Finite-Tree Analysis for Constraint Logic-Based Languages ⋆

Roberto Bagnara [a], Roberta Gori [b], Patricia M. Hill [c], Enea Zaffanella [a]

[a] *Department of Mathematics, University of Parma, Italy*
[b] *Department of Computer Science, University of Pisa, Italy*
[c] *School of Computing, University of Leeds, UK*

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 very efficient representations 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 correct only for those parts of programs working over finite trees. It is thus important to obtain, automatically, a knowledge of the 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, based on abstract interpretation, that captures such information. We present a parametric domain where a simple component for recording finite variables is coupled, in the style of the *open product* construction of Cortesi et al., 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. This finite-tree analysis domain is further enhanced by coupling it with a domain of Boolean functions, called *finite-tree dependencies*, that precisely captures how the finiteness of some variables influences the finiteness of other variables. We also summarize our experimental results showing how finite-tree analysis, enhanced with finite-tree dependencies, is a practical means of obtaining precise finiteness information.

Key words: static analysis, abstract interpretation, rational unification, occurs-check

1 Introduction

The intended computation domain of most logic-based languages [1] includes the algebra (or structure) of *finite trees*. Other (constraint) logic-based languages, such as Prolog II and its successors [1,2], SICStus Prolog [3], and Oz [4], refer to a computation domain of *rational trees*. [2] A rational tree is a possibly infinite tree with a finite number of distinct subtrees and where each node has a finite number of immediate descendants. These properties 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 [6], parser generation and grammar manipulation [1,7], and computing with finite-state automata [1]. Rational trees also constitute the basis of the abstract domain of *rigid type graphs*, which is used for type analysis of logic programs [8–10]. Other applications are described in [11] and [12]. Very recently, Manuel Carro has described a nice application of rational trees where they are used to represent imperative programs within interpreters. Taking a continuation-passing style approach, each instruction is coupled with a data structure representing the remaining part of the program to be executed so that sequences of instructions for realizing (backward) jumps, iterations and recursive calls give rise to cyclic structures in the form of rational trees. Compared to a naive interpreter for the same language, this threaded interpreter is faster and uses less memory, at the cost of a simple preliminary "compilation pass" to generate the rational tree representation for the program [13].

Going from Prolog to CLP, in [14] K. Mukai has combined constraints on rational trees and record structures, while the logic-based language *Oz* allows constraints over rational and feature trees [4]. The expressive power of rational trees is put to use, for instance, in several areas of natural language processing.

⋆ This work has been partly supported by MURST projects "Automatic Program Certification by Abstract Interpretation", "Abstract Interpretation, Type Systems and Control-Flow Analysis", and "Constraint Based Verification of Reactive Systems." Some of this work was done during visits of the fourth author to Leeds, funded by EPSRC under grant M05645.

Email addresses: `bagnara@cs.unipr.it` (Roberto Bagnara), `gori@di.unipi.it` (Roberta Gori), `hill@comp.leeds.ac.uk` (Patricia M. Hill), `zaffanella@cs.unipr.it` (Enea Zaffanella).

[1] That is, ordinary logic languages, (concurrent) constraint logic languages, functional logic languages and variations of the above.

[2] Support for rational trees is also provided as an option by the YAP Prolog system [5].

Rational trees are used in implementations of the HPSG formalism (Head-driven Phrase Structure Grammar) [15], in the ALE system (Attribute Logic Engine) [16], and in the ProFIT system (Prolog with Features, Inheritance and Templates) [17].

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 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 [3] and for complexity analysis [4] 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 term 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 interrupting the program. 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 with a finiteness 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 rational trees, the biggest effort concerns proper han-

[3] Such as TerminWeb [18,19], TermiLog [20], cTI [21], and LPTP [22,23].
[4] Systems like GAIA [24], CASLOG [25], and the Ciao-Prolog preprocessor [26].

dling 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: [5] failing, raising an exception, answering `Y = [a]` (if duplicates are deleted) and answering `Y = [a|Y]` (if duplicates are kept) can all 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 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. [6] 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 could be used 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 ensuring that such a predicate is never called with an illegal argument, will result in a robuster system. With the appropriate data-flow analyses, safe libraries do not have to be inefficient libraries.

Another serious problem is the following: the standard term ordering dictated by ISO Prolog [27] 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 stan-

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

[6] SICStus 3.11 still loops on `?- X = [97|X], name(Y, X).`

dard 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 any Prolog term-ordering predicate to terms where one or both of them is infinite 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 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.

Still on efficiency, it is worth noting that even for built-ins whose definition on rational trees is not problematic, there is often a performance penalty in catering for the possibility of infinite trees. Thus, for such predicates, which include rational unification provided by `=/2`, a compile-time knowledge of term finiteness can be beneficial. For instance, rational-tree implementations of the built-ins `ground/1`, `term_variables/2`, `copy_term/2`, `subsumes/2`, `variant/2` and `numbervars/3` need more expensive marking techniques to ensure they do not enter an infinite loop. With finiteness information it is possible to avoid this overhead.

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 (written with the initial of *Herbrand* and called 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 [28,29] 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 [30], it is possible to define and establish the correctness of abstract operators on the finite-tree domain independently from any particular domain for sharing analysis.

The information encoded by H is *attribute independent* [31], which means that each variable is considered in isolation. What this lacks is information about how finiteness of one variable affects the finiteness of other variables. This kind of information, usually called *relational information*, is not cap-

tured at all by H and is only partially captured by the composite domain $H \times P$. Moreover, $H \times P$ is designed to capture the "negative" aspect of term-finiteness, that is, the circumstances under which finiteness can be lost. However, term-finiteness has also a "positive" aspect: there are cases where a variable is granted to be bound to a finite term and this knowledge can be propagated to other variables. Guarantees of finiteness are provided by several built-ins like `unify_with_occurs_check/2`, `var/1`, `name/2`, all the arithmetic predicates, besides those explicitly provided to test for term-finiteness such as the `acyclic_term/1` predicate of SICStus Prolog. For these reasons $H \times P$ is coupled with a domain of Boolean functions that precisely captures how the finiteness of some variables influences the finiteness of other variables. This domain of *finite-tree dependencies* provides relational information that is important for the precision of the overall finite-tree analysis. It also combines obvious similarities, interesting differences and somewhat unexpected connections with classical domains for *groundness dependencies*. Finite-tree and groundness dependencies are similar in that they both track *covering* information (a term s covers t if all the variables in t also occur in s) and share several abstract operations. However, they are different because covering does not tell the whole story. Suppose x and y are free variables before either the unification $x = f(y)$ or the unification $x = f(x, y)$ are executed. In both cases, x will be ground if and only if y will be so. However, when $x = f(y)$ is the performed unification, this equivalence will also carry over to finiteness. In contrast, when the unification is $x = f(x, y)$, x will never be finite and will be totally independent, as far as finiteness is concerned, from y. Among the unexpected connections is the fact that finite-tree dependencies can improve the groundness information obtained by the usual approaches to groundness analysis.

The paper is structured as follows. The required notations and preliminary concepts are given in Section 2. The concrete domain for the analysis is presented in Section 3. The finite-tree domain is then introduced in Section 4: Section 4.1 provides the specification of the parameter domain P; Section 4.2 defines some computable operators that extract, from substitutions in rational solved form, properties of the denoted rational trees; Section 4.3 defines the abstraction function for the finiteness component H; Section 4.4 defines the abstract unification operator for $H \times P$. Section 5 introduces the use of Boolean functions for tracking finite-tree dependencies, whereas Section 6 illustrates the interaction between groundness and finite-tree dependencies. Our experimental results are presented in Section 7. We conclude the main body of the paper in Section 8. Appendix A specifies the sharing domain SFL defined in [32,33] as a possible instance of the parameter P.

This paper is a combination and improvement of [34] and [35]. As a result of an editorial requirement, the proofs of the stated results have been omitted from this version of the paper; the referees did check the proofs which were

part of the submitted version. The proofs can be found in [36], the unabridged version of this paper.

2 Preliminaries

2.1 Infinite Terms and Substitutions

The cardinality of a set S is denoted by $\#S$; $\wp(S)$ is the powerset of S, whereas $\wp_{\mathrm{f}}(S)$ 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 and Terms denote 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 non-leaf 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 $t \in \mathrm{Terms}$ then $\mathrm{vars}(t)$ and $\mathrm{mvars}(t)$ denote the set and the multiset of variables occurring in t, respectively. We will also write $\mathrm{vars}(o)$ to denote the set of variables occurring in an arbitrary syntactic object o.

Suppose $s, t \in \mathrm{Terms}$: s and t are *independent* if $\mathrm{vars}(s) \cap \mathrm{vars}(t) = \varnothing$; t is said to be *ground* if $\mathrm{vars}(t) = \varnothing$; t is *free* if $t \in \mathrm{Vars}$; if $y \in \mathrm{vars}(t)$ occurs exactly once in t, then we say that variable y *occurs linearly in* t, more briefly written using the predication $\mathrm{occ_lin}(y, t)$; t is *linear* if we have $\mathrm{occ_lin}(y, t)$ for all $y \in \mathrm{vars}(t)$; 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 GTerms, LTerms and HTerms, respectively. As we have specified that Sig contains function symbols of rank 0 and rank greater than 0, $\mathrm{GTerms} \cap \mathrm{HTerms} \neq \varnothing$ and $\mathrm{GTerms} \setminus \mathrm{HTerms} \neq \varnothing$.

A *substitution* is a total function $\sigma\colon \mathrm{Vars} \to \mathrm{HTerms}$ that is the identity almost everywhere; in other words, the *domain* of σ,

$$\mathrm{dom}(\sigma) \stackrel{\mathrm{def}}{=} \Big\{\, x \in \mathrm{Vars} \;\Big|\; \sigma(x) \neq x \,\Big\},$$

is finite. Given a substitution $\sigma\colon \mathrm{Vars} \to \mathrm{HTerms}$, we overload the symbol 'σ' so as to denote also the function $\sigma\colon \mathrm{HTerms} \to \mathrm{HTerms}$ defined as follows, for

each term $t \in \text{HTerms}$:

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

If $t \in \text{HTerms}$, we write $t\sigma$ to denote $\sigma(t)$ and $t\sigma\tau$ to denote $(t\sigma)\tau$.

If $x \in \text{Vars}$ and $t \in \text{HTerms} \setminus \{x\}$, 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

$$\Big\{ x \mapsto x\sigma \;\Big|\; x \in \text{dom}(\sigma) \Big\}.$$

We denote by $\text{vars}(\sigma)$ the set of variables occurring in the bindings of σ.

A substitution is said to be *circular* if, for $n > 1$, it has the form

$$\{x_1 \mapsto x_2, \ldots, x_{n-1} \mapsto x_n, x_n \mapsto x_1\},$$

where x_1, $\ldots$, 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 RSubst.

The composition of substitutions is defined in the usual way. Thus $\tau \circ \sigma$ is the substitution such that, for all terms $t \in \text{HTerms}$,

$$(\tau \circ \sigma)(t) = \tau\Big(\sigma(t)\Big) = t\sigma\tau$$

and has the formulation

$$\tau \circ \sigma = \Big\{ x \mapsto x\sigma\tau \;\Big|\; x \in \text{dom}(\sigma) \cup \text{dom}(\tau), x \neq x\sigma\tau \Big\}.$$

As usual, σ^0 denotes the identity function (i.e., the empty substitution) and, when $i > 0$, σ^i denotes the substitution $(\sigma \circ \sigma^{i-1})$.

Consider an infinite sequence of terms $t_0, t_1, t_2, \ldots$ with $t_i \in \text{HTerms}$ for each $i \in \mathbb{N}$. Suppose there exists $t \in \text{Terms}$ such that, for each $n \in \mathbb{N}$, there exists $m_0 \in \mathbb{N}$ such that, for each $m \in \mathbb{N}$ with $m \geq m_0$, the trees corresponding to the terms t and t_m coincide up to the first n levels. Then we say that *the sequence* $t_0, t_1, t_2, \ldots$ *converges to* t and we write $t = \lim_{i \to \infty} t_i$ [37].

For each $\sigma \in \text{RSubst}$ and $t \in \text{HTerms}$, the sequence of finite terms

$$\sigma^0(t), \sigma^1(t), \sigma^2(t), \ldots$$

converges [37,38]. Therefore, the function $\mathrm{rt}\colon \mathrm{HTerms} \times \mathrm{RSubst} \to \mathrm{Terms}$ such that

$$\mathrm{rt}(t, \sigma) \stackrel{\mathrm{def}}{=} \lim_{i \to \infty} \sigma^i(t)$$

is well defined.

