

Finite-Tree Analysis for Constraint Logic-Based Languages

Roberto Bagnara, Roberta Gori, Patricia M. Hill, Enea Zaffanella

Abstract

Logic languages based on the theory of rational, possibly infinite, trees have much appeal in that rational trees allow for faster unification (due to the safe omission of the occurs-check) and increased expressivity (cyclic terms can provide a very efficient representation of grammars and other useful objects). Unfortunately, the use of infinite rational trees has problems. For instance, many of the built-in and library predicates are ill-defined for such trees and need to be supplemented by run-time checks whose cost may be significant. Moreover, some widely-used program analysis and manipulation techniques are only correct for those parts of programs working over finite trees. It is thus important to obtain, automatically, a knowledge of those program variables (the *finite variables*) that, at the program points of interest, will always be bound to finite terms. For these reasons, we propose here a new data-flow analysis that captures such information. We present a parametric domain where a simple component for recording finite variables is coupled with a generic domain (the parameter of the construction) providing sharing information. The sharing domain is abstractly specified so as to guarantee the correctness of the combined domain and the generality of the approach.

Keywords: *Finite trees, rational trees, data-flow analysis, abstract interpretation, sharing analysis.*

1 Introduction

The intended computation domain of most logic-based languages includes the algebra (or structure) of *finite trees*. Other (constraint) logic-based languages, such as Prolog II and its successors [9, 11], SICStus Prolog [34], and Oz [32], refer to a

This work has been partly supported by MURST projects “Certificazione automatica di programmi mediante interpretazione astratta” and “Interpretazione astratta, sistemi di tipo e analisi Control-Flow”. Some of this work was done during visits of the fourth author to Leeds, funded by EPSRC under grant M05645.

Roberto Bagnara and Enea Zaffanella are with the Department of Mathematics of the University of Parma, Italy. Email: {bagnara,zaffanella}@cs.unipr.it.

Roberta Gori is with the Department of Computer Science of the University of Pisa, Italy. Email: gori@di.unipi.it.

Patricia M. Hill is with the School of Computing of the University of Leeds, United Kingdom. Email: hill@comp.leeds.ac.uk.

computation domain of *rational trees*. A rational tree is a tree (possibly infinite) with a finite number of distinct subtrees and where each node has a finite number of immediate descendants. For instance, the infinite list $[a, a, \dots]$ defined by $X = [a|X]$ is a rational tree, whose only subterms are a and X . These properties will ensure that rational trees, even though infinite in the sense that they admit paths of infinite length, can be finitely represented. One possible representation makes use of connected, rooted, directed and possibly cyclic graphs where nodes are labeled with variable and function symbols as is the case of finite trees.

Applications of rational trees in logic programming include graphics [17], parser generation and grammar manipulation [9, 20], and computing with finite-state automata [9]. Other applications are described in [19] and [22]. Going from Prolog to CLP, [30] combines constraints on rational trees and record structures, while the logic-based language *Oz* allows constraints over rational and feature trees [32]. The expressive power of rational trees is put to use, for instance, in several areas of natural language processing. Rational trees are used in implementations of the HPSG formalism (Head-driven Phrase Structure Grammar) [31], in the ALE system (Attribute Logic Engine) [8], and in the ProFIT system (Prolog with Features, Inheritance and Templates) [18].

While rational trees allow for increased expressivity, they also come equipped with a surprising number of problems. As we will see, some of these problems are so serious that rational trees must be used in a very controlled way, disallowing them in any context where they are “dangerous”. This, in turn, causes a secondary problem: in order to disallow rational trees in selected contexts one must first detect them, an operation that may be expensive.

The first thing to be aware of is that almost any semantics-based program manipulation technique developed in the field of logic programming —whether it be an analysis, a transformation, or an optimization— assumes a computation domain of *finite trees*. Some of these techniques might work with the rational trees but their correctness has only been proved in the case of finite trees. Others are clearly inapplicable. Let us consider a very simple Prolog program:

```
list([]).
list(_|T) :- list(T).
```

Most automatic and semi-automatic tools for proving program termination and for complexity analysis agree on the fact that `list/1` will terminate when invoked with a ground argument. Consider now the query `?- X = [a|X], list(X).` and note that, after the execution of the first rational unification, the variable X will be bound to a rational tree containing no variables, i.e., the predicate `list/1` will be invoked with X ground. However, if such a query is given to, say, SICStus Prolog, then the only way to get the prompt back is by pressing `^C`. The problem stems from the fact that the analysis techniques employed by these tools are only sound for finite trees: as soon as they are applied to a system where the creation of cyclic terms is possible, their results are inapplicable. The situation can be improved by combining these termination and/or complexity analyses by a finite-tree analysis providing the precondition for the applicability of the other techniques.

The implementation of built-in predicates is another problematic issue. Indeed, it is widely acknowledged that, for the implementation of a system that provides real support for the rational trees, the biggest effort concerns proper handling of built-ins. Of course, the meaning of ‘proper’ depends on the actual built-in. Built-ins such as `copy_term/2` and `==/2` maintain a clear semantics when passing from finite to rational trees. For others, like `sort/2`, the extension can be questionable:¹ for instance, both raising an exception and answering $Y = [a]$ can be argued to be “the right reaction” to the query `?- X = [a|X], sort(X, Y)`. Other built-ins do not tolerate infinite trees in some argument positions. A good implementation should check for finiteness of the corresponding arguments and make sure “the right thing” —failing or raising an appropriate exception— always happens. However, such behavior appears to be uncommon. A small experiment we conducted on six modern Prolog implementations with queries like

```
?- X = 1+X, Y is X.
?- X = [97|X], name(Y, X).
?- X = [X|X], Y =.. [f|X].
```

resulted in infinite loops, memory exhaustion and/or system thrashing, segmentation faults or other fatal errors. One of the implementations tested, SICStus Prolog, is a professional one and implements run-time checks to avoid most cases where built-ins can have catastrophic effects.² The remaining systems are a bit more than research prototypes, but will clearly have to do the same if they evolve to the stage of production tools. Again, a data-flow analysis aimed at the detection of those variables that are definitely bound to finite terms would allow to avoid a (possibly significant) fraction of the useless run-time checks. Note that what has been said for built-in predicates applies to libraries as well. Even though it may be argued that it is enough for programmers to know that they should not use a particular library predicate with infinite terms, it is clear that the use of a “safe” library, including automatic checks which ensure that such predicates are never called with an illegal argument, will result in more robust systems. With the appropriate data-flow analyses, safe libraries do not have to be inefficient libraries.

