Global Analysis of Constraint Logic Programs

M. GARCIA DE LA BANDA, M. HERMENEGILDO
Universidad Politécnica de Madrid
and
M. BRUYNOOGHE, V. DUMORTIER, G. JANSSENS and W. SIMOENS
Katholieke Universiteit Leuven

This paper presents and illustrates a practical approach to the dataflow analysis of constraint logic programming languages using abstract interpretation. It is first argued that, from the framework point of view, it suffices to propose relatively simple extensions of traditional analysis methods which have already been proved useful and practical and for which efficient fixpoint algorithms exist. This is shown by proposing a simple extension of Bruynooghe's traditional framework which allows it to analyze constraint logic programs. Then, and using this generalized framework, two abstract domains and their required abstract functions are presented: the first abstract domain approximates definiteness information and the second one freeness. Finally, an approach for combining those domains is proposed. The two domains and their combination have been implemented and used in the analysis of CLP($\Re$) and Prolog-III applications. Results from this implementation showing its performance and accuracy are also presented.

Categories and Subject Descriptors: D.1.6 [**Programming Techniques**]: Logic Programming; D.1.m [**Programming Techniques**]: Miscellaneous—*constraint logic programming*; D.3.2 [**Programming Languages**]: Language Classifications—*nonprocedural languages*; D.3.4 [**Programming Languages**]: Processors—*compilers*; F.3.1 [**Logics and Meaning of Programs**]: Specifying, Verifying and Reasoning about Programs

General Terms: Languages, Program Analysis

Additional Key Words and Phrases: Abstract Interpretation, Constraint Logic Programming, Global Program Analysis

1. INTRODUCTION

The constraint logic programming (CLP) paradigm [Jaffar and Lassez 1987] is a relatively recent proposal which has emerged as the natural combination of the

This research has been funded in part by the CEC through ESPRIT project 5246 PRINCE, by the Belgium National Fund for Scientific Research, by CICYT project TIC91-0106-CE and by HCM project CHRX-CT94-0624 (ABILE).

Authors' addresses:
M. García de la Banda and M. Hermenegildo, Universidad Politécnica de Madrid, Facultad de Informática, 28660-Boadilla del Monte, Madrid, Spain; email: {maria,herme}@fi.upm.es;
M. Bruynooghe, V. Dumortier, G. Janssens and W. Simoens, Katholieke Universiteit Leuven, Department of Computer Science, Celestijnenlaan 200A, 3001 Heverlee, Belgium; email: {maurice,veroniek,gerda,wims}@cs.kuleuven.ac.be.

Permission to make digital/hard copy of all or part of this material without fee is granted provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery, Inc. (ACM). To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

constraint solving and logic programming paradigms. This combination enhances the flexibility and expressiveness of conventional logic languages. In this context, traditional logic programming (LP) can be seen as an instance of CLP in which constraints are equations over terms and the constraint solving is done by the well known unification algorithm.

One of the main advantages of CLP languages is that they allow the programmer to specify the problem in a short, simple, and declarative way by means of high-level constraints; leaving the details of how these constraints are to be solved to the underlying constraint solver. When the execution of the program requires the full capabilities of the solver, the resulting efficiency is often quite good, in the sense that it would only be achievable in another language after an extensive and tedious programming effort. However, in the cases in which a simpler solver would suffice, the expressive power is paid in terms of efficiency. As it has recently been shown, efficiency can be recovered by performing several compile-time optimizations, mainly aimed at automatically specializing the program in order to reduce as much as possible the use of the general solver [Jørgensen et al. 1991; Jaffar et al. 1992; Marriott and Stuckey 1993; Jaffar and Maher 1994; Marriott et al. 1994; Dumortier 1994; García de la Banda 1994]. The significant speed-ups promised by these optimizations, and the fact that they need quite accurate compile-time information regarding the characteristics of the program, have motivated a growing interest in dataflow analysis of CLP languages and, in particular, in the application of abstract interpretation [Cousot and Cousot 1977].

Much work has been done using the abstract interpretation technique in the context of LP (e.g. [Mellish 1986; Debray 1989; Bruynooghe 1991; Marriott and Søndergaard 1989; Debray 1992b]). A number of systems have been built, some of which have shown the potential usefulness and practicality of this technique [Van Roy and Despain 1992; Warren et al. 1988; Muthukumar and Hermenegildo 1992; Debray 1992b; Le Charlier and Van Hentenryck 1994; Bueno et al. 1994]. Thus, it is natural to expect that this technique should also be useful in the context of CLP. A few general frameworks have already been defined for this purpose [Marriott and Søndergaard 1990; Codognet and Filé 1992; Bruynooghe and Janssens 1992; Giacobazzi et al. 1993]. However, one common characteristic of these frameworks is that they are either not implementation oriented or depart from the approaches that have been so far quite successful in the analysis of traditional logic programming (LP) languages.

This paper shows how some of the LP-based techniques already developed and implemented, can relatively easily be extended to the analysis of CLP programs. This point is illustrated by proposing a simple but quite powerful extension of Bruynooghe's traditional framework in order to make it applicable to the analysis of CLP programs. We also extend the framework to deal with passive constraints. Finally, we give correctness conditions for the resulting framework. The generalized description represents a fully specified algorithm for analysis of CLP programs.

Then, and using this generalized framework, two abstract domains and their required abstract functions are described. The abstract domain $\mathbf{Cons}^{\mathcal{D}}$ determines whether program variables are *definite*, i.e. constrained to a unique value. The abstract domain $\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$ determines whether program variables are *free*, i.e. whether they can still take any possible value (at least according to their type, e.g. a variable that is constrained to be numerical but still ranges over the complete domain of numbers is considered as free). Finally, an approach for combining those domains is proposed. The idea is to use the definiteness information provided by $\mathbf{Cons}^{\mathcal{D}}$ to obtain a more compact and efficient freeness abstraction while maintaining the precision of the original freeness abstraction. This combination leads to a *full* mode inference system which is, to the authors' knowledge, the first full mode system proposed for CLP.

The two abstract domains and their combination have been implemented within the abstract interpretation system PLAI [Muthukumar and Hermenegildo 1990; 1992]. This system is based on the framework of [Bruynooghe 1991], optimized with the specialized domain-independent fixpoint defined in [Muthukumar and Hermenegildo 1992], and generalized to support analysis of practical CLP languages, following the guidelines presented in this paper. Results from this implementation showing its performance and accuracy are also presented.

Parts of the work in this paper have already been presented previously. The generalization of abstract interpretation of LP towards CLP has been discussed in [García de la Banda and Hermenegildo 1993; Bruynooghe and Boulanger 1994]. A description of the definiteness analysis can be found in [García de la Banda and Hermenegildo 1993; García de la Banda 1994]. The freeness analysis and its optimizations have been described in [Dumortier et al. 1993; Dumortier and Janssens 1994; Dumortier 1994]. The paper [Dumortier and Janssens 1994] also explains how definiteness information can be exploited in order to improve the freeness abstraction.

2. BACKGROUND AND NOTATION

In this section we present some basic concepts of constraint logic programming, and abstract interpretation, as well as the notation which will be used throughout the paper. In doing this, we will follow mainly [Jaffar and Lassez 1987; Jaffar and Maher 1994] and [Cousot and Cousot 1977].

2.1 Constraint domains and programs

First, we introduce some notational conventions. Upper case letters generally denote collections of objects, while lower case letters generally denote individual objects. $\mathbf{u}, \mathbf{v}, \mathbf{w}, \mathbf{x}, \mathbf{y}, \mathbf{z}$ will denote variables, $\mathbf{t}$ will denote a term, $\mathbf{p}, \mathbf{q}$ will denote predicate symbols, $\mathbf{f}$ will denote a function symbol, $\mathbf{a}, \mathbf{h}$ will denote atoms, $\mathbf{c}$ will denote a constraint, ϵ will denote the empty constraint, $\mathbf{b}, \mathbf{g}$ will denote an atom or a constraint, ρ will denote a rule, $\mathbf{P}, \mathbf{Q}$ will denote programs and $\mathbf{B}, \mathbf{G}$ will denote goals, i.e. sequences of atoms and constraints. These symbols may be subscripted. $\tilde{\mathbf{x}}$ denotes a sequence of distinct variables and, by abuse of notation, also the corresponding set of variables. $\mathbf{vars}(\mathbf{o})$ denotes the set of variables occurring in the syntactic object $\mathbf{o}$. Finally, $\exists_{-\tilde{\mathbf{x}}}\phi$ denotes the existential closure of the formula ϕ except for the variables $\tilde{\mathbf{x}}$, and $\tilde{\exists}\phi$ denotes the full existential closure of the formula ϕ.

As an example of a simple CLP program, consider the following, adapted from [Jaffar and Maher 1994]: $\mathbf{sumto}(\mathbf{x}, \mathbf{y})$ expresses that $\mathbf{y}$ is the sum of the first $\mathbf{x}$ natural numbers.

$$\mathbf{sumto}(0,0).$$
$$\mathbf{sumto}(\mathbf{x},\mathbf{y}) \text{ :- } 1 \leq \mathbf{x}, \mathbf{x} \leq \mathbf{y}, \mathbf{x}' = \mathbf{x} - 1, \mathbf{y}' = \mathbf{y} - \mathbf{x}, \mathbf{sumto}(\mathbf{x}',\mathbf{y}').$$

This simple program can be used to compute $\mathbf{y}$ from $\mathbf{x}$ (e.g. $\mathbf{sumto}(3,\mathbf{y})$ returns $\mathbf{y} = 6$ and then terminates), compute the $\mathbf{x}$ from $\mathbf{y}$ (e.g. $\mathbf{sumto}(\mathbf{x},15)$ returns $\mathbf{x} = 5$ and then terminates), test whether a given $\mathbf{x}$ and $\mathbf{y}$ satisfy the relationship (e.g. $\mathbf{sumto}(5,15)$ succeeds and terminates, $\mathbf{sumto}(3,15)$ fails), or to answer more complex queries like ?- $\mathbf{y} \leq 3, \mathbf{sumto}(\mathbf{x},\mathbf{y})$ which gives rise to three answers $(\mathbf{x} = 0, \mathbf{y} = 0)$, $(\mathbf{x} = 1, \mathbf{y} = 1)$ and $(\mathbf{x} = 2, \mathbf{y} = 3)$ and then terminates. A direct translation of the above program into Prolog would require transforming each arithmetic equality into the `is/2` Prolog builtin:

$$\mathbf{sumto}(\mathbf{x},\mathbf{y}) \text{ :- } 1 \leq \mathbf{x}, \mathbf{x} \leq \mathbf{y}, \mathbf{x}' \textbf{ is } \mathbf{x} - 1, \mathbf{y} \textbf{ is } \mathbf{y}' + \mathbf{x}, \mathbf{sumto}(\mathbf{x}',\mathbf{y}').$$

However, since the Prolog arithmetic builtins $\mathbf{is}/2$ and $\leq /2$ require their second and both arguments, respectively, to be bound to a numerical value at run-time, the Prolog program would only execute queries in which both input arguments are constrained to a unique numerical value, such as ?- $\mathbf{sumto}(2,3)$ and ?- $\mathbf{sumto}(2,5)$. Carefully rewriting the second rule as

$$\mathbf{sumto}(\mathbf{x},\mathbf{y}) \text{ :- } 1 \leq \mathbf{x}, \mathbf{x}' \textbf{ is } \mathbf{x} - 1, \mathbf{sumto}(\mathbf{x}',\mathbf{y}'), \mathbf{y} \textbf{ is } \mathbf{y}' + \mathbf{x}, \mathbf{x} \leq \mathbf{y}.$$

will also allow executing queries in which the second input argument is an unconstrained variable, such as ?- $\mathbf{sumto}(3,\mathbf{y})$. While less general than the CLP program, this Prolog program will execute quite fast, because it is performing only simple arithmetic operations. Note also that rewriting the second rule as

$$\mathbf{sumto}(\mathbf{x},\mathbf{y}) \text{ :- } \mathbf{sumto}(\mathbf{x}',\mathbf{y}'), \mathbf{x} \textbf{ is } \mathbf{x}' + 1, \mathbf{y} \textbf{ is } \mathbf{y}' + \mathbf{x}, 1 \leq \mathbf{x}, \mathbf{x} \leq \mathbf{y}.$$

the query ?- $\mathbf{sumto}(3,\mathbf{y})$ will produce the answer $\mathbf{y} = 6$ but then go into an infinite loop. The same happens, after producing the three answers, for the query ?- $\mathbf{sumto}(\mathbf{x},\mathbf{y}), \mathbf{y} \leq 3$. In summary, although the functionality of the simple and elegant CLP program can only be obtained in Prolog by a more complex case-by-case program, the resulting Prolog program would probably execute faster than its CLP counterpart.

The example illustrates the expressiveness of CLP programming but also the challenge for the implementors. Ideally, it would be desirable for CLP systems to be strict generalizations of LP systems, not only from a functional point of view, but also from a performance point of view. I.e., in our example the general CLP program should offer a performance comparable with that of Prolog for queries such as ?- $\mathbf{sumto}(3,\mathbf{y})$. To achieve this, some form of program analysis and creation of code dedicated to queries in this class looks unavoidable. Furthermore, one would also like the program to perform as well as possible even when actual constraint solving is being performed by the program. As mentioned in the introduction, it has been shown that such optimizations often require information from global analysis.

In the example, we have a constraint domain that is based on the constants 0, 1, the function + and the predicates $=, <, \leq$ (3 is syntactic sugar for $1 + 1 + 1$, $\mathbf{x}' = \mathbf{x} - 1$ is syntactic sugar for $\mathbf{x} = \mathbf{x}' + 1$). In general, constants, functions and predicates make up the signature Σ underlying the constraint domain. The so called Σ-structure $\mathcal{D}$ consists of a domain, e.g. the domain of the real numbers, and an interpretation of constants, functions and predicates over this domain, e.g.

the standard arithmetic of the reals. A *primitive constraint* such as $1 \leq \mathbf{x}$ is built from a predicate in Σ and terms built from constants and functions in Σ and from variables. Using logical connectives and quantifiers, primitive constraints can be combined into expressions of a language $\mathcal{L}$, called *constraints.* The pair $(\mathcal{D}, \mathcal{L})$ defines the constraint domain. The interested reader should consult [Jaffar and Maher 1994] for a more formal and more detailed account, as well as for the assumptions that are usually made about $(\mathcal{D}, \mathcal{L})$.

As in the example above, a CLP program is a collection of rules of the form $\mathbf{h}$:- $\mathbf{B}$, where $\mathbf{h}$ (the *head*) is an atom built from a predicate (not in Σ), and $\mathbf{B}$ (the *body*) is a sequence $\mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}$ of atoms (not in Σ) and constraints. A goal $\mathbf{G}$ is also any sequence of atoms and constraints.

In [Jaffar and Maher 1994] four relevant operations on constraints are mentioned, the first one being almost obligatory in any implementation of a CLP language:

(1) *Consistency* or *satisfiability* of a constraint $\mathbf{c}$: $\mathcal{D} \models \tilde{\exists}\mathbf{c}$.
(2) *Implication* or *entailment* of a constraint $\mathbf{c}_1$ by another constraint $\mathbf{c}_0$: $\mathcal{D} \models \mathbf{c}_0 \rightarrow \mathbf{c}_1$.
(3) *Projection* of a constraint $\mathbf{c}$ onto variables $\tilde{\mathbf{x}}$: $\mathcal{D} \models \exists_{-\tilde{\mathbf{x}}}\mathbf{c}$.
(4) Detection that, given a constraint $\mathbf{c}$, there is only one value that a variable $\mathbf{x}$ can take that is consistent with $\mathbf{c}$: $\mathcal{D} \models \mathbf{c}(\mathbf{x}_1, \tilde{\mathbf{y}}) \wedge \mathbf{c}(\mathbf{x}_2, \tilde{\mathbf{y}}) \rightarrow \mathbf{x}_1 = \mathbf{x}_2$. We say that $\mathbf{x}$ is **definite** in $\mathbf{c}$ and denote by $\mathbf{def}(\mathbf{c})$ the set of definite variables in $\mathbf{c}$.

2.2 CLP operational semantics

In [Jaffar and Maher 1994], the interested reader can find a very general operational semantics which takes passive constraints into account, separates the generation of new constraints from the consistency check of the constraints accumulated in the constraint store, and is not tailored to any particular computation rule.

The work reported here concerns the analysis of programs under the widely used left-to-right computation rule (as in Prolog). In the first part of this paper, we focus on programs without passive (i.e. delayed) constraints. The treatment of passive constraints is deferred to Section 5.3. Another assumption is that the considered systems are quick-checking [Jaffar and Maher 1994], i.e. the addition of new constraints is immediately followed by a consistency check of the constraint store.

Under these simplifications, the state of the computation can be described by a pair $\langle \mathbf{G}; \mathbf{c} \rangle$, where $\mathbf{G}$ is the sequence of constraints and atoms yet to be executed and $\mathbf{c}$ is the constraint store containing the constraints accumulated so far. The operational behavior of a program can be described by a set of sequences of states (*SLD sequences*), each sequence starting with $\langle \mathbf{G}; \mathbf{true} \rangle$ where $\mathbf{G}$ is the query.

Such SLD sequences are manipulated by transitions whose behavior – given the left-to-right computation rule – is determined by the leftmost element in the goal of the last state in the sequence. Formally, an incomplete SLD sequence ending in a consistent state can be extended by the following transitions which are formulated as rewrite rules ($\mathbf{S}$ represents an SLD sequence, :: is used to concatenate SLD sequences):

—$\mathbf{S} :: \langle \mathbf{c}', \mathbf{G}; \mathbf{c} \rangle$
$\xrightarrow{\mathbf{c}} \mathbf{S} :: \langle \mathbf{c}', \mathbf{G}; \mathbf{c} \rangle :: \langle \mathbf{G}; \mathbf{c}' \wedge \mathbf{c} \rangle$ if $\mathbf{consistent}(\mathbf{c}' \wedge \mathbf{c})$ or

$\xrightarrow{\mathbf{c}}$ $\mathbf{S} :: \langle \mathbf{c'}, \mathbf{G}; \mathbf{c} \rangle :: \langle \mathbf{G}; \mathbf{false} \rangle$ if $\mathbf{inconsistent}(\mathbf{c'} \wedge \mathbf{c})$

—$\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle$ rewrites to a set of sequences, one for each rule of the program defining the predicate symbol of $\mathbf{a}$. Let $\rho : \mathbf{h}$:-$\mathbf{b}_1, \ldots, \mathbf{b_n}$ be such a (properly renamed) rule, then
$\xrightarrow{\mathbf{r}}$ $\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b_n}, \mathbf{G}; \mathbf{a} = \mathbf{h} \wedge \mathbf{c} \rangle$ if $\mathbf{consistent}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c})^1$ or
$\xrightarrow{\mathbf{r}}$ $\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b_n}, \mathbf{G}; \mathbf{false} \rangle$ if $\mathbf{inconsistent}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c})$

Note that an SLD sequence is very similar to a partial SLD derivation in [Lloyd 1987]. In particular, an SLD sequence represents a *complete* derivation if its last state

—contains a goal whose leftmost atom has a predicate symbol for which the program has no rules (a failed SLD derivation);

—has *false* in its constraint store (also a failed SLD derivation, but this case is distinguished from the previous one because some analyses can be interested in the question of at which points an inconsistent store can be introduced; such analyses will use abstractions that can distinguish an inconsistent store from consistent ones);

—contains an empty goal (a successful SLD derivation); the constraint store then provides the answer[2].

Ignoring the search rule (see [Le Charlier et al. 1994] for a framework taking the search rule into account), the operational semantics is given by the fixpoint of the operator which applies the above transitions on incomplete SLD sequences, starting from the initial sequence $\langle \mathbf{G}; \mathbf{true} \rangle$. (Alternatively, the set of all complete sequences can be collected in an SLD tree as in [Lloyd 1987]). The fixpoint of the operator – a set of complete SLD sequences – represents the operational semantics as it describes in sufficient detail the behavior of the program for the analyses considered in this paper. If desired, sequences could be instrumented with more detail (e.g. [Mulkers et al. 1994]).

2.3 Abstract interpretation

The most familiar framework for abstract interpretation is defined in terms of Galois connections and Galois insertions [Cousot and Cousot 1977; 1992a].

Definition 2.3.1. Galois connection
A Galois connection is a quadruple $(\mathbf{Dom}^{\mathcal{C}}, \alpha, \mathbf{Dom}^{\mathcal{A}}, \gamma)$ where:

(1) $(\mathbf{Dom}^{\mathcal{C}}, \leq^{\mathcal{C}})$ and $(\mathbf{Dom}^{\mathcal{A}}, \leq^{\mathcal{A}})$ are posets called *concrete* and *abstract domains* respectively;

(2) $\alpha : \mathbf{Dom}^{\mathcal{C}} \to \mathbf{Dom}^{\mathcal{A}}$ and $\gamma : \mathbf{Dom}^{\mathcal{A}} \to \mathbf{Dom}^{\mathcal{C}}$ are functions called *abstraction* and *concretization functions* respectively, satisfying that for every $\mathbf{d}^{\mathcal{A}} \in \mathbf{Dom}^{\mathcal{A}}$ and $\mathbf{d}^{\mathcal{C}} \in \mathbf{Dom}^{\mathcal{C}}$, $\alpha(\mathbf{d}^{\mathcal{C}}) \leq^{\mathcal{A}} \mathbf{d}^{\mathcal{A}}$ iff $\mathbf{d}^{\mathcal{C}} \leq^{\mathcal{C}} \gamma(\mathbf{d}^{\mathcal{A}})$.

[1] The label ρ on :: identifies the renamed rule used in solving $\mathbf{a}$. The expression $\mathbf{a} = \mathbf{h}$ is an abbreviation for the conjunction of the corresponding primitive equations.
[2] The answer can be conditional in case one allows passive constraints, as then the store may be inconsistent.

Definition 2.3.2. Galois insertion
A Galois insertion is a Galois connection satisfying $\alpha(\gamma(\mathbf{d}^{\mathcal{A}})) = \mathbf{d}^{\mathcal{A}}$.

The Galois connection corresponds to a perfect situation where each concrete property has a unique best abstract approximation. Thus, only one of $\{ \alpha, \gamma \}$ needs to be specified since if one exists, the other is determined by the properties of the definition. In addition a Galois insertion has no superfluous elements in the abstract domain. The following specifies the notion of approximation (in terms of γ) which is then extended from the primitive domains to function domains:

Definition 2.3.3. Approximation
Let $(\mathbf{Dom}^{\mathcal{C}}, \alpha, \mathbf{Dom}^{\mathcal{A}}, \gamma)$ be a Galois insertion and let $\mu^{\mathcal{C}} : \mathbf{Dom}^{\mathcal{C}} \rightarrow \mathbf{Dom}^{\mathcal{C}}$ and $\mu^{\mathcal{A}} : \mathbf{Dom}^{\mathcal{A}} \rightarrow \mathbf{Dom}^{\mathcal{A}}$ be monotonic functions. We say that $\mathbf{d}^{\mathcal{A}} \in \mathbf{Dom}^{\mathcal{A}}$ *γ-approximates* $\mathbf{d}^{\mathcal{C}} \in \mathbf{Dom}^{\mathcal{C}}$, denoted $\mathbf{d}^{\mathcal{A}} \propto_{\gamma} \mathbf{d}^{\mathcal{C}}$, if $\mathbf{d}^{\mathcal{C}} \leq^{\mathcal{C}} \gamma(\mathbf{d}^{\mathcal{A}})$. We say that $\mu^{\mathcal{A}}$ *γ-approximates* $\mu^{\mathcal{C}}$, denoted $\mu^{\mathcal{A}} \propto_{\gamma} \mu^{\mathcal{C}}$, if for every $\mathbf{d}^{\mathcal{A}} \in \mathbf{Dom}^{\mathcal{A}}$, $\mathbf{d}^{\mathcal{C}} \in \mathbf{Dom}^{\mathcal{C}}$ such that $\mathbf{d}^{\mathcal{A}} \propto_{\gamma} \mathbf{d}^{\mathcal{C}}$ then $\mu^{\mathcal{A}}(\mathbf{d}^{\mathcal{A}}) \propto_{\gamma} \mu^{\mathcal{C}}(\mathbf{d}^{\mathcal{C}})$.

As illustrated in Section 2.2, the information of interest about a program – in our case the operational semantics – can often be expressed as the least fixpoint of a function. Formally one writes $[[\mathbf{P}]] = \mathbf{lfp}(\mu^{\mathcal{C}})$ where $\mu^{\mathcal{C}} : \mathbf{Dom}^{\mathcal{C}} \rightarrow \mathbf{Dom}^{\mathcal{C}}$ is a monotonic operator on the concrete domain $\mathbf{Dom}^{\mathcal{C}}$ and $[[\mathbf{P}]]$ expresses the meaning of the program. Such a formalization provides the foundation for an abstract interpretation of the program. By introducing an appropriate Galois insertion $(\mathbf{Dom}^{\mathcal{C}}, \alpha, \mathbf{Dom}^{\mathcal{A}}, \gamma)$ and defining a monotonic function $\mu^{\mathcal{A}} : \mathbf{Dom}^{\mathcal{A}} \rightarrow \mathbf{Dom}^{\mathcal{A}}$, which approximates $\mu^{\mathcal{C}}$ and whose fixpoint can be computed or approximated by a finite computation, one can obtain information about the least fixpoint of $\mu^{\mathcal{C}}$. This is expressed by the following result [Cousot and Cousot 1992a]:

Theorem 2.3.4.
Let $(\mathbf{Dom}^{\mathcal{C}}, \alpha, \mathbf{Dom}^{\mathcal{A}}, \gamma)$ *be a Galois insertion and let* $\mu^{\mathcal{C}} : \mathbf{Dom}^{\mathcal{C}} \rightarrow \mathbf{Dom}^{\mathcal{C}}$ *and* $\mu^{\mathcal{A}} : \mathbf{Dom}^{\mathcal{A}} \rightarrow \mathbf{Dom}^{\mathcal{A}}$ *be monotonic functions such that* $\mu^{\mathcal{A}} \propto_{\gamma} \mu^{\mathcal{C}}$. *Then* $\mathbf{lfp}(\mu^{\mathcal{A}}) \propto_{\gamma} \mathbf{lfp}(\mu^{\mathcal{C}})$.