2.2 Equations

An *equation* is a statement of the form $s = t$ where $s, t \in \mathrm{HTerms}$. Eqs denotes the set of all equations. As usual, a system of equations (i.e., a conjuction of elements in Eqs) will be denoted by a subset of Eqs. A substitution σ may be regarded as a finite set of equations, that is, as the set $\{\, x = t \mid x \mapsto t \in \sigma \,\}$. A set of equations e is in *rational solved form* if $\left\{\, s \mapsto t \;\middle|\; (s = t) \in e \,\right\} \in \mathrm{RSubst}$. In the rest of the paper, we will often write a substitution $\sigma \in \mathrm{RSubst}$ 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 [1,39]. 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. For instance, $\{x = y, y = z\}$ and $\left\{x = f(y), y = f(x)\right\}$ are in rational solved form. We refer the interested reader to [40–42] for details on the subject.

Given a set of equations $e \in \wp_{\mathrm{f}}(\mathrm{Eqs})$ that is satisfiable in $\mathcal{RT}$, a substitution $\sigma \in \mathrm{RSubst}$ is called a *solution for e in $\mathcal{RT}$* if $\mathcal{RT} \vdash \forall(\sigma \to e)$, i.e., if theory $\mathcal{RT}$ entails the first order formula $\forall(\sigma \to e)$. If in addition $\mathrm{vars}(\sigma) \subseteq \mathrm{vars}(e)$, 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 solutions for e in $\mathcal{RT}$ will be denoted by $\mathrm{mgs}(e)$.

In the sequel, in order to model the constraint accumulation process of logic-based languages, we will need to characterize those sets of equations that are stronger than (that can be obtained by adding equations to) a given set of equations.

Definition 1 **($\downarrow(\cdot)$)** *The function $\downarrow(\cdot)\colon \mathrm{RSubst} \to \wp(\mathrm{RSubst})$ is defined, for*

each $\sigma \in \mathrm{RSubst}$*, by*

$$\downarrow \sigma \stackrel{\mathrm{def}}{=} \Big\{ \tau \in \mathrm{RSubst} \;\Big|\; \exists \sigma' \in \mathrm{RSubst} \,.\, \tau \in \mathrm{mgs}(\sigma \cup \sigma') \Big\}.$$

The next result shows that $\downarrow(\cdot)$ corresponds to the closure by entailment in $\mathcal{RT}$.

Proposition 2 *Let* $\sigma \in \mathrm{RSubst}$*. Then*

$$\downarrow \sigma = \Big\{ \tau \in \mathrm{RSubst} \;\Big|\; \mathcal{RT} \vdash \forall(\tau \to \sigma) \Big\}.$$

2.3 Boolean Functions

Boolean functions have already been extensively used for data-flow analysis of logic-based languages. An important class of these functions used for tracking groundness dependencies is Pos [43]. This domain was introduced in [44] under the name *Prop* and further refined and studied in [45,46].

The formal definition of the set of Boolean functions over a finite set of variables is based on the notion of Boolean valuation. Note that in all the following definitions we abuse notation by assuming that the finite set of variables V is clear from context, so as to avoid using it as a suffix everywhere.

Definition 3 (Boolean valuation and function.) *Let* $V \in \wp_{\mathrm{f}}(\mathrm{Vars})$ *and* $\mathrm{Bool} \stackrel{\mathrm{def}}{=} \{0, 1\}$*. The set of* Boolean valuations over V *is given by*

$$\mathrm{Bval} \stackrel{\mathrm{def}}{=} V \to \mathrm{Bool}.$$

The set of Boolean functions over V *is*

$$\mathrm{Bfun} \stackrel{\mathrm{def}}{=} \mathrm{Bval} \to \mathrm{Bool}.$$

Bfun *is partially ordered by the relation* $\models$ *where, for each* $\phi, \psi \in \mathrm{Bfun}$*,*

$$\phi \models \psi \quad \stackrel{\mathrm{def}}{\Longleftrightarrow} \quad \Big(\forall a \in \mathrm{Bval} : \phi(a) = 1 \implies \psi(a) = 1\Big).$$

Boolean functions are constructed from the elementary functions corresponding to variables and by means of the usual logical connectives. Thus, for each $x \in V$, x also denotes the Boolean function ϕ such that, for each $a \in \mathrm{Bval}$, $\phi(a) = 1$ if and only if $a(x) = 1$; for $\phi \in \mathrm{Bfun}$, we write $\neg\phi$ to denote the function ψ such that, for each $a \in \mathrm{Bval}$, $\psi(a) = 1$ if and only if $\phi(a) = 0$; for $\phi_1, \phi_2 \in \mathrm{Bfun}$, we write $\phi_1 \vee \phi_2$ to denote the function ϕ such that, for each

$a \in \mathrm{Bval}$, $\phi(a) = 0$ if and only if both $\phi_1(a) = 0$ and $\phi_2(a) = 0$. A variable is restricted away using Schröder's elimination principle [47]:

$$\exists x \,.\, \phi \stackrel{\text{def}}{=} \phi[1/x] \vee \phi[0/x]$$

where, for each $c \in \mathrm{Bool}$ and each $a \in \mathrm{Bval}$,

$$\begin{aligned} \phi[c/x](a) &\stackrel{\text{def}}{=} \phi\Big(a[c/x]\Big), \\ a[c/x](y) &\stackrel{\text{def}}{=} \begin{cases} c, & \text{if } x = y; \\ a(y), & \text{otherwise.} \end{cases} \end{aligned}$$

Note that existential quantification is both monotonic and extensive on Bfun. The other Boolean connectives and quantifiers are handled similarly. The distinguished elements $\bot, \top \in \mathrm{Bfun}$ are the functions defined by

$$\begin{aligned} \bot &\stackrel{\text{def}}{=} \lambda a \in \mathrm{Bval} \,.\, 0, \\ \top &\stackrel{\text{def}}{=} \lambda a \in \mathrm{Bval} \,.\, 1. \end{aligned}$$

For notational convenience, when $X \subseteq V$, we inductively define

$$\bigwedge X \stackrel{\text{def}}{=} \begin{cases} \top, & \text{if } X = \varnothing; \\ x \wedge \bigwedge\Big(X \setminus \{x\}\Big), & \text{if } x \in X. \end{cases}$$

The distinguished valuation $\mathbf{1} \stackrel{\text{def}}{=} \lambda x \in V \,.\, 1$ is also called the *everything-is-true* assignment. $\mathrm{Pos} \subset \mathrm{Bfun}$ consists precisely of those functions assuming the true value under the *everything-is-true* assignment, i.e.,

$$\mathrm{Pos} \stackrel{\text{def}}{=} \Big\{ \phi \in \mathrm{Bfun} \;\Big|\; \phi(\mathbf{1}) = 1 \Big\}.$$

For each $\phi \in \mathrm{Bfun}$, the *positive part of* ϕ, denoted $\mathrm{pos}(\phi)$, is the strongest Pos formula that is entailed by ϕ. Formally,

$$\mathrm{pos}(\phi) \stackrel{\text{def}}{=} \phi \vee \bigwedge V.$$

For each $\phi \in \mathrm{Bfun}$, the set of *variables necessarily true for* ϕ and the set of *variables necessarily false for* ϕ are given, respectively, by

$$\begin{aligned} \mathrm{true}(\phi) &\stackrel{\text{def}}{=} \Big\{ x \in V \;\Big|\; \forall a \in \mathrm{Bval} : \phi(a) = 1 \implies a(x) = 1 \Big\}, \\ \mathrm{false}(\phi) &\stackrel{\text{def}}{=} \Big\{ x \in V \;\Big|\; \forall a \in \mathrm{Bval} : \phi(a) = 1 \implies a(x) = 0 \Big\}. \end{aligned}$$

3 The Concrete Domain

A knowledge of the basic concepts of abstract interpretation theory [48,49] is assumed. In this paper, the concrete domain consists of pairs of the form (Σ, V), where V is a finite set of *variables of interest* [45] and Σ is a (possibly infinite) set of substitutions in rational solved form.

Definition 4 (The concrete domain.) *Let* $\mathcal{D}^\flat \stackrel{\text{def}}{=} \wp(\text{RSubst}) \times \wp_{\text{f}}(\text{Vars})$. *If* $(\Sigma, V) \in \mathcal{D}^\flat$, *then* (Σ, V) *represents the (possibly infinite) set of first-order formulas* $\left\{ \exists \Delta \,.\, \sigma \;\middle|\; \sigma \in \Sigma, \Delta = \text{vars}(\sigma) \setminus V \right\}$ *where* σ *is interpreted as the logical conjunction of the equations corresponding to its bindings.*

The operation of projecting $x \in \text{Vars}$ *away from* $(\Sigma, V) \in \mathcal{D}^\flat$ *is defined as follows:*

$$\exists x \,.\, (\Sigma, V) \stackrel{\text{def}}{=} \left\{ \sigma' \in \text{RSubst} \;\middle|\; \begin{array}{l} \sigma \in \Sigma, \overline{V} = \text{Vars} \setminus V, \\ \mathcal{RT} \vdash \forall\Big(\exists \overline{V} \,.\, (\sigma' \leftrightarrow \exists x \,.\, \sigma)\Big) \end{array} \right\}.$$

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 [50] for an example).

The concrete element $\Big(\Big\{\{x \mapsto f(y)\}\Big\}, \{x, y\}\Big)$ expresses a dependency between x and y. In contrast, $\Big(\Big\{\{x \mapsto f(y)\}\Big\}, \{x\}\Big)$ 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 $\text{VI} \in \wp_{\text{f}}(\text{Vars})$ that contains, at each stage of the analysis, the current variables of interest. [7] As a consequence, when the context makes it clear, we will write $\Sigma \in \mathcal{D}^\flat$ as a shorthand for $(\Sigma, \text{VI}) \in \mathcal{D}^\flat$.

[7] This 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.

4 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 [51,52] 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. [8] 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 5 (The finiteness component.) *The* finiteness component *is the set* $H \stackrel{\text{def}}{=} \wp(\mathrm{VI})$ *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 [28,29,53,54]. However, this would tie our specification to a particular sharing analysis domain, whereas the overall approach is 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 abstract domain construction. This approach can be formalized as an application of the *open product*

[8] Such a requirement is typically obtained by replacing the unification with a call to the standard predicate `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.

operator [30], where the interaction between the H and P components is modeled by defining a suite of generic query operators: thus, the overall accuracy of the finite-tree analysis will heavily depend on the accuracy with which any specific instance of the parameter P is able to answer these queries.

4.1 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 6 (The parameter component.) *The parameter component P is an abstract domain related to the concrete domain $\mathcal{D}^\flat$ by means of the* concretization function $\gamma_P \colon P \to \wp(\mathrm{RSubst})$ *such that, for all $p \in P$,*

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

The interface between H and P is provided by a set of abstract operators that satisfy suitable correctness criteria. We only specify those that are useful for defining abstract unification and projection on the combined domain $H \times P$. Other operations needed for a full description of the analysis, such as renaming and upper bound, are very simple and, as usual, do not pose any problems.