Another serious problem is the following: the ISO Prolog standard term ordering cannot be extended to rational trees [M. Carlsson, Personal communication, October 2000]. Consider the rational trees defined by $A = f(B, a)$ and $B = f(A, b)$. Clearly, $A == B$ does not hold. Since the standard term ordering is total, we must have either $A @< B$ or $B @< A$. Assume $A @< B$. Then $f(A, b) @< f(B, a)$, since the ordering of terms having the same principal functor is inherited by the ordering of subterms considered in a left-to-right fashion. Thus $B @< A$ must hold, which is a contradiction. A dual contradiction is obtained by assuming $B @< A$. As a consequence, applying one of the Prolog term-ordering predicates to one or two infinite terms may cause inconsistent results, giving rise to bugs that are exceptionally difficult to diagnose. For this reason, any system that extends ISO Prolog with rational trees ought to detect such situations and make sure they are not ignored (e.g., by

¹Even though `sort/2` is not required to be a built-in by the standard, it is offered as such by several implementations.

²SICStus 3.8.5 still loops on `?- X = [97|X], name(Y, X)`.

throwing an exception or aborting execution with a meaningful message). However, predicates such as the term-ordering ones are likely to be called a significant number of times, since they are often used to maintain structures implementing ordered collections of terms. This is another instance of the efficiency issue mentioned above.

In this paper, we present a parametric abstract domain for finite-tree analysis, denoted by $H \times P$. This domain combines a simple component H (the *finiteness* component), recording the set of definitely finite variables, with a generic domain P (the parameter of the construction), providing sharing information. The term “sharing information” is to be understood in its broader meaning, which includes variable aliasing, groundness, linearity, freeness and any other kind of information that can improve the precision on these components, such as explicit structural information. Several domain combinations and abstract operators, characterized by different precision/complexity trade-offs, have been proposed to capture these properties (see [5] for an account of some of them). By giving a generic specification for this parameter component, in the style of the *open product* construct proposed in [13], it is possible to define and establish the correctness of the abstract operators on the finite-tree domain independently from any particular domain for sharing analysis.

The paper is structured as follows. The required notations and preliminary concepts are given in Section 2. The finite-tree domain is then introduced in Section 3: Section 3.1 provides the specification of the parameter domain P ; Section 3.2 defines the abstraction function for the finiteness component H ; Section 3.3 defines the abstract unification operator for $H \times P$. A brief description of some further developments on the subject is given in Section 4. We conclude in Section 5.

A longer version of this paper with proofs of the results presented here is available as a technical report [2].

2 Preliminaries

2.1 Infinite Terms and Substitutions

For a set S , $\wp(S)$ is the powerset of S , whereas $\wp_f(S)$ is the set of all the *finite* subsets of S . Let Sig denote a possibly infinite set of function symbols, ranked over the set of natural numbers. It is assumed that Sig contains at least one function symbol having rank 0 and one having rank greater than 0. Let $Vars$ denote a denumerable set of variables, disjoint from Sig . Then $Terms$ denotes the free algebra of all (possibly infinite) terms in the signature Sig having variables in $Vars$. Thus a term can be seen as an ordered labeled tree, possibly having some infinite paths and possibly containing variables: every inner node is labeled with a function symbol in Sig with a rank matching the number of the node’s immediate descendants, whereas every leaf is labeled by either a variable in $Vars$ or a function symbol in Sig having rank 0 (a constant).

If $t \in Terms$ then $vars(t)$ and $mvars(t)$ denote the set and the multiset of variables occurring in t , respectively. We will also write $vars(o)$ to denote the set of variables occurring in an arbitrary syntactic object o . If a occurs more than once in a multiset

M we write $a \in M$.

Suppose $s, t \in \text{Terms}$: s and t are *independent* if $\text{vars}(s) \cap \text{vars}(t) = \emptyset$; if $y \in \text{vars}(t)$ and $\neg(y \in \text{mvars}(t))$ we say that variable y *occurs linearly in t* , more briefly written using the predication $\text{occ_lin}(y, t)$; t is said to be *ground* if $\text{vars}(t) = \emptyset$; t is *free* if $t \in \text{Vars}$; t is *linear* if, for all $y \in \text{vars}(t)$, we have $\text{occ_lin}(y, t)$; finally, t is a *finite term* (or *Herbrand term*) if it contains a finite number of occurrences of function symbols. The sets of all ground, linear and finite terms are denoted by $G\text{Terms}$, $L\text{Terms}$ and $H\text{Terms}$, respectively. As we have specified that Sig contains function symbols of rank 0 and rank greater than 0, $G\text{Terms} \cap H\text{Terms} \neq \emptyset$ and $G\text{Terms} \setminus H\text{Terms} \neq \emptyset$.

A *substitution* is a total function $\sigma: \text{Vars} \rightarrow H\text{Terms}$ that is the identity almost everywhere; in other words, the *domain of σ* , $\text{dom}(\sigma) \stackrel{\text{def}}{=} \{x \in \text{Vars} \mid \sigma(x) \neq x\}$, is finite. Given a substitution $\sigma: \text{Vars} \rightarrow H\text{Terms}$, we overload the symbol ‘ σ ’ so as to denote also the function $\sigma: H\text{Terms} \rightarrow H\text{Terms}$ defined as follows, for each term $t \in H\text{Terms}$:

$$\sigma(t) \stackrel{\text{def}}{=} \begin{cases} t, & \text{if } t \text{ is a constant symbol;} \\ \sigma(t), & \text{if } t \in \text{Vars;} \\ f(\sigma(t_1), \dots, \sigma(t_n)), & \text{if } t = f(t_1, \dots, t_n). \end{cases}$$

If $x \in \text{Vars}$ and $t \in H\text{Terms} \setminus \{x\}$, then $x \mapsto t$ is called a *binding*. The set of all bindings is denoted by Bind . Substitutions are denoted by the set of their bindings, thus a substitution σ is identified with the (finite) set $\{x \mapsto \sigma(x) \mid x \in \text{dom}(\sigma)\}$. We denote by $\text{vars}(\sigma)$ the set of variables occurring in the bindings of σ .