The construction of $\mu^{\mathcal{A}}$ often takes a systematic approach which involves replacing the basic operations in the concrete semantics operator $\mu^{\mathcal{C}}$ by the corresponding abstract operations in $\mu^{\mathcal{A}}$ (e.g. [Cousot and Cousot 1992a; Nielson 1988]). Given that the basic abstract operations approximate their concrete counterparts, it is generally straightforward to prove that $\mu^{\mathcal{A}}$ approximates $\mu^{\mathcal{C}}$.

3. TOWARDS A CLP ANALYSIS FRAMEWORK

There has been considerable interest in developing new abstract interpretation frameworks for CLP languages. To these authors' knowledge, at least four frameworks have been proposed previously or simultaneously with our work.[3] Marriott and Sondergaard [1990] present a general and elegant, semantics based framework. It is based on a definition-independent meta-language which can express the semantics of a wide variety of programming languages, including CLP languages. However, from a practical point of view, this framework does not provide much

[3] The ideas illustrated in this paper were first presented at the ICLP'91 Workshop on Constraint Logic Programming.

simplification to the developer of the abstract interpretation system, in the sense that many issues are left open.

In fact, one of the advantages of the most popular methods used in the analysis of conventional LP systems (for example Bruynooghe's method [Bruynooghe 1991] and the optimizations proposed for it [Muthukumar and Hermenegildo 1992]) is that they are "generic," in the sense that they specify much of what is needed leaving only the definition of the domain, domain dependent functions, and assurance of correctness criteria to be provided by the implementor. It is our intention to develop a framework for CLP program analysis at this level of specification.

Codognet and Filé [1992] also present a quite general framework for the description of both CLP languages and their static analyses, and an implementation approach. Although more concrete, this proposal is still more abstract than the level pointed out above as our objective. On the other hand this paper introduces the quite interesting idea of implementing the abstract functions actually using constraint solvers, to which we will return later.

The paper [Giacobazzi et al. 1993] formulates a general algebraic framework for constraint logic programming. It formulates the operational and fixpoint semantics within this framework and shows that abstract interpretation is simply another instance of the general framework which safely approximates the instance given by the concrete constraint system. Also, this work is in fairly general terms and does not offer much to the application developer.

Finally, Bruynooghe and Janssens [1992] present a specialized framework (which was developed in parallel with the proposal presented in this paper) which is based on the idea of adding complexity to the framework with the potential benefit of decreased complexity in the abstract domain. This is done by incorporating a local form of "suspension" so that some goals can be reconsidered if later execution in a different environment can provide further information. This extension is based on a particular view of the execution of a CLP program in which constraints are considered as goals which can suspend depending on the state of their arguments and on the particular constraint system.

The view of constraints as suspended goals could be interesting and worth pursuing. However this makes it more difficult to make the framework fully general. We prefer to take the more traditional notion presented in the CLP scheme (as introduced in the previous sections) in which constraints take the place of substitutions and goals always either succeed or fail, in the former case possibly *placing* new constraints[4].

One of the main points of this paper is to show that standard abstract interpretation frameworks for logic programs are useful for the analysis of constraint logic programs, provided the parts that relate to the abstraction of the Herbrand domain and unification functions are suitably generalized. Indeed, in this traditional view of CLP the role of goals and their control are basically identical to those in tradi-

[4] In fact, actual suspension, as is often used in the solving of non-linear arithmetic constraints or in programs with explicit coroutining can also be modeled in this way. However, we propose treating actual suspension directly using techniques such as those proposed for analyzing programs with delay declarations [Hanus 1993; Marriott et al. 1994]. This issue is discussed further in Section 5.3.

tional LP systems, the differences being essentially limited to replacing the notions of Herbrand domain, unification, and substitutions by those of constraint system, conjunction, and constraints.

In particular, we argue that the traditional framework of Bruynooghe and its extensions can be used for analyzing constraint logic programs by using the notions of abstract constraint and abstract conjunction and reformulating the safety conditions, but keeping the construction of the AND-OR graph, the implementation and optimizations of the fixpoint algorithm, the notions of projection and extension, etc. This has the advantage that the implementations based on this scheme or derivations thereof can be applied to CLP systems provided the safety conditions and other related requirements proposed herein are observed.

4. MODIFYING THE CLP OPERATIONAL SEMANTICS

The states appearing in the fixpoint of the concrete operational semantics are of the form $\langle \mathbf{g}, \mathbf{G}; \mathbf{c} \rangle$ where $\mathbf{c}$ is a constraint store over an unbounded number of variables. A basic insight underlying the framework of [Bruynooghe 1991] is that when optimizing a particular predicate, most optimizations only need information about the variables in the clauses defining such predicate. Therefore, when analyzing $\mathbf{g}$, the analysis is not interested in the properties of all program variables but only of the variables of the clause $\mathbf{g}$ belongs to. This information is collected by a slightly different operational semantics which is called LSLD (Local SLD) in [Bruynooghe and Boulanger 1994]. In our constraint setting, we can rephrase LSLD as an operator on LSLD sequences as follows:

—The $\mathbf{c}$-transition on $\mathbf{S} :: \langle \mathbf{c}', \mathbf{G}; \mathbf{c} \rangle$ is as before.

—The $\mathbf{r}$-transition on $\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle$ for **consistent**$(\mathbf{a} = \mathbf{h} \wedge \mathbf{c})$ becomes:
$\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \xrightarrow{\mathbf{r}} \mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}; \exists_{-\mathbf{vars}(\rho)}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c}) \rangle$, where ρ : $\mathbf{h}$:-$\mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}$.
The $\mathbf{r}$-transition for **inconsistent**$(\mathbf{a} = \mathbf{h} \wedge \mathbf{c})$ is unmodified.

Because the $\mathbf{r}$-transition computes a constraint store over the variables of ρ, it is called the *entry transition* in the future.

—In addition, an *exit transition* is introduced for states where the goal is the empty left-over of the body of a (uniquely renamed) rule ρ : $\mathbf{h}$:-$\mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}$ (denoted $\Box_\rho$). Note that the transition is not only based on the last state in the sequence, but also on the state prior to the application of the entry transition using ρ (marked by $\overset{\rho}{::}$):
$\mathbf{S}_1 :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}; \mathbf{c}_\mathbf{in} \rangle :: \mathbf{S}_2 :: \langle \Box_\rho; \mathbf{c}_\mathbf{out} \rangle \xrightarrow{\mathbf{exit}}$
$\mathbf{S}_1 :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b}_\mathbf{n}; \mathbf{c}_\mathbf{in} \rangle :: \mathbf{S}_2 :: \langle \Box_\rho; \mathbf{c}_\mathbf{out} \rangle :: \langle \mathbf{G}; \exists_{-\mathbf{vars}(\rho_0)}(\mathbf{c} \wedge \mathbf{a} = \mathbf{h} \wedge \mathbf{c}_\mathbf{out}) \rangle$.

ρ_0 is the rule containing $\mathbf{a}, \mathbf{G}$ as tail of its body. Note that, due to the renaming, there is a unique state $\langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle$ to which an entry transition using ρ was applied. Note also that there exists a constraint $\mathbf{c}_\mathbf{new}$ such that $\mathbf{c}_\mathbf{out} = \mathbf{c}_\mathbf{in} \wedge \mathbf{c}_\mathbf{new}$.

As before, the initial sequence is $\langle \mathbf{G}; \mathbf{true} \rangle$, with $\mathbf{G}$ being the query, and the operational semantics is given by the fixpoint of the operator applying transitions on incomplete sequences. Though the exit transitions introduce extra states

$\langle \square_\rho; \mathbf{c_{out}} \rangle$ in LSLD sequences, there is a strong equivalence between SLD sequences and LSLD sequences using the same renamed rules in the same order: for every state $\langle \mathbf{b}_1, \ldots, \mathbf{b_n}, \mathbf{G}; \mathbf{c} \rangle$ in an SLD sequence with $\mathbf{b}_1, \ldots, \mathbf{b_n}$ the tail of some renamed rule ρ, there is a state $\langle \mathbf{b}_1, \ldots, \mathbf{b_n}; \exists_{-\mathbf{vars}(\rho)} \mathbf{c} \rangle$ in the corresponding LSLD sequence. This can be proved by induction. Consequently, the fixpoint of the LSLD operator carries the same amount of relevant information (i.e. what are the properties of $\mathbf{vars}(\mathbf{b_i})$ of a state $\langle \mathbf{b_i}, \ldots; \mathbf{c} \rangle$) as the fixpoint of the original SLD operator.

An SLD sequence can be represented by an AND-tree (a proof tree, to be distinguished from an SLD tree which is a search tree). The children of the root are the atoms and constraints of the query. An atom $\mathbf{a}$ is paired with the head of the rule $\rho : \mathbf{h}\text{:-}\mathbf{b}_1, \ldots, \mathbf{b_n}$ that is used by the entry transition on $\mathbf{a}$ (the sequence contains $\langle \mathbf{a}, \ldots; \ldots \rangle \overset{\rho}{::} \langle \mathbf{b}_1, \ldots, \mathbf{b_n}; \mathbf{c}_1 \rangle$). The constraint store adorns the tree as shown in the fragment of Figure 1; $\mathbf{c_i}$ is the constraint store of the state $\langle \mathbf{b_i}, \ldots, \mathbf{b_n}; \mathbf{c_i} \rangle$. It contains the information about the variables of $\mathbf{b_i}$ at the point where $\mathbf{b_i}$ is to be processed. As described in [Bruynooghe 1991], the set of all AND-trees, which represents the operational semantics of the program, can be collected in an AND-OR tree where nodes are adorned with sets of constraint stores (this gives the collecting semantics). Using a tabulation technique, repeated computations can be avoided: there is no point in collecting states which are renamings of each other; therefore, states are tabled with $\langle \mathbf{b}_1, \ldots, \mathbf{b_n}; \mathbf{c}_1 \rangle$ as *key* and the corresponding $\langle \square_\rho; \mathbf{c_{out}} \rangle$ as *answer*. A sequence ending in a tabled state is extended with an exit operation which uses the tabled answer, thus avoiding the construction of a renaming of an already existing subsequence. The LSLD semantics is thus transformed into the LSLDT semantics ([Bruynooghe and Boulanger 1994]).

Tabulation allows abstracting the AND-OR tree, representing the concrete collecting semantics, by an AND-OR graph. In practice, however, abstract interpretation systems such as PLAI [Muthukumar and Hermenegildo 1990; 1992] and GAIA [Englebert et al. 1992; Le Charlier et al. 1991; Le Charlier and Van Hentenryck 1994] are based on a variant of the above tabulation technique, where the stored key is not $\langle \mathbf{b}_1, \ldots, \mathbf{b_n}; \mathbf{c}_1 \rangle$. Instead, with $\langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle$ as preceding state, $\langle \mathbf{a}; \exists_{-\mathbf{vars(a)}} \mathbf{c} \rangle$ is stored as key and $\exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{out}})$ as answer for an atom $\mathbf{a}$ if $\langle \square_\rho; \mathbf{c_{out}} \rangle$ has occurred when resolving $\mathbf{a}$ with ρ. If a state $\langle \mathbf{a'}, \mathbf{G'}; \mathbf{c'} \rangle$ is met such that $\langle \mathbf{a'}; \exists_{-\mathbf{vars(a')}} \mathbf{c'} \rangle$ is a renaming of $\langle \mathbf{a}; \exists_{-\mathbf{vars(a)}} \mathbf{c} \rangle$, then no entry transition with a renaming of ρ is performed. Instead, the sequence is extended with a state $\langle \mathbf{G'}; \exists_{-\mathbf{vars}(\rho'_0)}(\mathbf{c'} \wedge \mathbf{a'} = \mathbf{a} \wedge \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{out}})) \rangle$ for each tabled answer $\exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{out}}))$ (Table look-up transition), where ρ'_0 is the rule with $\mathbf{a'}, \mathbf{G'}$ as tail of its body, and $\mathbf{a'} = \mathbf{a}$ performs renaming. The advantage of this tabulation variant is to avoid an entry transition. However, some table look-ups can be missed because different states $\langle \mathbf{a}, \ldots; \mathbf{c} \rangle$ can give rise to $\langle \mathbf{b}_1, \ldots, \mathbf{b_n}; \mathbf{c_{in}} \rangle$ that are renamings of each other (and extra work will be done: a look-up transition for every atom $\mathbf{b_i}$, and a $\mathbf{c}$-transition for every constraint $\mathbf{b_i}$). With so-called *semi-normalized* programs where calls have the form $\mathbf{p}(\mathbf{x}_1, \ldots, \mathbf{x_n})$ (all $\mathbf{x_i}$ different), the disadvantage disappears. With the heads also of the form $\mathbf{p}(\mathbf{x}_1, \ldots, \mathbf{x_n})$ (*normalized* programs), $\exists_{-\mathbf{vars}(\rho)}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c})$ and $\exists_{-\mathbf{vars}(\rho_0)}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{out}})$, where $\mathbf{h}\text{:-}\ldots$ is used to resolve $\mathbf{a}$, reduce to simple renaming operations. The price for (semi-)normalization is that

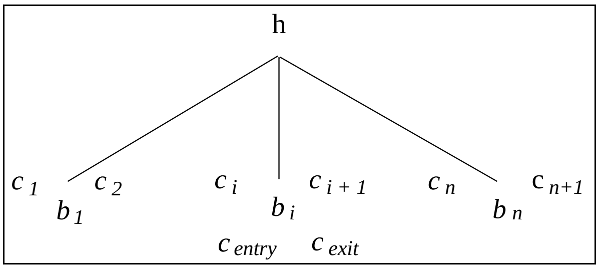

Fig. 1. Naming conventions for constraints

there are more constraints in rule bodies, and, more importantly, more variables. The latter can have a significant effect in applications where the size of an element in the abstract domain can be exponential in the number of rule variables. Also, for some applications, (semi-)normalization may result in loss of precision.

Within the proposed LSLD semantics, it is convenient to name constraint stores differently, depending on the point in a rule to which they correspond. The same conventions will be used for the abstract constraint stores. Consider, for example, the rule $\mathbf{h}$:- $\mathbf{b}_1, \cdots, \mathbf{b_n}$. Let $\mathbf{c_i}$ and $\mathbf{c_{i+1}}$ be the constraint stores to the left and right of the subgoal $\mathbf{b_i}, 1 \leq \mathbf{i} \leq \mathbf{n}$ in this rule. See Figure 1.

—$\mathbf{c_i}$ and $\mathbf{c_{i+1}}$ are, respectively, the *call constraint* and the *success constraint* for $\mathbf{b_i}$.

—$\mathbf{c}_1$ and $\mathbf{c_{n+1}}$ are, respectively, the *in constraint* and the *out constraint* of the rule (also denoted by $\mathbf{c_{in}}$ and $\mathbf{c_{out}}$). Note that $\mathbf{c}_1$ and $\mathbf{c_{n+1}}$ are also the call constraint for $\mathbf{b}_1$ and the success constraint for $\mathbf{b_n}$, respectively.

—$\mathbf{c_i}$ projected over the variables of $\mathbf{b_i}$ is the *entry constraint* (represented by $\mathbf{c_{entry}}$) of $\mathbf{b_i}$ and the answer constraint $\exists_{-\mathbf{vars(b_i)}}(\mathbf{b_i} = \mathbf{h}' \wedge \mathbf{c_{out}})$ for $\mathbf{b_i}$ is called *exit constraint* (represented by $\mathbf{c_{exit}}$). Note that these two constraints are defined over the variables in $\mathbf{b_i}$, instead of over the variables of the rule.

5. EXTENSION OF THE ANALYSIS FRAMEWORK

As mentioned in the previous section, the framework of [Bruynooghe 1991] provides an algorithm for safely abstracting an operational collecting semantics represented as an AND-OR tree by a finite AND-OR graph. The extension of the framework towards CLP is founded on the observation that the LSLD and LSLDT operational semantics are also valid for CLP. As a consequence, the extension replaces the set of substitutions adorning the AND-OR tree in the original framework by sets of constraint stores, and also replaces unification by conjunction. The algorithm is based on a number of primitive transitions which have to approximate transitions on states $\langle \mathbf{G}; \mathbf{C} \rangle$, where $\mathbf{G}$ is a sequence of constraints and atoms and $\mathbf{C}$ is a set of constraint stores belonging to $\mathbf{Cons}^{\mathcal{C}}_{\tilde{\mathbf{x}}}$ (denoting the set of all sets of constraint stores over the variables $\tilde{\mathbf{x}}$). The abstract transitions operate on states $\langle \mathbf{G}; \mathbf{AC} \rangle$ with the abstract constraint $\mathbf{AC}$, a description of a set of constraint stores, belonging to $\mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$ (denoting the set of all descriptions of sets of constraint stores over $\tilde{\mathbf{x}}$). The extension of the framework also involves a reformulation of the safety conditions of the primitive transitions in the constraint setting.

In some program points, the set of constraint stores $\mathbf{C}$ to be abstracted as $\mathbf{AC}$

is the fixpoint of a sequence $\mathbf{C}_1 \subseteq \mathbf{C}_2 \subseteq \mathbf{C}_3 \subseteq \ldots$. Here, the standard theory of abstract interpretation comes in with the Galois insertion as the most popular approach for linking $\mathbf{Cons}^{\mathbf{C}}_{\tilde{\mathbf{x}}}$ with $\mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$ (see [Marriott 1993] for others). The theory provides a method for safely approximating the fixpoint of the sequence $\mathbf{C}_1 \subseteq \mathbf{C}_2 \subseteq \ldots$. Having for each $\mathbf{AC}$ in the AND-OR graph a Galois insertion between $\mathbf{Cons}^{\mathbf{C}}_{\tilde{\mathbf{x}}}$ and $\mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$, a Galois insertion is induced between the set of all AND-OR trees representing the collecting semantics and the set of all abstract AND-OR graphs.

5.1 The abstract domain

The elements to be abstracted in the collecting semantics are sets of constraint stores, a constraint store being built from primitive constraints through conjunction and projection. All constraint stores in the same set are over some set of variables $\tilde{\mathbf{x}}$. Thus, the concrete domain is $(\mathbf{Cons}^{\mathbf{C}}_{\tilde{\mathbf{x}}}, \leq^{\mathbf{C}})$ where $\leq^{\mathbf{C}}$ is the subset relation. The concrete domain is a lattice whose minimal element is $\emptyset$ and whose maximal element is the set of all possible constraint stores over $\tilde{\mathbf{x}}$. Whether *false* is also considered as a constraint store depends on the kind of analysis one is interested in.

The abstract domain $\mathbf{Cons}^{\mathbf{A}}_{\tilde{\mathbf{x}}}$ consists of descriptions (denoted AC) and is equipped with an order relation $\leq^{\mathbf{A}}$. Descriptions are given a meaning by the concretization function γ. For the analyses considered in this paper, the meaning of descriptions are sets closed under equivalence. For example, if $\mathbf{x}+\mathbf{y} = 2 \wedge \mathbf{x}-\mathbf{y} = 0$ is in $\gamma(\mathbf{AC})$, then so will be $\mathbf{x} = 1 \wedge \mathbf{y} = 1$.

A special class of descriptions are those where the represented sets are closed under *anti-entailment*: a description representing a constraint also represents all stronger constraints. Formally: if $\mathbf{c} \in \gamma(\mathbf{AC})$ and $\mathbf{c}' \rightarrow \mathbf{c}$ then $\mathbf{c}' \in \gamma(\mathbf{AC})$[5]. This class is the CLP counterpart of substitution closed (downward closed) descriptions in abstract interpretation of logic programs [Debray 1992a]. Such domains have the special property that, if AC is a valid description of the computation at state $\mathbf{s_i}$ in the collecting semantics, then it is also a valid description (though usually rather imprecise) of the state $\mathbf{s_{i+1}}$. Indeed, the standard semantics can only strengthen the constraints by adding constraints to the store. The definiteness domain developed in Section 6 is such a domain. If a variable is constrained to a unique value by some constraint then it is certainly so under stronger constraints.

Another special class of descriptions represents sets closed under entailment (upward closed) (at least if unsatisfiable constraints are discarded): a description representing a constraint ($\neq$ **false**) represents also all weaker constraints, formally: if $\mathbf{c} \in \gamma(\mathbf{AC})$ ($\mathbf{c} \neq \mathbf{false}$) and $\mathbf{c} \rightarrow \mathbf{c}'$, then $\mathbf{c}' \in \gamma(\mathbf{AC})$. The freeness domain developed in Section 7 is such a domain. If a variable can still take all possible values under some constraint, then it can do so under weaker constraints.

As stated in Section 2.3, the most familiar setting for abstract interpretation is the Galois insertion. The concretization function γ and the abstraction function α provide a tight linkage between the concrete and the abstract domain. As a consequence one can specify the safety conditions of the functions used in formulating

[5] Note that each description that is closed under anti-entailment automatically represents the constraint **false**.

the abstract semantics as well in terms of the concretization function as in terms of the abstraction function. As discussed in [Marriott 1993], abstract interpretation has also been studied in settings with a weaker linkage between abstract and concrete domain. Here we follow the weaker setting of the original framework of [Bruynooghe 1991] where only a concretization function γ is assumed. However, the formulation is modified in one aspect. To ensure termination for abstract domains allowing for infinite ascending chains $\mathbf{AC}_1 <^{\mathcal{A}} \mathbf{AC}_2 <^{\mathcal{A}} \mathbf{AC}_3 <^{\mathcal{A}} \ldots$, the standard notion of a widening operator [Cousot and Cousot 1977; 1992b] is used.

Let $(\mathbf{Cons}^{\mathbf{C}}_{\tilde{\mathbf{x}}}, \leq^{\mathbf{C}})$ be the powerset of the set of all constraint stores, ordered by set inclusion. The minimal requirements on $(\mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}, \leq^{\mathcal{A}})$ are:

(1) A pre-order $\leq^{\mathcal{A}}$ satisfying that $\forall \mathbf{AC}_1, \mathbf{AC}_2 \in \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$ if $\mathbf{AC}_1 \leq^{\mathcal{A}} \mathbf{AC}_2$ then $\gamma(\mathbf{AC}_1) \subseteq \gamma(\mathbf{AC}_2)$. The pre-order allows to define an equivalence relation: $\mathbf{AC}_1 \equiv^{\mathcal{A}} \mathbf{AC}_2$ iff $\mathbf{AC}_1 \leq^{\mathcal{A}} \mathbf{AC}_2$ and $\mathbf{AC}_2 \leq^{\mathcal{A}} \mathbf{AC}_1$. The relation $\equiv^{\mathcal{A}}$ has the property $\mathbf{AC}_1 \equiv^{\mathcal{A}} \mathbf{AC}_2 \rightarrow \gamma(\mathbf{AC}_1) = \gamma(\mathbf{AC}_2)$.

(2) An upper bound operator $\mathbf{upp} : \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}} \times \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}} \rightarrow \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$ such that $\mathbf{AC_i} \leq^{\mathcal{A}} \mathbf{upp}(\mathbf{AC}_1, \mathbf{AC}_2)$ $(\mathbf{i} = 1, 2)$.

(3) A maximal element named $\top$ such that $\gamma(\top) =$ the set of all constraints over $\tilde{\mathbf{x}}$ and $\forall \mathbf{AC} \in \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}},\ \mathbf{AC} \leq^{\mathcal{A}} \top$.

(4) A minimal element $\perp$ such that $\gamma(\perp) = \emptyset$ and $\forall \mathbf{AC} \in \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}},\ \perp \leq^{\mathcal{A}} \mathbf{AC}$.

(5) A widening operator $\mathcal{W} : \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}} \times \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}} \rightarrow \mathbf{Cons}^{\mathcal{A}}_{\tilde{\mathbf{x}}}$ such that $\mathbf{AC_i} \leq^{\mathcal{A}} \mathcal{W}(\mathbf{AC}_1, \mathbf{AC}_2)(\mathbf{i} = 1, 2)$ and such that there does not exist an infinite chain $\mathbf{ACa}_1, \mathbf{ACb}_1, \mathbf{ACa}_2, \mathbf{ACb}_2, \mathbf{ACa}_3, \ldots$ such that, for all i: $\mathbf{not}(\mathbf{ACb_i} \leq^{\mathcal{A}} \mathbf{ACa_i})$ and, for all $\mathbf{i} > 1 : \mathbf{ACa_i} = \mathcal{W}(\mathbf{ACa_{i-1}}, \mathbf{ACb_{i-1}})$.