Definition 7 (Abstract operators on P.) *Let $s, t \in$ HTerms be finite terms. For each $p \in P$, we specify the following predicates:*

s and t are independent in *p if and only if* $\mathrm{ind}_p \colon \mathrm{HTerms}^2 \to \mathrm{Bool}$ *holds for (s,t), where*

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

s and t share linearly in *p if and only if* $\mathrm{share_lin}_p \colon \mathrm{HTerms}^2 \to \mathrm{Bool}$ *holds for (s,t), where*

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

t is ground in *p if and only if* $\mathrm{ground}_p \colon \mathrm{HTerms} \to \mathrm{Bool}$ *holds for t, where*

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

t is ground-or-free in *p if and only if* $\mathrm{gfree}_p\colon \mathrm{HTerms} \to \mathrm{Bool}$ *holds for t, where*

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

s is linear in *p if and only if* $\mathrm{lin}_p\colon \mathrm{HTerms} \to \mathrm{Bool}$ *holds for s, where*

$$\mathrm{lin}_p(s) \implies \forall \sigma \in \gamma_P(p) : \mathrm{rt}(s,\sigma) \in \mathrm{LTerms};$$

s and t are or-linear in *p if and only if* $\mathrm{or_lin}_p\colon \mathrm{HTerms}^2 \to \mathrm{Bool}$ *holds for (s,t), where*

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

For each $p \in P$, the following functions compute subsets of the set of variables of interest:

the function $\mathrm{share_same_var}_p\colon \mathrm{HTerms} \times \mathrm{HTerms} \to \wp(\mathrm{VI})$ *returns a set of variables that may share with the given terms via the same variable. For each pair of terms $s,t \in \mathrm{HTerms}$,*

$$\mathrm{share_same_var}_p(s,t) \supseteq \left\{ y \in \mathrm{VI} \;\middle|\; \begin{array}{l} \exists \sigma \in \gamma_P(p)\ . \\ \quad \exists z \in \mathrm{vars}\Big(\mathrm{rt}(y,\sigma)\Big)\ . \\ \qquad z \in \mathrm{vars}\Big(\mathrm{rt}(s,\sigma)\Big) \cap \mathrm{vars}\Big(\mathrm{rt}(t,\sigma)\Big) \end{array} \right\};$$

the function $\mathrm{share_with}_p\colon \mathrm{HTerms} \to \wp(\mathrm{VI})$ *yields a set of variables that may share with the given term. For each $t \in \mathrm{HTerms}$,*

$$\mathrm{share_with}_p(t) \stackrel{\mathrm{def}}{=} \Big\{ y \in \mathrm{VI} \;\Big|\; y \in \mathrm{share_same_var}_p(y,t) \Big\}.$$

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

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

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

The function $\mathrm{proj}_P\colon P \times \mathrm{VI} \to P$ *correctly captures the operation of projecting away a variable from an element of P. For each $x \in \mathrm{VI}$, $p \in P$ and $\sigma \in \gamma_P(p)$, if $\tau \in \exists\, x\ .\ \{\sigma\}$, then $\tau \in \gamma_P\Big(\mathrm{proj}_P(p,x)\Big)$.*

As it will be shown in Appendix A, some of these generic operators can be directly mapped to the corresponding abstract operators defined for well-known

sharing analysis domains. However, the specification given in Definition 7, besides being more general than a particular implementation, also allows for a modular approach when proving correctness results.

4.2 Operators on Substitutions in Rational Solved Form

There are cases when an analysis tries to capture properties of the particular substitutions computed by a specific (ordinary or rational) unification algorithm. This is the case, for example, when the analysis needs to track structure sharing for the purpose of compile-time garbage collection, or provide upper bounds on the amount of memory needed to perform a given computation. More often the interest is on properties of the (finite or rational) trees that are denoted by such substitutions.

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 (as usual, this is modulo the possible renaming of variables). 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

$$\sigma_n = \left\{ x \mapsto \overbrace{f(\cdots f(}^{n} x) \cdots) \right\}$$

for $n = 1, 2, \ldots$, all map the variable x to 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 $\sigma \in \mathrm{RSubst}$ can be obtained by computing the limit function

$$\sigma^\infty \stackrel{\mathrm{def}}{=} \lambda t \in \mathrm{HTerms} \,.\, \mathrm{rt}(t, \sigma),$$

obtained by fixing the substitution parameter of 'rt'. The problem is that, in general, σ^∞ is not a substitution: while having a finite domain, its "bindings" $x \mapsto \lim_{i \to \infty} \sigma^i(x)$ can map a domain variable x to an infinite rational term. 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. Of course, it is important that the properties under investigation are exactly captured, so as to avoid any unnecessary precision loss. Pursuing this goal requires an ability to observe properties of (infinite) rational trees while just dealing with one of

their finite representations. This is not always an easy task since even simple properties can be "hidden" when using non-idempotent substitutions. For instance, when σ^∞ maps variable x to an infinite and ground rational tree (i.e., when $\mathrm{rt}(x,\sigma) \in \mathrm{GTerms} \setminus \mathrm{HTerms}$), all of its finite representations in RSubst (i.e., all the $\tau \in \mathrm{RSubst}$ such that $\mathcal{RT} \models \forall(\sigma \leftrightarrow \tau)$) will map the variable x into a finite term that is not ground. These are the motivations behind the introduction of the following computable operators on substitutions.

The groundness operator 'gvars' captures the set of variables that are mapped to ground rational trees by rt. We define it by means of the *occurrence operator* 'occ'. This was introduced in [55] as a replacement for the sharing-group operator 'sg' of [56]. In [55] the 'occ' operator is used to define a new abstraction function for set-sharing analysis that, differently from the classical ones [57,56], maps equivalent substitutions in rational solved form to the same abstract element.

Definition 8 (Occurrence and groundness operators.) *For each $n \in \mathbb{N}$, the* occurrence function $\mathrm{occ}_n \colon \mathrm{RSubst} \times \mathrm{Vars} \to \wp_{\mathrm{f}}(\mathrm{Vars})$ *is defined, for each $\sigma \in \mathrm{RSubst}$ and each $v \in \mathrm{Vars}$, by*

$$\mathrm{occ}_n(\sigma, v) \stackrel{\mathrm{def}}{=} \begin{cases} \{v\} \setminus \mathrm{dom}(\sigma), & \text{if } n = 0; \\ \Big\{\, y \in \mathrm{Vars} \;\Big|\; \mathrm{vars}(y\sigma) \cap \mathrm{occ}_{n-1}(\sigma, v) \neq \varnothing \,\Big\}, & \text{if } n > 0. \end{cases}$$

The occurrence operator $\mathrm{occ} \colon \mathrm{RSubst} \times \mathrm{Vars} \to \wp_{\mathrm{f}}(\mathrm{Vars})$ *is given, for each $\sigma \in \mathrm{RSubst}$ and $v \in \mathrm{Vars}$, by $\mathrm{occ}(\sigma, v) \stackrel{\mathrm{def}}{=} \mathrm{occ}_\ell(\sigma, v)$, where $\ell = \#\,\sigma$.*

The groundness operator $\mathrm{gvars} \colon \mathrm{RSubst} \to \wp_{\mathrm{f}}(\mathrm{Vars})$ *is given, for each substitution $\sigma \in \mathrm{RSubst}$, by*

$$\mathrm{gvars}(\sigma) \stackrel{\mathrm{def}}{=} \Big\{\, y \in \mathrm{dom}(\sigma) \;\Big|\; \forall v \in \mathrm{vars}(\sigma) : y \notin \mathrm{occ}(\sigma, v) \,\Big\}.$$

Example 9 *Let*

$$\sigma = \Big\{ x \mapsto f(y,z), y \mapsto g(z,x), z \mapsto f(a) \Big\}.$$

Then $\mathrm{gvars}(\sigma) = \{x, y, z\}$, *although* $\mathrm{vars}(x\sigma^i) \neq \varnothing$ *and* $\mathrm{vars}(y\sigma^i) \neq \varnothing$, *for all* $0 \leq i < \infty$.

The *finiteness operator* is defined, like 'occ', by means of a fixpoint construction.

Definition 10 (Finiteness functions.) *For each $n \in \mathbb{N}$, the* finiteness function $\mathrm{hvars}_n \colon \mathrm{RSubst} \to \wp(\mathrm{Vars})$ *is defined, for each $\sigma \in \mathrm{RSubst}$, by*

$$\mathrm{hvars}_0(\sigma) \stackrel{\mathrm{def}}{=} \mathrm{Vars} \setminus \mathrm{dom}(\sigma)$$

and, for $n > 0$, *by*

$$\mathrm{hvars}_n(\sigma) \stackrel{\text{def}}{=} \mathrm{hvars}_{n-1}(\sigma) \cup \Big\{\, y \in \mathrm{dom}(\sigma) \;\Big|\; \mathrm{vars}(y\sigma) \subseteq \mathrm{hvars}_{n-1}(\sigma) \,\Big\}.$$

For each $\sigma \in \mathrm{RSubst}$ and each $i \geq 0$, we have $\mathrm{hvars}_i(\sigma) \subseteq \mathrm{hvars}_{i+1}(\sigma)$ and also that $\mathrm{Vars} \setminus \mathrm{hvars}_i(\sigma) \subseteq \mathrm{dom}(\sigma)$ is a finite set. By these two properties, the chain $\mathrm{hvars}_0(\sigma) \subseteq \mathrm{hvars}_1(\sigma) \subseteq \cdots$ is stationary and finitely computable. In particular, if $\ell = \#\,\sigma$, then, for all $n \geq \ell$, $\mathrm{hvars}_\ell(\sigma) = \mathrm{hvars}_n(\sigma)$.

Definition 11 (Finiteness operator.) *For each* $\sigma \in \mathrm{RSubst}$, *the* finiteness operator $\mathrm{hvars}\colon \mathrm{RSubst} \to \wp(\mathrm{Vars})$ *is given by* $\mathrm{hvars}(\sigma) \stackrel{\text{def}}{=} \mathrm{hvars}_\ell(\sigma)$ *where* $\ell \stackrel{\text{def}}{=} \ell(\sigma) \in \mathbb{N}$ *is such that* $\mathrm{hvars}_\ell(\sigma) = \mathrm{hvars}_n(\sigma)$ *for all* $n \geq \ell$.

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

Proposition 12 *If* $\sigma \in \mathrm{RSubst}$ *and* $x \in \mathrm{Vars}$ *then*

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

Example 13 *Consider* $\sigma \in \mathrm{RSubst}$, *where*

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

Then,

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

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

The following proposition states how 'gvars' and 'hvars' behave with respect to the further instantiation of variables.

Proposition 14 *Let* $\sigma, \tau \in \mathrm{RSubst}$, *where* $\tau \in \downarrow \sigma$. *Then*

$$\mathrm{hvars}(\sigma) \supseteq \mathrm{hvars}(\tau), \tag{14a}$$

$$\mathrm{gvars}(\sigma) \cap \mathrm{hvars}(\sigma) \subseteq \mathrm{gvars}(\tau) \cap \mathrm{hvars}(\tau). \tag{14b}$$

4.3 The Abstraction Function for H

A Galois connection between the concrete domain $\wp(\mathrm{RSubst})$ and the finiteness component $H = \wp(\mathrm{VI})$ can now be defined naturally.

Definition 15 (The Galois connection between $\wp(\mathrm{RSubst})$ and H.) *The abstraction function $\alpha_H\colon \mathrm{RSubst} \to H$ is defined, for each $\sigma \in \mathrm{RSubst}$, by*

$$\alpha_H(\sigma) \stackrel{\mathrm{def}}{=} \mathrm{VI} \cap \mathrm{hvars}(\sigma).$$

The concrete domain $\mathcal{D}^\flat$ is related to H by means of the abstraction function *$\alpha_H\colon \mathcal{D}^\flat \to H$ such that, for each $\Sigma \in \wp(\mathrm{RSubst})$,*

$$\alpha_H(\Sigma) \stackrel{\mathrm{def}}{=} \bigcap \Big\{\, \alpha_H(\sigma) \;\Big|\; \sigma \in \Sigma \,\Big\}.$$

Since the abstraction function α_H is additive, the concretization function is given by its adjoint [48]*: whenever $h \in H$,*

$$\begin{aligned} \gamma_H(h) &\stackrel{\mathrm{def}}{=} \Big\{\, \sigma \in \mathrm{RSubst} \;\Big|\; \alpha_H(\sigma) \supseteq h \,\Big\} \\ &\stackrel{\mathrm{def}}{=} \Big\{\, \sigma \in \mathrm{RSubst} \;\Big|\; \mathrm{hvars}(\sigma) \supseteq h \,\Big\}. \end{aligned}$$

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

Theorem 16 *If $\sigma, \tau \in \mathrm{RSubst}$ and $\mathcal{RT} \vdash \forall(\sigma \leftrightarrow \tau)$, then $\alpha_H(\sigma) = \alpha_H(\tau)$.*

4.4 Abstract Unification and Projection 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 17 (Abstract unification on $H \times P$.) *A term $t \in \mathrm{HTerms}$ is a* finite tree *in $h \in H$ if and only if the predicate $\mathrm{hterm}_h\colon \mathrm{HTerms} \to \mathrm{Bool}$ holds for t, where*

$$\mathrm{hterm}_h(t) \stackrel{\mathrm{def}}{=} \Big(\mathrm{vars}(t) \subseteq h\Big).$$

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

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

where h' is given by the first case that applies in

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

The abstract unification function $\mathrm{amgu}\colon (H \times P) \times \mathrm{Bind} \to H \times P$*, for any* $\langle h, p\rangle \in H \times P$ *and* $(x \mapsto t) \in \mathrm{Bind}$*, is given by*

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

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 [51,52]. Accordingly, it is required that both x and t describe finite terms that do not share. The use of the implicitly disjunctive predicate $\mathrm{or_lin}_p$ allows for the application of this case even when neither x nor t are known to be definitely linear. For instance, as observed in [51], this may happen when the component P embeds the

domain Pos for groundness analysis. [9]

(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 [28,29] 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.