A substitution is said to be *circular* if and only if, for some $n > 1$, it has the form $\{x_1 \mapsto x_2, \dots, x_{n-1} \mapsto x_n, x_n \mapsto x_1\}$, where x_1, \dots, x_n are distinct variables. A substitution is in *rational solved form* if it has no circular subset. The set of all substitutions in rational solved form is denoted by $R\text{Subst}$.

If $t \in H\text{Terms}$, we write $t\sigma$ to denote $\sigma(t)$ and $t[x/s]$ to denote $t\{x \mapsto s\}$.

The composition of substitutions is defined in the usual way. Thus $\tau \circ \sigma$ is the substitution such that, for all terms $t \in H\text{Terms}$, $(\tau \circ \sigma)(t) = \tau(\sigma(t))$ and has the formulation

$$\tau \circ \sigma = \{x \mapsto x\sigma\tau \mid x \in \text{dom}(\sigma), x \neq x\sigma\tau\} \cup \{x \mapsto x\tau \mid x \in \text{dom}(\tau) \setminus \text{dom}(\sigma)\}.$$

As usual, σ^0 denotes the identity function (i.e., the empty substitution) and, when $i > 0$, σ^i denotes the substitution $(\sigma \circ \sigma^{i-1})$. For each $\sigma \in R\text{Subst}$ and $s \in H\text{Terms}$, the sequence of finite terms $\sigma^0(s), \sigma^1(s), \sigma^2(s), \dots$ converges to a (possibly infinite) term, denoted $\sigma^\infty(s)$ [24, 28]. Thus, the function $\text{rt}: H\text{Terms} \times R\text{Subst} \rightarrow \text{Terms}$ such that $\text{rt}(s, \sigma) = \sigma^\infty(s)$ is well defined. Note that, in general, this function is not a substitution: while having a finite domain, its “bindings” $x \mapsto t$ can map a domain variable x into a term $t \in \text{Terms} \setminus H\text{Terms}$.

2.2 Equations

An *equation* is of the form $s = t$ where $s, t \in H\text{Terms}$. Eqs denotes the set of all equations. A substitution σ may be regarded as a finite set of equations, that is, as

the set $\{x = t \mid x \mapsto t \in \sigma\}$. We say that a set of equations e is in *rational solved form* if $\{s \mapsto t \mid (s = t) \in e\} \in RSubst$. In the rest of the paper, we will often write a substitution $\sigma \in RSubst$ to denote a set of equations in rational solved form (and vice versa).

Languages such as Prolog II, SICStus and Oz are based on \mathcal{RT} , the theory of rational trees [9, 10]. This is a syntactic equality theory (i.e., a theory where the function symbols are uninterpreted), augmented with a *uniqueness axiom* for each substitution in rational solved form. Informally speaking these axioms state that, after assigning a ground rational tree to each non-domain variable, the substitution uniquely defines a ground rational tree for each of its domain variables. Thus, any set of equations in rational solved form is, by definition, satisfiable in \mathcal{RT} . Note that being in rational solved form is a very weak property. Indeed, unification algorithms returning a set of equations in rational solved form are allowed to be much more “lazy” than one would usually expect. We refer the interested reader to [26, 27, 29] for details on the subject.

Given a set of equations $e \in \wp_f(Eqs)$ that is satisfiable in \mathcal{RT} , a substitution $\sigma \in RSubst$ is called a *solution for e in \mathcal{RT}* if $\mathcal{RT} \vdash \forall(\sigma \rightarrow e)$, i.e., if every model of the theory \mathcal{RT} is also a model of the first order formula $\forall(\sigma \rightarrow e)$. If in addition $\text{vars}(\sigma) \subseteq \text{vars}(e)$, then σ is said to be a *relevant solution for e* . Finally, σ is a *most general solution for e in \mathcal{RT}* if $\mathcal{RT} \vdash \forall(\sigma \leftrightarrow e)$. In this paper, the set of all the relevant most general solution for e in \mathcal{RT} will be denoted by $\text{mgs}(e)$.

2.3 The Concrete Domain

Throughout the paper, we assume a knowledge of the basic concepts of abstract interpretation theory [14, 15].

For the purpose of this paper, we assume a concrete domain constituted by pairs of the form (Σ, V) , where V is a finite set of *variables of interest* and Σ is a (possibly infinite) set of substitutions in rational solved form.

Definition 2.1 (The concrete domain.) *Let $\mathcal{D}^b \stackrel{\text{def}}{=} \wp(RSubst) \times \wp_f(Vars)$. If $(\Sigma, V) \in \mathcal{D}^b$, then (Σ, V) represents the (possibly infinite) set of first-order formulas $\{\exists \Delta . \sigma \mid \sigma \in \Sigma, \Delta = \text{vars}(\sigma) \setminus V\}$ where σ is interpreted as the logical conjunction of the equations corresponding to its bindings.*

Concrete domains for constraint languages would be similar. If the analyzed language allows the use of constraints on various domains to restrict the values of the variable leaves of rational trees, the corresponding concrete domain would have one or more extra components to account for the constraints (see [3] for an example).

The concrete element $(\{\{x \mapsto f(y)\}\}, \{x, y\})$ expresses a dependency between x and y . In contrast, $(\{\{x \mapsto f(y)\}\}, \{x\})$ only constrains x . The same concept can be expressed by saying that in the first case the variable name ‘ y ’ matters, but it does not in the second case. Thus, the set of variables of interest is crucial for defining the meaning of the concrete and abstract descriptions. Despite this, always specifying the set of variables of interest would significantly clutter the presentation. Moreover, most of the needed functions on concrete and abstract descriptions preserve the

set of variables of interest. For these reasons, we assume the existence of a set $VI \in \wp_f(\text{Vars})$ that contains, at each stage of the analysis, the current variables of interest.³ As a consequence, when the context makes it clear that $\Sigma \in \wp(\text{RSubst})$, we will write $\Sigma \in \mathcal{D}^b$ as a shorthand for $(\Sigma, VI) \in \mathcal{D}^b$.