Condition 1 allows different descriptions that are not equivalent to represent the same set of constraints. However this is better avoided as it can decrease precision and increase computation time. Condition 2 states that there must be an upper bound operator, i.e. that it must be possible to approximate two or more descriptions by a single one. Of course, it is desirable to define **upp** as precise as possible. With the abstract domain a complete partial order, the optimal **upp** is the least upper bound. Condition 3 implies the existence of a maximal element, it is a convention to name it $\top$. Condition 3 also states that it must represent the set of all constraints. This assures that every set of constraints has an abstraction. Condition 4 imposes a minimal element $\perp$ representing the empty set of constraints. This provides a precise abstraction for states in unreachable program points. Also, it provides the initial value for computing a fixpoint of a function over the abstract domain. Finally, condition 5 ensures existence of a widening operator which can enforce a safe approximation of a fixpoint in a finite number of steps ($\mathbf{ACa}_1, \mathbf{ACa}_2, \ldots$ are the successive approximations of the fixpoint, $\mathbf{ACb}_1, \mathbf{ACb}_2, \ldots$ the values resulting from the new iterations). Notice that **upp** can be used as widening operator in domains without infinite ascending chains.

5.2 The abstract operations

The algorithm computes an AND-OR graph adorned with abstract constraints (elements of the abstract domain). It also computes Table, an initially empty table, with elements of the form $(\langle \mathbf{a}, \mathbf{AC_{entry}} \rangle, \mathbf{AC_{exit}})$. In these entries $\mathbf{a}$ is an atom and $\mathbf{AC_{entry}}$ (the entry constraint) and $\mathbf{AC_{exit}}$ (the exit constraint) are abstract constraints over $\mathbf{vars(a)}$. The pair $\langle \mathbf{a}, \mathbf{AC_{entry}} \rangle$ is the key of the table element and $\mathbf{AC_{exit}}$ is the (current) answer for the call $\mathbf{a}$ with abstract entry constraint $\mathbf{AC_{entry}}$. $\mathbf{AC_{exit}}$ is used by table look-ups[6]. The graph is initialized with an AND-node having one child for each atom or constraint in the query ?-$\mathbf{g}_1, \ldots, \mathbf{g_n}$ and an abstract call constraint $\mathbf{AC}$ for $\mathbf{g}_1$. This initialization represents the set of initial LSLDT sequences $\langle \mathbf{g}_1, \ldots, \mathbf{g_n}; \mathbf{c} \rangle$ where $\mathbf{c} \in \gamma(\mathbf{AC})$. The algorithm builds a complete graph by applying transitions in a controlled way.

Below, we make use of abstract projection, denoted $\exists^{\mathcal{A}}_{-\tilde{\mathbf{x}}}$, and abstract conjunction $\wedge^{\mathcal{A}}$. They are intended to approximate projection and conjunction respectively. More formally:

—The abstract projection $\exists^{\mathcal{A}}_{-\tilde{\mathbf{x}}}$ is a safe approximation of the concrete projection if for any constraint $\mathbf{c}$ and for any abstract constraint $\mathbf{AC}$ such that $\mathbf{c} \in \gamma(\mathbf{AC})$ it holds that $\exists_{-\tilde{\mathbf{x}}}\mathbf{c} \in \gamma(\exists^{\mathcal{A}}_{-\tilde{\mathbf{x}}}\mathbf{AC})$.

—The abstract conjunction $\wedge^{\mathcal{A}}$ is a safe approximation of the concrete conjunction if for any two constraints $\mathbf{c}_1, \mathbf{c}_2$ and for any two abstract constraints $\mathbf{AC}_1, \mathbf{AC}_2$ such that $\mathbf{c}_1 \in \gamma(\mathbf{AC}_1)$ and $\mathbf{c}_2 \in \gamma(\mathbf{AC}_2)$ it holds that $\mathbf{c}_1 \wedge \mathbf{c}_2 \in \gamma(\mathbf{AC}_1 \wedge^{\mathcal{A}} \mathbf{AC}_2)$.

With the concrete and abstract domains linked by a Galois connection or insertion, the safety condition can also be formulated in terms of the abstraction function:

—The abstract projection $\exists^{\mathcal{A}}_{-\tilde{\mathbf{x}}}$ is a safe approximation of the concrete projection if for any set of constraints $\mathbf{C}$ and for any abstract constraint $\mathbf{AC}$ such that $\alpha(\mathbf{C}) \leq^{\mathcal{A}} \mathbf{AC}$ it holds that $\alpha(\exists_{-\tilde{\mathbf{x}}}\mathbf{C}) \leq^{\mathcal{A}} \exists^{\mathcal{A}}_{-\tilde{\mathbf{x}}}\mathbf{AC}$.

—The abstract conjunction $\wedge^{\mathcal{A}}$ is a safe approximation of the concrete conjunction if for any two sets of constraints $\mathbf{C}_1, \mathbf{C}_2$ and for any two abstract constraints $\mathbf{AC}_1, \mathbf{AC}_2$ such that $\alpha(\mathbf{C}_1) \leq^{\mathcal{A}} \mathbf{AC}_1$, $\alpha(\mathbf{C}_2) \leq^{\mathcal{A}} \mathbf{AC}_2$ it holds that $\alpha(\mathbf{C}_1 \wedge \mathbf{C}_2) \leq^{\mathcal{A}} \mathbf{AC}_1 \wedge^{\mathcal{A}} \mathbf{AC}_2$ where $\mathbf{C}_1 \wedge \mathbf{C}_2$ is the collecting conjunction, i.e. $\mathbf{C}_1 \wedge \mathbf{C}_2 = \{\mathbf{c}_1 \wedge \mathbf{c}_2 \mid \mathbf{c}_1 \in \mathbf{C}_1, \mathbf{c}_2 \in \mathbf{C}_2\}$.

Let $\mathbf{a}$ be a leaf atom of the AND-OR graph, and $\mathbf{AC}$ be its abstract call constraint. Also, let $\rho_1, \ldots, \rho_{\mathbf{m}}$ be the rules of the program $\mathbf{P}$ defining the predicate of $\mathbf{a}$ with the $\mathbf{j^{th}}$ rule $\rho_{\mathbf{j}}$ of the form $\mathbf{h_j}$:-$\mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}$. Basic transitions on the AND-OR graph are:

—**abstract_entry**$(\mathbf{a}, \mathbf{AC})$

This abstract transition has to approximate all entry transitions over LSLDT sequences $\mathbf{S} :: \langle \mathbf{a}; \mathbf{c} \rangle$ with $\mathbf{c} \in \gamma(\mathbf{AC})$. As explained in Section 4, for each rule $\rho_{\mathbf{j}}$, the entry transition extends the sequence $\mathbf{S} :: \langle \mathbf{a}; \mathbf{c} \rangle$ with the state $\langle \mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}; \mathbf{c^j_{in}} \rangle$, where $\mathbf{c^j_{in}} = \exists_{-\mathbf{vars}(\rho_{\mathbf{j}})}(\mathbf{a} = \mathbf{h_j} \wedge \mathbf{c})$ (and creates an entry in Table with key $\langle \mathbf{a}, \mathbf{c_{entry}} \rangle$, where $\mathbf{c_{entry}} = \exists_{-\mathbf{vars(a)}}\mathbf{c}$). Therefore, in this transition the leaf node $\mathbf{a}$ becomes an OR-node, with the nodes $\mathbf{h_j}$ as children. A node $\mathbf{h_j}$ becomes an AND-node with the atoms/constraints $\mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}$ as children. The abstract call constraints $\mathbf{AC^j_{in}}$ of $\mathbf{b_{j1}}$, for all $\mathbf{j}$, are computed. Finally, the transition computes $\mathbf{AC_{entry}}$, an intermediate abstract constraint over $\mathbf{vars(a)}$ approximating

[6] A similar table is used in the concrete LSLDT semantics, but a key is then associated with a set of answers.

$\mathbf{c_{entry}}$. The pair $\langle \mathbf{a}, \mathbf{AC_{entry}} \rangle$ will be a key in Table. This gives the following safety conditions:

—for $\mathbf{AC_{entry}}$: $\mathbf{c} \in \gamma(\mathbf{AC}) \rightarrow \mathbf{c_{entry}} \in \gamma(\mathbf{AC_{entry}})$.

—for $\mathbf{AC^j_{in}}$: $\mathbf{c} \in \gamma(\mathbf{AC}) \rightarrow \mathbf{c^j_{in}} \in \gamma(\mathbf{AC^j_{in}})$.

—**extension_from_table**$(\mathbf{a}, \mathbf{AC}, \mathbf{a^{tab}}, \mathbf{AC^{tab}})$

This abstract transition has to approximate all Table look-up transitions on LSLDT sequences of the form $\mathbf{S}_1 :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle$ for which there is an element in Table with key $\langle \mathbf{a^{tab}}, \mathbf{c^{tab}_{entry}} \rangle$ such that $\langle \mathbf{a^{tab}}, \mathbf{c^{tab}_{entry}} \rangle \delta = \langle \mathbf{a}, \mathbf{c_{entry}} \rangle$ for some renaming δ. For each stored answer $\mathbf{c^{tab}_{exit}}$, LSLDT extends such a sequence with $\langle \mathbf{G}; \mathbf{c} \wedge \mathbf{c^{tab}_{exit}} \delta \rangle$.

The abstract transition has to compute the abstract success constraint $\mathbf{AC'}$ of $\mathbf{a}$. With $(\langle \mathbf{a^{tab}}, \mathbf{AC^{tab}_{entry}} \rangle, \mathbf{AC^{tab}_{exit}})$ the table entry such that $\langle \mathbf{a^{tab}}, \mathbf{AC^{tab}_{entry}} \rangle \delta = \langle \mathbf{a}, \mathbf{AC_{entry}} \rangle$ for some renaming δ, the safety condition is:

—$\mathbf{c} \in \gamma(\mathbf{AC})$, $\mathbf{c^{tab}_{exit}} \in \gamma(\mathbf{AC^{tab}_{exit}})$, $(\mathbf{c^{tab}_{exit}} \delta \rightarrow \mathbf{c_{entry}}) \rightarrow \mathbf{c} \wedge \mathbf{c^{tab}_{exit}} \delta \in \gamma(\mathbf{AC'})$.

—**abstract_exit**$(\mathbf{a}, \mathbf{AC}, \{\mathbf{h}_1, \ldots, \mathbf{h_m}\}, \{\mathbf{AC^1_{out}}, \ldots, \mathbf{AC^m_{out}}\})$

Let $\mathbf{AC^j_{out}}$ be the abstract out constraint of the rule $\rho_\mathbf{j}$. This abstract transition has to approximate all exit transitions over LSLDT sequences of the form $\mathbf{S}_1 :: \langle \mathbf{a}, \mathbf{G}; \mathbf{c} \rangle \overset{\rho_\mathbf{j}}{::} \langle \mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}; \mathbf{c^j_{in}} \rangle :: \mathbf{S}_2 :: \langle \Box_{\rho_\mathbf{j}}; \mathbf{c^j_{out}} \rangle$ where $\mathbf{c} \in \gamma(\mathbf{AC})$, $\mathbf{c^j_{out}} \in \gamma(\mathbf{AC^j_{out}})$ and $\mathbf{c^j_{out}} \rightarrow \exists_{-\mathbf{vars}(\rho_\mathbf{j})}(\mathbf{c} \wedge \mathbf{a} = \mathbf{h_j})$. Such an exit transition computes $\mathbf{c_{exit}} = \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h_j} \wedge \mathbf{c^j_{out}})$ to be stored in Table as an answer for the key $\langle \mathbf{a}, \mathbf{c_{entry}} \rangle$ and extends the sequence with the state $\langle \mathbf{G}; \mathbf{c} \wedge \mathbf{c_{exit}} \rangle$.

The abstract transition has to compute $\mathbf{AC_{exit}}$, the abstract constraint over $\mathbf{vars(a)}$ to be stored as answer in Table for the element with key $\langle \mathbf{a}, \mathbf{AC_{entry}} \rangle$, and the abstract success constraint $\mathbf{AC'}$ of $\mathbf{a}$. The safety conditions are:

—for $\mathbf{AC_{exit}}$: $\mathbf{c} \in \gamma(\mathbf{AC})$, $\mathbf{c^j_{out}} \in \gamma(\mathbf{AC^j_{out}})$, $(\mathbf{c^j_{out}} \rightarrow \exists_{-\mathbf{vars}(\rho_\mathbf{j})}(\mathbf{a} = \mathbf{h_j} \wedge \mathbf{c})) \rightarrow \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h_j} \wedge \mathbf{c^j_{out}}) \in \gamma(\mathbf{AC_{exit}})$.

—for $\mathbf{AC'}$: $\mathbf{c} \in \gamma(\mathbf{AC})$, $\mathbf{c_{exit}} \in \gamma(\mathbf{AC_{exit}})$, $(\mathbf{c_{exit}} \rightarrow \mathbf{c_{entry}}) \rightarrow \mathbf{c} \wedge \mathbf{c_{exit}} \in \gamma(\mathbf{AC'})$.

Alternatively, the condition for $\mathbf{AC'}$ can be formulated without relying on $\mathbf{AC_{exit}}$:
$\mathbf{c} \in \gamma(\mathbf{AC})$, $\mathbf{c^j_{out}} \in \gamma(\mathbf{AC^j_{out}})$, $(\mathbf{c^j_{out}} \rightarrow \exists_{-\mathbf{vars}(\rho_\mathbf{j})}(\mathbf{a} = \mathbf{h_j} \wedge \mathbf{c})) \rightarrow \exists_{-\mathbf{vars}(\rho_0)}(\mathbf{c} \wedge \mathbf{a} = \mathbf{h_j} \wedge \mathbf{c^j_{out}}) \in \gamma(\mathbf{AC'})$

A straightforward definition in terms of abstract projection, abstract conjunction and constraint abstraction for the abstractions mentioned above, which satisfies the safety requirements, is:

—$\mathbf{AC_{entry}} = \exists^{\mathcal{A}}_{-\mathbf{vars(a)}} \mathbf{AC}$,

—$\mathbf{AC^j_{in}} = \exists^{\mathcal{A}}_{-\mathbf{vars}(\rho_\mathbf{j})}(\mathbf{AC_{entry}} \wedge^{\mathcal{A}} \alpha(\mathbf{a} = \mathbf{h_j}))$,

—$\mathbf{AC_{exit}} = \mathbf{upp}(\mathbf{AC^1_{exit}}, \ldots, \mathbf{AC^m_{exit}})$, where $\mathbf{AC^j_{exit}} = \exists^{\mathcal{A}}_{-\mathbf{vars(a)}}(\mathbf{AC^j_{out}} \wedge^{\mathcal{A}} \alpha(\mathbf{a} = \mathbf{h_j}))$,

—$\mathbf{AC'} = \mathbf{AC} \wedge^{\mathcal{A}} \mathbf{AC_{exit}}$, and in **extension_from_table** $\mathbf{AC'} = \mathbf{AC} \wedge^{\mathcal{A}} \mathbf{AC^{tab}} \delta$ where δ is a renaming such that $\mathbf{a^{tab}} \delta = \mathbf{a}$. This computation is often called *extension.*

However, other definitions are feasible. As we will see later, different definitions can yield more accurate results, depending on the characteristics of the particular abstract domain considered.

Now we can describe the $\mathbf{call\_to\_success}(\mathbf{g}, \mathbf{AC})$ procedure which controls a succession of transitions of which **abstract_entry** and **abstract_exit** are the most important ones. Assuming, without loss of generality, that a query consists of a single atom or constraint $\mathbf{g}$ with abstract call constraint $\mathbf{AC}$, the abstract operational semantics (the AND-OR graph) is computed by $\mathbf{call\_to\_success}(\mathbf{g}, \mathbf{AC})$.

If $\mathbf{g}$ is a constraint, then an $\mathbf{AC}'$ satisfying $\mathbf{c} \in \gamma(\mathbf{AC}) \rightarrow \mathbf{c} \wedge \mathbf{g} \in \gamma(\mathbf{AC}')$ has to be computed. Defining $\mathbf{AC}'$ as $\mathbf{AC} \wedge^{\mathcal{A}} \alpha(\mathbf{g})$ satisfies this condition. However, other definitions (e.g. not relying on α and $\wedge^{\mathcal{A}}$) are feasible.

If $\mathbf{g}$ is an atom, the procedure is as follows:

(1) Compute $\mathbf{AC_{entry}}$.

(2) If Table has an entry $(\langle \mathbf{g^{tab}}, \mathbf{AC^{tab}_{entry}}\rangle, \mathbf{AC^{tab}_{exit}})$ such that $\mathbf{g^{tab}}$ is a renaming of $\mathbf{g}$ and $\mathbf{AC^{tab}_{entry}} \equiv^{\mathcal{A}} \mathbf{AC_{entry}}$[7] (*table look-up*), then $\mathbf{AC}'$ is computed by $\mathbf{extension\_from\_table}(\mathbf{g}, \mathbf{AC}, \mathbf{g^{tab}}, \mathbf{AC^{tab}_{exit}})$.

(3) Else if there is an ancestor node $\mathbf{g^{anc}}$ with associated entry $(\langle \mathbf{g^{anc}}, \mathbf{AC^{anc}_{entry}}\rangle, \mathbf{AC^{anc}_{exit}})$ in Table such that $\mathbf{g}$ is a renaming of $\mathbf{g^{anc}}$ and for which $\mathbf{similar}(\mathbf{AC^{anc}_{entry}}, \mathbf{AC_{entry}})$ holds (*table look-up*), then
   —If $\mathbf{AC_{entry}} \leq^{\mathcal{A}} \mathbf{AC^{anc}_{entry}}$ then $\mathbf{AC}' = \mathbf{extension\_from\_table}(\mathbf{g}, \mathbf{AC}, \mathbf{g^{anc}}, \mathbf{AC^{anc}_{exit}})$.
   —Else backtrack to $\mathbf{g^{anc}}$ and restart with $\mathbf{call\_to\_success}(\mathbf{g^{anc}}, \mathbf{AC^{anc}})$ but with $\mathbf{AC^{anc}_{entry}} = \mathcal{W}(\mathbf{AC^{anc}_{entry}}, \mathbf{AC_{entry}})$.
   (The original computation of $\mathbf{call\_to\_succes}(\mathbf{g^{anc}}, \mathbf{AC^{anc}})$ becomes obsolete.)

(4) Else
   —Create an entry[8] $(\langle \mathbf{g}, \mathbf{AC_{entry}}\rangle, \bot)$ in Table.
   —Apply $\mathbf{abstract\_entry}(\mathbf{g}, \mathbf{AC})$ obtaining the set of abstract *in constraints* $\mathbf{AC^1_{in}}, \ldots, \mathbf{AC^m_{in}}$, one for each rule $\mathbf{h_j} \leftarrow \mathbf{B_j}$ $(1 \leq \mathbf{j} \leq \mathbf{m})$.
   —The states $\langle \mathbf{B_j}, \mathbf{AC^j_{in}}\rangle$ are analyzed, applying, from left to right, **call_to_success** on the subgoals of the $\mathbf{B_j}$. Eventually one obtains the abstract *out constraints* $\mathbf{AC^1_{out}}, \ldots, \mathbf{AC^m_{out}}$.
   —Apply $\mathbf{abstract\_exit}(\mathbf{g}, \mathbf{AC}, \{\mathbf{h}_1, \ldots, \mathbf{h_m}\}, \{\mathbf{AC^1_{out}}, \ldots, \mathbf{AC^m_{out}}\})$. The intermediate result $\mathbf{AC_{exit}}$ is used to update the entry $(\langle \mathbf{g}, \mathbf{AC_{entry}}\rangle, \mathbf{AC^{tab}_{exit}})$ of Table as follows.
   If $\mathbf{AC_{exit}} \leq^{\mathbf{A}} \mathbf{AC^{tab}_{exit}}$ then no update
   Else if $\mathbf{AC^{tab}_{exit}}$ has already been used in a *table look-up* (this implies that $\mathbf{g}$ is a recursive predicate) then
   —The new value is $\mathcal{W}(\mathbf{AC^{tab}_{exit}}, \mathbf{AC_{exit}})$.
   —*Redo all computations* whose outcome depends directly or indirectly on the value $\mathbf{AC^{tab}_{exit}}$ which was used in the *table look-up*s (again part of the

[7] Here and in the sequel we implicitly assume proper renaming of formulas.

[8] It can sometimes be preferable to enlarge $\mathbf{AC_{entry}}$, for example because it contains uninteresting details or because there are already too many different entry patterns for $\mathbf{g}$. If the enlarged $\langle \mathbf{g}, \mathbf{AC^{enl}_{entry}}\rangle$ can be solved by *table look-up*, then $\mathbf{AC}'$ is computed as in step 2.

computations becomes obsolete). These are the "iterations" mentioned below. A crude way is to backtrack and to restart $\mathbf{call\_to\_success}(\mathbf{g}, \mathbf{AC})$.

Else the new value is $\mathbf{upp}(\mathbf{AC}^{\mathbf{tab}}_{\mathbf{exit}}, \mathbf{AC}_{\mathbf{exit}})$.

The test $\mathbf{similar}(\mathbf{AC}^{\mathbf{anc}}_{\mathbf{entry}}, \mathbf{AC}_{\mathbf{entry}})$ must be such that no infinite chain of similar ancestors $\langle \mathbf{g}, \mathbf{AC}_{\mathbf{entry}}\rangle, \langle \mathbf{g}^{\mathbf{anc}}, \mathbf{AC}^{\mathbf{anc}}_{\mathbf{entry}}\rangle, \langle \mathbf{g}^{\mathbf{anc}^2}, \mathbf{AC}^{\mathbf{anc}^2}_{\mathbf{entry}}\rangle, \ldots$ is created. A straightforward method is to put an arbitrary bound on the length of such chains. A more intelligent way would be to judge whether the differences between $\mathbf{AC}_{\mathbf{entry}}$, $\mathbf{AC}^{\mathbf{anc}}_{\mathbf{entry}}$ and $\mathcal{W}(\mathbf{AC}_{\mathbf{entry}}, \mathbf{AC}^{\mathbf{anc}}_{\mathbf{entry}})$ are significant with regard to the properties of interest (i.e. whether specialization for the different calls is worthwhile).

The relevant information about the analysis (the atoms with their abstract entry and exit constraint) are all collected in Table. Abstract interpretation systems such as PLAI [Muthukumar and Hermenegildo 1990; 1992], GAIA [Englebert et al. 1992; Le Charlier et al. 1991] and AMAI [Janssens et al. 1995] do not construct the AND-OR graph explicitly. The systems use a more compact dependency structure which is sufficient to control the order of the transitions to be performed on the implicit AND-OR graph. Major difference between the systems is in the way they organize to "redo all computations dependent on an invalid *table look-up*": the way they attempt to minimize the number of transitions to be redone and how they attempt to make the best use of what has already been computed. Our PLAI implementation of the fixpoint algorithm [Muthukumar and Hermenegildo 1989; 1990; 1992; Hermenegildo et al. 1995] is performed as follows. The program is preprocessed in order to determine recursive predicates and recursive rules. This allows analyzing non-recursive predicates in one pass without checking whether there is an ancestor node. For the recursive predicates, non-recursive rules are analyzed first and once, and the result is taken as a first approximation of the answer. Then, the analysis for the recursive rules starts. The number of iterations performed in this computation is reduced by keeping track of the dependencies among nodes in the abstract AND-OR graph and the state of the information being computed. In some cases the fixpoint algorithm is able to finish in a single iteration.

5.3 Passive constraints

The extended analysis framework proposed in the previous sections does not consider passive constraints. Integrating passive constraints in the concrete operational semantics can be done by using a more general representation of a state as a tuple $\langle \mathbf{G}, \mathbf{c}, \mathbf{s}\rangle$ ($\mathbf{s}$ being a conjunction of constraints whose consistency has not been checked), modifying the conjunction operation so that it adds the constraints to $\mathbf{s}$ instead of to $\mathbf{c}$, and including an $\mathbf{infer}(\mathbf{c}, \mathbf{s}) = (\mathbf{c}', \mathbf{s}')$ step after each conjunction operation. This step moves active constraints from $\mathbf{s}$ to $\mathbf{c}$, and is immediately followed by a test for consistency [Jaffar and Maher 1994], at least if the considered CLP system is quick-checking[9].

[9] Special care is needed to perform safe analyses of systems that are not quick-checking. The problem that the analysis recognizes a state as a failure while the actual computation would proceed several steps, visiting several states that are not described by the output of the analysis, can be avoided by transforming the analyzed program so that it fails at run time at the same point. Assuming absence of side effects, this is a transformation that cannot modify the observable behavior of the program and that always reduces runtime.

When considering the modifications needed at the abstract level, the fundamental question is what kind of information is required from the analysis and at what level of accuracy. Assume that gathering information regarding *which* constraints are passive and *when* they become active is not required from the analysis, and that we prefer to lose accuracy rather than complicate the abstract operations. Then, the simplest method is to abstract both active and passive constraints by a single abstract component, without distinguishing between the information regarding passive constraints and that provided by the active constraints. This abstraction has to be safe with respect to all possible (future) activations of the passive constraint, and therefore it is possible to lose accuracy. However, this method significantly simplifies the abstract operations and allows such analysis to be integrated in the framework described above. We adopted this simple approach in the implementation of the analyzers presented in the following sections.