The following result, together with the assumption on amgu_P as specified in Definition 7, ensures that abstract unification on the combined domain $H \times P$ is correct.

Theorem 18 *Let $\langle h, p\rangle \in H \times P$ and $(x \mapsto t) \in \mathrm{Bind}$, where $\{x\} \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$. Let also $\sigma \in \gamma_H(h) \cap \gamma_P(p)$ and $h' = \mathrm{amgu}_H\Big(\langle h, p\rangle, x \mapsto t\Big)$. Then*

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

Abstract projection on the composite domain $H \times P$ is much simpler than abstract unification, because in this case there is no interaction between the two components of the abstract domain.

Definition 19 (Abstract projection on $H \times P$.) *The function $\mathrm{proj}_H \colon H \times \mathrm{VI} \to H$ captures the effects, on the H component, of projecting away a variable. For each $h \in H$ and $x \in \mathrm{VI}$,*

$$\mathrm{proj}_H(h, x) \stackrel{\mathrm{def}}{=} h \cup \{x\}.$$

The abstract variable projection function $\mathrm{proj} \colon (H \times P) \times \mathrm{VI} \to H \times P$, for any $\langle h, p\rangle \in H \times P$ and $x \in \mathrm{VI}$, is given by

$$\mathrm{proj}\Big(\langle h, p\rangle, x\Big) \stackrel{\mathrm{def}}{=} \Big\langle \mathrm{proj}_H(h, x), \mathrm{proj}_P(p, x)\Big\rangle.$$

[9] Let t be y. Let also P be Pos. Then, given the Pos formula $\phi \stackrel{\mathrm{def}}{=} (x \vee y)$, both $\mathrm{ind}_\phi(x, y)$ and $\mathrm{or_lin}_\phi(x, y)$ 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.

As a consequence, as far as the H component is concerned, the correctness of the projection function does not depend on the assumption on proj_P as specified in Definition 7.

Theorem 20 *Let $x \in \mathrm{VI}$, $h \in H$ and $\sigma \in \gamma_H(h)$. Then*

$$\tau \in \exists x \,.\, \{\sigma\} \implies \tau \in \gamma_H\Big(\mathrm{proj}_H(h, x)\Big).$$

We do not consider the disjunction and conjunction operations here. The implementation (and therefore proof of correctness) for disjunction is straightforward and omitted. The implementation of *independent conjunction* where the descriptions are renamed apart is also straightforward. On the other hand, full conjunction, which is only needed for a top-down analysis framework, can be approximated by combining unification and independent conjunction, obtaining a correct (although possibly less precise) analysis.

Several abstract domains for sharing analysis can be used to implement the parameter component P. As a basic implementation, one could consider the well-known set-sharing domain of Jacobs and Langen [56]. In such a case, most of the required correctness results have already been established in [55]. Note however that, since no freeness and linearity information is recorded in the plain set-sharing domain, some of the predicates of Definition 7 need to be grossly approximated. For instance, the predicate gfree_p will provide useful information only when applied to an argument that is known to be definitely ground. Another possibility would be to use the domain based on pair-sharing, definite groundness and definite linearity described in [38]. A more precise choice is constituted by the SFL domain (an acronym standing from Set-sharing plus Freeness plus Linearity) introduced in [58,33]. Even in this case, all the non-trivial correctness results have already been proved. In particular, in [32,33] it is shown that the abstraction function satisfies the requirement of Definition 6 and that the abstract unification operator is correct with respect to rational-tree unification. In order to better highlight the generality of our specification of the sharing component P, the instantiation of P to SFL is presented in Appendix A. Notice that the quest for more precision does not end with SFL: a number of possible precision improvements are presented and discussed in [28,29].

5 Finite-Tree Dependencies

The precision of the finite-tree analysis based on $H \times P$ is highly dependent on the precision of the generic component P. As explained before, the information provided by P on groundness, freeness, linearity, and sharing of variables is

exploited, in the combination $H \times P$, to circumscribe as much as possible the creation and propagation of cyclic terms. However, finite-tree analysis can also benefit from other kinds of relational information. In particular, we now show how *finite-tree dependencies* allow a positive propagation of finiteness information.

Let us consider the finite terms $t_1 = f(x)$, $t_2 = g(y)$, and $t_3 = h(x, y)$: it is clear that, for each assignment of rational terms 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 $t_3 \leftrightarrow (t_1 \wedge t_2)$.[10] The reasoning is based on the following facts:

(1) t_1, t_2, and t_3 are finite terms, so that the finiteness of their instances depends only on the finiteness of the terms that take the place of x and y.
(2) $\mathrm{vars}(t_3) \supseteq \mathrm{vars}(t_1) \cup \mathrm{vars}(t_2)$, that is, t_3 *covers* both t_1 and t_2; this means that, if an assignment to the variables of t_3 produces a finite instance of t_3, that very assignment will necessarily result in finite instances of t_1 and t_2. Conversely, an assignment producing non-finite instances of t_1 or t_2 will forcibly result in a non-finite instance of t_3.
(3) Similarly, t_1 and t_2, taken together, cover t_3.

The important point to notice is that this dependency will keep holding for any further simultaneous instantiation of t_1, t_2, and t_3. In other words, such dependencies are preserved by forward computations (which proceed by consistently instantiating program variables).

Consider $x \mapsto t \in \mathit{Bind}$ where $t \in \mathrm{HTerms}$ and $\mathrm{vars}(t) = \{y_1, \ldots, y_n\}$. After this binding has been successfully applied, the destinies of x and t concerning term-finiteness are tied together: forever. This tie can be described by the dependency formula

$$x \leftrightarrow (y_1 \wedge \cdots \wedge y_n), \tag{2}$$

meaning that x will be bound to a finite term if and only if y_i is bound to a finite term, for each $i = 1, \ldots, n$. While the dependency expressed by (2) is a correct description of any computation state following the application of the binding $x \mapsto t$, it is not as precise as it could be. Suppose that x and y_k are indeed the same variable. Then (2) is logically equivalent to

$$x \rightarrow (y_1 \wedge \cdots \wedge y_{k-1} \wedge y_{k+1} \wedge \cdots \wedge y_n). \tag{3}$$

Although this is correct —whenever x is bound to a finite term, all the other variables will be bound to finite terms— it misses the point that x has just been bound, irrevocably, to a non-finite term: no forward computation can change this. Thus, the implication (3) holds vacuously. A more precise and

[10] The introduction of such Boolean formulas, called *dependency formulas*, is originally due to P. W. Dart [59].

correct description for the state of affairs caused by the cyclic binding is, instead, the negated atom $\neg x$, whose intuitive reading is "x is not (and never will be) finite."

We are building an abstract domain for finite-tree dependencies where we are making the deliberate choice of including only information that cannot be withdrawn by forward computations. The reason for this choice is that we want the concrete constraint accumulation process to be paralleled, at the abstract level, by another constraint accumulation process: logical conjunction of Boolean formulas. For this reason, it is important to distinguish between *permanent* and *contingent* information. Permanent information, once established for a program point p, maintains its validity in all points that follow p in any forward computation. Contingent information, instead, does not carry its validity beyond the point where it is established. An example of contingent information is given by the h component of $H \times P$: having $x \in h$ in the description of some program point means that x is definitely bound to a finite term *at that point*; nothing is claimed about the finiteness of x at later program points and, in fact, unless x is ground, x can still be bound to a non-finite term. However, if at some program point x is finite and ground, then x will remain finite. In this case we will ensure our Boolean dependency formula entails the positive atom x.

At this stage, we already know something about the abstract domain we are designing. In particular, we have positive and negated atoms, the requirement of describing program predicates of any arity implies that arbitrary conjunctions of these atomic formulas must be allowed and, finally, it is not difficult to observe that the merge-over-all-paths operation [48] will be logical disjunction, so that the domain will have to be closed under this operation. This means that the carrier of our domain must be able to express any Boolean function over the finite set VI of the variables of interest: Bfun is the carrier.

Definition 21 ($\gamma_F \colon \mathrm{Bfun} \to \wp(\mathrm{RSubst})$**.)** *The function* $\mathrm{hval} \colon \mathrm{RSubst} \to \mathrm{Bval}$ *is defined, for each* $\sigma \in \mathrm{RSubst}$ *and each* $x \in \mathrm{VI}$*, by*

$$\mathrm{hval}(\sigma)(x) = 1 \quad \stackrel{\mathrm{def}}{\Longleftrightarrow} \quad x \in \mathrm{hvars}(\sigma).$$

The concretization function $\gamma_F \colon \mathrm{Bfun} \to \wp(\mathrm{RSubst})$ *is defined, for* $\phi \in \mathrm{Bfun}$*, by*

$$\gamma_F(\phi) \stackrel{\mathrm{def}}{=} \Big\{ \sigma \in \mathrm{RSubst} \;\Big|\; \forall \tau \in {\downarrow}\sigma : \phi\big(\mathrm{hval}(\tau)\big) = 1 \Big\}.$$

The domain of positive Boolean functions Pos used, among other things, for groundness analysis is so popular that our use of the domain Bfun deserves some further comments. For the representation of finite-tree dependencies, the presence in the domain of negative functions such as $\neg x$, meaning that x is bound to an infinite term, is an important feature. One reason why it is so

is that knowing about definite non-finiteness can improve the information on definite finiteness. The easiest example goes as follows: if we know that either x or y is finite (i.e., $x \vee y$) and we know that x is *not* finite (i.e., $\neg x$), then we can deduce that y must be finite (i.e., y). It is important to observe that this reasoning can be applied, verbatim, to groundness: a knowledge of non-groundness may improve groundness information. The big difference is that non-finiteness is information of the permanent kind while non-groundness is only contingent. As a consequence, a knowledge of finiteness and non-finiteness can be monotonically accumulated along computation paths by computing the logical conjunction of Boolean formulae. An approach where groundness and non-groundness information is represented by elements of Bfun would need to use a much more complex operation and significant extra information to correctly model the constraint accumulation process.

The other reason why the presence of negative functions in the domain is beneficial is efficiency. The most efficient implementations of Pos and Bfun, such as the ones described in [43,60], are based on Reduced Ordered Binary Decision Diagrams (ROBDD) [61]. While an ROBDD representing the imprecise information given by the formula (3) has a worst case complexity that is exponential in n, the more precise formula $\neg x$ has constant complexity.

The following theorem shows how most of the operators needed to compute the concrete semantics of a logic program can be correctly approximated on the abstract domain Bfun. Notice how the addition of equations is modeled by logical conjunction and projection of a variable is modeled by existential quantification.

Theorem 22 *Let* $\Sigma, \Sigma_1, \Sigma_2 \in \wp(\mathrm{RSubst})$ *and* $\phi, \phi_1, \phi_2 \in \mathrm{Bfun}$ *be such that* $\gamma_F(\phi) \supseteq \Sigma$, $\gamma_F(\phi_1) \supseteq \Sigma_1$, *and* $\gamma_F(\phi_2) \supseteq \Sigma_2$. *Let also* $(x \mapsto t) \in \mathrm{Bind}$, *where* $\{x\} \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$. *Then the following hold:*

$$\gamma_F\Big(x \leftrightarrow \bigwedge \mathrm{vars}(t)\Big) \supseteq \Big\{\{x \mapsto t\}\Big\}; \tag{22a}$$

$$\gamma_F(\neg x) \supseteq \Big\{\{x \mapsto t\}\Big\}, \ \textit{if } x \in \mathrm{vars}(t); \tag{22b}$$

$$\gamma_F(x) \supseteq \Big\{\, \sigma \in \mathrm{RSubst} \;\Big|\; x \in \mathrm{gvars}(\sigma) \cap \mathrm{hvars}(\sigma) \,\Big\}; \tag{22c}$$

$$\gamma_F(\phi_1 \wedge \phi_2) \supseteq \Big\{\, \mathrm{mgs}(\sigma_1 \cup \sigma_2) \;\Big|\; \sigma_1 \in \Sigma_1, \sigma_2 \in \Sigma_2 \,\Big\}; \tag{22d}$$

$$\gamma_F(\phi_1 \vee \phi_2) \supseteq \Sigma_1 \cup \Sigma_2; \tag{22e}$$

$$\gamma_F(\exists x \,.\, \phi) \supseteq \exists x \,.\, \Sigma. \tag{22f}$$

Cases (22a), (22b), and (22d) of Theorem 22 ensure that the following definition of amgu_F provides a correct approximation on Bfun of the concrete unification of rational trees.