3 An Abstract Domain for Finite-Tree Analysis

Finite-tree analysis applies to logic-based languages computing over a domain of rational trees where cyclic structures are allowed. In contrast, analyses aimed at occurs-check reduction [16, 33] apply to programs that are meant to compute on a domain of finite trees only, but have to be executed over systems that are either designed for rational trees or intended just for the finite trees but omit the occurs-check for efficiency reasons. Despite their different objectives, finite-tree and occurs-check analyses have much in common: in both cases, it is important to detect all program points where cyclic structures can be generated.

Note however that, when performing occurs-check reduction, one can take advantage of the following invariant: all data structures generated so far are finite. This property is maintained by transforming the program so as to force finiteness whenever it is possible that a cyclic structure could have been built.⁴ In contrast, a finite-tree analysis has to deal with the more general case when some of the data structures computed so far may be cyclic. It is therefore natural to consider an abstract domain made up of two components. The first one simply represents the set of variables that are guaranteed not to be bound to infinite terms. We will denote this *finiteness component* by H (from *Herbrand*).

Definition 3.1 (The finiteness component.) *The finiteness component is the set $H \stackrel{\text{def}}{=} \wp(VI)$ partially ordered by reverse subset inclusion.*

The second component of the finite-tree domain should maintain any kind of information that may be useful for computing finiteness information.

It is well-known that sharing information as a whole, therefore including possible variable aliasing, definite linearity, and definite freeness, has a crucial role in occurs-check reduction so that, as observed before, it can be exploited for finite-tree analysis too. Thus, a first choice for the second component of the finite-tree domain would be to consider one of the standard combinations of sharing, freeness and linearity as defined, e.g., in [5, 6, 21]. However, this would tie our specification to a particular sharing analysis domain, whereas the overall approach seems to be inherently more general. For this reason, we will define a finite-tree analysis based on the abstract domain schema $H \times P$, where the generic *sharing component* P is a parameter of the

³This parallels what happens in the efficient implementation of data-flow analyzers. In fact, almost all the abstract domains currently in use do not need to represent explicitly the set of variables of interest. In contrast, this set is maintained externally and in a unique copy, typically by the fixpoint computation engine.

⁴Such a requirement is typically obtained by replacing the problematic unification with a call to `unify_with_occurs_check/2`. As an alternative, in some systems based on rational trees it is possible to insert, after each problematic unification, a finiteness test for the generated term.

abstract domain construction. This approach can be formalized as an application of the *open product* operator [13].

3.1 The parameter component P

Elements of P can encode any kind of information. We only require that substitutions that are equivalent in the theory \mathcal{RT} are identified in P .

Definition 3.2 (The parameter component.) *The parameter component P is an abstract domain related to the concrete domain \mathcal{D}^b by means of the concretization function $\gamma_P: P \rightarrow \wp(\mathcal{RSubst})$ such that, for all $p \in P$,*

$$\left(\sigma \in \gamma_P(p) \wedge (\mathcal{RT} \vdash \forall(\sigma \leftrightarrow \tau)) \right) \implies \tau \in \gamma_P(p).$$

The interface between H and P is provided by a set of predicates and functions that satisfy suitable correctness criteria. Note that, for space limitations, we will only specify those abstract operations that are useful to define abstract unification on the combined domain $H \times P$. The other operations needed for a full description of the analysis, such as renamings, upper bound operators and projections, are very simple and, as usual, do not pose any problems.

Definition 3.3 (Abstract operators on P .) *Let $s, t \in HTerms$ be finite terms. For each $p \in P$, we define the following predicates:
 s and t are independent in p if and only if $\text{ind}_p: HTerms^2 \rightarrow Bool$ holds for (s, t) , where*

$$\text{ind}_p(s, t) \implies \forall \sigma \in \gamma_P(p) : \text{vars}(\text{rt}(s, \sigma)) \cap \text{vars}(\text{rt}(t, \sigma)) = \emptyset;$$

s and t share linearly in p if and only if $\text{share_lin}_p: HTerms^2 \rightarrow Bool$ holds for (s, t) , where

$$\begin{aligned} \text{share_lin}_p(s, t) \implies \forall \sigma \in \gamma_P(p) : \\ \forall y \in \text{vars}(\text{rt}(s, \sigma)) \cap \text{vars}(\text{rt}(t, \sigma)) : \\ \text{occ_lin}(y, \text{rt}(s, \sigma)) \wedge \text{occ_lin}(y, \text{rt}(t, \sigma)); \end{aligned}$$

t is ground in p if and only if $\text{ground}_p: HTerms \rightarrow Bool$ holds for t , where

$$\text{ground}_p(t) \implies \forall \sigma \in \gamma_P(p) : \text{rt}(t, \sigma) \in GTerms;$$

t is ground-or-free in p if and only if $\text{gfree}_p: HTerms \rightarrow Bool$ holds for t , where

$$\text{gfree}_p(t) \implies \forall \sigma \in \gamma_P(p) : \text{rt}(t, \sigma) \in GTerms \vee \text{rt}(t, \sigma) \in Vars;$$

s and t are or-linear in p if and only if $\text{or_lin}_p: HTerms^2 \rightarrow Bool$ holds for (s, t) , where

$$\text{or_lin}_p(s, t) \implies \forall \sigma \in \gamma_P(p) : \text{rt}(s, \sigma) \in LTerms \vee \text{rt}(t, \sigma) \in LTerms;$$

s is linear in p if and only if $\text{lin}_p: H\text{Terms} \rightarrow \text{Bool}$ holds for s , where

$$\text{lin}_p(s) \stackrel{\text{def}}{\iff} \text{or_lin}_p(s, s).$$

For each $p \in P$, the following functions compute subsets of the set of variables of interest: the function $\text{share_same_var}_p: H\text{Terms} \times H\text{Terms} \rightarrow \wp(VI)$ returns a set of variables that may share with the given terms via the same variable. For each $s, t \in H\text{Terms}$,

$$\text{share_same_var}_p(s, t) \supseteq \left\{ y \in VI \mid \begin{array}{l} \exists \sigma \in \gamma_P(p) . \\ \exists z \in \text{vars}(\text{rt}(y, \sigma)) . \\ z \in \text{vars}(\text{rt}(s, \sigma)) \cap \text{vars}(\text{rt}(t, \sigma)) \end{array} \right\};$$

the function $\text{share_with}_p: H\text{Terms} \rightarrow \wp(VI)$ yields a set of variables that may share with the given term. For each $t \in H\text{Terms}$,

$$\text{share_with}_p(t) \stackrel{\text{def}}{=} \{ y \in VI \mid y \in \text{share_same_var}_p(y, t) \}.$$