If the information provided by the analysis is aimed at detecting program points at which all constraints are definitely active, then we have to abstract in some way the **infer** function. In order to do this, the abstract domain should be able to approximate the information used by **infer** to decide if a constraint is definitely active. Then, for each constraint **c** analyzed, the abstract **infer** function must decide if under the current abstract constraint store, **c** is definitely active, and, if it is not the case, it must abstract the fact that a passive constraint may appear (possibly without identifying which particular constraint it is) and the properties needed for such passive constraint to be definitely woken. Note that, if the domain is closed under anti-entailment, as is the case for definiteness analysis, then the approximation remains safe when a constraint is active before being recognized as such, so there is no need to deal with *possible* wake-ups of passive constraints. This option is closely related to the work in [Hanus 1993] which presents an abstract domain for detecting CLP($\Re$) programs for which all passive non-linear constraints eventually become linear at run-time. Otherwise, as for the freeness analysis, we must take possible wake-ups into account.

Finally, if the information is aimed at accurately modeling the delay and wake-up behavior, and we want to be able to determine which are the passive constraints, when they become passive and when they are woken, we should split up the abstraction in two parts, an active part representing the active constraints and a passive part representing the passive constraints. In this case, the abstract projection function has to preserve enough information to ensure the correct wake-up behavior. A possible technique is to project only the abstract active constraints and to keep the passive part. Then an abstract constraint is no longer restricted to a finite number of variables (the variables of the rule, goal or query) as it is in the original abstract interpretation framework. As a consequence, termination is not guaranteed and some new kind of widening should be introduced. This is related to the work of Marriott et al [Marriott et al. 1994], which gives a simple denotational semantics and a generic global data-flow analysis algorithm which is based on the semantics sketched above, for languages in which the computation generally proceeds left-to-right but in which some calls are dynamically delayed until their arguments are sufficiently instantiated, a very similar case to that of the passive constraints. An alternative technique which is able to project both active and passive components while maintaining accuracy, has been recently described in [García de la Banda et al. 1995].

6. INFERENCE OF DEFINITENESS INFORMATION

In this section we present the abstract domain $\mathbf{Cons}^{\mathcal{D}}$, which approximates definiteness information in CLP programs, and the corresponding abstract functions as required for the extended framework developed above. The abstraction is based on a high-level description of definiteness dependencies which are easy to obtain for each particular type of constraint in an actual system. We have attempted to give intuitively comprehensible definitions of the different operations, rather than algorithmic versions. The algorithms can be found in [García de la Banda 1994], where proofs of correctness for such algorithms are also provided.

6.1 Abstract domain and abstraction function

Let $\wp(\mathbf{S})$ denote the powerset of a set $\mathbf{S}$ and $\wp_\emptyset(\mathbf{S})$ denote $\wp(\mathbf{S}) \setminus \{\emptyset\}$. Also, let $\mathbf{Var}$ denote a denumerable set of variables and $\mathbf{Pvar} \subseteq \mathbf{Var}$ a distinguished (denumerable) set of variables which may occur in programs. An abstract constraint $\mathbf{AC}^{\mathcal{D}} = (\mathbf{D}, \mathbf{R})$ of the abstract domain $\mathbf{Cons}^{\mathcal{D}}$ is an element of $\wp(\mathbf{Pvar}) \times \wp(\mathbf{Pvar} \times \wp_\emptyset(\wp_\emptyset(\mathbf{Pvar})))$ which is in simplified form. A variable $\mathbf{x}$ in $\mathbf{D}$ represents a variable that is known to be definite, which can be represented by the propositional formula $\mathbf{x} \leftarrow \mathbf{true}$. An element $(\mathbf{x}, \{\mathbf{S_1}, \ldots, \mathbf{S_n}\}) \in \mathbf{R}$ with $\mathbf{S_i} = \{\mathbf{x_{i1}}, \ldots, \mathbf{x_{im_i}}\}$ represents known dependencies between variables. These dependencies can be expressed by the propositional formula $\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S_1}) \vee \ldots \vee \mathbf{conj}(\mathbf{S_n})$ where $\mathbf{conj}(\mathbf{S_i}) = \mathbf{x_{i1}} \wedge \ldots \wedge \mathbf{x_{im_i}}$ (this is equivalent with $(\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S_1})) \wedge \ldots \wedge (\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S_n}))$), where a formula $\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S_i})$ expresses that $\mathbf{x}$ is definite if $\mathbf{x_{i1}}$ up to $\mathbf{x_{im_i}}$ are. An element $(\mathbf{D}, \mathbf{R})$ is in simplified form if it encodes at most one formula $\mathbf{x} \leftarrow \ldots$ for each variable $\mathbf{x}$ and has an explicit representation of all implied non-redundant formulas of the form $\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S})$. A formula $\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S})$ is considered redundant if it is a tautology (i.e. $\mathbf{x} \in \mathbf{S}$) or if it is implied by another formula $\mathbf{x} \leftarrow \mathbf{conj}(\mathbf{S'})$ (i.e. $\mathbf{S'} \subseteq \mathbf{S}$). Putting formulas in simplified form gives a more compact representation and reduces the cost of key operations, such as testing for equivalence and performing abstract projection. A simplified form can be obtained by applying the following rewrite rules:

(1) $(\mathbf{D}, \{(\mathbf{x}, \mathbf{SS_1})\} \cup \{(\mathbf{x}, \mathbf{SS_2})\} \cup \mathbf{R}) \Rightarrow (\mathbf{D}, \{(\mathbf{x}, \mathbf{SS_1} \cup \mathbf{SS_2})\} \cup \mathbf{R})$.

(2) $(\mathbf{D}, \{(\mathbf{x}, \{\mathbf{S_1}\} \cup \{\mathbf{S_2}\} \cup \mathbf{SS})\} \cup \mathbf{R}) \Rightarrow (\mathbf{D}, \{(\mathbf{x}, \{\mathbf{S_1}\} \cup \mathbf{SS})\} \cup \mathbf{R})$ if $\mathbf{S_1} \subset \mathbf{S_2}$.

(3) $(\mathbf{D}, \{(\mathbf{x}, \mathbf{SS})\} \cup \mathbf{R}) \Rightarrow (\mathbf{D}, \mathbf{R})$ if $\mathbf{x} \in \mathbf{D}$.

(4) $(\mathbf{D}, \{(\mathbf{x}, \{\{\mathbf{y}\} \cup \mathbf{S}\} \cup \mathbf{SS})\} \cup \mathbf{R}) \Rightarrow (\mathbf{D}, \{(\mathbf{x}, \{\mathbf{S}\} \cup \mathbf{SS})\} \cup \mathbf{R})$ if $\mathbf{y} \in \mathbf{D}$.

(5) $(\mathbf{D}, \{(\mathbf{x}, \{\emptyset\} \cup \mathbf{SS})\} \cup \mathbf{R}) \Rightarrow (\{\mathbf{x}\} \cup \mathbf{D}, \mathbf{R})$.

(6) $(\mathbf{D}, \{(\mathbf{x}, \{\{\mathbf{y}\} \cup \mathbf{S_1}\} \cup \mathbf{SS_1})\} \cup \{(\mathbf{y}, \{\mathbf{S_2}\} \cup \mathbf{SS_2})\} \cup \mathbf{R}) \Rightarrow (\mathbf{D}, \{(\mathbf{x}, \{\mathbf{S_1} \cup \mathbf{S_2}\} \cup \{\{\mathbf{y}\} \cup \mathbf{S_1}\} \cup \mathbf{SS_1})\} \cup \{(\mathbf{y}, \{\mathbf{S_2}\} \cup \mathbf{SS_2})\} \cup \mathbf{R})$ if $\mathbf{x} \notin \mathbf{S_2}$ and $\nexists \mathbf{S} \in \mathbf{SS_1}$ such that $\mathbf{S} \subseteq (\mathbf{S_1} \cup \mathbf{S_2})$.

Rule 1 merges several definite dependencies approximated for the same variable. Knowing that the definiteness of $\mathbf{x}$ can be derived from the definiteness of a set $\mathbf{S_2}$ of variables is useless once the definiteness of $\mathbf{x}$ is known to be derived from a subset $\mathbf{S_1}$ of $\mathbf{S_2}$. Rule 2 eliminates those useless $\mathbf{S_2}$ sets. Approximating that the definiteness of $\mathbf{x}$ can be derived from the definiteness of any other set of variables is useless once $\mathbf{x}$ is in $\mathbf{D}$. Rule 3 performs such simplification. If a variable $\mathbf{y}$ in a set $\mathbf{S} \in \mathbf{SS}$ of $(\mathbf{x}, \mathbf{SS})$ is in $\mathbf{D}$, $\mathbf{y}$ can be removed from $\mathbf{S}$ without losing information. Rule 4 removes those variables. The element $(\mathbf{x}, \{\emptyset\} \cup \mathbf{SS})$ is obtained once $\mathbf{x}$ is known to be definite. Rule 5 eliminates $(\mathbf{x}, \{\emptyset\} \cup \mathbf{SS})$ from $\mathbf{R}$ and adds $\mathbf{x}$ to $\mathbf{D}$. If the definiteness of $\mathbf{y}$ can be inferred from that of the variables in $\mathbf{S}_2$, and the definiteness of $\mathbf{x}$ can in turn be inferred from that of $\{\mathbf{y}\} \cup \mathbf{S}_1$, we can conclude that the definiteness of $\mathbf{x}$ can also be inferred from that of $\mathbf{S}_2 \cup \mathbf{S}_1$. This propagation of definiteness dependencies is performed by rule 6. Note that the condition "$\nexists \mathbf{S} \in \mathbf{SS}_1$ such that $\mathbf{S} \subseteq (\mathbf{S}_1 \cup \mathbf{S}_2)$" avoids infinite applications of rule 6 (if $\mathbf{S} = \mathbf{S}_1 \cup \mathbf{S}_2$) or infinite alternate applications of rules 6 and 2 (if $\mathbf{S} \subset \mathbf{S}_1 \cup \mathbf{S}_2$).

Let $\mathbf{simplify}(\mathbf{D}, \mathbf{R})$ denote the abstract constraint obtained by applying the rewrite rules to $(\mathbf{D}, \mathbf{R})$ until no rule can be applied. We can now formally define $\mathbf{Cons}^{\mathcal{D}}$ as $\{\bot\} \cup \{(\mathbf{D}, \mathbf{R}) \in \wp(\mathbf{Pvar}) \times \wp(\mathbf{Pvar} \times \wp_{\emptyset}(\wp_{\emptyset}(\mathbf{Pvar}))) \mid \mathbf{simplify}(\mathbf{D}, \mathbf{R}) = (\mathbf{D}, \mathbf{R})\}$[10]. For convenience, in the rest of the section we will denote by $\mathbf{min}\mathcal{D}(\mathbf{SS})$ the set of sets obtained by applying rule 2 to a particular $\mathbf{SS}$ of $(\mathbf{x}, \mathbf{SS}) \in \mathbf{R}$.

Definition 6.1.1. Abstraction of a constraint: $\alpha^{\mathbf{d}}$
Let $\mathbf{c}$ be a constraint.
Then $\alpha^{\mathbf{d}}(\mathbf{c}) = \bot$ if $\neg\mathbf{consistent}(\mathbf{c})$, otherwise $\alpha^{\mathbf{d}}(\mathbf{c}) = (\mathbf{D}, \mathbf{R})$ where
1. $\mathbf{D} = \mathbf{def}(\mathbf{c})$[11]
2. $\mathbf{R} = \{(\mathbf{x}, \mathbf{SS}) \mid \mathbf{x} \in \mathbf{vars}(\mathbf{c}), \mathbf{SS} = \mathbf{min}\mathcal{D}(\mathbf{gr\_dep}(\mathbf{c}, \mathbf{x})), \mathbf{SS} \neq \emptyset, \mathbf{SS} \neq \{\emptyset\}\}$
3. $\mathbf{gr\_dep}(\mathbf{c}, \mathbf{x}) = \{\tilde{\mathbf{y}} \subseteq \mathbf{vars}(\mathbf{c}) \setminus \{\mathbf{x}\} \mid$ **for all sequences of values** $\tilde{\mathbf{v}}$ **s.t.** $\mathbf{consistent}(\mathbf{c} \wedge \tilde{\mathbf{y}} = \tilde{\mathbf{v}})$, **holds that** $\mathbf{x} \in \mathbf{def}(\mathbf{c} \wedge \tilde{\mathbf{y}} = \tilde{\mathbf{v}})\}$

Note that $\emptyset \in \mathbf{gr\_dep}(\mathbf{c}, \mathbf{x})$ for any $\mathbf{x} \in \mathbf{def}(\mathbf{c})$ and for any $\mathbf{x}$ such that no definite dependency can be found. In such cases $\mathbf{min}\mathcal{D}(\mathbf{gr\_dep}(\mathbf{c}, \mathbf{x})) = \{\emptyset\}$.

Example 6.1.2.

$$\begin{aligned}
&\alpha^{\mathbf{d}}(\mathbf{x} = 3) &&= (\{\mathbf{x}\}, \emptyset)\\
&\alpha^{\mathbf{d}}(\mathbf{x} = \mathbf{f}(\mathbf{y}, \mathbf{z})) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{x}\}\}), (\mathbf{z}, \{\{\mathbf{x}\}\})\})\\
&\alpha^{\mathbf{d}}(\mathbf{x} = 3\mathbf{y} + 2\mathbf{z}) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{x}, \mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{x}, \mathbf{y}\}\})\})\\
&\alpha^{\mathbf{d}}(\mathbf{x} > \mathbf{y}) &&= (\emptyset, \emptyset)\\
&\alpha^{\mathbf{d}}(\mathbf{x} \neq \mathbf{y}) &&= (\emptyset, \emptyset)\\
&\alpha^{\mathbf{d}}(\mathbf{x} = \mathbf{y} * \mathbf{z}) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{z}\}\})\})\\
&\alpha^{\mathbf{d}}(\mathbf{x} =< \mathbf{y} > .\mathbf{z}) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{x}\}\}), (\mathbf{z}, \{\{\mathbf{x}\}\})\})\\
&\alpha^{\mathbf{d}}(\mathbf{x} =< \mathbf{y} > . < \mathbf{w} > .\mathbf{z}) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{w}, \mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{x}\}\}), (\mathbf{w}, \{\{\mathbf{x}\}\}), (\mathbf{z}, \{\{\mathbf{x}\}\})\})\\
&\alpha^{\mathbf{d}}(\mathbf{x} =< \mathbf{y} > .\mathbf{w}.\mathbf{z}) &&= (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}, \mathbf{w}, \mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{x}\}\}), (\mathbf{w}, \{\{\mathbf{x}, \mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{x}, \mathbf{w}\}\})\})
\end{aligned}$$

Note that the symbol "." stands for concatenation of PrologIII lists, and "$< \mathbf{y} >$" is a list with one element.

Definition 6.1.3. Order relation
Let $(\mathbf{D}_1, \mathbf{R}_1), (\mathbf{D}_2, \mathbf{R}_2) \in \mathbf{Cons}^{\mathcal{D}}$.
Then $(\mathbf{D}_1, \mathbf{R}_1) \leq^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2)$ iff:

[10] For reasons of readability most of the following definitions and operations do not explicitly deal with $\bot$. Their extensions are trivial.
[11] As mentioned in Section 2.1, $\mathbf{def}(\mathbf{c})$ denotes the set of definite variables in $\mathbf{c}$.

1. $\mathbf{D}_2 \subseteq \mathbf{D}_1$
2. $\forall (\mathbf{x}, \mathbf{SS}_2) \in \mathbf{R}_2 : \mathbf{x} \in \mathbf{D}_1$
   or $\exists (\mathbf{x}, \mathbf{SS}_1) \in \mathbf{R}_1$ such that $\forall \mathbf{S}_2 \in \mathbf{SS}_2 : \exists \mathbf{S}_1 \in \mathbf{SS}_1,\ \mathbf{S}_1 \subseteq \mathbf{S}_2$

Intuitively, this means that for every formula represented by $(\mathbf{D}_2, \mathbf{R}_2)$, there is a formula in $(\mathbf{D}_1, \mathbf{R}_1)$ which is at least as strong.

Definition 6.1.4. Equivalence
Let $(\mathbf{D}_1, \mathbf{R}_1), (\mathbf{D}_2, \mathbf{R}_2) \in \mathbf{Cons}^{\mathcal{D}}$.
Then $(\mathbf{D}_1, \mathbf{R}_1) \equiv^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2)$ iff:
1. $\mathbf{D}_1 = \mathbf{D}_2$
2. $\mathbf{R}_1 = \mathbf{R}_2$

Definition 6.1.5. Least upper bound
Let $(\mathbf{D}_1, \mathbf{R}_1), (\mathbf{D}_2, \mathbf{R}_2) \in \mathbf{Cons}^{\mathcal{D}}$.
Then $\mathbf{upp}^{\mathcal{D}}((\mathbf{D}_1, \mathbf{R}_1), (\mathbf{D}_2, \mathbf{R}_2)) = (\mathbf{D}, \mathbf{R})$ where
1. $\mathbf{D} = \mathbf{D}_1 \cap \mathbf{D}_2$
2. $\mathbf{R} = \{(\mathbf{x}, \mathbf{SS}) \in \mathbf{R_i} \mid \mathbf{x} \in \mathbf{D_j}, \mathbf{i}, \mathbf{j} \in \{1,2\}, \mathbf{i} \neq \mathbf{j}\} \cup$
   $\{(\mathbf{x}, \mathbf{min}\mathcal{D}(\mathbf{SS'})) \mid \mathbf{SS'} = \{\mathbf{S}_1 \cup \mathbf{S}_2 \mid (\mathbf{x}, \mathbf{SS}_1) \in \mathbf{R}_1, \mathbf{S}_1 \in \mathbf{SS}_1,$
   $(\mathbf{x}, \mathbf{SS}_2) \in \mathbf{R}_2, \mathbf{S}_2 \in \mathbf{SS}_2\}\}$

The definition can easily be extended to compute the least upper bound of m (m $>$ 2) abstractions. In the following we will assume that the function **upp** applies to a set of abstract constraints.

Definition 6.1.6. Abstraction of a set of constraints: $\alpha^{\mathcal{D}}$
Let $\mathbf{C} \in \mathbf{Cons}^{\mathcal{C}}$.
Then $\alpha^{\mathcal{D}}(\mathbf{C}) = \bot$ if $\mathbf{C} = \emptyset$, otherwise $\alpha^{\mathcal{D}}(\mathbf{C}) = \mathbf{upp}(\{\alpha^{\mathbf{d}}(\mathbf{c}) \mid \mathbf{c} \in \mathbf{C}\})$.

Definition 6.1.7. Maximal and minimal elements
The maximal element is $\top = (\emptyset, \emptyset)$.
The minimal element is $\bot$, denoting the empty set of constraints.

The concretization function $\gamma^{\mathcal{D}}$ can be defined based upon $\alpha^{\mathcal{D}}$ as described in [Cousot and Cousot 1992a]: $\gamma^{\mathcal{D}}(\mathbf{AC}) = \bigcup\{\mathbf{C} \in \mathbf{Cons}^{\mathcal{C}} \mid \alpha^{\mathcal{D}}(\mathbf{C}) \leq^{\mathcal{D}} \mathbf{AC}\}$. Then $(\mathbf{Cons}^{\mathcal{C}}, \subseteq, \mathbf{Cons}^{\mathcal{D}}, \leq^{\mathcal{D}})$ is a Galois insertion [García de la Banda and Hermenegildo 1993].

6.2 Abstract projection and abstract conjunction functions

Definition 6.2.1. Abstract projection
Let $(\mathbf{D}_1, \mathbf{R}_1) \in \mathbf{Cons}^{\mathcal{D}}$ and $\tilde{\mathbf{x}}$ be a set of variables.
Then $\exists^{\mathcal{D}}_{-\tilde{\mathbf{x}}}(\mathbf{D}_1, \mathbf{R}_1) = (\mathbf{D}, \mathbf{R})$ where
1. $\mathbf{D} = \mathbf{D}_1 \cap \tilde{\mathbf{x}}$
2. $\mathbf{R} = \{(\mathbf{x}, \mathbf{SS}) \mid (\mathbf{x}, \mathbf{SS}_1) \in \mathbf{R}_1, \mathbf{x} \in \tilde{\mathbf{x}}, \mathbf{SS} = \{\mathbf{S} \in \mathbf{SS}_1 \mid \mathbf{S} \subseteq \tilde{\mathbf{x}}\}, \mathbf{SS} \neq \emptyset\}$

The propositional formula represented by $(\mathbf{D}, \mathbf{R})$ is the projection of the formula represented by $(\mathbf{D}_1, \mathbf{R}_1)$. Intuitively, $\mathbf{D}$ is the subset of variables in $\tilde{\mathbf{x}}$ which are known to be definite in $\mathbf{D}_1$, and $\mathbf{R}$ contains the definiteness dependencies (if any) approximated by $\mathbf{R}_1$ for the possibly non-definite variables in $\tilde{\mathbf{x}}$. Since only the variables in $\tilde{\mathbf{x}}$ are taken into account, any element $(\mathbf{y}, \mathbf{SS}_1) \in \mathbf{R}_1$ approximating the dependencies for a variable which is not in $\tilde{\mathbf{x}}$ (i.e., $\mathbf{y} \notin \tilde{\mathbf{x}}$) is eliminated. Furthermore, the dependency sets in $\mathbf{SS}_1$ of the elements $(\mathbf{x}, \mathbf{SS}_1) \in \mathbf{R}_1, \mathbf{x} \in \tilde{\mathbf{x}}$ which are not subsets of $\tilde{\mathbf{x}}$ are also eliminated as groundness of *all* variables in a dependency set is required to ground $\mathbf{x}$, yielding $\mathbf{SS}$. Note that if as a result $\mathbf{SS}$ becomes empty, there is no information for the definiteness dependencies of $\mathbf{x}$ w.r.t. the variables in $\tilde{\mathbf{x}}$ and no $(\mathbf{x}, \mathbf{SS})$ will appear in $\mathbf{R}$.

Definition 6.2.2. Abstract conjunction
Let $(\mathbf{D}_1, \mathbf{R}_1), (\mathbf{D}_2, \mathbf{R}_2) \in \mathbf{Cons}^{\mathcal{D}}$.
Then $(\mathbf{D}_1, \mathbf{R}_1) \wedge^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2) = \mathbf{simplify}(\mathbf{D}_1 \cup \mathbf{D}_2, \mathbf{R}_1 \cup \mathbf{R}_2)$.

A more implementation oriented definition of the abstract conjunction function would state that we should first apply rule 1, then rules 3, 4 and 5 (thus propagating definiteness), and finally rule 6 (propagating definite dependencies). Note that we may also need to apply rule 2 immediately after the application of rules 1, 4 or 6. The order in which those steps are performed has been chosen to increase efficiency, but they can be performed in any order affecting neither correctness nor accuracy. For a more implementation oriented definition of this operation, see [García de la Banda 1994].

Example 6.2.3.
Consider the abstract constraints:

$(\mathbf{D}_1, \mathbf{R}_1)$	$(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\})\})$
$(\mathbf{D}_2, \mathbf{R}_2)$	$(\{\mathbf{z}\}, \{(\mathbf{y}, \{\{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\})$

Then $(\mathbf{D}_1, \mathbf{R}_1) \wedge^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2)$ yields the abstract constraint $(\mathbf{D}, \mathbf{R})$ as follows:

$$\begin{array}{l}
\mathbf{simplify}(\{\mathbf{z}\}, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^1 \\
\mathbf{simplify}(\{\mathbf{z}\}, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}, \{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^4 \\
\mathbf{simplify}(\{\mathbf{z}\}, \{(\mathbf{x}, \{\{\mathbf{y}\}, \emptyset\}), (\mathbf{y}, \{\{\mathbf{z}\}, \{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^4 \\
\mathbf{simplify}(\{\mathbf{z}\}, \{(\mathbf{x}, \{\{\mathbf{y}\}, \emptyset\}), (\mathbf{y}, \{\emptyset, \{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^5 \\
\mathbf{simplify}(\{\mathbf{x}, \mathbf{z}\}, \{(\mathbf{y}, \{\emptyset, \{\mathbf{w}\}\}), (\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^5 \\
\mathbf{simplify}(\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}, \{(\mathbf{w}, \{\{\mathbf{y}\}\})\}) \rightarrow^4 \\
\mathbf{simplify}(\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}, \{(\mathbf{w}, \{\emptyset\})\}) \rightarrow^5 \\
\mathbf{simplify}(\{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}\}, \emptyset) = (\{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}\}, \emptyset)
\end{array}$$

where $\rightarrow^{\mathbf{n}}$ represents the application of the nth rule. Thus $(\mathbf{D}_1, \mathbf{R}_1) \wedge^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2) = (\{\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}\}, \emptyset)$

Consider now the abstract constraints:

$(\mathbf{D}_1, \mathbf{R}_1)$	$(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}, \mathbf{w}\}\}), (\mathbf{y}, \{\{\mathbf{z}, \mathbf{w}\}\})\})$
$(\mathbf{D}_2, \mathbf{R}_2)$	$(\emptyset, \{(\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\})$.