Definition 23 *The function* $\mathrm{amgu}_F \colon \mathrm{Bfun} \times \mathrm{Bind} \to \mathrm{Bfun}$ *captures the effects*

of a binding on a finite-tree dependency formula. Let $\phi \in$ Bfun *and* $(x \mapsto t) \in$ Bind *be such that* $\{x\} \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$. *Then*

$$\mathrm{amgu}_F(\phi, x \mapsto t) \stackrel{\mathrm{def}}{=} \begin{cases} \phi \wedge \Big(x \leftrightarrow \bigwedge \mathrm{vars}(t)\Big), & \text{if } x \notin \mathrm{vars}(t); \\ \phi \wedge \neg x, & \text{otherwise.} \end{cases}$$

Other semantic operators, such as the consistent renaming of variables, are very simple and omitted for the sake of brevity.

The next result shows how finite-tree dependencies may improve the finiteness information encoded in the h component of the domain $H \times P$.

Theorem 24 *Let* $h \in H$ *and* $\phi \in$ Bfun. *Let also* $h' \stackrel{\mathrm{def}}{=} \mathrm{true}\Big(\phi \wedge \bigwedge h\Big)$. *Then*

$$\gamma_H(h) \cap \gamma_F(\phi) = \gamma_H(h') \cap \gamma_F(\phi).$$

Example 25 *Consider the following program, where it is assumed that the only "external" query is '*`?- r(X, Y)`*':*

```
p(X, Y) :- X = f(Y, _).
q(X, Y) :- X = f(_, Y).
r(X, Y) :- p(X, Y), q(X, Y), acyclic_term(X).
```

Then the predicate `p/2` *in the clause defining* `r/2` *will be called with* `X` *and* `Y` *both unbound. Computing on the abstract domain* $H \times P$ *gives us the finiteness description* $h_p = \{x, y\}$, *expressing the fact that both* `X` *and* `Y` *are bound to finite terms. Computing on the finite-tree dependencies domain* Bfun, *gives us the Boolean formula* $\phi_p = x \to y$ *(*`Y` *is finite if* `X` *is so).*

Considering now the call to the predicate `q/2`, *we note that, since variable* `X` *is already bound to a non-variable term sharing with* `Y`, *all the finiteness information encoded by* H *will be lost (i.e.,* $h_q = \varnothing$*). So, both* `X` *and* `Y` *are detected as possibly cyclic. However, the finite-tree dependency information is preserved, since we have* $\phi_q = (x \to y) \wedge (x \to y) = x \to y$.

Finally, consider the effect of the abstract evaluation of `acyclic_term(X)`. *On the* $H \times P$ *domain we can only infer that variable* `X` *cannot be bound to an infinite term, while* `Y` *will be still considered as possibly cyclic, so that* $h_r = \{x\}$. *On the domain* Bfun *we can just confirm that the finite-tree dependency computed so far still holds, so that* $\phi_r = x \to y$ *(no stronger finite-tree dependency can be inferred, since the finiteness of* `X` *is only contingent). Thus, by applying the result of* Theorem 24, *we can recover the finiteness of* `Y`*:*

$$h'_r = \mathrm{true}\Big(\phi_r \wedge \bigwedge h_r\Big) = \mathrm{true}\Big((x \to y) \wedge x\Big) = \mathrm{true}(x \wedge y) = \{x, y\}.$$

Information encoded in $H \times P$ and Bfun is not completely orthogonal and the following result provides a kind of consistency check.

Theorem 26 *Let $h \in H$ and $\phi \in$ Bfun. Then*

$$\gamma_H(h) \cap \gamma_F(\phi) \neq \varnothing \quad \Longrightarrow \quad h \cap \mathrm{false}\Big(\phi \wedge \bigwedge h\Big) = \varnothing.$$

Note however that, provided the abstract operators are correct, the computed descriptions will always be mutually consistent, unless $\phi = \bot$.

6 Groundness Dependencies

Since information about the groundness of variables is crucial for many applications, it is natural to consider a static analysis domain including both a finite-tree and a groundness component. In fact, any reasonably precise implementation of the parameter component P of the abstract domain specified in Section 4 will include some kind of groundness information.[11] We highlight similarities, differences and connections relating the domain Bfun for finite-tree dependencies to the abstract domain Pos for groundness dependencies. Note that these results also hold when considering a combination of Bfun with the groundness domain Def [43].

We first define how elements of Pos represent sets of substitutions in rational solved form.

Definition 27 **($\gamma_G\colon$ Pos $\to \wp$(RSubst).)** *The function* gval$\colon$ RSubst $\to$ Bval *is defined as follows, for each $\sigma \in$ RSubst and each $x \in$ VI:*

$$\mathrm{gval}(\sigma)(x) = 1 \quad \stackrel{\mathrm{def}}{\Longleftrightarrow} \quad x \in \mathrm{gvars}(\sigma).$$

The concretization function $\gamma_G\colon$ Pos $\to \wp$(RSubst) is defined, for each $\psi \in$ Pos,

$$\gamma_G(\psi) \stackrel{\mathrm{def}}{=} \Big\{ \sigma \in \mathrm{RSubst} \;\Big|\; \forall \tau \in \downarrow \sigma : \psi\big(\mathrm{gval}(\tau)\big) = 1 \Big\}.$$

The following is a simple variant of the standard abstract unification operator for groundness analysis over finite-tree domains: the only difference concerns the case of cyclic bindings [65].

[11] One could define P so that it explicitly contains the abstract domain Pos. Even when this is not the case, it should be noted that, as soon as the parameter P includes the set-sharing domain of Jacobs and Langen [62], then it will subsume the groundness information captured by the domain Def [63,64].

Definition 28 *The function* $\mathrm{amgu}_G\colon \mathrm{Pos} \times \mathrm{Bind} \to \mathrm{Pos}$ *captures the effects of a binding on a groundness dependency formula. Let* $\psi \in \mathrm{Pos}$ *and* $(x \mapsto t) \in \mathrm{Bind}$ *be such that* $\{x\} \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$. *Then*

$$\mathrm{amgu}_G(\psi, x \mapsto t) \stackrel{\mathrm{def}}{=} \psi \wedge \Big(x \leftrightarrow \bigwedge\big(\mathrm{vars}(t) \setminus \{x\}\big)\Big).$$

The next result shows how, by exploiting the finiteness component H, the finite-tree dependencies (Bfun) component and the groundness dependencies (Pos) component can improve each other.

Theorem 29 *Let* $h \in H$, $\phi \in \mathrm{Bfun}$ *and* $\psi \in \mathrm{Pos}$. *Let also* $\phi' \in \mathrm{Bfun}$ *and* $\psi' \in \mathrm{Pos}$ *be defined as* $\phi' = \exists \mathrm{VI} \setminus h \,.\, \psi$ *and* $\psi' = \mathrm{pos}(\exists \mathrm{VI} \setminus h \,.\, \phi)$. *Then*

$$\gamma_H(h) \cap \gamma_F(\phi) \cap \gamma_G(\psi) = \gamma_H(h) \cap \gamma_F(\phi) \cap \gamma_G(\psi \wedge \psi'); \tag{29a}$$
$$\gamma_H(h) \cap \gamma_F(\phi) \cap \gamma_G(\psi) = \gamma_H(h) \cap \gamma_F(\phi \wedge \phi') \cap \gamma_G(\psi). \tag{29b}$$

Moreover, even without any knowledge of the H component, combining Theorem 24 and Eq. (29a), the groundness dependencies component can be improved.

Theorem 30 *Let* $\phi \in \mathrm{Bfun}$ *and* $\psi \in \mathrm{Pos}$. *Then*

$$\gamma_F(\phi) \cap \gamma_G(\psi) = \gamma_F(\phi) \cap \gamma_G\Big(\psi \wedge \bigwedge \mathrm{true}(\phi)\Big).$$

The following example shows that, when computing on rational trees, finite-tree dependencies may provide groundness information that is not captured by the usual approaches.

Example 31 *Consider the program:*

```
p(a, Y).
p(X, a).
q(X, Y) :- p(X, Y), X = f(X, Z).
```

The abstract semantics of `p/2`*, for both finite-tree and groundness dependencies, is* $\phi_p = \psi_p = x \vee y$. *The finite-tree dependency for* `q/2` *is* $\phi_q = (x \vee y) \wedge \neg x = \neg x \wedge y$. *Using* Definition 28*, the groundness dependency for* `q/2` *is*

$$\psi_q = \exists z \,.\, \Big((x \vee y) \wedge (x \leftrightarrow z)\Big) = x \vee y.$$

This can be improved, using Theorem 30*, to*

$$\psi'_q = \psi_q \wedge \bigwedge \mathrm{true}(\phi_q) = y.$$

It is worth noticing that the groundness information can be improved regardless of whether, like Pos, the groundness domain captures disjunctive information: groundness information represented by the less expressive domain Def [43] can be improved as well. The next example illustrates this point.

Example 32 *Consider the following program:*

```
p(a, a).
p(X, Y) :- X = f(X, _).
q(X, Y) :- p(X, Y), X = a.
```

Consider the predicate `p/2`*. Concerning finite-tree dependencies, the abstract semantics of* `p/2` *is expressed by the Boolean formula* $\phi_p = (x \wedge y) \vee \neg x = x \rightarrow y$ *(*`Y` *is finite if* `X` *is so). In contrast, the* Pos*-groundness abstract semantics of* `p/2` *is a plain "don't know": the Boolean formula* $\psi_p = (x \wedge y) \vee \top = \top$*. In fact, the groundness of* `X` *and* `Y` *can be completely decided by the call-pattern of* `p/2`*.*

Consider now the predicate `q/2`*. The finiteness semantics of* `q/2` *is given by* $\phi_q = (x \rightarrow y) \wedge x = x \wedge y$*, whereas the* Pos *formula expressing groundness dependencies is* $\psi_q = \top \wedge x = x$*. By* Theorem 30*, we obtain*

$$\psi'_q = \psi_q \wedge \bigwedge \mathrm{true}(\phi_q) = x \wedge y,$$

therefore recovering the groundness of variable y*.*

Since better groundness information, besides being useful in itself, may also improve the precision of many other analyses such as sharing [28,29,63], the reduction steps given by Theorems 29 and 30 can trigger improvements to the precision of other components. Theorem 29 can also be exploited to recover precision after the application of a widening operator on either the groundness dependencies or the finite-tree dependencies component.

7 Experimental Results

The work described here has been experimentally evaluated in the framework provided by China [65], a data-flow analyzer for constraint logic languages

(i.e., ISO Prolog, CLP($\mathcal{R}$), `clp(FD)` and so forth). CHINA performs bottom-up analysis deriving information on both call-patterns and success-patterns by means of program transformations and optimized fixpoint computation techniques. [12] An abstract description is computed for the call- and success-patterns for each predicate defined in the program.

We implemented and compared the three domains Pattern(P), Pattern($H \times P$) and Pattern(Bfun $\times$ H $\times$ P), [13] where the parameter component P has been instantiated to the domain $\mathrm{Pos} \times \mathrm{SFL}_2$ [28,32,33] for tracking groundness, freeness, linearity and (non-redundant) set-sharing information. The Pattern($\cdot$) operator [50] further upgrades the precision of its argument by adding explicit structural information. Note that the analyzer tracks the finiteness of the terms that can be bound to those abstract variables occurring as leaves in the acyclic term structure computed by the Pattern($\cdot$) component; therefore, in order to show that an abstract variable is definitely bound to a finite term, the basic domain Pattern(P) has to prove that this variable is definitely free. [14]

Concerning the Bfun component, the implementation was straightforward, since all the techniques described in [60] (and almost all the code, including the widenings) was reused unchanged, obtaining comparable efficiency. As a consequence, most of the implementation effort was in the coding of the abstract operators on the H component and in the reduction processes between the different components. A key choice, in this sense, is *when* the reduction steps given in Theorems 24 and 29 should be applied. When striving for maximum precision, a trivial strategy is to perform reductions immediately after any application of any abstract operator. This is how predicates like `acyclic_term/1` should be handled: after adding the variables of the argument to the H component, the reduction process is applied to propagate the new information to all domain components. However, such an approach turns out to be unnecessarily inefficient. In fact, the next result shows that Theorems 24 and 29 cannot lead to a precision improvement if applied just after the abstract evaluation of the merge-over-all-paths or the existential quantification operations (provided the initial descriptions are already reduced).

[12] More precisely, CHINA uses a variation of the *Magic Templates* algorithm [66], in order to obtain goal-dependent information, and a sophisticated chaotic iteration strategy proposed in [67,68] (recursive fixpoint iteration on the weak topological ordering defined by partitioning of the call graph into strongly-connected subcomponents).