The function $\text{amgu}_P: P \times \text{Bind} \rightarrow P$ correctly captures the effects of a binding on an element of P . For each $(x \mapsto t) \in \text{Bind}$ and $p \in P$, let

$$p' \stackrel{\text{def}}{=} \text{amgu}_P(p, x \mapsto t).$$

For all $\sigma \in \gamma_P(p)$, if $\tau \in \text{mgs}(\sigma \cup \{x = t\})$, then $\tau \in \gamma_P(p')$.

Some of these generic operators can be directly mapped into the corresponding abstract operators defined for well-known sharing analysis domains. However, the specification given in Definition 3.3, besides being more general than a particular implementation, also allows for a modular approach when proving correctness results.

Example 3.1 Let $p \in P$ and assume $\sigma \in \gamma_P(p)$, where

$$\sigma \stackrel{\text{def}}{=} \{x \mapsto f(x_1, x_2), y \mapsto g(x_1, x_3), w \mapsto h(x_1), z \mapsto g(x_2, x_3)\}.$$

Then we have $w \in \text{share_same_var}_p(x, y)$, but $z \notin \text{share_same_var}_p(x, y)$. Note that z shares with both x and y , but through different variables. Formally,

$$\text{vars}(\text{rt}(w, \sigma)) \cap \text{vars}(\text{rt}(x, \sigma)) \cap \text{vars}(\text{rt}(y, \sigma)) = \{x_1\},$$

but

$$\text{vars}(\text{rt}(z, \sigma)) \cap \text{vars}(\text{rt}(x, \sigma)) \cap \text{vars}(\text{rt}(y, \sigma)) = \emptyset.$$

3.2 The abstraction function for H

When the concrete domain is based on the theory of finite trees, idempotent substitutions provide a finitely computable *strong normal form* for domain elements, meaning that different substitutions describe different sets of finite trees.⁵ In contrast, when working on a concrete domain based on the theory of rational trees, substitutions in rational solved form, while being finitely computable, no longer satisfy this property: there can be an infinite set of substitutions in rational solved form all describing the same set of rational trees (i.e., the same element in the “intended” semantics). For instance, the substitutions

$$\sigma_n = \{x \mapsto \overbrace{f(\dots f(x)\dots)}^n\}$$

for $n = 1, 2, \dots$, all map the variable x into the same rational tree (which is usually denoted by f^ω).

Ideally, a strong normal form for the set of rational trees described by a substitution $\sigma \in RSubst$ can be obtained by computing the limit σ^∞ . The problem is that we may end up with $\sigma^\infty \notin RSubst$, as σ^∞ can map domain variables to infinite rational trees.

This poses a non-trivial problem when trying to define a “good” abstraction function, since it would be really desirable for this function to map any two equivalent concrete elements to the same abstract element. As shown in [23], the classical abstraction function for set-sharing analysis [12, 25], which was defined for idempotent substitutions only, does not enjoy this property when applied, as it is, to arbitrary substitutions in rational solved form. A possibility is to look for a more general abstraction function that allows to obtain the desired property. For example, in [23] the sharing-group operator sg of [25] is replaced by an occurrence operator, occ , defined by means of a fixpoint computation. We now provide a similar fixpoint construction defining the finiteness operator.

Definition 3.4 (Finiteness functions.) *For each $n \in \mathbb{N}$, the finiteness function $hvars_n: RSubst \rightarrow \wp(Vars)$ is defined, for each $\sigma \in RSubst$, by*

$$hvars_0(\sigma) \stackrel{\text{def}}{=} Vars \setminus \text{dom}(\sigma)$$

and, for $n > 0$, by

$$hvars_n(\sigma) \stackrel{\text{def}}{=} hvars_{n-1}(\sigma) \cup \{y \in \text{dom}(\sigma) \mid \text{vars}(y\sigma) \subseteq hvars_{n-1}(\sigma)\}.$$

For each $\sigma \in RSubst$ and each $i \geq 0$, we have $hvars_i(\sigma) \subseteq hvars_{i+1}(\sigma)$ and also that $Vars \setminus hvars_i(\sigma) \subseteq \text{dom}(\sigma)$ is a finite set. By these two properties, the following fixpoint computation is well defined and finitely computable.

Definition 3.5 (Finiteness operator.) *For each $\sigma \in RSubst$, the finiteness operator $hvars: RSubst \rightarrow \wp(Vars)$ is given by $hvars(\sigma) \stackrel{\text{def}}{=} hvars_\ell(\sigma)$ where $\ell \stackrel{\text{def}}{=} \ell(\sigma) \in \mathbb{N}$ is such that $hvars_\ell(\sigma) = hvars_n(\sigma)$ for all $n \geq \ell$.*

⁵As usual, this is modulo the possible renaming of variables.

The following proposition shows that the hvars operator precisely captures the intended property.

Proposition 3.1 *If $\sigma \in RSubst$ and $x \in Vars$ then*

$$x \in \text{hvars}(\sigma) \iff \text{rt}(x, \sigma) \in HTerms.$$

Example 3.2 *Consider $\sigma \in RSubst$, where*

$$\sigma \stackrel{\text{def}}{=} \{x_1 \mapsto f(x_2), x_2 \mapsto g(x_5), x_3 \mapsto f(x_4), x_4 \mapsto g(x_3)\}.$$

Then,

$$\begin{aligned} \text{hvars}_0(\sigma) &= Vars \setminus \{x_1, x_2, x_3, x_4\}, \\ \text{hvars}_1(\sigma) &= Vars \setminus \{x_1, x_3, x_4\}, \\ \text{hvars}_2(\sigma) &= Vars \setminus \{x_3, x_4\} = \text{hvars}(\sigma). \end{aligned}$$

Thus, $x_1 \in \text{hvars}(\sigma)$, although $\text{vars}(x_1\sigma) \subseteq \text{dom}(\sigma)$.