Then $(\mathbf{D}_1, \mathbf{R}_1) \wedge^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2)$ yields the abstract constraint $(\mathbf{D}, \mathbf{R})$ as follows:

$$\mathbf{simplify}(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}, \mathbf{w}\}\}), (\mathbf{y}, \{\{\mathbf{z}, \mathbf{w}\}\}), \{(\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\}) \rightarrow^1$$

$\mathbf{simplify}(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}, \mathbf{w}\}\}), (\mathbf{y}, \{\{\mathbf{z}, \mathbf{w}\}, \{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\}) \rightarrow^2$
$\mathbf{simplify}(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}, \mathbf{w}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\}) \rightarrow^6$
$\mathbf{simplify}(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}, \{\mathbf{z}, \mathbf{w}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\}) \rightarrow^2$
$\mathbf{simplify}(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\}) =$
$(\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\})$

Thus $(\mathbf{D}_1, \mathbf{R}_1) \wedge^{\mathcal{D}} (\mathbf{D}_2, \mathbf{R}_2) = (\emptyset, \{(\mathbf{x}, \{\{\mathbf{y}\}, \{\mathbf{z}\}\}), (\mathbf{y}, \{\{\mathbf{z}\}\}), (\mathbf{z}, \{\{\mathbf{y}\}\})\})$

Let us now present how the abstractions required by the framework are computed. Let $\mathbf{g}$ be a constraint or an atom and $\mathbf{AC}$ be its abstract call constraint. If $\mathbf{g}$ is a constraint, then $\mathbf{AC}'$ is defined as $\mathbf{AC} \wedge^{\mathcal{D}} \alpha^{\mathbf{d}}(\mathbf{g})$. If $\mathbf{g}$ is an atom, let $\rho_1, \ldots, \rho_{\mathbf{m}}$ be the rules of the program $\mathbf{P}$ defining the predicate of $\mathbf{g}$, $\rho_{\mathbf{j}}$ be $\mathbf{h_j}$:-$\mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}$, and let $\mathbf{AC}^1_{\mathbf{out}}, \ldots, \mathbf{AC}^{\mathbf{m}}_{\mathbf{out}}$ be the abstract out constraints of rules $\rho_1, \ldots, \rho_{\mathbf{m}}$ respectively. Then, the abstract entry, in, exit and success constraints are defined as follows:

—$\mathbf{AC_{entry}} = \exists^{\mathcal{D}}_{-\mathbf{vars(g)}} \mathbf{AC}$,

—$\mathbf{AC}^{\mathbf{j}}_{\mathbf{in}} = \exists^{\mathcal{D}}_{-\mathbf{vars}(\rho_{\mathbf{j}})} (\mathbf{AC_{entry}} \wedge^{\mathcal{D}} \alpha^{\mathbf{d}}(\mathbf{g} = \mathbf{h_j}))$,

—$\mathbf{AC_{exit}} = \mathbf{upp}^{\mathcal{D}}(\mathbf{AC}^1_{\mathbf{exit}}, \ldots, \mathbf{AC}^{\mathbf{m}}_{\mathbf{exit}})$, where $\mathbf{AC}^{\mathbf{j}}_{\mathbf{exit}} = \exists^{\mathcal{D}}_{-\mathbf{vars(g)}} (\mathbf{AC}^{\mathbf{j}}_{\mathbf{out}} \wedge^{\mathcal{D}} \mathbf{AC_{entry}} \wedge^{\mathcal{D}} \alpha^{\mathbf{d}}(\mathbf{g} = \mathbf{h_j}))$,

—$\mathbf{AC}' = \mathbf{AC} \wedge^{\mathcal{D}} \mathbf{AC_{exit}}$, and in **extension_from_table** $\mathbf{AC}' = \mathbf{AC} \wedge^{\mathcal{D}} \mathbf{AC^{tab}} \delta_{\mathbf{a}=\mathbf{a^{tab}}}$.

It is clear that all definitions satisfy the safety requirements imposed by the framework. However, two important issues must be pointed out. The first issue is related to one of the three properties of the abstract operations identified in [Jacobs and Langen 1992]: *additivity.* This property requires that precision should not be lost when commuting the least upper bound with an abstract operation. Additive upper bounds are not common, and $\mathbf{upp}^{\mathcal{D}}$ is not an exception. As a result, it is possible to obtain a more accurate $\mathbf{AC}'$ by computing $\mathbf{AC}'$ as $\mathbf{upp}(\mathbf{AC} \wedge^{\mathcal{D}} \mathbf{AC}^1_{\mathbf{exit}}, \ldots, \mathbf{AC} \wedge^{\mathcal{D}} \mathbf{AC}^{\mathbf{m}}_{\mathbf{exit}})$. However, the price is $\mathbf{m}$ applications of $\wedge^{\mathcal{D}}$ instead of one. As for this analysis abstract conjunction is an expensive computation, this approach is not taken.

The second issue is related to the definition of $\mathbf{AC_{exit}}$ and, in particular, to the appearance of $\mathbf{AC_{entry}}$ in the definition of each $\mathbf{AC}^{\mathbf{j}}_{\mathbf{exit}}$[12]. This redundant constraint is added in order to avoid a loss of precision caused by the interaction among approximating a property that is closed under anti-entailment (downwards closed), non-normalization, a loss of precision in the abstract projection function, and the tabulation method. Let us illustrate the problem with a simple example.

Example 6.2.4. Assume we have a program $\mathbf{P}$ with only one rule $\rho_1 : \mathbf{p}(\mathbf{z})$. (i.e., a fact). The computation of $\mathbf{call\_to\_success}(\mathbf{p}(\mathbf{f}(\mathbf{x}, \mathbf{y})), \mathbf{AC})$, where $\mathbf{AC} = (\{\mathbf{x}\}, \emptyset)$, will proceed as follows:

(1) $\mathbf{abstract\_entry}(\mathbf{p}(\mathbf{f}(\mathbf{x}, \mathbf{y})), \mathbf{AC})$. Following the definitions above, we will obtain $\mathbf{AC_{entry}} = (\{\mathbf{x}\}, \emptyset)$ and $\mathbf{AC}^1_{\mathbf{in}} = (\emptyset, \emptyset)$.
(2) Since the body of ρ_1 is empty, $\mathbf{AC}^1_{\mathbf{out}} = \mathbf{AC}^1_{\mathbf{in}} = (\emptyset, \emptyset)$.

[12] Recall that $\mathbf{AC}^{\mathbf{j}}_{\mathbf{exit}}$ can be defined in a simpler way, such as $\exists^{\mathcal{D}}_{-\mathbf{vars(g)}} (\mathbf{AC}^{\mathbf{j}}_{\mathbf{out}} \wedge^{\mathcal{D}} \alpha^{\mathbf{d}}(\mathbf{g} = \mathbf{h_j}))$, while still satisfying the safety conditions.

(3) $\mathbf{abstract\_exit}(\mathbf{p}(\mathbf{f}(\mathbf{x},\mathbf{y})),\mathbf{AC},\{\mathbf{p}(\mathbf{z})\},\{\mathbf{AC}^1_{\mathbf{out}}\})$. If we compute $\mathbf{AC}^1_{\mathbf{exit}}$ as $\exists^{\mathcal{D}}_{-\mathbf{vars(a)}}(\mathbf{AC}^{\mathbf{j}}_{\mathbf{out}} \wedge^{\mathcal{D}} \alpha^{\mathbf{d}}(\mathbf{f}(\mathbf{x},\mathbf{y}) = \mathbf{p}(\mathbf{z})))$, we will obtain $\mathbf{AC}^1_{\mathbf{exit}} = (\emptyset,\emptyset)$. Then, $\mathbf{AC}_{\mathbf{exit}} = (\emptyset,\emptyset)$ and $\mathbf{AC}' = (\{\mathbf{x}\},\emptyset)$. On the other hand, if we include $\mathbf{AC}_{\mathbf{entry}}$ in the definition of $\mathbf{AC}^1_{\mathbf{exit}}$ (as proposed in the above definitions), we obtain $\mathbf{AC}_{\mathbf{exit}} = \mathbf{AC}^1_{\mathbf{exit}} = (\{\mathbf{x}\},\emptyset)$, thus avoiding a loss of precision.

Although accuracy is always recovered when computing $\mathbf{AC}'$, the difference can have an adverse effect on memory (tabulation) and time consumption (computing $\mathbf{AC}'$). Also, for some applications it is convenient that $\mathbf{AC}_{\mathbf{exit}}$ provides accurate information about the success state (for example, the output mode of the predicate). Finally, the loss of accuracy in $\mathbf{AC}_{\mathbf{exit}}$ could imply a greater number of fixpoint iterations, since they depend on the value of $\mathbf{AC}_{\mathbf{exit}}$. Regarding the extra cost introduced by our definition, note that since $\mathbf{AC}_{\mathbf{entry}} \wedge^{\mathcal{A}} \alpha^{\mathbf{d}}(\mathbf{g} = \mathbf{h}_{\mathbf{j}})$ is already computed during the **abstract_entry** operation, the alternative computation does not introduce a significant overhead.

As a last remark, we use $\mathbf{upp}^{\mathcal{D}}$ as a widening operator, since $\mathbf{Cons}^{\mathcal{D}}$ (when considered over a finite set of variables) does not have infinite ascending chains.

There are at least two other domains which are closely related to ours. One is the domain proposed by Hanus [1995], and originally used for detecting situations in which the residuation rule[13] can be guaranteed to never be activated in a given program (this is similar in some ways to a "non-suspension" analysis). The non-residuation requirements imply groundness requirements for the arguments of certain functions and a domain similar to the one defined in this section is used for inferring such groundness.

The second related domain is the domain of positive Boolean functions which are closed under intersection. This domain was defined early on by Dart [1988] under the name of *dependency formulae*, and applied to the inference of groundness in deductive databases. Our domain can be seen as a compact representation of this domain (including a formulation of efficient operations for it). The main difference is that, for efficiency reasons, we require the abstraction to be in a particular simplified form. Recently, the different possible subsets of the Boolean functions which can be used for tracking dependencies in program analysis and their representations have been studied and greatly clarified [Armstrong et al. 1994]. Our domain corresponds essentially to the **Def** domain in this taxonomy. The work developed in [Armstrong et al. 1994] also illustrates that the representation that we have proposed is closely related to the **CDF** representation which is shown therein to offer an advantageous cost-performance tradeoff.

7. INFERENCE OF FREENESS INFORMATION

The definiteness analysis infers whether variables are *definite*, i.e. constrained to a unique value. The analysis takes into account *definite* dependencies between variables in order to perform accurate definiteness propagation. The freeness analysis derives whether variables are *free*, i.e. whether they can range over the whole domain specified by their type: e.g. a variable that is constrained to be numerical but

[13] Residuation is an operational mechanism for the integration of functions into logic programming. The residuation rule delays the evaluation of functions during the unification process until the arguments are sufficiently instantiated.

still ranges over the complete domain of numbers is considered as free. It keeps track of *possible* dependencies between variables to take care of non-freeness propagation: in order to obtain definite freeness information we must trace all possible dependencies. These dependencies are established via the constraints in the program either directly or through entailment. The derived information is useful for example to perform constraint reordering (cf. [Dumortier 1994]).

The most closely related work to our freeness analysis is the **LSign** abstraction of Marriott and Stuckey [1993] that describes sets of linear equations and inequalities. In [Marriott and Stuckey 1994], this domain is further elaborated and extended towards the treatment of non-linear constraints and unification constraints. The major advantage of the abstraction compared with ours is its enhanced precision, especially for inequalities but also for equations (it keeps track of the constraint symbol and the sign of the coefficients, which are discarded in our analysis). However, the main deficiencies are that (1) no implementation is reported, such that the efficiency (especially with respect to the increased precision) cannot be judged and (2) some aspects that are relevant in order to obtain a complete analyzer are not (sufficiently) elaborated (such as procedure-exit, the upper bound operation, the order relation and the interaction between the unification and the numerical part). Recently [Ramachandran and Van Hentenryck 1995] described some improvements.

7.1 Abstract domain and abstraction function

Let $\mathbf{c}$ denote a constraint. A set of variables $\{\mathbf{x}_1, \ldots, \mathbf{x_n}\} \subseteq \mathbf{vars(c)}$ is *constrained* by $\mathbf{c}$ iff there exists a set of values $\{\mathbf{v}_1, \ldots, \mathbf{v_n}\}$, with each $\mathbf{v_i}$ in the domain of $\mathbf{x_i}$, such that $\mathbf{c} \wedge \mathbf{x}_1 = \mathbf{v}_1 \wedge \ldots \wedge \mathbf{x_n} = \mathbf{v_n}$ is inconsistent while for any $\{\mathbf{i}_1, \ldots, \mathbf{i_m}\} \subset \{1, \ldots, \mathbf{n}\}$ it holds that $\mathbf{c} \wedge \mathbf{x_{i_1}} = \mathbf{v_{i_1}} \wedge \ldots \wedge \mathbf{x_{i_m}} = \mathbf{v_{i_m}}$ is consistent.

Example 7.1.1.
Let $\mathbf{c}$ be $\mathbf{x} = \mathbf{f}(\mathbf{y}_1, \ldots, \mathbf{y_n})$ $(\mathbf{n} \geq 0)$. The sets constrained by $\mathbf{c}$ are $\{\mathbf{x}\}$, $\{\mathbf{x}, \mathbf{y}_1\}, \ldots,$ $\{\mathbf{x}, \mathbf{y_n}\}$ (e.g. for $\{\mathbf{x}, \mathbf{y}_1\}$, $\mathbf{c} \ \wedge \ \mathbf{x} = \mathbf{f}(1, \ldots, \mathbf{n}) \ \wedge \ \mathbf{y}_1 = 3$ is inconsistent while any subpart of the conjunction is consistent).

Let $\mathbf{c}$ be $\mathbf{a}_1\mathbf{x}_1 + \ldots + \mathbf{a_n x_n} = \mathbf{b}$ $(\mathbf{n} \geq 1)$ where the $\mathbf{a_i}$ and $\mathbf{b}$ are numbers $(\mathbf{a_i} \neq 0)$. Then $\mathbf{c}$ constrains the set $\{\mathbf{x}_1, \ldots, \mathbf{x_n}\}$.

Let $\mathbf{c}$ be $\mathbf{x} > \mathbf{y}$. Then $\mathbf{c}$ constrains $\{\mathbf{x}, \mathbf{y}\}$.

Let $\mathbf{c}$ be $\mathbf{x} = \mathbf{y} * \mathbf{z}$. Then $\mathbf{c}$ constrains $\{\mathbf{x}, \mathbf{y}\}$, $\{\mathbf{x}, \mathbf{z}\}$ and $\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$ (for $\{\mathbf{x}, \mathbf{z}\}$, $\mathbf{c} \ \wedge \ \mathbf{x} = 1 \ \wedge \ \mathbf{z} = 0$ is inconsistent while $\mathbf{c} \ \wedge \ \mathbf{x} = 1$ and $\mathbf{c} \ \wedge \ \mathbf{z} = 0$ are consistent; for $\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$, $\mathbf{c} \ \wedge \ \mathbf{x} = 2 \ \wedge \ \mathbf{y} = 1 \ \wedge \ \mathbf{z} = 1$ is inconsistent while any subpart of the conjunction is consistent).

Let $\mathbf{c}$ be $\mathbf{x} =< \mathbf{y} > .\mathbf{w}.\mathbf{z}$. The sets constrained by $\mathbf{c}$ are $\{\mathbf{x}, \mathbf{y}\}, \{\mathbf{x}, \mathbf{w}\}$ and $\{\mathbf{x}, \mathbf{z}\}$.

A variable $\mathbf{x}$ is *free* in $\mathbf{c}$ iff $\{\mathbf{x}\}$ is not constrained by $\mathbf{c}$, so freeness can be derived by safely approximating all possible constrained sets. A constrained set $\{\mathbf{x}_1, \ldots, \mathbf{x_n}\}$ with $\mathbf{n} > 1$ indicates a possible dependency between those variables in the sense that constraining all variables but for example $\mathbf{x_i}$ can constrain $\mathbf{x_i}$ (can cause non-freeness of $\mathbf{x_i}$). Such constrained sets are the key concept used to perform non-freeness propagation. The formal development in [Dumortier 1994] (which is too long to include) shows that constrained sets that can be obtained as union of others (e.g. the set $\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$ in the last example), and also unions of constrained sets, are redundant with respect to non-freeness propagation (the sub-dependencies impose stronger restrictions). These *non-minimal* sets can therefore be omitted in the abstraction[14].

Definition 7.1.2. Minimal set
Let $\mathbf{SS} \in \wp(\wp_\emptyset(\mathbf{Pvar}))$. Then $\mathbf{S} \in \mathbf{SS}$ is minimal in $\mathbf{SS}$ iff $\nexists \mathbf{S}_1, \ldots, \mathbf{S_m} \in \mathbf{SS} \setminus \{\mathbf{S}\}$ $(\mathbf{m} \geq 2)$ such that $\mathbf{S} = \mathbf{S}_1 \cup \ldots \cup \mathbf{S_m}$.

Definition 7.1.3. $\min\mathcal{F}$
Let $\mathbf{SS} \in \wp(\wp_\emptyset(\mathbf{Pvar}))$. Then $\mathbf{min}\mathcal{F}(\mathbf{SS}) = \{\mathbf{S} \in \mathbf{SS} \mid \mathbf{S\ is\ a\ minimal\ set\ in\ SS}\}$.

Definition 7.1.4. Abstraction of a constraint: $\alpha^{\mathbf{f}}$
Let $\mathbf{c}$ be a constraint.
Then $\alpha^{\mathbf{f}}(\mathbf{c}) = \perp$ if $\neg\mathbf{consistent}(\mathbf{c})$, otherwise $\alpha^{\mathbf{f}}(\mathbf{c}) = \mathbf{min}\mathcal{F}(\{\{\mathbf{x}_1, \ldots, \mathbf{x_n}\} \subseteq \mathbf{vars}(\mathbf{c}) \mid \{\mathbf{x}_1, \ldots, \mathbf{x_n}\}\ \mathbf{is\ constrained\ by\ c}\})$.

The abstract domain $\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$ can now be formally defined as $\{\perp\} \cup \{\mathbf{AC} \in \wp(\wp_\emptyset(\mathbf{Pvar})) \mid \mathbf{min}\mathcal{F}(\mathbf{AC}) = \mathbf{AC}\}$[15].

While it is rather straightforward to derive the abstraction of *primitive* constraints, it is more involved for a *conjunction* of primitive constraints. Let us consider some examples.

Example 7.1.5.
Let $\mathbf{c}$ be $\mathbf{y} = \mathbf{f}(\mathbf{g}(\mathbf{x})) \ \wedge\ \mathbf{z} = \mathbf{x}$. Constrained sets are $\{\mathbf{y}\}$ and $\{\mathbf{y}, \mathbf{x}\}$ (from the first primitive constraint) and $\{\mathbf{z}, \mathbf{x}\}$ (from the second primitive constraint), but also $\{\mathbf{y}, \mathbf{z}\}$ from the entailed primitive constraint $\mathbf{y} = \mathbf{f}(\mathbf{g}(\mathbf{z}))$.

Let $\mathbf{c}$ be $\mathbf{x} + \mathbf{y} = 3 \ \wedge\ \mathbf{y} - \mathbf{z} = 2$. Constrained sets are $\{\mathbf{x}, \mathbf{y}\}$ and $\{\mathbf{y}, \mathbf{z}\}$ but also $\{\mathbf{x}, \mathbf{z}\}$ as there is an entailed primitive constraint $\mathbf{x} + \mathbf{z} = 1$.

This suggests that it is sufficient to consider the constrained sets for all entailed primitive constraints. However, this does not suffice for conjunctions composed of constraints of different constraint domains, as shown by the following example.

Example 7.1.6.
Let $\mathbf{c}$ be $\mathbf{x} = \mathbf{f}(\mathbf{u}, \mathbf{v}) \ \wedge\ \mathbf{u} - \mathbf{v} + \mathbf{t} = 3$. Besides the constrained sets of the first conjunct ($\{\mathbf{x}\}$, $\{\mathbf{x}, \mathbf{u}\}$ and $\{\mathbf{x}, \mathbf{v}\}$) and of the second conjunct ($\{\mathbf{u}, \mathbf{v}, \mathbf{t}\}$), there is also a constrained set $\{\mathbf{x}, \mathbf{t}\}$. Indeed, e.g. $\mathbf{c} \ \wedge\ \mathbf{x} = \mathbf{f}(1, 2) \ \wedge\ \mathbf{t} = 1$ is inconsistent while any subpart of the conjunction is consistent.

In our implementation, we have not attempted to compute constrained sets of non-primitive constraints, but rather use abstract conjunction to obtain their abstraction from the abstractions of the composing conjuncts. It is recommended to first put the conjunction in solved form as the presence of redundant conjuncts will

[14] The non-minimal freeness abstraction of a constraint $\mathbf{c}$ as developed in [Dumortier et al. 1993] exhaustively enumerates not only minimal constrained sets in $\mathbf{c}$ but also all possible unions of these. These unions are needed at abstract conjunction (cf. Definition 7.2.2). Adding the unions at once instead of computing them at abstract conjunction contributes to the precision of the analysis. However, it also limits its practical use as the size of the abstractions is in the worst case exponential with respect to the number of variables.

[15] For reasons of readability most of the following definitions and operations do not explicitly deal with $\perp$. Their extensions are trivial.

severely affect precision[16]. Even in the absence of redundancy, one can obtain a more precise result when starting from the solved form, as will be illustrated below. For the Herbrand domain, the solved form can be obtained by applying the Martelli-Montanari unification algorithm [Martelli and Montanari 1982]; for generalized linear constraints, a solved form can be obtained by the algorithm of Lassez and McAloon [1992].

Before discussing abstract conjunction, let us first further develop the abstract domain.

Definition 7.1.7. Order relation
Let $\mathbf{AC}_1$, $\mathbf{AC}_2 \in \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$. Then $\mathbf{AC}_1 \leq^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}_2$ iff $\mathbf{AC}_1 \subseteq \mathbf{close}(\mathbf{AC}_2)$ where $\mathbf{close}(\mathbf{AC})$ is the closure under union of $\mathbf{AC}$.

Definition 7.1.8. Equivalence
Let $\mathbf{AC}_1$, $\mathbf{AC}_2 \in \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$. Then $\mathbf{AC}_1 \equiv^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}_2$ iff $\mathbf{AC}_1 = \mathbf{AC}_2$.

Definition 7.1.9. Least upper bound
Let $\mathbf{AC}_1$, $\mathbf{AC}_2 \in \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$. Then $\mathbf{upp}^{\mathcal{F}^{\mathbf{m}}}(\mathbf{AC}_1, \mathbf{AC}_2) = \mathbf{min}\mathcal{F}(\mathbf{AC}_1 \cup \mathbf{AC}_2)$.

This definition can easily be extended to compute the least upper bound of m (m > 2) abstractions. In the following we will assume that **upp** applies to a set of abstract constraints.

To abstract a set of constraints, ideally $\wp(\wp(\wp_{\emptyset}(\mathbf{Pvar})))$ should be the abstract domain. However, this may give rise to impractically large abstractions. Therefore, the abstraction of a set of constraints is approximated by the least upper bound of the abstractions of the individual constraints in the set.

Definition 7.1.10. Abstraction of a set of constraints: $\alpha^{\mathcal{F}^{\mathbf{m}}}$
Let $\mathbf{C} \in \mathbf{Cons}^{\mathcal{C}}$.
Then $\alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}) = \perp$ if $\mathbf{C} = \emptyset$, otherwise $\alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}) = \mathbf{upp}(\{\alpha^{\mathbf{f}}(\mathbf{c}) \mid \mathbf{c} \in \mathbf{C}\})$.

Definition 7.1.11. Maximal and minimal elements
The maximal element is $\mathbf{min}\mathcal{F}(\wp(\wp_{\emptyset}(\mathbf{Pvar}))) = \{\{\mathbf{x}\} \mid \mathbf{x} \in \mathbf{Pvar}\}$.
The minimal element is $\perp$.

The concretization function $\gamma^{\mathcal{F}^{\mathbf{m}}}$ can be defined based upon $\alpha^{\mathcal{F}^{\mathbf{m}}}$ as described in [Cousot and Cousot 1992a]: $\gamma^{\mathcal{F}^{\mathbf{m}}}(\mathbf{AC}) = \bigcup\{\mathbf{C} \in \mathbf{Cons}^{\mathcal{C}} \mid \alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}) \leq^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}\}$. Then $(\mathbf{Cons}^{\mathcal{C}}, \subseteq, \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}, \leq^{\mathcal{F}^{\mathbf{m}}})$ is a Galois insertion [Dumortier 1994].

7.2 Abstract projection and abstract conjunction functions

Definition 7.2.1. Abstract projection
Let $\mathbf{AC} \in \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$ and $\tilde{\mathbf{x}}$ be a sequence of variables.
Then $\exists^{\mathcal{F}^{\mathbf{m}}}_{-\tilde{\mathbf{x}}} \mathbf{AC} = \{\mathbf{S} \in \mathbf{AC} \mid \mathbf{S} \subseteq \tilde{\mathbf{x}}\}$.

[16] This is not done in the actual implementation based on the PLAI system, which is written in Prolog. In this case the only highly efficient solved form algorithm readily available in the system itself is the one for unification constraints inherited from the Prolog implementation. However, as pointed out in [Codognet and Filé 1992], implementing the system in the CLP language to be analyzed would allow to use all built-in solved form algorithms. On the other hand it should also be noted that for the actual benchmarks analyzed in Section 9 not applying the solved form algorithm does not affect precision.