[13] For ease of notation, the domain names are shortened to P, H and B, respectively.

[14] Put in other words, by considering just the variables occurring inside the pattern structure, we systematically disregard those cases when the basic domain is able to prove that a particular argument position is definitely bound to a finite and ground term such as $f(a)$. Clearly, the same approach is consistently adopted when considering the more accurate analysis domains.

Prec. class	P	H	B
$p = 100$	2	84	86
$80 \leq p < 100$	1	31	36
$60 \leq p < 80$	7	26	23
$40 \leq p < 60$	6	41	40
$20 \leq p < 40$	47	47	46
$0 \leq p < 20$	185	19	17

Prec. improvement	P $\rightarrow$ H	H $\rightarrow$ B
$i > 20$	185	4
$10 < i \leq 20$	31	3
$5 < i \leq 10$	11	6
$2 < i \leq 5$	4	10
$0 < i \leq 2$	2	24
no improvement	15	201

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

Theorem 33 *Let* $x \in \mathrm{VI}$, $h, h' \in H$ $\phi, \phi' \in \mathrm{Bfun}$ *and* $\psi, \psi' \in \mathrm{Pos}$ *and suppose that* $\gamma_H(h) \cap \gamma_F(\phi) \neq \varnothing$. *Let*

$$h_1 \stackrel{\mathrm{def}}{=} h \cap h', \qquad \phi_1 \stackrel{\mathrm{def}}{=} \phi \vee \phi', \qquad \psi_1 \stackrel{\mathrm{def}}{=} \psi \vee \psi',$$
$$h_2 \stackrel{\mathrm{def}}{=} \mathrm{proj}_H(h, x), \qquad \phi_2 \stackrel{\mathrm{def}}{=} \exists x \,.\, \phi, \qquad \psi_2 \stackrel{\mathrm{def}}{=} \exists x \,.\, \psi.$$

Let also

$$h \supseteq \mathrm{true}\Big(\phi \wedge \bigwedge h\Big), \qquad \phi \models (\exists \mathrm{VI} \setminus h \,.\, \psi), \qquad \psi \models \mathrm{pos}(\exists \mathrm{VI} \setminus h \,.\, \phi),$$
$$h' \supseteq \mathrm{true}\Big(\phi' \wedge \bigwedge h'\Big), \qquad \phi' \models (\exists \mathrm{VI} \setminus h' \,.\, \psi'), \qquad \psi' \models \mathrm{pos}(\exists \mathrm{VI} \setminus h' \,.\, \phi').$$

Then, for $i = 1, 2$,

$$h_i \supseteq \mathrm{true}\Big(\phi_i \wedge \bigwedge h_i\Big), \qquad \phi_i \models (\exists \mathrm{VI} \setminus h_i \,.\, \psi_i), \qquad \psi_i \models \mathrm{pos}(\exists \mathrm{VI} \setminus h_i \,.\, \phi_i).$$

A goal-dependent analysis was run for all the programs in our benchmark

suite. [15] For 116 of them, the analyzer detects that the program in not amenable to goal-dependent analysis, either because the entry points are unknown or because the program uses builtins in a way that every predicate can be called with any call-pattern, so that the analysis provides results that are so imprecise to be irrelevant. The precision results for the remaining 248 programs are summarized in Table 1. Here, the precision is measured as the percentage of the total number of variables that the analyser can show to be finite. Two alternative views are provided.

In the first view, 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 finite using the domain H is between 80% and 100%.*"

The second view provides a better picture of the precision *improvements* obtained when moving from P to H (in the column '$\mathrm{P} \rightarrow \mathrm{H}$') and from H to B (in the column '$\mathrm{H} \rightarrow \mathrm{B}$'). For instance, the value '10' at the intersection of column '$\mathrm{H} \rightarrow \mathrm{B}$' and row '$2 < i \leq 5$' is to be read as "*when moving from H to B, for 10 benchmarks the improvement i in the percentage of the total number of variables shown to be finite was between 2% and 5%.*"

It can be seen from Table 1 that, even though the H domain is remarkably precise, the inclusion of the Bfun component allows for a further, and sometimes significant, precision improvement for a number of benchmarks. It is worth noting that the current implementation of China does not yet fully exploit the finite-tree dependencies arising when evaluating many of the built-in predicates, therefore incurring an avoidable precision loss. We are working on this issue and we expect that the specialized implementation of the abstract evaluation of some built-ins will result in more and better precision improvements. The experimentation has also shown that, in practice, the Bfun component does not improve the groundness information.

Concerning efficiency, our experimentation shown that the techniques we propose are really practical. The total analysis time for the 248 programs for which we give precision results in Table 1 is 596 seconds for P, 602 seconds for

[15] The suite comprises all the logic programs we have access to (including everything we could find by systematically dredging the Internet): 364 programs, 24 MB of code, 800 K lines. Besides classical benchmarks, several real programs of respectable size are included, the largest one containing 10063 clauses in 45658 lines of code. The suite also comprises a few synthetic benchmarks, which are artificial programs explicitly constructed to stress the capabilities of the analyzer and of its abstract domains with respect to precision and/or efficiency. The interested reader can find more information at the URI `http://www.cs.unipr.it/China/`.

H, and 1211 seconds for B. [16] It should be stressed that, as mentioned before, the implementation of Bfun was derived in a straightforward way from the one of Pos described in [60]. We believe that a different tuning of the widenings we employ in that component could reduce the gap between the efficiency of H and the one of B.

8 Conclusion

Several modern logic-based languages offer a computation domain based on rational trees. On the one hand, the use of such trees is encouraged by the possibility of using efficient and correct unification algorithms and by an increase in expressivity. On the other hand, these gains are countered by the extra problems rational trees bring with themselves and that can be summarized as follows: several built-ins, library predicates, program analysis and manipulation techniques are only well-defined for program fragments working with finite trees.

As a consequence, those applications that exploit rational trees tend to do so in a very controlled way, that is, most program variables can only be bound to finite terms. By detecting the program variables that may be bound to infinite terms with a good degree of accuracy, we can significantly reduce the disadvantages of using rational trees.

In this paper we have proposed an abstract-interpretation based solution to this problem, where the composite abstract domain $H \times P$ allows tracking of the creation and propagation of infinite terms. Even though this information is crucial to any finite-tree analysis, propagating the guarantees of finiteness that come from several built-ins (including those that are explicitly provided to test term-finiteness) is also important. Therefore, we have introduced a domain of Boolean functions Bfun for finite-tree dependencies which, when coupled to the domain $H \times P$, can enhance its expressive power. Since Bfun has many similarities with the domain Pos used for groundness analysis, we have investigated how these two domains relate to each other and, in particular, the synergy arising from their combination in the "global" domain of analysis.

[16] On a PC system equipped with an Athlon XP 2800 CPU, 1 GB of RAM memory and running GNU/Linux.

Acknowledgment

We would like to express our gratitude to the Journal referees for their useful comments that have helped improve the final versions of the paper.

References

[1] A. Colmerauer, Prolog and infinite trees, in: K. L. Clark, S. Å. Tärnlund (Eds.), Logic Programming, APIC Studies in Data Processing, Vol. 16, Academic Press, New York, 1982, pp. 231–251.

[2] A. Colmerauer, An introduction to Prolog-III, Communications of the ACM 33 (7) (1990) 69–90.

[3] Swedish Institute of Computer Science, Intelligent Systems Laboratory, SICStus Prolog User's Manual, release 3.9 Edition (2002).

[4] G. Smolka, R. Treinen, Records for logic programming, Journal of Logic Programming 18 (3) (1994) 229–258.

[5] V. Santos Costa, L. Damas, R. Reis, R. Azevedo, YAP User's Manual, Universidade do Porto, version 4.3.20 Edition (2001).

[6] P. R. Eggert, K. P. Chow, Logic programming, graphics and infinite terms, Tech. Rep. UCSB DoCS TR 83-02, Department of Computer Science, University of California at Santa Barbara (1983).

[7] F. Giannesini, J. Cohen, Parser generation and grammar manipulation using Prolog's infinite trees, Journal of Logic Programming 3 (1984) 253–265.

[8] P. Cousot, R. Cousot, Formal language, grammar and set-constraint-based program analysis by abstract interpretation, in: Proceedings of the Seventh ACM Conference on Functional Programming Languages and Computer Architecture, ACM Press, La Jolla, California, 1995, pp. 170–181.

[9] G. Janssens, M. Bruynooghe, Deriving descriptions of possible values of program variables by means of abstract interpretation, Journal of Logic Programming 13 (2&3) (1992) 205–258.

[10] P. Van Hentenryck, A. Cortesi, B. Le Charlier, Type analysis of Prolog using type graphs, Journal of Logic Programming 22 (3) (1995) 179–209,.

[11] M. Filgueiras, A Prolog interpreter working with infinite terms, in: Campbell [69], pp. 250–258.

[12] S. Haridi, D. Sahlin, Efficient implementation of unification of cyclic structures, in: Campbell [69], pp. 234–249.

[13] M. Carro, An application of rational trees in a logic programming interpreter for a procedural language, Tech. Rep. arXiv:cs.DS/0403028, School of Computer Science, Technical University of Madrid (UPM), available from http://arxiv.org/ (2004).

[14] K. Mukai, Constraint logic programming and the unification of information, Ph.D. thesis, Department of Computer Science, Faculty of Engineering, Tokio Institute of Technology (1991).

[15] C. Pollard, I. A. Sag, Head-Driven Phrase Structure Grammar, University of Chicago Press, Chicago, 1994.

[16] B. Carpenter, The Logic of Typed Feature Structures with Applications to Unification-based Grammars, Logic Programming and Constraint Resolution, Vol. 32 of Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, New York, 1992.

[17] 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, Dublin, Ireland, 1995, pp. 180–187.

[18] M. Codish, C. Taboch, A semantic basis for termination analysis of logic programs and its realization using symbolic norm constraints, in: M. Hanus, J. Heering, K. Meinke (Eds.), Algebraic and Logic Programming, 6th International Joint Conference, Vol. 1298 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Southampton, U.K., 1997, pp. 31–45.

[19] M. Codish, C. Taboch, A semantic basis for the termination analysis of logic programs, Journal of Logic Programming 41 (1) (1999) 103–123.

[20] N. Lindenstrauss, Y. Sagiv, A. Serebrenik, TermiLog: A system for checking termination of queries to logic programs, in: O. Grumberg (Ed.), Computer Aided Verification: Proceedings of the 9th International Conference, Vol. 1250 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Haifa, Israel, 1997, pp. 444–447.

[21] F. Mesnard, R. Bagnara, cTI: A constraint-based termination inference tool for ISO-Prolog, Theory and Practice of Logic Programming 5 (1&2), to appear.

[22] R. F. Stärk, Total correctness of pure Prolog programs: A formal approach, in: R. Dyckhoff, H. Herre, P. Schroeder-Heister (Eds.), Extensions of Logic Programming: Proceedings of the 5th International Workshop, Vol. 1050 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Leipzig, Germany, 1996, pp. 237–254.

[23] R. F. Stärk, The theoretical foundations of LPTP (a Logic Program Theorem Prover), Journal of Logic Programming 36 (3) (1998) 241–269.

[24] A. Cortesi, B. Le Charlier, S. Rossi, Specification-based automatic verification of Prolog programs, in: J. P. Gallagher (Ed.), Logic Program Synthesis and Transformation: Proceedings of the 6th International Workshop, Vol. 1207 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Stockholm, Sweden, 1997, pp. 38–57.

[25] S. Debray, N.-W. Lin, Cost analysis of logic programs, ACM Transactions on Programming Languages and Systems 15 (5) (1993) 826–875.

[26] M. V. Hermenegildo, F. Bueno, G. Puebla, P. López, Program analysis, debugging, and optimization using the ciao system preprocessor, in: D. De Schreye (Ed.), Logic Programming: The 1999 International Conference, MIT Press Series in Logic Programming, The MIT Press, Las Cruces, New Mexico, 1999, pp. 52–66.

[27] ISO/IEC, ISO/IEC 13211-1: 1995 Information technology — Programming languages — Prolog — Part 1: General core, International Standard Organization (1995).

[28] R. Bagnara, E. Zaffanella, P. M. Hill, Enhanced sharing analysis techniques: A comprehensive evaluation, in: M. Gabbrielli, F. Pfenning (Eds.), Proceedings of the 2nd International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, Association for Computing Machinery, Montreal, Canada, 2000, pp. 103–114.