The abstraction function for H can then be defined in the obvious way.

Definition 3.6 (The abstraction function for H .) *The abstraction function $\alpha_H: RSubst \rightarrow H$ is defined, for each $\sigma \in RSubst$, by $\alpha_H(\sigma) \stackrel{\text{def}}{=} VI \cap \text{hvars}(\sigma)$.*

The concrete domain \mathcal{D}^b is related to H by means of the abstraction function $\alpha_H: \mathcal{D}^b \rightarrow H$ such that, for each $\Sigma \in \wp(RSubst)$, $\alpha_H(\Sigma) \stackrel{\text{def}}{=} \bigcap \{ \alpha_H(\sigma) \mid \sigma \in \Sigma \}$.

Since the abstraction function α_H is additive, the concretization function is given by its adjoint [14]: $\gamma_H(h) \stackrel{\text{def}}{=} \{ \sigma \in RSubst \mid \alpha_H(\sigma) \supseteq h \}$.

With these definitions, we have the desired result: equivalent substitutions in rational solved form have the same finiteness abstraction.

Theorem 3.2 *If $\sigma, \tau \in RSubst$ and $\mathcal{RT} \vdash \forall(\sigma \leftrightarrow \tau)$, then $\alpha_H(\sigma) = \alpha_H(\tau)$.*

3.3 Abstract unification on $H \times P$

The abstract unification for the combined domain $H \times P$ is defined by using the abstract predicates and functions as specified for P as well as a new finiteness predicate for the domain H .

Definition 3.7 (Abstract unification on $H \times P$.) *A term $t \in HTerms$ is a finite tree in h if and only if the predicate $\text{hterm}_h: HTerms \rightarrow \text{Bool}$ holds for t , where $\text{hterm}_h(t) \stackrel{\text{def}}{=} \text{vars}(t) \subseteq h$.*

The function $\text{amgu}_H: (H \times P) \times \text{Bind} \rightarrow H$ captures the effects of a binding on an H element. Let $\langle h, p \rangle \in H \times P$ and $(x \mapsto t) \in \text{Bind}$. Then

$$\text{amgu}_H(\langle h, p \rangle, x \mapsto t) \stackrel{\text{def}}{=} h',$$

where

$$h' \stackrel{\text{def}}{=} \begin{cases} h \cup \text{vars}(t), & \text{if } \text{hterm}_h(x) \wedge \text{ground}_p(x); \\ h \cup \{x\}, & \text{if } \text{hterm}_h(t) \wedge \text{ground}_p(t); \\ h, & \text{if } \text{hterm}_h(x) \wedge \text{hterm}_h(t) \\ & \wedge \text{ind}_p(x, t) \wedge \text{or_lin}_p(x, t); \\ h, & \text{if } \text{hterm}_h(x) \wedge \text{hterm}_h(t) \\ & \wedge \text{gfree}_p(x) \wedge \text{gfree}_p(t); \\ h \setminus \text{share_same_var}_p(x, t), & \text{if } \text{hterm}_h(x) \wedge \text{hterm}_h(t) \\ & \wedge \text{share_lin}_p(x, t) \\ & \wedge \text{or_lin}_p(x, t); \\ h \setminus \text{share_with}_p(x), & \text{if } \text{hterm}_h(x) \wedge \text{lin}_p(x); \\ h \setminus \text{share_with}_p(t), & \text{if } \text{hterm}_h(t) \wedge \text{lin}_p(t); \\ h \setminus (\text{share_with}_p(x) \cup \text{share_with}_p(t)), & \text{otherwise.} \end{cases}$$

The abstract unification function $\text{amgu}: (H \times P) \times \text{Bind} \rightarrow H \times P$ is given, for each $\langle h, p \rangle \in H \times P$ and each $(x \mapsto t) \in \text{Bind}$, by

$$\text{amgu}(\langle h, p \rangle, x \mapsto t) \stackrel{\text{def}}{=} \langle \text{amgu}_H(\langle h, p \rangle, x \mapsto t), \text{amgu}_P(p, x \mapsto t) \rangle.$$

In the computation of h' (the new finiteness component resulting from the abstract evaluation of a binding) there are eight cases based on properties holding for the concrete terms described by x and t .

1. In the first case, the concrete term described by x is both finite and ground. Thus, after a successful execution of the binding, any concrete term described by t will be finite. Note that t could have contained variables which may be possibly bound to cyclic terms just before the execution of the binding.
2. The second case is symmetric to the first one. Note that these are the only cases when a “positive” propagation of finiteness information is correct. In contrast, in all the remaining cases, the goal is to limit as much as possible the propagation of “negative” information, i.e., the possible cyclicity of terms.
3. The third case exploits the classical results proved in research work on occurs-check reduction [16, 33]. Accordingly, it is required that both x and t describe finite terms that do not share. The use of the implicitly disjunctive predicate or_lin_p allows for the application of this case even when neither x nor t are known to be definitely linear. For instance, as observed in [16], this may happen when the component P embeds the domain Pos for groundness analysis.⁶

⁶Let t be y . Let also P be Pos . Then, given the Pos formula $\phi \stackrel{\text{def}}{=} (x \vee y)$, both $\text{ind}_\phi(x, y)$ and $\text{or_lin}_\phi(x, y)$ satisfy the conditions in Definition 4. Note that from ϕ we cannot infer that x is definitely linear and neither that y is definitely linear.