The abstract conjunction of two abstract constraints $\mathbf{AC}_1$ and $\mathbf{AC}_2$, denoted $\mathbf{AC}_1 \wedge^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}_2$, must safely approximate the constrained sets of all constraints $\mathbf{c}_1 \wedge \mathbf{c}_2$ where $\mathbf{c}_1$ and $\mathbf{c}_2$ are abstracted by $\mathbf{AC}_1$ and $\mathbf{AC}_2$ respectively. It is obvious that constrained sets of $\mathbf{c}_1$ resp. $\mathbf{c}_2$ are also constrained sets of $\mathbf{c}_1 \wedge \mathbf{c}_2$. Actually, if $\mathbf{c}_1$ and $\mathbf{c}_2$ do not share variables, these are the only ones. The hard case is when $\mathbf{c}_1$ and $\mathbf{c}_2$ do share variables. Consider a simple example in the numerical domain. Let $\mathbf{c}_1$ be $\mathbf{x} = \mathbf{y} \wedge \mathbf{u} = \mathbf{v}$ and $\mathbf{c}_2$ be $\mathbf{y} + \mathbf{v} = \mathbf{z}$. Constrained sets of $\mathbf{c}_1$ are $\{\mathbf{x}, \mathbf{y}\}$ and $\{\mathbf{u}, \mathbf{v}\}$; $\{\mathbf{y}, \mathbf{v}, \mathbf{z}\}$ is the only constrained set of $\mathbf{c}_2$. The conjunction $\mathbf{c}_1 \wedge \mathbf{c}_2$ entails constraints $\mathbf{x} + \mathbf{v} = \mathbf{z}$, $\mathbf{y} + \mathbf{u} = \mathbf{z}$ and $\mathbf{x} + \mathbf{u} = \mathbf{z}$, giving rise to the constrained sets $\{\mathbf{x}, \mathbf{v}, \mathbf{z}\}$, $\{\mathbf{y}, \mathbf{u}, \mathbf{z}\}$ and $\{\mathbf{x}, \mathbf{u}, \mathbf{z}\}$. At the concrete level, the key operation in obtaining entailed constraints is variable elimination. At the abstract level, the operation is mimicked by taking the union of an element of $\mathbf{close}(\mathbf{AC}_1)$ (which, to abstract $\mathbf{c}_1$, must contain $\{\mathbf{x}, \mathbf{y}\}$, $\{\mathbf{u}, \mathbf{v}\}$ and $\{\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{v}\}$) and an element of $\mathbf{close}(\mathbf{AC}_2)$ (which, to abstract $\mathbf{c}_2$, must contain $\{\mathbf{y}, \mathbf{v}, \mathbf{z}\}$) and removing some elements from the intersection: removing $\mathbf{y}$ and $\mathbf{v}$ from $\{\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{v}\} \cup \{\mathbf{y}, \mathbf{v}, \mathbf{z}\}$ yields $\{\mathbf{x}, \mathbf{u}, \mathbf{z}\}$, removing $\mathbf{y}$ from $\{\mathbf{x}, \mathbf{y}\} \cup \{\mathbf{y}, \mathbf{v}, \mathbf{z}\}$ yields $\{\mathbf{x}, \mathbf{v}, \mathbf{z}\}$ and deleting $\mathbf{v}$ from $\{\mathbf{u}, \mathbf{v}\} \cup \{\mathbf{y}, \mathbf{v}, \mathbf{z}\}$ yields $\{\mathbf{y}, \mathbf{u}, \mathbf{z}\}$. Notice that one should not only remove the complete intersection, as shown by the following example. Consider $\{\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{v}\}$ as element of $\mathbf{AC}_1$ which abstracts for example $\mathbf{c}_1 \equiv \mathbf{x} + \mathbf{y} = \mathbf{u} + \mathbf{v}$, and $\{\mathbf{x}, \mathbf{y}, \mathbf{t}\}$ as element of $\mathbf{AC}_2$ which abstracts for example $\mathbf{c}_2 \equiv \mathbf{x} + \mathbf{y} = \mathbf{t}$ but also $\mathbf{c}_2' \equiv \mathbf{x} + 2\mathbf{y} = \mathbf{t}$. Now $\mathbf{c}_1 \wedge \mathbf{c}_2$ entails $\mathbf{t} = \mathbf{u} + \mathbf{v}$ with constrained set $\{\mathbf{t}, \mathbf{u}, \mathbf{v}\}$, while $\mathbf{c}_1 \wedge \mathbf{c}_2'$ entails $\mathbf{y} = \mathbf{t} - \mathbf{u} - \mathbf{v}$ and $\mathbf{x} = 2\mathbf{u} + 2\mathbf{v} - \mathbf{t}$ with constrained sets $\{\mathbf{t}, \mathbf{u}, \mathbf{v}, \mathbf{y}\}$ and $\{\mathbf{t}, \mathbf{u}, \mathbf{v}, \mathbf{x}\}$. This also illustrates that computing the abstraction of $\mathbf{c}_1 \wedge \mathbf{c}_2$ by abstract conjunction of the abstractions of $\mathbf{c}_1$ and $\mathbf{c}_2$ can be less precise than directly determining the constrained sets of the conjunction (which can be done by first transforming the conjunction to solved form).

In [Dumortier 1994] it is shown how a similar reasoning applies for Herbrand constraints, PrologIII tuple constraints and mixed constraints (over more than one constraint domain) and that abstract conjunction as defined below always yields a safe approximation (the proof is too long to be included here).

Definition 7.2.2. Abstract conjunction
Let $\mathbf{AC}_1, \mathbf{AC}_2 \in \mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$.
Then $\mathbf{AC}_1 \wedge^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}_2 = \mathbf{min}\mathcal{F}(\mathbf{AC}_1 \cup \mathbf{AC}_2 \cup (\mathbf{close}(\mathbf{AC}_1) \oplus \mathbf{close}(\mathbf{AC}_2)))$ where $\mathbf{SS}_1 \oplus \mathbf{SS}_2 = \left\{ (\mathbf{S}_1 \cup \mathbf{S}_2) \setminus \mathbf{D} \,\middle|\, \mathbf{S}_1 \in \mathbf{SS}_1, \mathbf{S}_2 \in \mathbf{SS}_2, \mathbf{D} \subseteq \mathbf{S}_1 \cap \mathbf{S}_2, \mathbf{D} \neq \emptyset \right\} \setminus \{\emptyset\}$ and $\mathbf{close}(\mathbf{AC})$ is the closure under union of $\mathbf{AC}$[17].

An equivalent but more efficient algorithm corresponding to Definition 7.2.2 is obtained by closing only the necessary parts of $\mathbf{AC}_1$ and $\mathbf{AC}_2$ (i.e. those parts containing common variables) and by taking care of not generating non-minimal sets when combining the two.

[17] $\mathbf{AC}_1$ and $\mathbf{AC}_2$ are abstractions of sets of constraints, that are obtained by joining the abstractions of the individual constraints in the set (Definition 7.1.10). Thus, closing $\mathbf{AC}_1$ and $\mathbf{AC}_2$ at abstract conjunction implies that also constrained sets originating from different (independent) constraints are combined. This results in a possible loss of precision. The non-minimal freeness abstraction of [Dumortier et al. 1993], however, exhaustively represents all combinations of constrained sets when abstracting each constraint (instead of computing these combinations at abstract conjunction) and hence prevents the loss of precision.

Let us now present how the abstract operations required by the framework are computed. One can take the same approach as in Section 6, defining $\mathbf{AC_{entry}}$, $\mathbf{AC^{j}_{in}}$, $\mathbf{AC_{exit}}$ and $\mathbf{AC'}$ in terms of abstract projection $\exists^{\mathcal{F}^{\mathbf{m}}}$, abstract conjunction $\wedge^{\mathcal{F}^{\mathbf{m}}}$ and abstraction $\alpha^{\mathbf{f}}$. However, this results in a very poor precision. The reason is that it is disastrous to precision to add a numerical constraint to an abstract constraint store which already describes that constraint. For example, let $\mathbf{c}$ be the constraint $\mathbf{a_1x_1 + \ldots + a_nx_n = a_{n+1}}$ and $\mathbf{AC}$ an abstract constraint store containing its abstraction, i.e. $\{\mathbf{x_1, \ldots, x_n}\} \in \mathbf{AC}$. Performing $\alpha^{\mathbf{f}}(\mathbf{c}) \wedge^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC}$ creates an abstract constraint store $\mathbf{AC'}$ which includes the singleton $\{\mathbf{x_i}\}$ for each of the variables $\mathbf{x_i}$; hence, $\mathbf{AC'}$ indicates that each $\mathbf{x_i}$ is possibly non-free. This computation reflects that $\mathbf{AC}$ abstracts an equation $\mathbf{c'}$, $\mathbf{b_1x_1 + \ldots + b_nx_n = b_{n+1}}$. With an appropriate choice of values for $\mathbf{b_1, \ldots, b_n}$, the constraint $\mathbf{c} \wedge \mathbf{c'}$ entails a constraint $\mathbf{dx_i = e}$ which is abstracted as $\{\{\mathbf{x_i}\}\}$, so it is required that $\{\mathbf{x_i}\} \in \mathbf{AC'}$. When **abstract_entry** passes an abstraction of a constraint to the entered procedure, then **abstract_exit** returns it and the computation of $\mathbf{AC'}$ as suggested above destroys the freeness of all the involved variables.

To overcome this problem, we slightly revise the concrete semantics: a constraint $\mathbf{c}$ is represented as a pair $(\mathbf{c_{old}}, \mathbf{c_{new}})$ such that $\mathbf{c} = \mathbf{c_{old}} \wedge \mathbf{c_{new}}$. The corresponding rewrite rules are:

—The c-transition (if consistent):
$\mathbf{S} :: \langle \mathbf{c}, \mathbf{G}; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle \xrightarrow{\mathbf{c}} \mathbf{S} :: \langle \mathbf{c}, \mathbf{G}; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle :: \langle \mathbf{G}; (\mathbf{c_{old}}, \mathbf{c_{new}} \wedge \mathbf{c})\rangle$.

—The r-transition:
$\mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle \xrightarrow{\mathbf{r}} \mathbf{S} :: \langle \mathbf{a}, \mathbf{G}; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle :: \langle \mathbf{b_1}, \ldots, \mathbf{b_n};\ (\exists_{-\mathbf{vars}(\rho)}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{old}} \wedge \mathbf{c_{new}}), \mathbf{true})\rangle$
where $\rho : \mathbf{h}\text{:-}\mathbf{b_1}, \ldots, \mathbf{b_n}$.

—The exit transition:
$\mathbf{S_1} :: \langle \mathbf{a}, \mathbf{G}; (\mathbf{c_o}, \mathbf{c_n})\rangle \overset{\rho}{::} \langle \mathbf{b_1}, \ldots, \mathbf{b_n}; (\mathbf{c_{old}}, \mathbf{true})\rangle :: \mathbf{S_2} :: \langle \Box_\rho; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle \xrightarrow{\mathbf{exit}} \mathbf{S_1} :: \langle \mathbf{a}, \mathbf{G}; (\mathbf{c_o}, \mathbf{c_n})\rangle \overset{\rho}{::} \langle \mathbf{b_1}, \ldots, \mathbf{b_n}; (\mathbf{c_{old}}, \mathbf{true})\rangle :: \mathbf{S_2} :: \langle \Box_\rho; (\mathbf{c_{old}}, \mathbf{c_{new}})\rangle :: \langle \mathbf{G}; (\mathbf{c_o}, \mathbf{c_n} \wedge \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{new}}))\rangle$.
The modification of the exit transition is valid because $\mathbf{c_o} \wedge \mathbf{c_n} \wedge \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{old}} \wedge \mathbf{c_{new}})$ where $\mathbf{c_{old}} = \exists_{-\mathbf{vars}(\rho)}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_o} \wedge \mathbf{c_n})$ is equivalent with $\mathbf{c_o} \wedge \mathbf{c_n} \wedge \exists_{-\mathbf{vars(a)}}(\mathbf{a} = \mathbf{h} \wedge \mathbf{c_{new}})$.

Now abstract constraint stores are also represented by a pair $(\mathbf{AC_{old}}, \mathbf{AC_{new}})$. The idea is that $(\mathbf{c_{old}}, \mathbf{c_{new}}) \in \gamma((\mathbf{AC_{old}}, \mathbf{AC_{new}}))$ iff $\mathbf{c_{new}} \in \gamma(\mathbf{AC_{new}})$ and $\mathbf{c_{old}} \wedge \mathbf{c_{new}} \in \gamma(\mathbf{AC_{old}} \cup \mathbf{AC_{new}})$. Reformulating the safety conditions of the framework for these modifications is a rather straightforward task and is omitted. The abstract operations can be defined as follows. With $\mathbf{g}$ a constraint and $(\mathbf{AC_{old}}, \mathbf{AC_{new}})$ its abstract call constraint, $\mathbf{AC'}$ is defined as $(\mathbf{AC_{old}}, \mathbf{AC_{new}} \wedge^{\mathcal{F}^{\mathbf{m}}} \alpha^{\mathbf{f}}(\mathbf{g}))$. With $\mathbf{g}$ an atom, $(\mathbf{AC_{old}}, \mathbf{AC_{new}})$ its abstract call constraint, $\rho_1, \ldots, \rho_{\mathbf{m}}$ the rules of the program $\mathbf{P}$ defining the predicate of $\mathbf{g}$, and $\rho_{\mathbf{j}} : \mathbf{h_j}\text{:-}\mathbf{b_{j1}}, \ldots, \mathbf{b_{jn_j}}$, $\mathbf{AC_{entry}}$ is defined as $\exists^{\mathcal{F}^{\mathbf{m}}}_{-\mathbf{vars(g)}}(\mathbf{AC_{old}} \cup \mathbf{AC_{new}})$ and $\mathbf{AC^{j}_{in}}$ is defined as $(\exists^{\mathcal{F}^{\mathbf{m}}}_{-\mathbf{vars}(\rho_{\mathbf{j}})}(\mathbf{AC_{entry}} \wedge^{\mathcal{F}^{\mathbf{m}}} \alpha^{\mathbf{f}}(\mathbf{g} = \mathbf{h_j})), \emptyset)$. Finally, with $(\mathbf{AC^{1}_{old}}, \mathbf{AC^{1}_{new}}), \ldots, (\mathbf{AC^{m}_{old}}, \mathbf{AC^{m}_{new}})$ the abstract out constraints of rules $\rho_1, \ldots, \rho_{\mathbf{m}}$, $\mathbf{AC_{exit}}$ is defined as $\mathbf{upp}(\mathbf{AC^{1}_{exit}}, \ldots, \mathbf{AC^{m}_{exit}})$ where $\mathbf{AC^{j}_{exit}} = \exists^{\mathcal{F}^{\mathbf{m}}}_{-\mathbf{vars(g)}}(\mathbf{AC^{j}_{new}} \wedge^{\mathcal{F}^{\mathbf{m}}} \alpha^{\mathbf{f}}(\mathbf{g} = \mathbf{h_j}))$ and $\mathbf{AC'}$ is defined as $(\mathbf{AC_{old}}, \mathbf{AC_{new}} \wedge^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC_{exit}})$ and, in **extension_from_table** as $(\mathbf{AC_{old}}, \mathbf{AC_{new}} \wedge^{\mathcal{F}^{\mathbf{m}}} \mathbf{AC^{tab}}\delta_{\mathbf{g}=\mathbf{g^{tab}}})$.

Notice that $\mathbf{AC_{entry}}$, $\mathbf{AC_{exit}}$ and all entries in Table are not pairs but elements of $\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$ and that $\mathbf{AC^{j}_{old}}$ does not contribute to $\mathbf{AC^{j}_{exit}}$. For further details, the reader is referred to [Dumortier 1994].

Making the distinction between new and old information in the analysis of logic programs has been applied previously by Plaisted [1984] and also by Mulkers [1993], Mulkers et al. [1990] and Mulkers et al. [1994].

Example 7.2.3. $\mathcal{F}^{\mathbf{m}}$ analysis for the sumlist program
The initial call pattern of $\mathbf{sumlist}(\mathbf{A},\mathbf{B})$ is $\{\ \{\mathbf{A}\}\ \}$, which is also the call pattern of the recursive call (the abstract information written out is the union of the old and new components of the compound abstract constraints).

```
sumlist(x, w)  :-          % { {x} }
       {x = [ ],           % { {x} }
       w = 0}.             % { {x}, {w} }
sumlist(x, w) :-           % { {x} }
       {x = [y | z],       % { {x}, {y}, {z} }
       w = y + w'},        % { {x}, {y}, {z}, {w, w'} }
       sumlist(z,w').      % { {x}, {y}, {z}, {w}, {w'} }
```

The analysis indicates that at the end of each rule, $\mathbf{x}$ and $\mathbf{w}$ are possibly non-free. In the second rule, $\mathbf{w}$ and $\mathbf{w}'$ are free before the recursive call and depend on each other.

8. COMBINING THE TWO DOMAINS

The information inferred by the definiteness analysis and the freeness analysis of the previous two sections is enough to obtain a full mode system: the former provides modes **d** and **a** and the latter modes **f** and **a**. In a combination along the lines of the paper [Cousot and Cousot 1979] (applied in [Codish et al. 1995]), the abstract domains and the original components of the basic operations remain the same, while during analysis interactions between the computed abstractions occur to refine them. This results in a precise combined analysis, in particular when the analyses being composed contain a sufficient degree of "overlapping" information. As our domains are in a sense complementary, we present another kind of combination. Essentially, the **D** part of the definiteness analysis can be used as additional knowledge for the freeness abstraction. In this section we briefly present the improved freeness abstraction $\mathbf{Cons}^{\mathcal{D}\mathcal{F}^{\mathbf{m}}}$ which is based on the minimal freeness abstraction and which uses additional knowledge about *definiteness* of program variables. In section 9.1, we discuss how such an interaction between analyzers can be realized in a practical abstract interpretation system such as PLAI.

Definite variables occur in the minimal freeness abstraction as possibly non-free variables. The presence of their corresponding singletons implies that the abstract operations have to take them into account – for example when computing the closure under union – although they play a very specific role in the propagation of possible non-freeness. Efficiency of the analysis can be improved by separating out the definite variables. The assumption that the definite variables are known is reasonable as the definiteness analysis computes a safe approximation (denoted by $\mathbf{defvars}(\alpha^{\mathcal{D}}(\mathbf{C}))$).

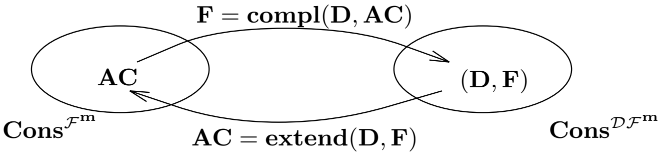
$\mathbf{F} = \mathbf{compl}(\mathbf{D}, \mathbf{AC})$
$\mathbf{AC}$
$(\mathbf{D}, \mathbf{F})$
$\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$
$\mathbf{AC} = \mathbf{extend}(\mathbf{D}, \mathbf{F})$
$\mathbf{Cons}^{\mathcal{DF}^{\mathbf{m}}}$

Fig. 2. Relation between $\mathcal{F}^{\mathbf{m}}$ and $\mathcal{DF}^{\mathbf{m}}$ abstraction (for a given $\mathbf{D}$)

Given the set of definite variables $\mathbf{D}$, $\alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C})$ can be split into a set of singletons containing definite variables and a set of sets containing no definite variables, namely $\mathbf{compl}(\mathbf{D}, \alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C})) = \{\mathbf{S} \in \alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}) \mid \mathbf{S} \cap \mathbf{D} = \emptyset\}$. The $\mathcal{DF}^{\mathbf{m}}$ abstraction is based on the observation that the minimal freeness abstraction can be expressed in terms of $\mathbf{compl}(\mathbf{D}, \alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}))$ and $\mathbf{D}$ (without loss of precision). The abstract domain $\mathbf{Cons}^{\mathcal{DF}^{\mathbf{m}}}$ is a set of pairs $(\mathbf{D}, \mathbf{F})$ where $\mathbf{D} \subseteq \mathbf{Pvar}$ and $\mathbf{F} \in \wp(\wp_{\emptyset}(\mathbf{Pvar} \backslash \mathbf{D}))$ such that $\mathbf{min}\mathcal{F}(\mathbf{F}) = \mathbf{F}$, to which $\bot$ is added as minimal element.

Definition 8.0.1. Abstraction of a set of constraints: $\alpha^{\mathcal{DF}^{\mathbf{m}}}$
Let $\mathbf{C} \in \mathbf{Cons}^{\mathcal{C}}$. Then $\alpha^{\mathcal{DF}^{\mathbf{m}}}(\mathbf{C}) = \bot$ if $\mathbf{C} = \emptyset$, otherwise $\alpha^{\mathcal{DF}^{\mathbf{m}}}(\mathbf{C}) = (\mathbf{D}, \mathbf{F})$ where $\mathbf{D}$ could be given by $\mathbf{defvars}(\alpha^{\mathcal{D}}(\mathbf{C}))$ and $\mathbf{F} = \mathbf{compl}(\mathbf{D}, \alpha^{\mathcal{F}^{\mathbf{m}}}(\mathbf{C}))$.

There is a 1-to-1 correspondence between the abstractions in $\mathbf{Cons}^{\mathcal{DF}^{\mathbf{m}}}$ and $\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$ and vice-versa, for a given $\mathbf{D}$ (see Figure 2).

Definition 8.0.2. extend
Let $(\mathbf{D}, \mathbf{F}) \in \mathbf{Cons}^{\mathcal{DF}^{\mathbf{m}}}$. Then $\mathbf{extend}(\mathbf{D}, \mathbf{F}) = \mathbf{F} \cup \{\{\mathbf{x}\} \mid \mathbf{x} \in \mathbf{D}\}$.

The operations on $\mathbf{Cons}^{\mathcal{DF}^{\mathbf{m}}}$ are based on the corresponding operations on $\mathbf{Cons}^{\mathcal{D}}$ and $\mathbf{Cons}^{\mathcal{F}^{\mathbf{m}}}$. For their exact definitions we refer to [Dumortier and Janssens 1994; Dumortier 1994]. Concerning abstract conjunction of two abstract constraints $(\mathbf{D}_1, \mathbf{F}_1)$ and $(\mathbf{D}_2, \mathbf{F}_2)$, an efficient operation is obtained as follows: the $\mathbf{D}$ parts are joined first and then the obtained definiteness information is propagated onto the freeness parts $\mathbf{F}_1$ and $\mathbf{F}_2$, thus reducing them considerably, before these are joined. Consequently, the $\mathcal{DF}^{\mathbf{m}}$ abstract conjunction is much more efficient than if one would perform the $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ abstract conjunctions on the $\mathbf{D}$ and $\mathbf{F}$ parts separately, afterwards deleting the definite variables from the resulting $\mathbf{F}$ part.

Again, as for the $\mathcal{F}^{\mathbf{m}}$ analysis, an abstract constraint should be split into an old component, containing the information passed down from a calling environment, and a new component, containing the information that is gathered during local analysis of the rule body. Otherwise, too much precision would be lost at **abstract_exit**.

Example 8.0.3. $\mathcal{DF}^{\mathbf{m}}$ analysis for the sumlist program
The initial call pattern is $\mathbf{sumlist}(\mathbf{d}, \mathbf{f})$, which is also the call pattern of the recursive call. The definiteness information is as in [García de la Banda and Hermenegildo 1993]; we obtain the same freeness information as in Example 7.2.3 but in a more compact form (old and new components of the freeness part are put together).

```
sumlist(x, w) :-          % ({x}, ∅)
    { x = [ ],            % ({x}, ∅)
    w = 0 }.              % ({x, w}, ∅)
sumlist(x, w) :-          % ({x}, ∅)
    { x = [y | z],        % ({x, y, z}, ∅)
    w = y + w′ },         % ({x, y, z}, {{w, w′}})
    sumlist(z, w′).       % ({x, w, y, z, w′}, ∅)
```

9. EXPERIMENTAL RESULTS

In this section we present the results of the experiments that we have performed in order to evaluate the efficiency and accuracy of the analyses. We start by describing the implementation and the benchmarks used. Our attention then first focuses on the issue of efficiency, and, thus, of the feasibility and scalability of the approach. This is an important issue since it has been shown that even relatively simple analyses of LP programs have worst case exponential behavior [Debray 1995]. On the other hand, it has also been shown experimentally that average case behaviors have much better characteristics for typical analyses [Warren et al. 1988; Van Roy and Despain 1992; Muthukumar and Hermenegildo 1992; Debray 1992b; Le Charlier and Van Hentenryck 1994; Bueno et al. 1994]. It is obviously interesting to explore if this practical behavior carries over to our CLP analyses, both when analyzing CLP programs and also when analyzing traditional LP programs (for comparison with LP analyzers). To study this point, we present a summary of the analysis times for a set of benchmarks which includes CLP programs (both relatively small and larger ones) and also LP programs. The larger CLP programs are the largest size programs available to us and they should be instrumental in giving an idea of the scalability of the results in the new application area.

We then focus on the effect of an important technique related to the scalability issue: the application of widening operations in order to trade precision for efficiency. We investigate the effects on the efficiency and precision of our analyses of the introduction of widening in the freeness abstraction.

Finally, we perform a more detailed evaluation, focused on a representative set of CLP-programs, in order to gain insight into the potential of the analyses, the main causes for loss of accuracy, and the advantages and disadvantages of the combined analyses. We do not address herein the obviously interesting issue of how the derived information can be used to optimize CLP programs, which we consider to be outside the scope of the paper. However, and as mentioned previously, this subject has recently been addressed by several authors and their results show that, if information from global analysis such as that obtained by our analyses is available, it can in fact be used to perform optimizations which result in significant speedups [Jørgensen et al. 1991; Jaffar et al. 1992; Marriott and Stuckey 1993; Jaffar and Maher 1994; Marriott et al. 1994; Dumortier 1994; García de la Banda 1994].

9.1 Implementation issues

The abstract domains described in Sections 6, 7, and 8 have been implemented within the PLAI abstract interpretation system [Muthukumar and Hermenegildo 1990; 1992] which is an incarnation of the framework presented. The resulting analyses can deal with CLP(H,N) programs and also with some of the PrologIII specific features, namely tuples and size relations.