[29] R. Bagnara, E. Zaffanella, P. M. Hill, Enhanced sharing analysis techniques: A comprehensive evaluation, Theory and Practice of Logic Programming 5 (1&2), to appear.

[30] A. Cortesi, B. Le Charlier, P. Van Hentenryck, Combinations of abstract domains for logic programming: Open product and generic pattern construction, Science of Computer Programming 38 (1–3) (2000) 27–71.

[31] P. Cousot, R. Cousot, Abstract interpretation and applications to logic programs, Journal of Logic Programming 13 (2&3) (1992) 103–179.

[32] P. M. Hill, E. Zaffanella, R. Bagnara, A correct, precise and efficient integration of set-sharing, freeness and linearity for the analysis of finite and rational tree languages, Theory and Practice of Logic Programming 4 (3) (2004) 289–323, to appear.

[33] E. Zaffanella, Correctness, precision and efficiency in the sharing analysis of real logic languages, Ph.D. thesis, School of Computing, University of Leeds, Leeds, U.K., available at `http://www.cs.unipr.it/~zaffanella/` (2001).

[34] R. Bagnara, R. Gori, P. M. Hill, E. Zaffanella, Finite-tree analysis for constraint logic-based languages, in: P. Cousot (Ed.), Static Analysis: 8th International Symposium, SAS 2001, Vol. 2126 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Paris, France, 2001, pp. 165–184.

[35] R. Bagnara, E. Zaffanella, R. Gori, P. M. Hill, Boolean functions for finite-tree dependencies, in: R. Nieuwenhuis, A. Voronkov (Eds.), Proceedings of the 8th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR 2001), Vol. 2250 of Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, Havana, Cuba, 2001, pp. 579–594.

[36] R. Bagnara, R. Gori, P. M. Hill, E. Zaffanella, Finite-tree analysis for constraint logic-based languages: The complete unabridged version, Quaderno 363,

Dipartimento di Matematica, Università di Parma, Italy, available at `http://www.cs.unipr.it/Publications/`. Also published as `arXiv:cs.PL/0404055`, available from `http://arxiv.org/` (2004).

[37] A. Berarducci, M. Venturini Zilli, Generalizations of unification, Journal of Symbolic Computation 15 (1993) 479–491.

[38] A. King, Pair-sharing over rational trees, Journal of Logic Programming 46 (1–2) (2000) 139–155.

[39] A. Colmerauer, Equations and inequations on finite and infinite trees, in: Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS'84), ICOT, Tokyo, Japan, 1984, pp. 85–99.

[40] J. Jaffar, J.-L. Lassez, M. J. Maher, Prolog-II as an instance of the logic programming scheme, in: M. Wirsing (Ed.), Formal Descriptions of Programming Concepts III, North-Holland, Amsterdam, 1987, pp. 275–299.

[41] T. Keisu, Tree constraints, Ph.D. thesis, The Royal Institute of Technology, Stockholm, Sweden, also available in the SICS Dissertation Series: SICS/D–16–SE (May 1994).

[42] M. J. Maher, Complete axiomatizations of the algebras of finite, rational and infinite trees, in: Proceedings, Third Annual Symposium on Logic in Computer Science, IEEE Computer Society Press, Edinburgh, Scotland, 1988, pp. 348–357.

[43] T. Armstrong, K. Marriott, P. Schachte, H. Søndergaard, Two classes of Boolean functions for dependency analysis, Science of Computer Programming 31 (1) (1998) 3–45.

[44] K. Marriott, H. Søndergaard., Notes for a tutorial on abstract interpretation of logic programs, North American Conference on Logic Programming, Cleveland, Ohio, USA (1989).

[45] A. Cortesi, G. Filé, W. Winsborough, *Prop* revisited: Propositional formula as abstract domain for groundness analysis, in: Proceedings, Sixth Annual IEEE Symposium on Logic in Computer Science, IEEE Computer Society Press, Amsterdam, The Netherlands, 1991, pp. 322–327.

[46] K. Marriott, H. Søndergaard, Precise and efficient groundness analysis for logic programs, ACM Letters on Programming Languages and Systems 2 (1–4) (1993) 181–196.

[47] E. Schröder, Der Operationskreis des Logikkalkuls, B. G. Teubner, Leibzig, 1877.

[48] P. Cousot, 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, ACM Press, New York, 1977, pp. 238–252.

[49] P. Cousot, R. Cousot, Abstract interpretation frameworks, Journal of Logic and Computation 2 (4) (1992) 511–547.

[50] R. Bagnara, P. M. Hill, E. Zaffanella, Efficient structural information analysis for real CLP languages, in: M. Parigot, A. Voronkov (Eds.), Proceedings of the 7th International Conference on Logic for Programming and Automated Reasoning (LPAR 2000), Vol. 1955 of Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, Réunion Island, France, 2000, pp. 189–206.

[51] L. Crnogorac, A. D. Kelly, H. Søndergaard, A comparison of three occur-check analysers, in: R. Cousot, D. A. Schmidt (Eds.), Static Analysis: Proceedings of the 3rd International Symposium, Vol. 1145 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Aachen, Germany, 1996, pp. 159–173.

[52] H. Søndergaard, An application of abstract interpretation of logic programs: Occur check reduction, in: B. Robinet, R. Wilhelm (Eds.), Proceedings of the 1986 European Symposium on Programming, Vol. 213 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Saarbrücken, Federal Republic of Germany, 1986, pp. 327–338.

[53] M. Bruynooghe, M. Codish, A. Mulkers, Abstract unification for a composite domain deriving sharing and freeness properties of program variables, in: F. S. de Boer, M. Gabbrielli (Eds.), Verification and Analysis of Logic Languages, Proceedings of the W2 Post-Conference Workshop, International Conference on Logic Programming, Santa Margherita Ligure, Italy, 1994, pp. 213–230.

[54] W. Hans, S. Winkler, Aliasing and groundness analysis of logic programs through abstract interpretation and its safety, Tech. Rep. 92–27, Technical University of Aachen (RWTH Aachen) (1992).

[55] P. M. Hill, R. Bagnara, E. Zaffanella, Soundness, idempotence and commutativity of set-sharing, Theory and Practice of Logic Programming 2 (2) (2002) 155–201.

[56] D. Jacobs, A. Langen, Accurate and efficient approximation of variable aliasing in logic programs, in: E. L. Lusk, R. A. Overbeek (Eds.), Logic Programming: Proceedings of the North American Conference, MIT Press Series in Logic Programming, The MIT Press, Cleveland, Ohio, USA, 1989, pp. 154–165.

[57] A. Cortesi, G. Filé, Sharing is optimal, Journal of Logic Programming 38 (3) (1999) 371–386.

[58] P. M. Hill, E. Zaffanella, R. Bagnara, A correct, precise and efficient integration of set-sharing, freeness and linearity for the analysis of finite and rational tree languages, Quaderno 273, Dipartimento di Matematica, Università di Parma, Italy, available at `http://www.cs.unipr.it/Publications/`. Also published as technical report No. 2001.22, School of Computing, University of Leeds, U.K. (2001).

[59] P. W. Dart, On derived dependencies and connected databases, Journal of Logic Programming 11 (1&2) (1991) 163–188.

[60] R. Bagnara, P. Schachte, Factorizing equivalent variable pairs in ROBDD-based implementations of *Pos*, in: A. M. Haeberer (Ed.), Proceedings of the "Seventh International Conference on Algebraic Methodology and Software Technology (AMAST'98)", Vol. 1548 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Amazonia, Brazil, 1999, pp. 471–485.

[61] R. E. Bryant, Symbolic boolean manipulation with ordered binary-decision diagrams, ACM Computing Surveys 24 (3) (1992) 293–318.

[62] D. Jacobs, A. Langen, Static analysis of logic programs for independent AND parallelism, Journal of Logic Programming 13 (2&3) (1992) 291–314.

[63] M. Codish, H. Søndergaard, P. J. Stuckey, Sharing and groundness dependencies in logic programs, ACM Transactions on Programming Languages and Systems 21 (5) (1999) 948–976.

[64] A. Cortesi, G. Filé, W. Winsborough, The quotient of an abstract interpretation for comparing static analyses, Theoretical Computer Science 202 (1&2) (1998) 163–192.

[65] R. Bagnara, Data-flow analysis for constraint logic-based languages, Ph.D. thesis, Dipartimento di Informatica, Università di Pisa, Pisa, Italy, printed as Report TD-1/97 (Mar. 1997).

[66] R. Ramakrishnan, Magic Templates: A spellbinding approach to logic programs, in: R. A. Kowalski, K. A. Bowen (Eds.), Logic Programming: Proceedings of the Fifth International Conference and Symposium on Logic Programming, MIT Press Series in Logic Programming, The MIT Press, Seattle, USA, 1988, pp. 140–159.

[67] F. Bourdoncle, Efficient chaotic iteration strategies with widenings, in: D. Bjørner, M. Broy, I. V. Pottosin (Eds.), Proceedings of the International Conference on "Formal Methods in Programming and Their Applications", Vol. 735 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Academgorodok, Novosibirsk, Russia, 1993, pp. 128–141.

[68] F. Bourdoncle, Sémantiques des langages impératifs d'ordre supérieur et interprétation abstraite, PRL Research Report 22, DEC Paris Research Laboratory (1993).

[69] J. A. Campbell (Ed.), Implementations of Prolog, Ellis Horwood/Halsted Press/Wiley, 1984.

[70] R. Bagnara, P. M. Hill, E. Zaffanella, Set-sharing is redundant for pair-sharing, Theoretical Computer Science 277 (1-2) (2002) 3–46.

[71] E. Zaffanella, P. M. Hill, R. Bagnara, Decomposing non-redundant sharing by complementation, Theory and Practice of Logic Programming 2 (2) (2002) 233–261.

[72] M. Codish, D. Dams, E. Yardeni, Derivation and safety of an abstract unification algorithm for groundness and aliasing analysis, in: K. Furukawa

(Ed.), Logic Programming: Proceedings of the Eighth International Conference on Logic Programming, MIT Press Series in Logic Programming, The MIT Press, Paris, France, 1991, pp. 79–93.

[73] F. Scozzari, Abstract domains for sharing analysis by optimal semantics, in: J. Palsberg (Ed.), Static Analysis: 7th International Symposium, SAS 2000, Vol. 1824 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, Santa Barbara, CA, USA, 2000, pp. 397–412.

A An Instance of the Parameter Domain P

As discussed in Section 4, several abstract domains for sharing analysis can be used to implement the parameter component P. We here consider the abstract domain SFL [32,33], integrating the set-sharing domain of Jacobs and Langen with definite freeness and linearity information.

Definition 34 (The set-sharing domain SH**.)** *The set* SH *is defined by* $\mathrm{SH} \stackrel{\mathrm{def}}{=} \wp(\mathrm{SG})$*, where* $\mathrm{SG} \stackrel{\mathrm{def}}{=} \wp(\mathrm{VI}) \setminus \{\varnothing\}$ *is the set of* sharing groups*.* SH *is ordered by subset inclusion.*

The information about definite freeness and linearity is encoded by two sets of variables, one for each property.

Definition 35 (The domain SFL**.)** *Let* $F \stackrel{\mathrm{def}}{=} \wp(\mathrm{VI})$ *and* $L \stackrel{\mathrm{def}}{=} \wp(\mathrm{VI})$ *be partially ordered by reverse subset inclusion. The domain* SFL *is defined by the Cartesian product* $\mathrm{SFL} \stackrel{\mathrm{def}}{=} \mathrm{SH} \times F \times L$ *ordered by '$\leq_S$', the component-wise extension of the orderings defined on the sub-domains; the bottom element is* $\bot_S \stackrel{\mathrm{def}}{=} \langle \varnothing, \mathrm{VI}, \mathrm{VI} \rangle$*.*

In the next definition we introduce a few well-known operations on the set-sharing domain SH. These will be used to define the operations on the domain SFL.

