuth. Semantics of context-free languages: Correction. *Mathematical Systems Theory*, 5:95–96, 1971.
14. M. Kuiper and S. D. Swierstra. Using attribute grammars to derive efficient functional programs. In *Computing Science in the Netherlands CSN '87*, 1987.
15. Mathematics of Program Construction Group, Eindhoven University of Technology. Fixed-point calculus. *Information Processing Letters Special Issue on The Calculational Method*, 53:131–136, 1995.
16. S. L. Peyton Jones. *The Implementation of Functional Programming Languages*. Series in Computer Science. Prentice Hall, 1987.
17. J. C. Reynolds. Types, abstraction and parametric polymorphism. In R. E. A. Mason, editor, *Proceedings 9th IFIP World Computer Congress, Information Processing '83, Paris, France, 19–23 Sept 1983*, pages 513–523. North-Holland, Amsterdam, 1983.
18. M. Rosendahl. Strictness analysis for attribute grammars. In *PLILP'92*, volume 631 of *LNCS*, pages 145–157. Springer-Verlag, 1992.
19. E. Van Wyk, O. de Moor, K. Backhouse, and P. Kwiatkowski. Forwarding in attribute grammars for modular language design. *Compiler Construction 2002*.
20. P. Wadler. Theorems for free! In *FPCA '89, London, England*, pages 347–359. ACM Press, September 1989.

A Definitions and Theorems

Pairs $fst(x, y) = x$ $snd(x, y) = y$ $(f \times g)(x, y) = (f\ x, g\ y)$

Sequences. Throughout this paper, we use sequences rather than lists. Our motivation is that this simplifies the use of the abstraction rule (see below). The type of sequences is:

type $Seq\ \alpha = Nat^+ \rightarrow \alpha$

Every sequence xs in this document has a finite length n , which means that $\forall\langle k > n \bullet xs_k = \perp \rangle$. (Hence, sequences should be over a domain with a \perp element.) We sometimes use the notation $[a, b, c]$ to denote a finite length sequence and xs_k is equivalent to $xs\ k$. Some useful operations on sequences are:

$$\begin{aligned}
\text{map} & :: (\alpha \rightarrow \beta) \rightarrow \text{Seq } \alpha \rightarrow \text{Seq } \beta \\
\text{map } f \text{ } xs & = f \circ xs \\
\text{appSeq} & :: \text{Seq } (\alpha \rightarrow \beta) \rightarrow \text{Seq } \alpha \rightarrow \text{Seq } \beta \\
\text{appSeq } fs \text{ } xs \text{ } k & = fs \text{ } k \text{ } (xs \text{ } k) \\
\text{zipWith} & :: (\alpha \rightarrow \beta \rightarrow \gamma) \rightarrow \text{Seq } \alpha \rightarrow \text{Seq } \beta \rightarrow \text{Seq } \gamma \\
\text{zipWith } f \text{ } xs \text{ } ys \text{ } k & = f \text{ } (xs \text{ } k) \text{ } (ys \text{ } k)
\end{aligned}$$

Note that *appSeq* and *zipWith* are equivalent to the *S* and *S'* combinators, respectively (see Peyton Jones [16, pages 260 and 270]). Therefore, the claim on page 148 can be stated as, for all c_1, c_2, xs, ys and k :

$$S' \ c_1 \ xs \ (S' \ c_2 \ ys \ zs) \ k = S \ (S' \ (\circ) \ (c_1 \circ \text{xs}) \ (c_2 \circ \text{ys})) \ zs \ k$$

The proof is a simple manipulation of the definitions of *S* and *S'*.

Rose Trees. Our rose trees differ from Bird's [4, page 195], because we use sequences, rather than lists:

data *Rose* $\alpha = \text{Node } \alpha \ (\text{Seq } (\text{Rose } \alpha))$

All our rose trees are finite in size. Some useful functions are:

$$\begin{aligned}
\text{split} & :: \text{Rose } \alpha \rightarrow (\alpha, \text{Forest } \alpha) & \text{merge} & :: (\alpha, \text{Forest } \alpha) \rightarrow \text{Rose } \alpha \\
\text{split } (\text{Node } x \ xs) & = (x, xs) & \text{merge } (x, xs) & = \text{Node } x \ xs \\
\text{mapRose} & :: (\alpha \rightarrow \beta) \rightarrow \text{Rose } \alpha \rightarrow \text{Rose } \beta & \text{root} & :: \text{Rose } \alpha \rightarrow \alpha \\
\text{mapRose } f \ (\text{Node } x \ xs) & = & \text{root} & = \text{fst} \circ \text{split} \\
& \text{Node } (f \ x) \ (\text{map } (\text{mapRose } f) \ xs) & & \\
\text{appRose} & :: \text{Rose } (\alpha \rightarrow \beta) \rightarrow \text{Rose } \alpha \rightarrow \text{Rose } \beta \\
\text{appRose } (\text{Node } f \ fs) \ (\text{Node } x \ xs) & = \text{Node } (f \ x) \ (\text{zipWith } \text{appRose } fs \ xs) \\
\text{foldRose} & :: ((\alpha, \text{Seq } \beta) \rightarrow \beta) \rightarrow \text{Rose } \alpha \rightarrow \beta \\
\text{foldRose } r \ (\text{Node } x \ xs) & = r \ (x, \text{map } (\text{foldRose } r) \ xs)
\end{aligned}$$

The hylomorphism theorem [5, page 142] on rose trees states that, for all r, s :

$$(\text{foldrose } r) \circ (\text{foldrose } s)^\circ = \mu \langle h \bullet r \circ (\text{id} \times \text{map } h) \circ s^\circ \rangle$$

Fixpoint Theorems. A proper introduction to these theorems is given by the Eindhoven MPC group [15]. The theorems assume that the fixpoints exist.

1. Rolling Rule. *If f is a monotonic function of type $A \rightarrow B$ and g is a monotonic function of type $B \rightarrow A$ then: $\mu (f \circ g) = f (\mu (g \circ f))$.*
2. Abstraction Rule. *Suppose f is a function of type $A \rightarrow B \rightarrow B$. If f is monotonic in both its arguments then: $\mu (S \ f) = \langle x \bullet \mu (f \ x) \rangle$. In the context of sequences, this means that for all xs : $\mu (\text{appSeq } xs) = \text{map } \mu \ xs$.*

3. Mutual Recursion Rule. *Suppose f is a monotonic function of type $(A, B) \rightarrow (A, B)$. If we define $f_1 = \text{fst} \circ f$ and $f_2 = \text{snd} \circ f$, then:*

$$\mu f = (\mu \langle x \bullet f_1 (x, p x) \rangle, \mu \langle y \bullet f_2 (q y, y) \rangle)$$

$$\mathbf{where} \quad p x = \mu \langle v \bullet f_2 (x, v) \rangle$$

$$q y = \mu \langle u \bullet f_1 (u, y) \rangle$$