A few details of PLAI are worth mentioning since they are instrumental in understanding the results obtained during our evaluation. PLAI in principle assumes finite abstract domains and analyzes each predicate for each distinct key (the pair $\langle \mathbf{a}, \mathbf{AC_{entry}} \rangle$). This implies that PLAI performs a quite detailed analysis and can obtain several annotations for the same predicate (versions). The current implementation allows the user to choose between obtaining a transformed program in which the different versions of the predicates appear explicitly and are each annotated with their corresponding inferred information or, alternatively, obtaining essentially the original program where predicates are annotated with the upper bound of the annotations of the different versions of that predicate. In our experiments the former approach was selected (exceptions are indicated).

It is important to note that the only modification that was needed for extending PLAI to CLP languages was the elimination of a "unifiability" test performed before executing the abstract entry function. This test is performed in the analysis of traditional LP languages in order to avoid analyzing rules whose head does not (syntactically) unify with the current subgoal. Naturally, the domain-dependent abstract functions had to be implemented and incorporated into the system, but almost all the existing implementation was reused. We argue that this strongly supports our claim regarding the practical usefulness of the approach that we propose, specially considering that, as we believe our measurements show, the resulting system can analyze reasonably sized programs in quite reasonable times.

Finally, the integration of the $\mathcal{DF}^{\mathbf{m}}$ analyzer has been performed as follows. Since the $\mathcal{DF}^{\mathbf{m}}$ analysis uses definiteness information provided by the $\mathcal{D}$ analysis, the $\mathcal{D}$ and $\mathcal{DF}^{\mathbf{m}}$ analysis are executed in a coroutining fashion. At each point of the analysis (i.e. at the application of one of the higher-level abstract operations), the definiteness operation is called first. Afterwards, the set of definite variables is extracted from the result of that operation and passed as an extra parameter to the freeness operation. If the definiteness operation results in the abstract constraint $\perp$, the freeness operation proceeds with $\perp$. Thus, information is always passed from the definiteness to the freeness analysis; information passing in the other direction is restricted to the passing of $\perp$ information: if a freeness operation yields $\perp$ where the preceding definiteness operation did not give $\perp$, the subsequent definiteness analysis continues with $\perp$ (thus computation of useless information is avoided).

The effect of such combination can be quite subtle. On the one hand, the efficiency (both in terms of memory and time) of the $\mathcal{DF}^{\mathbf{m}}$ analyzer can be better than that of simply running both the $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyzers. This can be due to several factors. First, the potential reduction in the size of the $\mathcal{DF}^{\mathbf{m}}$ abstractions can reduce the memory consumption, which in turn affects the analysis times. Second, reductions in the size of the abstract constraints can also reduce the cost of the abstract operations. Finally, the combination has a "loop merging" effect — a single pass over the program is sufficient for $\mathcal{DF}^{\mathbf{m}}$ instead of the two passes needed otherwise.

On the other hand, if one of the analyses requires more fixpoint iterations than the other, this may have a negative effect on the efficiency of the combined execution. If, for example, the definiteness analysis reaches the fixpoint first, the extra iterations will imply some unnecessary *table look-ups*, projections, and extensions for this analyzer. If the freeness analyzer is the one who first reaches the fixpoint, the overhead may be more substantial. This is because part of the definiteness abstraction is included in the freeness abstraction, and therefore all abstract operations may be redone. Such extra iterations could be avoided by first performing the definiteness analysis by itself, and then using the programs annotated by the definiteness analysis as input for the freeness analysis. The detailed evaluation for a subset of CLP programs discusses the interaction in depth.

9.2 Benchmarks

Table I. Properties of the CLP benchmarks

Program	Pr	Rl	R	TR	NR	MaxV	AvgV
dnf	3	32	2	1	0	7	2.3
vecmat1	8	15	0	7	1	8	2.9
laplace1	2	4	0	2	0	12	6.0
fib	1	3	1	0	0	3	1.0
meal	6	11	0	0	6	6	0.9
listlength	1	2	0	1	0	4	2.0
sumlist	1	2	0	1	0	4	2.0
mining	25	50	4	10	11	18	2.5
power	18	42	0	9	9	19	3.3
rectangle	5	10	2	2	1	9	3.3
vecmat2	8	15	0	7	1	8	3.1
num	17	97	0	0	17	10	2.4
laplace2	2	4	0	2	0	16	10.8
sendmm	4	7	0	3	1	11	2.7
trap	4	5	0	1	3	9	6.4
runkut	4	5	0	1	3	9	6.2
mortgage1	1	2	0	1	0	5	4.0
mortgage3	1	2	0	1	0	5	4.0
mortgage2	1	2	0	1	0	5	4.0
bridge	29	90	0	14	15	13	1.6
color4	8	21	0	3	5	9	2.2
color4F	8	110	0	3	5	9	0.4
cutstock	50	77	3	19	28	21	3.9
magic	7	14	0	6	1	5	2.4
magicC	5	9	1	3	1	8	2.4
periodic	3	5	0	1	2	11	4.0
perm	11	20	0	7	4	6	2.6
triangle	34	47	0	5	29	24	6.3
warehouse	12	38	1	4	7	21	2.

The global set of benchmarks used contains 29 CLP programs and 25 LP programs. The CLP programs solve typical CLP problems and include from small to relatively large programs (i.e., programs with 1 to 50 predicates and with 2 to 110 rules). Part of them are taken from the CLP($\mathcal{R}$) distribution, the PrologIII distribution, and from the CLP literature [Van Hentenryck 1989; Van Hentenryck and Ramachandran 1994; Colmerauer 1990]. Others have been obtained from the

Table II. Properties of the LP benchmarks

Program	Pr	Rl	R	TR	NR	MaxV	AvgV
akl	10	18	2	4	4	10	3.6
akl_old	7	12	0	4	3	10	3.6
ann	53	187	19	13	21	17	2.5
append	1	2	0	1	0	4	2.5
bid	22	53	0	7	15	7	2.2
boyer	28	138	3	1	24	8	2.3
browse	16	32	1	11	4	12	3.7
deriv	15	62	4	3	8	6	3.0
grammar	7	15	0	0	7	6	1.9
icomp	71	170	19	18	36	20	5.0
kalah	41	78	9	10	22	12	3.8
mapcolor	8	12	0	4	4	6	3.1
peephole	16	134	10	3	3	8	2.8
pg	10	18	0	6	4	10	3.6
plan	16	29	0	4	12	6	2.7
qplan	44	148	16	11	17	16	3.1
qsort	3	6	1	1	1	7	3.5
queens	5	9	0	4	1	5	2.4
rdtok	18	55	9	6	3	7	3.3
read	25	91	7	4	14	13	3.9
serialize	6	12	2	2	2	7	3.8
tarjan	37	90	12	14	11	20	4.9
vlok	46	137	0	17	29	12	2.6
vlokgr	46	137	0	17	29	12	2.6
witt	79	163	22	23	34	18	4.5

partners in the PRINCE ESPRIT project, from P. Van Hentenryck, and from the vendor of Prolog III and Prolog IV, PrologIA. We have also included a large collection of LP benchmarks, ranging from relatively simple to quite complex programs, which has been used previously in the literature to evaluate analyzers for LP programs [Mulkers et al. 1995; Codish et al. 1995]. The number of predicates in these benchmarks ranges from 1 to 79 and the number of rules from 2 to 187. Since all LP programs are also CLP programs, the latter set of benchmarks adds another dimension to the benchmark suite which allows us to expand our study of the scalability issue. A brief description of all the benchmarks is given in the Appendix. Here we include Table I and Table II which list properties of the benchmark programs to which the complexity of the analysis is related. The size of the programs is indicated by means of the number of user-defined predicates (Pr), and the number of rules (Rl). The recursiveness of the programs is indicated by means of the number of recursive predicates that are not tail-recursive (R), the number of tail-recursive predicates (TR), and the number of non-recursive predicates (NR). Programs containing recursive predicates lead to a more complex analysis than non-recursive programs, especially if they are not tail-recursive. The tables also list the maximum and average number of variables in the program rules (MaxV and AvgV). The number of variables in a rule typically affects the size of the abstract constraints for the different program points in the rule, which in turn influences the cost of the abstract operations.

9.3 Efficiency Results

In order to get an idea of the feasibility of the analyses proposed in this paper Table III and Table IV list the total analysis times for the CLP and the LP programs respectively. The figures include the time for garbage collection and stack shifts, and are averaged over 10 runs. All measurements have been done on a SUN Sparc 2 using SICStus 2.1 with the "fastcode" option. "–" indicates that the analyzer did not produce a result (because it ran out of memory). The last column in Tables III and IV gives the ratio of time taken by the combined analysis $\mathcal{DF}^{\mathbf{m}}$ to the sum of $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$. "Inf" indicates that the combined analysis is definitely better, since in these cases $\mathcal{F}^{\mathbf{m}}$ does not produce a result. The average of Table III does not take into account **laplace**1.

Table III. Timings of the analyses for the CLP programs

	Analysis times (seconds)			Comparison
Program	$\mathcal{D}$	$\mathcal{F}^{\mathbf{m}}$	$\mathcal{DF}^{\mathbf{m}}$	$\frac{\mathcal{DF}^{\mathbf{m}}}{\mathcal{D}+\mathcal{F}^{\mathbf{m}}}$
dnf	0.960	5.158	3.492	0.57
vecmat1	0.116	0.438	0.298	0.54
laplace1	0.060	–	0.134	Inf
fib	0.044	0.082	0.092	0.73
meal	0.032	0.074	0.078	0.74
listlength	0.010	0.022	0.020	0.63
sumlist	0.014	0.028	0.020	0.48
mining	1.852	7.990	9.700	0.99
power	2.192	10.558	4.688	0.37
rectangle	6.960	3.172	11.184	1.10
vecmat2	0.324	0.544	0.928	1.07
num	0.976	1.680	2.000	0.75
laplace2	0.776	0.262	1.222	1.18
sendmm	1.378	3.198	4.154	0.91
trap	2.472	0.494	3.718	1.25
runkut	0.052	0.394	0.142	0.32
mortgage1	0.124	0.148	0.358	1.32
mortgage3	0.030	0.148	0.122	0.69
mortgage2	0.125	0.086	0.206	0.98
bridge	3.316	9.470	33.508	2.62
color4	0.226	1.346	0.922	0.59
color4F	0.416	1.798	1.684	0.76
cutstock	1.250	5.552	2.940	0.43
magic	0.172	0.406	0.346	0.60
magicC	0.522	0.290	0.906	1.12
periodic	0.748	0.320	1.430	1.34
perm	0.250	0.760	0.958	0.95
triangle	52.030	216.662	256.734	0.96
warehouse	1.044	0.742	1.834	1.03
Average	2.706	9.708	11.856	0.89

Table IV. Timings of the analyses for the LP programs

Program	Analysis times (seconds)			Comparison
	$\mathcal{D}$	$\mathcal{F}^{\mathbf{m}}$	$\mathcal{DF}^{\mathbf{m}}$	$\frac{\mathcal{DF}^{\mathbf{m}}}{\mathcal{D}+\mathcal{F}^{\mathbf{m}}}$
akl	0.264	2.850	0.654	0.21
akl_old	0.182	2.626	0.454	0.16
ann	6.944	6.936	14.754	1.06
append	0.018	0.108	0.030	0.24
bid	0.348	1.304	0.708	0.43
boyer	3.850	9.822	14.756	1.08
browse	1.004	1.508	1.912	0.76
deriv	1.174	2.938	1.774	0.43
grammar	0.044	0.172	0.122	0.56
icomp	6.914	30.480	33.512	0.90
kalah	1.002	4.374	1.700	0.32
mapcolor	0.554	33.490	31.444	0.92
peephole	2.658	14.382	10.510	0.62
pg	0.192	0.868	0.386	0.36
plan	0.222	1.544	0.560	0.32
qplan	1.262	17.060	3.376	0.18
qsort	0.096	0.312	0.274	0.67
queens	0.066	0.284	0.140	0.40
rdtok	1.740	4.434	4.616	0.75
read	2.162	9.818	4.222	0.35
serialize	0.784	0.966	1.742	1.00
tarjan	1.900	124.340	6.126	0.05
vlok	1.236	34.452	3.886	0.11
vlokgr	0.964	33.826	2.688	0.08
witt	2.354	17.376	4.552	0.23
Average	1.517	14.251	5.796	0.37

For most benchmarks the analysis times are acceptable. The average $\mathcal{D}$ and $\mathcal{DF}^{\mathbf{m}}$ analysis times of the LP programs are better than for the CLP benchmarks. This is to be expected since constraints in LP programs (unification constraints) are in general less complex than typical CLP constraints, leading to smaller constrained sets and smaller abstractions. Also, there is usually more definiteness information to be exploited: LP programs are frequently "generate and test," whereas CLP programs are often of the "constrain and generate" type, which implies that definiteness information is only derived towards the end of the program. For most programs (46 out of 54) the $\mathcal{F}^{\mathbf{m}}$ analysis takes longer than the $\mathcal{D}$ analysis. This can be partly explained by the different natures of the abstractions. $\mathcal{D}$ propagates definiteness information and collects definite dependencies between non-definite variables. $\mathcal{F}^{\mathbf{m}}$ propagates possible non-freeness and collects possible dependencies between all variables. The freeness analysis inherently has a larger time and space complexity than the definiteness analysis. But also the abstract query[18] which is analyzed for the benchmark programs plays an important role (e.g laplace with **laplace**1(**d**) and **laplace**2(**M**) where **M** is a matrix of free

[18] More details are in the Appendix.

variables, and mortgage with **mortgage**1(**a**, **a**, **a**, **a**,**f**), **mortgage**2(**a**, **a**, **a**, **f**, **a**) and **mortgage**3(**f**, **d**, **d**, **d**, **d**)).

For some benchmarks the execution time of $\mathcal{F}^{\mathbf{m}}$ becomes quite large (**triangle** and **tarjan**) or even infinite (**laplace**1). The combined $\mathcal{DF}^{\mathbf{m}}$ analysis seems to offer a partial solution. The size of the $\mathcal{F}^{\mathbf{m}}$ abstractions can be reduced (sometimes substantially) when definiteness information is available. For **laplace**1 and **tarjan**, $\mathcal{DF}^{\mathbf{m}}$ has very good performance, but **triangle** does not really seem to benefit from the combination. The ratio $\frac{\mathcal{DF}^{\mathbf{m}}}{(\mathcal{D}+\mathcal{F}^{\mathbf{m}})}$ shows that also for the other programs the $\mathcal{DF}^{\mathbf{m}}$ analysis performs quite well. Due to the previously discussed interactions with the fixpoint computations $\mathcal{DF}^{\mathbf{m}}$ sometimes introduces overhead, but this remains acceptable (see the programs with a ratio > 1). The average for the ratio $\frac{\mathcal{DF}^{\mathbf{m}}}{(\mathcal{D}+\mathcal{F}^{\mathbf{m}})}$ is 0.89 for the CLP programs (not taking into account **laplace**1) and 0.36 for the LP programs. The execution times for $\mathcal{DF}^{\mathbf{m}}$ vary between 0.020 and 256.7 seconds, with an average of 5.796.

9.4 Effects of Widening

As mentioned before, we have also studied the effect of the application of widening operations in order to trade precision for efficiency, which is an important technique related to the scalability issue. For the current set of benchmarks, scalability seems to be a problem of the freeness abstraction but not of the definiteness abstraction. Therefore, we decided to apply widening in the freeness analysis and in the freeness part of the $\mathcal{DF}^{\mathbf{m}}$ analysis. The idea is to avoid the **close** operation (used for example during abstract conjunction) for large freeness abstractions as in those cases this operation is too expensive. An abstraction is considered to be too large if it contains a number of non-singleton sets above some bound. We experimented with two different bounds **B**: 10 (referred to as wid10 in Table V) and 8 (wid8 in Table V). If an abstraction contains **B** or more non-singletons, widening is applied. A strong form of widening was used: that all variables involved in the abstraction are given mode "any."

Table V indicates the influence of widening on the timings and precision of the $\mathcal{F}^{\mathbf{m}}$ analysis and freeness part of the $\mathcal{DF}^{\mathbf{m}}$ analysis. As mentioned before, the column "Wid" indicates whether the analyzer includes widening or not and, if so, what the bound is on the number of non-singleton sets in the abstraction ("wid10" or "wid8"). The column "Free" indicates the number of free variable annotations derived by the analysis.[19] It shows when precision is lost due to widening. The following two columns contain the timings (in seconds) for the $\mathcal{F}^{\mathbf{m}}$ and $\mathcal{DF}^{\mathbf{m}}$ analyses respectively. The last two columns indicate the memory consumption (in Mbyte, giving maximum amount of memory allocated by the UNIX system to the PLAI process). The table contains only those benchmarks where widening is actually applied: for wid8, 10 out of the 54 benchmarks effectively apply widening in the case of $\mathcal{F}^{\mathbf{m}}$, and 5 in the case of $\mathcal{DF}^{\mathbf{m}}$. For wid10, widening happens only in, respectively, 6 and 2 benchmarks. "≡" indicates that widening is not triggered for

[19] These figures are the same for the $\mathcal{F}^{\mathbf{m}}$ and $\mathcal{DF}^{\mathbf{m}}$ analysis and they are computed selecting the analysis output option that returns an annotated version of the original program with each predicate annotated with the upper bound of the analysis results for all the different entry-exit patterns (the different version) inferred during the analysis.

Table V. Widening information

Program	Wid	Free	Time $\mathcal{F}^{\mathbf{m}}$	Time $\mathcal{DF}^{\mathbf{m}}$	Mem $\mathcal{F}^{\mathbf{m}}$	Mem $\mathcal{DF}^{\mathbf{m}}$
laplace1	nowid	0	–	0.134	–	4.298
	wid10	0	0.890	≡	4.298	≡
	wid8	0	0.610	≡	4.298	≡
mining	nowid	143	7.990	9.700	4.888	4.950
	wid8	143	2.500	4.660	4.829	4.892
power	nowid	191	10.558	4.688	4.888	4.856
	wid8	191	3.610	≡	4.829	≡
rectangle	nowid	18	3.172	11.184	4.794	4.825
	wid8	18	2.360	10.250	4.798	4.798
sendmm	nowid	66	3.198	4.154	4.888	4.888
	wid10	57	0.310	1.670	4.298	4.798
	wid8	57	0.290	1.710	4.298	4.798
bridge	nowid	204	9.470	33.508	4.888	6.075
	wid8	204	5.490	9.380	4.829	4.892
triangle	nowid	1664	216.662	256.734	13.263	13.950
	wid10	1508	44.250	106.370	6.392	6.329
	wid8	1508	6.450	98.710	5.204	6.329
tarjan	nowid	439	124.340	6.126	6.138	5.075
	wid10	439	89.280	≡	6.142	≡
	wid8	439	32.020	≡	6.142	≡
vlok	nowid	216	34.452	3.886	8.513	5.138
	wid10	216	13.370	≡	5.142	≡
	wid8	216	8.320	≡	5.142	≡
vlokgr	nowid	216	33.826	2.688	8.513	5.138
	wid10	216	13.400	≡	5.142	≡
	wid8	216	8.440	≡	5.142	≡

the particular program and analysis (time and memory figures then correspond to the nowid figures). Notice that widening is not triggered in the case of the $\mathcal{DF}^{\mathbf{m}}$ analysis of the LP benchmarks. Applying the widening operation (wid8) allows analyzing the **laplace**1 benchmark in 0.610 seconds using 4.298 Mbyte, where the original $\mathcal{F}^{\mathbf{m}}$ analysis (without widening) did not produce a result within reasonable time and memory bounds (indicated by "–" in the table).

In general, if widening is applied it improves considerably the execution times and memory consumption. For some programs (**sendmm** and **triangle**) the difference is an order of magnitude. The maximum analysis time for $\mathcal{F}^{\mathbf{m}}$ (nowid) is 216.7 s for **triangle** and with wid8 it goes down to 6.5 s, while for $\mathcal{DF}^{\mathbf{m}}$ it goes down from 256.7 s to 98.7 s. The impact of widening is not the same for $\mathcal{F}^{\mathbf{m}}$ and $\mathcal{DF}^{\mathbf{m}}$. In the case of **triangle** this is due to the $\mathcal{D}$ part of the analysis. For the $\mathcal{D}$ analysis, the execution time of **triangle** (52 s) differs one order of magnitude with respect to the other execution times. The $\mathcal{D}$ analysis infers large definite dependency sets since **triangle** uses the CLP "constrain and generate" programming technique. A widening for the $\mathcal{D}$ analysis could lessen this kind of inefficiencies.

The effect on the precision is acceptable as only in two cases (**sendmm** and **triangle**) precision is lost when considering for each predicate the single upper bound which is computed from the analysis results for the different entry-exit patterns. These experiments suggest that widening allows to avoid exponentional time and memory consumption and to keep the loss of precision very small.

9.5 A More Detailed Evaluation

In addition to the more general study reported above, in order to gain more insight into the behaviour of the analyses in typical CLP programs we performed a more detailed evaluation on a subset of the benchmarks (namely, the first 19 CLP programs of Table I), which we consider a representative selection of typical CLP programs.

9.6 A Closer look at the efficiency results. **Measurements.**

Our aim is (1) to study the time and memory consumption of each of the analyzers (only taking into account the kernel of the analysis, i.e. the execution of the fixpoint algorithm, and not the pre- and post-processing phase) and (2) to evaluate the effectiveness of the combined analysis with respect to the individual $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyses. For this purpose the following figures have been collected:

- Regarding the fixpoint computation: the number of entry-exit patterns for all predicates (EE) and for the recursive predicates only (EEr), and the number of fixpoint iterations (FIX). They are presented in Table VI. These numbers largely determine the complexity of the analyses and will be used to explain the time and memory figures of the combined analysis, when compared to those of the individual $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyses. Table VI also lists the number of syntactically different calls modulo renaming (DCls), which is a lower bound on the number of entry-exit patterns that will be computed by the analyzer under the assumption that the program does not have dead code.
- Regarding time consumption: the total analysis times (including the time for garbage collection and stack shifts) averaged over 10 runs. They are given in Table III, the last column provides the execution time comparison.
- Regarding memory consumption:
  (1) the maximal amount of memory allocated by the UNIX system to the PLAI process. It falls between 4.293808 Mbytes and 4.918808 Mbytes (except for the $\mathcal{F}^{\mathbf{m}}$ analysis of **laplace**1 which runs out of memory). Note that the sum of the memory allocated by the $\mathcal{D}$ and the $\mathcal{F}^{\mathbf{m}}$ analyzers is always larger than the memory allocated by $\mathcal{DF}^{\mathbf{m}}$;
  (2) the total amount of global stack space (Glob) and program space (Prog)[20] used during the actual analysis. This is given in Table VII. The last column in Table VII compares the sum of the global stack and program space used by the $\mathcal{DF}^{\mathbf{m}}$ analysis with the maximum of the sum of global stack and program space used in the $\mathcal{D}$ and the $\mathcal{F}^{\mathbf{m}}$ analyses, i.e. the memory consumption[21] comparison figure.

In order to aid in the interpretation of the results we divide the programs into two classes:

[20] The global stack stores the compound terms. Program space refers to the amount of memory allocated for compiled and interpreted rules, symbol tables, the record database, and the like.

[21] In what follows, we refer to the global stack and program space consumption simply as memory consumption.

Table VI. Number of fixpoint iterations and entry-exit patterns

		$\mathcal{D}$			$\mathcal{F}^{\mathbf{m}}$			$\mathcal{DF}^{\mathbf{m}}$		
Program	DCls	EE	EEr	FIX	EE	EEr	FIX	EE	EEr	FIX
dnf	14	14	14	14	26	26	26	26	26	26
vecmat1	9	9	8	8	11	10	13	11	10	10
laplace1	4	4	4	4	-	-	-	4	4	4
fib	3	3	3	3	3	3	3	3	3	3
meal	6	6	0	0	6	0	0	6	0	0
listlength	1	1	1	1	1	1	1	1	1	1
sumlist	1	1	1	1	1	1	1	1	1	1
mining	32	36	25	39	32	20	28	36	25	39
power	24	28	19	27	24	15	18	28	19	27
rectangle	8	11	10	21	9	8	14	11	10	22
vecmat2	10	12	11	16	13	12	15	16	15	22
num	17	20	0	0	17	0	0	20	0	0
laplace2	3	3	3	4	3	3	4	3	3	4
sendmm	6	6	5	7	6	5	6	6	5	7
trap	5	11	4	5	5	2	3	11	4	5
runkut	6	6	1	1	6	1	2	6	1	2
mortgage1	2	2	2	2	3	3	3	3	3	3
mortgage3	2	2	2	2	3	3	3	3	3	3
mortgage2	2	2	2	2	2	2	2	2	2	2

Table VII. Global stack space and program space used during analysis (in Mbyte)

	$\mathcal{D}$		$\mathcal{F}^{\mathbf{m}}$		$\mathcal{DF}^{\mathbf{m}}$		Comparison
Program	Glob	Prog	Glob	Prog	Glob	Prog	Gl+Pr $\frac{\mathcal{DF}^{\mathbf{m}}}{\mathbf{max}(\mathcal{D},\mathcal{F}^{\mathbf{m}})}$
dnf	0.303	0.065	1.737	0.133	1.081	0.149	0.66
vecmat1	0.049	0.008	0.194	0.011	0.110	0.013	0.60
laplace1	0.033	0.004	–	–	0.060	0.006	Inf
fib	0.022	0.003	0.045	0.003	0.040	0.004	0.92
meal	0.018	0.004	0.041	0.004	0.038	0.005	0.96
listlength	0.014	0.001	0.019	0.001	0.017	0.001	0.90
sumlist	0.014	0.001	0.021	0.001	0.017	0.001	0.82
mining	0.805	0.054	1.312	0.055	2.121	0.083	1.61
power	1.062	0.063	2.151	0.052	1.825	0.090	0.87
rectangle	3.054	0.019	0.856	0.016	4.127	0.028	1.35
vecmat2	0.147	0.012	0.216	0.014	0.401	0.022	1.84
num	0.216	0.064	0.532	0.061	0.541	0.084	1.05
laplace2	0.444	0.008	0.105	0.006	0.535	0.011	1.21
sendmm	0.484	0.012	0.792	0.010	1.250	0.017	1.58
trap	1.280	0.017	0.198	0.007	1.629	0.024	1.27
runkut	0.026	0.005	0.145	0.007	0.061	0.008	0.45
mortgage1	0.076	0.002	0.092	0.004	0.217	0.004	2.30
mortgage3	0.020	0.002	0.092	0.004	0.064	0.004	0.71
mortgage2	0.076	0.002	0.063	0.003	0.135	0.003	1.77

(1) Programs that, for the given entry patterns, constrain many variables from the start to definite values and related dependencies (dnf, vecmat1, laplace1, fib, meal, listlength and sumlist).