4. The fourth case exploits the observation that cyclic terms cannot be created when unifying two finite terms that are either ground or free. Ground-or-freeness [5] is a safe, more precise and inexpensive replacement for the classical freeness property when combining sharing analysis domains.
5. The fifth case applies when unifying a linear and finite term with another finite term possibly sharing with it, provided they can only share linearly (namely, all the shared variables occur linearly in the considered terms). In such a context, only the shared variables can introduce cycles.
6. In the sixth case, we drop the assumption about the finiteness of the term described by t . As a consequence, all variables sharing with x become possibly cyclic. However, provided x describes a finite and linear term, all finite variables independent from x preserve their finiteness.
7. The seventh case is symmetric to the sixth one.
8. The last case states that term finiteness is preserved for all variables that are independent from both x and t . Note that this case is only used when none of the other cases apply.

The following result, together with the assumption on amgu_P as specified in Definition 3.3, ensures that abstract unification on the combined domain $H \times P$ is correct.

Theorem 3.3 *Let $\langle h, p \rangle \in H \times P$ and $(x \mapsto t) \in \text{Bind}$, where $\{x\} \cup \text{vars}(t) \subseteq VI$. Let also $\sigma \in \gamma_H(h) \cap \gamma_P(p)$ and $h' = \text{amgu}_H(\langle h, p \rangle, x \mapsto t)$. Then*

$$\tau \in \text{mgs}(\sigma \cup \{x = t\}) \implies \tau \in \gamma_H(h').$$

4 Further Developments

The finite-tree analysis proposed in this paper has been experimentally evaluated in the framework provided by the CHINA analyzer [1]. We have developed an implementation of $H \times P$, where the P component is the *SFL* domain as presented in [5], including groundness, freeness, linearity and (non-redundant) set-sharing information. In order to exploit the precision provided by information about the actual structure of terms, this domain has been further upgraded by using the generic $\text{Pattern}(\cdot)$ construction [3]. This not only improves the precision of the computation on the *SFL* component, but also allows for a better identification of where cyclic structures may appear.

A goal-dependent analysis was run for all the programs in our benchmark suite and a comparison was made between the precision results obtained using the two domains $\text{Pattern}(SFL)$ and $\text{Pattern}(H \times SFL)$. The results for finite-tree analysis are summarized in Table 1, where the domain names are shortened to P and H , respectively. Precision is measured as the percentage of the total number of variables that the analyzer can show to be Herbrand. Since the domain $\text{Pattern}(SFL)$ has no explicit Herbrand component, in this case the number of finite variables has been

| Prec. class | P | H |
|-------------------|-----|----|
| $p = 100$ | 2 | 84 |
| $80 \leq p < 100$ | 1 | 31 |
| $60 \leq p < 80$ | 7 | 26 |
| $40 \leq p < 60$ | 6 | 41 |
| $20 \leq p < 40$ | 47 | 47 |
| $0 \leq p < 20$ | 185 | 19 |

Table 1: The precision on finite variables when using P and H.

obtained by counting the variables that are shown to be free. In the table, each column is labeled by an analysis domain and each row is labeled by a precision interval. For instance, the value ‘31’ at the intersection of column ‘H’ and row ‘ $80 \leq p < 100$ ’ is to be read as “for 31 benchmarks, the percentage p of the total number of variables that the analyzer can show to be Herbrand using the domain H is between 80% and 100%.”

As can be seen from Table 1, the domain $\text{Pattern}(H \times SFL)$ is remarkably precise. This domain, as any other domain obtained as an instance of $H \times P$, only captures the negative aspect of term-finiteness, that is, the circumstances under which finiteness can be lost. When a binding has the potential for creating one or more rational terms, the operator amgu_H removes from h all the variables that may be bound to non-finite terms. However, term-finiteness has also a positive aspect: there are cases where, provided a variable is guaranteed to be bound to a finite term, this knowledge can be propagated to other variables. This kind of information, termed *relational information*, is usually modeled by using *dependencies*.

Consider the terms $t_1 \stackrel{\text{def}}{=} f(x)$, $t_2 \stackrel{\text{def}}{=} g(y)$, and $t_3 \stackrel{\text{def}}{=} h(x, y)$: it is clear that, for each assignment of rational trees to x and y , t_3 is finite if and only if t_1 and t_2 are so. We can capture this by the Boolean formula $t_3 \leftrightarrow (t_1 \wedge t_2)$. The important point to notice is that the indicated dependency will continue to hold for any further simultaneous instantiation of t_1 , t_2 , and t_3 . In other words, such dependencies are preserved by forward computations (since they proceed by consistently instantiating program variables).

Following the intuition outlined above, in [4] we have studied a domain, whose carrier is the set of all Boolean functions, for representing and propagating finiteness dependencies. Coupling this new domain with $H \times P$ further improves the precision of the analysis.

5 Conclusion

Several modern logic-based languages offer a computation domain based on rational trees. On the one hand, such trees allow an increase in expressivity while having efficient (and correct) unification algorithms. On the other hand, these gains are countered by the extra problems rational trees bring; these being summarized as follows: several built-ins, library predicates, program analysis and manipulation techniques are only well-defined for program fragments working with finite trees.

In this paper we have proposed a solution, based on an abstract interpretation, to the problem of detecting program variables that can only be bound to finite terms. The rationale behind this is that applications exploiting rational trees tend to do so in a very controlled way. With the proposed analysis we have a practical way of taking advantage of rational trees while minimizing the impact of their disadvantages.