Definition 36 (Abstract operators on SH**.)** *For each* $\mathrm{sh} \in \mathrm{SH}$ *and each* $V \subseteq \mathrm{VI}$*, the extraction of the* relevant component of sh with respect to V *is given by the function* $\mathrm{rel}\colon \wp(\mathrm{VI}) \times \mathrm{SH} \to \mathrm{SH}$ *defined as*

$$\mathrm{rel}(V, \mathrm{sh}) \stackrel{\mathrm{def}}{=} \{\, S \in \mathrm{sh} \mid S \cap V \neq \varnothing \,\}.$$

For each $\mathrm{sh} \in \mathrm{SH}$ *and each* $V \subseteq \mathrm{VI}$*, the function* $\overline{\mathrm{rel}}\colon \wp(\mathrm{VI}) \times \mathrm{SH} \to \mathrm{SH}$ *gives the* irrelevant component of sh with respect to V*. It is defined as*

$$\overline{\mathrm{rel}}(V, \mathrm{sh}) \stackrel{\mathrm{def}}{=} \mathrm{sh} \setminus \mathrm{rel}(V, \mathrm{sh}).$$

The function $(\cdot)^\star \colon \mathrm{SH} \to \mathrm{SH}$*, called* star-union*, is given, for each* $\mathrm{sh} \in \mathrm{SH}$*, by*

$$\mathrm{sh}^\star \stackrel{\text{def}}{=} \left\{ S \in \mathrm{SG} \;\middle|\; \exists n \geq 1 \,.\, \exists T_1, \ldots, T_n \in \mathrm{sh} \,.\, S = \bigcup_{i=1}^{n} T_i \right\}.$$

For each $\mathrm{sh}_1, \mathrm{sh}_2 \in \mathrm{SH}$*, the function* $\mathrm{bin} \colon \mathrm{SH} \times \mathrm{SH} \to \mathrm{SH}$*, called* binary union*, is given by*

$$\mathrm{bin}(\mathrm{sh}_1, \mathrm{sh}_2) \stackrel{\text{def}}{=} \{\, S_1 \cup S_2 \mid S_1 \in \mathrm{sh}_1, S_2 \in \mathrm{sh}_2 \,\}.$$

For each $\mathrm{sh} \in \mathrm{SH}$ *and each* $(x \mapsto t) \in \mathrm{Bind}$*, the function* $\mathrm{cyclic}_x^t \colon \mathrm{SH} \to \mathrm{SH}$ *strengthens the sharing set* sh *by forcing the coupling of* x *with* t*:*

$$\mathrm{cyclic}_x^t(\mathrm{sh}) \stackrel{\text{def}}{=} \overline{\mathrm{rel}}\Big(\{x\} \cup \mathrm{vars}(t), \mathrm{sh}\Big) \cup \mathrm{rel}\Big(\mathrm{vars}(t) \setminus \{x\}, \mathrm{sh}\Big).$$

For each $\mathrm{sh} \in \mathrm{SH}$ *and each* $x \in \mathrm{VI}$*, the function* $\mathrm{proj}_{\mathrm{SH}} \colon \mathrm{SH} \times \mathrm{VI} \to \mathrm{SH}$ *projects away variable* x *from* sh*:*

$$\mathrm{proj}_{\mathrm{SH}}(\mathrm{sh}, x) \stackrel{\text{def}}{=} \Big\{\{x\}\Big\} \cup \Big\{\, S \setminus \{x\} \;\Big|\; S \in \mathrm{sh}, S \neq \{x\} \,\Big\}.$$

It is now possible to define the implementation, on the domain SFL, of all the predicates and functions specified in Definition 7.

Definition 37 (Abstract operators on SFL**.)** *For each* $\mathrm{d} \in \mathrm{SFL}$ *and* $s, t \in \mathrm{HTerms}$*, where* $\mathrm{d} = \langle \mathrm{sh}, f, l \rangle$ *and* $\mathrm{vars}(s) \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$*, let* $\mathrm{sh}_s =$

$\mathrm{rel}\Big(\mathrm{vars}(s), \mathrm{sh}\Big)$ *and* $\mathrm{sh}_t = \mathrm{rel}\Big(\mathrm{vars}(t), \mathrm{sh}\Big)$. *Then*

$$
\begin{aligned}
\mathrm{ind_d}(s,t) &\stackrel{\mathrm{def}}{=} \Big(\mathrm{sh}_s \cap \mathrm{sh}_t = \varnothing\Big);\\
\mathrm{ground_d}(t) &\stackrel{\mathrm{def}}{=} \Big(\mathrm{vars}(t) \subseteq \mathrm{VI} \setminus \mathrm{vars}(\mathrm{sh})\Big);\\
\mathrm{occ_lin_d}(y,t) &\stackrel{\mathrm{def}}{=} \mathrm{ground_d}(y) \vee \Big(\mathrm{occ_lin}(y,t) \wedge (y \in l)\\
&\qquad\qquad \wedge \forall z \in \mathrm{vars}(t) : \Big(y \neq z \implies \mathrm{ind_d}(y,z)\Big)\Big);\\
\mathrm{share_lin_d}(s,t) &\stackrel{\mathrm{def}}{=} \forall y \in \mathrm{vars}(\mathrm{sh}_s \cap \mathrm{sh}_t) :\\
&\qquad\qquad y \in \mathrm{vars}(s) \implies \mathrm{occ_lin_d}(y,s)\\
&\qquad\qquad \wedge\, y \in \mathrm{vars}(t) \implies \mathrm{occ_lin_d}(y,t);\\
\mathrm{free_d}(t) &\stackrel{\mathrm{def}}{=} \exists y \in \mathrm{VI}\,.\,(y = t) \wedge (y \in f);\\
\mathrm{gfree_d}(t) &\stackrel{\mathrm{def}}{=} \mathrm{ground_d}(t) \vee \mathrm{free_d}(t);\\
\mathrm{lin_d}(t) &\stackrel{\mathrm{def}}{=} \forall y \in \mathrm{vars}(t) : \mathrm{occ_lin_d}(y,t);\\
\mathrm{or_lin_d}(s,t) &\stackrel{\mathrm{def}}{=} \mathrm{lin_d}(s) \vee \mathrm{lin_d}(t);\\
\mathrm{share_same_var_d}(s,t) &\stackrel{\mathrm{def}}{=} \mathrm{vars}(\mathrm{sh}_s \cap \mathrm{sh}_t);\\
\mathrm{share_with_d}(t) &\stackrel{\mathrm{def}}{=} \mathrm{vars}(\mathrm{sh}_t).
\end{aligned}
$$

The function $\mathrm{amgu}_S\colon \mathrm{SFL} \times \mathrm{Bind} \to \mathrm{SFL}$ *captures the effects of a binding on an element of* SFL. *Let* $\mathrm{d} = \langle \mathrm{sh}, f, l\rangle \in \mathrm{SFL}$ *and* $(x \mapsto t) \in \mathrm{Bind}$, *where* $\{x\} \cup \mathrm{vars}(t) \subseteq \mathrm{VI}$. *Let also*

$$\mathrm{sh}' \stackrel{\mathrm{def}}{=} \mathrm{cyclic}_x^t(\mathrm{sh}_- \cup \mathrm{sh}''),$$

where

$$
\begin{aligned}
\mathrm{sh}_x &\stackrel{\mathrm{def}}{=} \mathrm{rel}\Big(\{x\}, \mathrm{sh}\Big), & \mathrm{sh}_t &\stackrel{\mathrm{def}}{=} \mathrm{rel}\Big(\mathrm{vars}(t), \mathrm{sh}\Big),\\
\mathrm{sh}_{xt} &\stackrel{\mathrm{def}}{=} \mathrm{sh}_x \cap \mathrm{sh}_t, & \mathrm{sh}_- &\stackrel{\mathrm{def}}{=} \overline{\mathrm{rel}}\Big(\{x\} \cup \mathrm{vars}(t), \mathrm{sh}\Big),
\end{aligned}
$$

$$
\mathrm{sh}'' \stackrel{\mathrm{def}}{=} \begin{cases}
\mathrm{bin}(\mathrm{sh}_x, \mathrm{sh}_t), & \textit{if } \mathrm{free_d}(x) \vee \mathrm{free_d}(t);\\
\mathrm{bin}\Big(\mathrm{sh}_x \cup \mathrm{bin}(\mathrm{sh}_x, \mathrm{sh}^\star_{xt}),\\
\qquad \mathrm{sh}_t \cup \mathrm{bin}(\mathrm{sh}_t, \mathrm{sh}^\star_{xt})\Big), & \textit{if } \mathrm{lin_d}(x) \wedge \mathrm{lin_d}(t);\\
\mathrm{bin}(\mathrm{sh}^\star_x, \mathrm{sh}_t), & \textit{if } \mathrm{lin_d}(x);\\
\mathrm{bin}(\mathrm{sh}_x, \mathrm{sh}^\star_t), & \textit{if } \mathrm{lin_d}(t);\\
\mathrm{bin}(\mathrm{sh}^\star_x, \mathrm{sh}^\star_t), & \textit{otherwise.}
\end{cases}
$$

Letting $S_x \stackrel{\text{def}}{=} \text{share_with}_\text{d}(x)$ *and* $S_t \stackrel{\text{def}}{=} \text{share_with}_\text{d}(t)$*, we also define*

$$f' \stackrel{\text{def}}{=} \begin{cases} f, & \textit{if } \text{free}_\text{d}(x) \wedge \text{free}_\text{d}(t); \\ f \setminus S_x, & \textit{if } \text{free}_\text{d}(x); \\ f \setminus S_t, & \textit{if } \text{free}_\text{d}(t); \\ f \setminus (S_x \cup S_t), & \textit{otherwise;} \end{cases}$$
$$l' \stackrel{\text{def}}{=} \Big(\text{VI} \setminus \text{vars}(\text{sh}')\Big) \cup f' \cup l'',$$

where

$$l'' \stackrel{\text{def}}{=} \begin{cases} l \setminus (S_x \cap S_t), & \textit{if } \text{lin}_\text{d}(x) \wedge \text{lin}_\text{d}(t); \\ l \setminus S_x, & \textit{if } \text{lin}_\text{d}(x); \\ l \setminus S_t, & \textit{if } \text{lin}_\text{d}(t); \\ l \setminus (S_x \cup S_t), & \textit{otherwise.} \end{cases}$$

Then

$$\text{amgu}_S\Big(\text{d}, x \mapsto t\Big) \stackrel{\text{def}}{=} \langle \text{sh}', f', l' \rangle.$$

The function $\text{proj}_S \colon \text{SFL} \times \text{VI} \to \text{SFL}$ *correctly captures the operation of projecting away a variable from an element of* SFL*. For each* $\text{d} \in \text{SFL}$ *and* $x \in \text{VI}$*,*

$$\text{proj}_S(\text{d}, x) \stackrel{\text{def}}{=} \begin{cases} \bot_S, & \textit{if } \text{d} = \bot_S; \\ \Big\langle \text{proj}_\text{SH}(\text{sh}, x), f \cup \{x\}, l \cup \{x\} \Big\rangle, & \textit{if } \text{d} = \langle \text{sh}, f, l \rangle \neq \bot_S. \end{cases}$$

Observe that a set-sharing domain such as SFL is strictly more precise for term finiteness information than a pair-sharing domain such as SFL_2 [32,33] (where the set-sharing component SH in SFL is replaced by the domain PSD as defined in [70,71]). To see this, consider the abstract evaluation of the binding $x \mapsto y$ and the description $\langle h, \text{d} \rangle \in H \times \text{SFL}$, where $h = \{x, y, z\}$ and $\text{d} = \langle \text{sh}, f, l \rangle$ is such that $\text{sh} = \Big\{\{x, y\}, \{x, z\}, \{y, z\}\Big\}$, $f = \varnothing$ and $l = \{x, y, z\}$. Then $z \notin \text{share_same_var}_\text{d}(x, y)$ so that we have $h' = \{z\}$. In contrast, when using a pair sharing domain such as SFL_2 the element d is equivalent to $\text{d}' = \langle \text{sh}', f, l \rangle$, where $\text{sh}' = \text{sh} \cup \Big\{\{x, y, z\}\Big\}$. Hence we have $z \in \text{share_same_var}_{\text{d}'}(x, y)$ and $h' = \varnothing$. Thus, in sh the information provided by the sharing group $\{x, y, z\}$ is redundant for the pair-sharing and groundness properties, but not redundant for term finiteness. Note that the above observation holds regardless of the pair-sharing variant considered, so that similar examples can be obtained for ASub [72,52] and $\mathsf{Sh}^{\mathsf{PSh}}$ [73].

Although the domain SFL described here is very precise and used to implement the parameter component P for computing our experimental results, it is not intended as the target of the generic specification given in Definition 7; more powerful sharing domains can also satisfy this schema, including

all the enhanced combinations considered in [28,29]. For instance, as the predicate gfree_{d} defined on SFL does not fully exploit the disjunctive nature of its generic specification gfree_p, the precision of the analysis may be improved by adding a domain component explicitly tracking ground-or-freeness, as proposed in [28,29]. The same argument applies to the predicate or_lin_{d}, with respect to or_lin_p, when considering the combination with the groundness domain Pos.