References

- [1] R. Bagnara. *Data-Flow Analysis for Constraint Logic-Based Languages*. PhD thesis, Dipartimento di Informatica, Università di Pisa, Pisa, Italy, 1997. Printed as Report TD-1/97.
- [2] R. Bagnara, R. Gori, P. M. Hill, and E. Zaffanella. Finite-tree analysis for constraint logic-based languages. Quaderno 251, Dipartimento di Matematica, Università di Parma, 2001. Available at <http://www.cs.unipr.it/~bagnara/>.
- [3] R. Bagnara, P. M. Hill, and E. Zaffanella. Efficient structural information analysis for real CLP languages. In M. Parigot and A. Voronkov, editors, *Proceedings of the 7th International Conference on Logic for Programming and Automated Reasoning (LPAR 2000)*, volume 1955 of *Lecture Notes in Computer Science*, pages 189–206, Réunion Island, France, 2000. Springer-Verlag, Berlin.
- [4] R. Bagnara, E. Zaffanella, R. Gori, and P. M. Hill. Boolean functions for finite-tree dependencies. Quaderno 252, Dipartimento di Matematica, Università di Parma, 2001. Available at <http://www.cs.unipr.it/~bagnara/>.
- [5] R. Bagnara, E. Zaffanella, and P. M. Hill. Enhanced sharing analysis techniques: A comprehensive evaluation. In M. Gabbriellini and F. Pfenning, editors, *Proceedings of the 2nd International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, pages 103–114, Montreal, Canada, 2000. Association for Computing Machinery.
- [6] M. Bruynooghe, M. Codish, and A. Mulkers. A composite domain for freeness, sharing, and compoundness analysis of logic programs. Technical Report CW 196, Department of Computer Science, K.U. Leuven, Belgium, July 1994.
- [7] J. A. Campbell, editor. *Implementations of Prolog*. Ellis Horwood/Halsted Press/Wiley, 1984.
- [8] B. Carpenter. *The Logic of Typed Feature Structures with Applications to Unification-based Grammars, Logic Programming and Constraint Resolution*, volume 32 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, New York, 1992.
- [9] A. Colmerauer. Prolog and infinite trees. In K. L. Clark and S. Å. Tärnlund, editors, *Logic Programming, APIC Studies in Data Processing*, volume 16, pages 231–251. Academic Press, New York, 1982.

- [10] A. Colmerauer. Equations and inequations on finite and infinite trees. In *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS'84)*, pages 85–99, Tokyo, Japan, 1984. ICOT.
- [11] A. Colmerauer. An introduction to Prolog-III. *Communications of the ACM*, 33(7):69–90, 1990.
- [12] A. Cortesi and G. Filé. Sharing is optimal. *Journal of Logic Programming*, 38(3):371–386, 1999.
- [13] A. Cortesi, B. Le Charlier, and P. Van Hentenryck. Combinations of abstract domains for logic programming: Open product and generic pattern construction. *Science of Computer Programming*, 38(1–3), 2000.
- [14] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Fourth Annual ACM Symposium on Principles of Programming Languages*, pages 238–252, 1977.
- [15] P. Cousot and R. Cousot. Abstract interpretation frameworks. *Journal of Logic and Computation*, 2(4):511–547, 1992.
- [16] L. Crnogorac, A. D. Kelly, and H. Søndergaard. A comparison of three occur-check analysers. In R. Cousot and D. A. Schmidt, editors, *Static Analysis: Proceedings of the 3rd International Symposium*, volume 1145 of *Lecture Notes in Computer Science*, pages 159–173, Aachen, Germany, 1996. Springer-Verlag, Berlin.
- [17] P. R. Eggert and K. P. Chow. Logic programming, graphics and infinite terms. Technical Report UCSB DoCS TR 83-02, Department of Computer Science, University of California at Santa Barbara, 1983.
- [18] G. Erbach. ProFIT: Prolog with Features, Inheritance and Templates. In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics*, pages 180–187, Dublin, Ireland, 1995.
- [19] M. Filgueiras. A Prolog interpreter working with infinite terms. In Campbell [7], pages 250–258.
- [20] F. Giannesini and J. Cohen. Parser generation and grammar manipulation using Prolog's infinite trees. *Journal of Logic Programming*, 3:253–265, 1984.
- [21] W. Hans and S. Winkler. Aliasing and groundness analysis of logic programs through abstract interpretation and its safety. Technical Report 92–27, Technical University of Aachen (RWTH Aachen), 1992.
- [22] S. Haridi and D. Sahlin. Efficient implementation of unification of cyclic structures. In Campbell [7], pages 234–249.

- [23] P. M. Hill, R. Bagnara, and E. Zaffanella. Soundness, idempotence and commutativity of set-sharing. *Theory and Practice of Logic Programming*, 2001. To appear. Available at <http://arXiv.org/abs/cs.PL/0102030>.
- [24] B. Intrigila and M. Venturini Zilli. A remark on infinite matching vs infinite unification. *Journal of Symbolic Computation*, 21(3):2289–2292, 1996.
- [25] D. Jacobs and A. Langen. Accurate and efficient approximation of variable aliasing in logic programs. In E. L. Lusk and R. A. Overbeek, editors, *Logic Programming: Proceedings of the North American Conference*, MIT Press Series in Logic Programming, pages 154–165, Cleveland, Ohio, USA, 1989. The MIT Press.
- [26] J. Jaffar, J-L. Lassez, and M. J. Maher. Prolog-II as an instance of the logic programming scheme. In M. Wirsing, editor, *Formal Descriptions of Programming Concepts III*, pages 275–299. North-Holland, 1987.
- [27] T. Keisu. *Tree Constraints*. PhD thesis, The Royal Institute of Technology, Stockholm, Sweden, May 1994. Also available in the SICS Dissertation Series: SICS/D-16-SE.
- [28] A. King. Pair-sharing over rational trees. *Journal of Logic Programming*, 46(1–2):139–155, 2000.
- [29] M. J. Maher. Complete axiomatizations of the algebras of finite, rational and infinite trees. In *Proceedings, Third Annual Symposium on Logic in Computer Science*, pages 348–357, Edinburgh, Scotland, 1988. IEEE Computer Society.
- [30] K. Mukai. *Constraint Logic Programming and the Unification of Information*. PhD thesis, Department of Computer Science, Faculty of Engineering, Tokio Institute of Technology, 1991.
- [31] C. Pollard and I. A. Sag. *Head-Driven Phrase Structure Grammar*. University of Chicago Press, Chicago, 1994.
- [32] Gert Smolka and Ralf Treinen. Records for logic programming. *Journal of Logic Programming*, 18(3):229–258, 1994.
- [33] H. Søndergaard. An application of abstract interpretation of logic programs: Occur check reduction. In B. Robinet and R. Wilhelm, editors, *Proceedings of the 1986 European Symposium on Programming*, volume 213 of *Lecture Notes in Computer Science*, pages 327–338. Springer-Verlag, Berlin, 1986.
- [34] Swedish Institute of Computer Science, Programming Systems Group. *SICStus Prolog User's Manual*, release 3 #0 edition, 1995.