(2) Programs that do not allow inferring much definiteness information or where it is inferred only towards the end of the program (mining, power, rectangle, vecmat2, num, laplace2, sendmm, trap, runkut, mortgage1, mortgage3, and mortgage2). They create and handle large sets of possible dependencies.

In each class the benchmarks are ordered starting with the highest estimate for the size of the AND-OR graph.[22]

Evaluation. For the first class of programs, the time and memory figures correspond quite well with the complexity estimate used for ordering them. The number and the size of the dependencies is small and hence has not much influence on the figures. For the second class of programs, however, the complexity estimate is no longer adequate to predict the time and memory consumption. In this case, the number and the size of the dependencies can have a more important impact. Note that for the $\mathcal{F}^{\mathbf{m}}$ analysis, the programs with large time and memory figures (**power**, **mining**, **sendmm** and **rectangle**) have relatively many variables in their rules (i.e. large MaxV and AvgV numbers in Table I and II). This trend can also be observed for the $\mathcal{D}$ analysis (**rectangle**, **trap**, **power**, **mining** and **sendmm**). This trend is also confirmed by the actual output of the analysis and by the correlation between analysis time and global stack consumption, since the latter is dominated by the size of the abstract constraints built during the analysis.

The $\mathcal{DF}^{\mathbf{m}}$ analysis yields quite satisfactory results for the considered benchmarks, both concerning time and memory consumption. The execution times vary between 0.020 and 11.184 seconds. For most benchmarks (14 out of 19), the execution time comparison figure of Table III is smaller than 1. Also, for 10 of the 19 benchmarks the memory consumption comparison figure of Table VII is smaller than 1, and only for 1 benchmark it is larger than 2. This provides evidence that combining the $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyzers indeed results in a practical full mode analysis system. We now perform a more detailed evaluation of these figures, based upon the classes of programs and their complexity in terms of entry-exit patterns and fixpoint iterations (Table VI).

As mentioned before, the first class of programs yields many definite variables right at the beginning of the execution. For these programs, the definiteness information is effectively used in the freeness part. This is reflected both in the timings (upper part of Table III) and the memory consumption (upper part of Table VII). In some cases (**dnf**, **vecmat**1), the $\mathcal{D}$ analyzer has to iterate along with the $\mathcal{F}^{\mathbf{m}}$ analyzer when combining the two, i.e. the FIX and EE numbers of the $\mathcal{DF}^{\mathbf{m}}$ analyzer correspond to the ones of the $\mathcal{F}^{\mathbf{m}}$ analyzer and are larger than those of the $\mathcal{D}$ analyzer (Table VI). But even then, this overhead is more than compensated by the benefit of exploiting definiteness information, so there is still a considerable improvement of $\mathcal{DF}^{\mathbf{m}}$ with respect to $\mathcal{D} + \mathcal{F}^{\mathbf{m}}$.

[22] We use the formula $(\mathbf{Rl}/\mathbf{Pr}) * \mathbf{AvgV} * (\mathbf{NR} + \mathbf{FIX} * (\mathbf{TR} + 3 * \mathbf{R}))$ with **FIX** of $\mathcal{DF}^{\mathbf{m}}$ given in Table VI and the rest in Table I and II.

For the second class of programs, the combination does not always pay off. Clearly, it depends on whether or not the gain obtained by exploiting definiteness information in the freeness part outweighs the overhead caused by extra fixpoint iterations (FIX) and entry-exit patterns (EE) in $\mathcal{DF}^{\mathbf{m}}$ compared to $\mathcal{D}$ and/or $\mathcal{F}^{\mathbf{m}}$. Four situations can be distinguished concerning the FIX and EE numbers of the $\mathcal{DF}^{\mathbf{m}}$ analysis with respect to those of $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$. First of all, for the **mining**, **power**, **num**, **rectangle**, **trap** and **sendmm** programs, the EE and FIX figures for the $\mathcal{F}^{\mathbf{m}}$ analyzer are smaller than the ones for the $\mathcal{D}$ analyzer. Thus, when combining the two analyses, the $\mathcal{DF}^{\mathbf{m}}$ analysis has to perform at least as many iterations as the $\mathcal{D}$ analysis. However, it now not only computes the definite part but it also takes into account the (reduced) freeness part. In the case of **mining**, **power**, **sendmm** and **num**, this overhead is outweighed by the gain obtained from exploiting definiteness information (the execution time comparison figures are smaller than 1 and the memory consumption comparison figures fall between 0.869 and 1.612), whereas for **rectangle** and **trap**, the freeness part cannot benefit much from the definiteness information (note that in those cases the time and memory consumption for the $\mathcal{D}$ analysis is large –both by itself and compared with the $\mathcal{F}^{\mathbf{m}}$ analysis–, which indicates that mostly definite dependencies are derived rather than definite variables). Secondly, for the **runkut**, **mortgage**1 and **mortgage**3 benchmarks, the EE and FIX numbers of the $\mathcal{DF}^{\mathbf{m}}$ analysis correspond to those of the $\mathcal{F}^{\mathbf{m}}$ analysis and are larger than the $\mathcal{D}$ ones. In the case of **runkut** and **mortgage**3, the time and memory figures show that the freeness part can benefit quite well from the definite information. Also, the $\mathcal{D}$ time and memory consumption is small compared to that of $\mathcal{F}^{\mathbf{m}}$, so the extra iterations of $\mathcal{D}$ (when forced to execute along with $\mathcal{F}^{\mathbf{m}}$) are not outweighing the gain. For **mortgage**1 however, the situation is just the opposite: the execution time comparison figure is larger than 1 and the memory consumption comparison figure is larger than 2. Thirdly, the **laplace**2 and **mortgage**2 benchmarks have the same EE and FIX numbers for all analyses. Although there are no extra iterations, there is almost no definiteness information to be exploited. The combination may cause a slight overhead due to the extra operations dealing with the (in this case useless) communication between the two analyses (cf. time for **laplace**2). Finally, for the **vecmat**2 benchmark, the $\mathcal{DF}^{\mathbf{m}}$ analysis performs more iterations and has more entry-exit patterns than either one of the $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyses. This results in a slight overhead in memory consumption and analysis time.

9.7 Accuracy results. **Measurements.** The accuracy of the analyzers is determined by comparing the outcome of concrete executions of the benchmarks with the results obtained by the analyses. More precisely, the correct (concrete) modes of the variables at each program point are compared with the modes derived by the analyzers. If specialized versions of a predicate arise during concrete execution, these are considered separately. The predicate versions produced by the analyzers are mapped onto the concrete versions (usually, there is a one-to-one correspondence between the concrete and abstract predicate versions; in some cases however, several abstract versions map onto one concrete version or vice-versa). The figures for the $\mathcal{DF}^{\mathbf{m}}$ analysis are presented in Table VIII (a similar study could be made for the individual $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyses which may infer a different number of predicate versions; herein we approximate these figures by considering the $\mathcal{D}$ part and the $\mathcal{F}^{\mathbf{m}}$ part of the combined $\mathcal{DF}^{\mathbf{m}}$ analysis separately). Column "Annot" gives the total number of variable annotations (summed up over the predicate versions and the program points). "ImpD" and "ImpF" give the number of imprecise variable annotations (derivation of mode **a** instead of **d**, respectively mode **a** instead of **f**). The columns "PrecD" and "PrecF" give the percentages of variable modes that are correctly inferred by the $\mathcal{D}$ part of the analysis and the $\mathcal{F}^{\mathbf{m}}$ part respectively. "PrecD+F" is the percentage of correct variable modes derived by the combined $\mathcal{DF}^{\mathbf{m}}$ analysis.[23] The average precision is shown at the bottom of the table. The last column indicates the cause of imprecision.

Table VIII. Accuracy of the analyzers (only w.r.t. variable modes)

Program	Annot	ImpD	ImpF	PrecD	PrecF	PrecD+F	Imp. cause
dnf	3772	0	33	100.0	99.1	99.1	(1)
vecmat1	125	0	6	100.0	95.2	95.2	(2)
laplace1	112	0	0	100.0	100.0	100.0	
fib	12	0	0	100.0	100.0	100.0	
meal	56	0	0	100.0	100.0	100.0	
listlength	12	0	0	100.0	100.0	100.0	
sumlist	12	0	0	100.0	100.0	100.0	
mining	1064	105	80	90.1	92.5	82.6	(1)
power	1256	126	73	89.9	94.2	84.1	(1,2)
rectangle	343	0	4	100.0	98.8	98.8	(1)
vecmat2	208	0	13	100.0	93.7	93.7	(1,2)
num	1402	13	0	99.1	100.0	99.1	
laplace2	124	0	60	100.0	51.6	51.6	(1)
sendmm	141	0	0	100.0	100.0	100.0	
trap	135	73	8	46.0	94.0	40.0	(2)
runkut	119	0	14	100.0	88.2	88.2	(2)
mortgage1	54	0	8	100.0	85.0	85.0	(2)
mortgage3	36	3	2	91.6	94.4	86.0	(2)
mortgage2	36	0	0	100.0	100.0	100.0	
Average				95.6	94.0	89.6	

Besides the accuracy of mode annotations, one can additionally consider the accuracy of the dependency information.[24] In the case of possible dependencies, the same precision is obtained with the $\mathcal{F}^{\mathbf{m}}$ and $\mathcal{DF}^{\mathbf{m}}$ analyzers. Even if correct modes are inferred at a particular program point, the inferred possible dependencies may not occur in the concrete case or may be too strong compared to the concrete

[23] This number is lower than or equal to the corresponding PrecD and PrecF number, as it takes into account *both* imprecision due to deriving mode **a** instead of **d** and that due to deriving mode **a** instead of **f**, whereas in the $\mathcal{D}$ part only imprecision of the type "mode **a** instead of **d**" is taken into account (as mode **a** is the most precise abstraction for free variables in the $\mathcal{D}$ analysis), and in the $\mathcal{F}^{\mathbf{m}}$ part only imprecision of the type "mode **a** instead of **f**" is taken into account (as mode **a** is the most precise abstraction for definite variables in the $\mathcal{F}^{\mathbf{m}}$ analysis).

[24] We only consider the *possible* dependency information. A similar study could be made for the definite dependencies.

dependencies, thus possibly leading to imprecise mode annotations at subsequent program points. Imprecise dependency information not affecting the precision of the mode information (not visible in Table VIII) is derived when analyzing the **sendmm** program (about 40% of the dependencies are too strong), and, to a lesser extent, also in the **power** and **runkut** benchmarks.

Evaluation. The average precision for the $\mathcal{D}$ part is 95.6%, 94% for the $\mathcal{F}^{\mathbf{m}}$ part, and 89.6% for the combined $\mathcal{DF}^{\mathbf{m}}$ analyzer. For the $\mathcal{D}$ part and the combined $\mathcal{DF}^{\mathbf{m}}$ analysis, the worst case occurs for the **trap** benchmark (resp. 46% and 40%). For the $\mathcal{F}^{\mathbf{m}}$ part, the worst results are for the **laplace**2 benchmark (51.6%).

There are three sources of inaccuracy: (1) the lack of information about term structures, (2) the treatment of non-linear constraints, and (3) the abstraction of primitive constraints instead of the abstraction of conjunctions of primitive constraints. The first is mainly related to inaccuracy of modes. The third mainly affects the accuracy of the dependency information. The second influences both.

Regarding the lack of information about term structure, when selecting a component of a partially instantiated term having mode **a**, the definiteness analysis cannot discover the definiteness of a definite subterm. Similarly, the freeness analysis cannot recognize free variables within the term. Consider a program scheme of the form

build_structure(**Data**), **constrain**(**Data**), **instantiate**(**Data**)

where one first builds a data-structure, then imposes constraints on that structure and finally instantiates it. Such a scheme is used quite frequently within CLP (e.g. **mining**, **power**, **rectangle**, **sendmm**,...). It gives rise to a loss of precision when selecting components of the structure within **constrain**/1 and **instantiate**/1. Imprecision due to the absence of structure information occurs in case of the **dnf**, **rectangle**, **laplace**2 and **mining** benchmarks, and is also causing part of the imprecision in **power** and **vecmat**2. Although adding structure information [Janssens and Bruynooghe 1992; Le Charlier and Van Hentenryck 1994; Mulkers et al. 1995] could clearly improve precision, it also complicates the analysis in the sense that, when changing the abstract representation for the unification part, the interaction between the unification and numerical part has to be revised. The second source of imprecision is the abstraction of non-linear constraints. This is the cause of inaccuracy in **runkut**, **trap**, **mortgage**1, **mortgage**3, **vecmat**1 and also partly in **power** and **vecmat**2. Finally, in the **sendmm** benchmark, imprecise possible dependency information is derived due to abstracting primitive constraints and joining their abstraction via abstract conjunction, instead of abstracting a conjunction of primitive constraints at once. In theory, loss of precision (at least for the freeness part) is also possible due to the imprecise abstraction of disequations and inequalities. However, it did not occur in the benchmarks considered. Also, minimization could lead to loss of precision, by combining via union dependencies that are unrelated (i.e. which result from different OR-branches in the computation). Note that this is not as bad as applying transitivity on dependency relations, as is done for some LP mode analyses, but it may nevertheless lead to imprecise results. Again, no such imprecision was found for the benchmarks. It might be argued that in practical programs, different predicate rules usually establish the same or comparable dependencies between the variables of a call to the predicate.

9.8 Conclusion

The detailed experimental evaluation provides good insight regarding the potential efficiency and accuracy of the analyzers, the main causes for loss of accuracy, and the advantages and disadvantages of the combined analyzer. It shows that the combination of the $\mathcal{D}$ and $\mathcal{F}^{\mathbf{m}}$ analyzers results in a practical full mode analysis system. Moreover, our experiments indicate that the analyses scale up quite well for larger programs. Problems – if any – have not so much to do with the size of the program but with the number of variables in a clause and can be overcome with the use of a widening operator. Our results provide evidence of the feasibility of abstract interpretation as a powerful tool for the analysis of CLP programs.

10. CONCLUSIONS AND DISCUSSION

The generalization of analysis frameworks for logic programs (based on abstract interpretation) has been presented as a practical approach to the dataflow analysis of constraint logic programming languages. In particular, we have proposed an extension of Bruynooghe's traditional framework which allows it to analyze constraint logic programs. Using this generalized framework, two analyses have been proposed for approximating definiteness and freeness information respectively, as well as a combined analysis inferring both properties. We have also reported on the implementations of the framework and the domains and on the study of these implementations. Finally, we have shown that simple widening operators are adequate for controlling the analysis time of large or complex programs. The experimental results support our claim that with the approach proposed it is possible to obtain practical, accurate, and efficient analyses, while reusing much of the framework technology developed for traditional logic programming.

We believe that, given the adaptability of traditional frameworks to CLP analysis, future work might concentrate on accurately approximating the new properties needed for effectively applying the different optimizations relevant to the CLP paradigm. Encouraging examples in this direction are [García de la Banda et al. 1993; Marriott and Stuckey 1994; Macdonald et al. 1993]. The difficulties in this task come from many sources. First, it requires a good abstraction of (possibly many) constraint solver algorithms which are typically more complex than the well known unification. This in turn implies abstracting enough information for simulating the way in which the solver propagates the property of interest. This information seems to be closely related to the abstraction of the entailment relation. The problem is then to determine which constraints from all those entailed are relevant to the property being abstracted. It is interesting to note how correctness problems encountered by early analyzers for LP in the context of variable "aliasing" can be reinterpreted in this context. After analyzing the goals $\mathbf{X} = \mathbf{Y}$, $\mathbf{Y} = \mathbf{Z}$, and $\mathbf{Z} = \mathbf{a}$, $\mathbf{X}$ can be inferred (incorrectly) to be a free variable or (inaccurately) to be $\top$. This problem can now be seen as related to not taking into account the entailed relation $\mathbf{X} = \mathbf{Z}$ which is relevant to the propagation of non-freeness and groundness information.

Second, most CLP languages are defined over several constraint systems and in most cases the theoretical separation among the objects (functors, constraint predicates, domain variables, etc.) of each constraint system is not maintained. Therefore, one must take into account the effects that the conjunction of a particular constraint can produce with respect to any of the other constraint systems in the language. The abstract domains proposed in this paper handle this directly. However, it may be preferable to be able to specify the abstraction for each constraint domain separately and then deal with the interactions. This suggests organizing the domains and analyses as a hierarchy where there is a top-level domain applicable to all constraint systems and some lower level domains which are constraint system specific. The top-level domain would be used for performing the transfer of information among the lower level domains that is necessary in order to preserve correctness and achieve reasonable efficiency. Alternatively, rather than having a top-level domain, transfer functions between all domains can be specified. A negative aspect of the separation of domains and of the explicit interaction between them is that the same information could be represented several times. Also, for some abstractions, it could be difficult to define interaction rules such that there is no loss of precision.

Finally, and from a practical point of view, one must consider the vehicle to be used for implementing the abstract operations. As mentioned before, Codognet and Filé [1992] propose the direct use of CLP solvers in specifying the abstract solving algorithms. The use of the constraint solving capabilities of the implementation language is a very elegant solution and has the advantage that the abstract algorithm can be specified in a declarative way. On the other hand, one favorable aspect of formulating analyses so that they can be executed using only equalities over the Herbrand domain is generality: it will be quite simple to implement them on a large number of CLP systems (and traditional logic programming systems!), given that in general all CLP systems include the Herbrand domain and a unification algorithm.

Appendix: Benchmark programs

The CLP benchmark programs solve typical CLP programs. A representative subset of the CLP programs are used in our detailed experiments. The programs solve typical CLP problems. Most of them are taken from the CLP($\mathcal{R}$) distribution, the PrologIII distribution, or the PRINCE project benchmarks. For this subset, we specify the abstract query. This query is given in the simplified "mode" format available to the user. Modes **d**, **f** and **a** mean that the argument is definite, free or any term respectively. This specification is translated into the appropriate representation for each domain.

—**dnf:** converts a propositional formula into disjunctive normal form; entry pattern $\mathbf{dnf}(\mathbf{d}, \mathbf{f})$.

—**fib:** $\mathbf{fib}(\mathbf{N}, \mathbf{F})$ expresses that $\mathbf{F}$ is the $\mathbf{N^{th}}$ Fibonacci number; entry pattern $\mathbf{fib}(\mathbf{d}, \mathbf{f})$.

—**laplace:** solves the Dirichlet problem for Laplace's equation in two dimensions using Leibman's five-point finite-difference approximation; entry patterns $\mathbf{laplace}1(\mathbf{d})$ and $\mathbf{laplace}2(\mathbf{M})$ where $\mathbf{M}$ is a matrix of free variables.

—**listlength:** specifies the relation between a list and its length; entry pattern **listlength**(**d**, **f**).

—**meal:** computes a balanced meal; entry pattern **lightMeal**(**f**, **f**, **f**).

—**mining:** optimizes the revenue of an open mine; entry pattern **mining**(**f**, **f**).

—**mortgage:** well-known mortgage program; entry patterns **mortgage**1(**a**, **a**, **a**, **a**,**f**), **mortgage**2(**a**, **a**, **a**, **f**, **a**) and **mortgage**3(**f**, **d**, **d**, **d**, **d**).

—**num:** transforms numbers into a sequence of letters and phonemes; entry pattern **nombre**(**d**, **f**, **f**).

—**power:** minimizes the production cost of power stations; entry pattern **pow**(**f**).

—**rectangle:** fills a rectangle with squares; entry pattern **fillRectangle**(**f**, **a**).

—**runkut:** first-order ordinary differential equation solving, using the Runge-Kutta method; entry pattern **solve**(**d**, **d**, **f**).

—**sendmm:** **send** + **more** = **money** puzzle; entry pattern **solution**(**f**, **f**, **f**).

—**sumlist:** specifies the relation between a list of numbers and the sum of its elements; entry pattern **sumlist**(**d**, **f**).

—**trap:** first-order ordinary differential equation solving, using the trapezoidal method; entry pattern **solve**([**d**, **d**], **d**, [**d**, **f**]).

—**vecmat:** performs vector and matrix operations (vector addition **vecadd**, multiplication of a matrix and a vector **matvecmul** and matrix multiplication **matmul**); entry patterns **vecmat**1 which gives rise to **matvecmul**(**d**, **d**, **f**), **vecadd**(**f**, **d**, **d**) and **matmul**(**d**, **d**, **f**), and **vecmat**2 which gives rise to **matvecmul**(**f**, **d**, **d**), **vecadd**(**f**, **d**,**a**) and **matmul**(**d**, **f**, **d**).

The other CLP benchmarks are obtained from P. Van Hentenryck (**bridge**, **cutstock**, **warehouse**), from the book of P. Van Hentenryck [Van Hentenryck 1989] (**magic** (p. 155), **perm** (p. 152)) and his article [Van Hentenryck and Ramachandran 1994] (**periodic** (p. 350)), from PrologIA (**color**4, **color**4**F**, **triangle**) and from an article of A. Colmerauer [Colmerauer 1990] (**magicC**).

Most of the LP benchmarks are used in [Mulkers et al. 1994], from which we borrow the following brief description of the programs. **akl** (called **init_vars** in [Mulkers et al. 1994]) initializes two abstract substitutions to have the same set of variables; **akl_old** is a slightly modified version of **akl**; **ann** is a simplified version of &-Prolog's parallelizing annotator [Hermenegildo and Greene 1990]; **bid** computes an opening bid for a bridge hand; **boyer** is a Boyer-Moore theorem prover from the Gabriel benchmarks (as translated by E. Tick); **browse** is a program for pattern matching also taken from the Gabriel benchmarks (as translated by T. Dobry and H. Touati); **deriv** performs symbolic differentiation of an equation; **grammar** is a program that generates and recognizes a small set of English; **icomp** is a code generator for the WAM, written by Demoen; **kalah** is the Kalah playing program from [Sterling and Shapiro 1994] which uses alpha-beta pruning; **mapcolor** is a map coloring program for a map representation of six countries; **peephole** is the optimizer of SB-Prolog, written by Debray; **rdtok** is O'Keefe's public domain Prolog tokenizer; **read** is Warren and O'Keefe's public domain Prolog parser; **serialize** is a program manipulating lists of numbers; **tarjan** is a program for computing strongly connected components written by Gallagher; **vlokgr** is a consistency checker for a lectures-administration database, written by Janssens; **vlok** is the same program but using an open-ended list for the list of lectures to be checked. The remaining benchmarks are the following: **append** is the well-known append program; **pg** is a program written by W. Older to solve a specific mathematical problem; **plan** is a simple planner in the blocks world; **qsort** implements the quicksort algorithm; **queens** is a generate-and-test program to solve the n-queens problem; **qplan** is part of CHAT, a natural language query interpreter; **witt** is a conceptual clustering system (written by Manuel Carro).

ACKNOWLEDGMENTS

The authors would like to thank Francisco Bueno for his help in improving the implementation of the analysis framework and German Puebla for modifying the system in order to output the inferred information for all the analyzed versions. Thanks are due to Pascal Van Hentenryck, PrologIA and the other PRINCE project partners for providing the benchmark programs. The authors are also grateful to the anonymous referees for their suggestions which have improved both the contents and the presentation of the paper.

REFERENCES

Armstrong, T., Marriott, K., Schachte, P., and Søndergaard, H. 1994. Boolean Functions for Dependency Analysis: Algebraic Properties and Efficient Representation. In *Proceedings of the Static Analysis Symposium*, B. Le Charlier, Ed. Number 864 in Lecture Notes in Computer Science. Springer-Verlag, Namur, Belgium, 266–280.

Bruynooghe, M. 1991. A Practical Framework for the Abstract Interpretation of Logic Programs. *J. Logic Program. 10,* 2 (Feb.), 91–124.

Bruynooghe, M. and Boulanger, D. 1994. Abstract Interpretation for (Constraint) Logic Programming. In *Constraint Programming, NATO ASI Series, Vol. F/131*, B. Mayoh, E. Tyugu, and J. Penjam, Eds. Springer-Verlag, 228–258.

Bruynooghe, M. and Janssens, G. 1992. Propagation: A New Operation in a Framework for Abstract Interpretation of Logic Programs. In *Proceedings of the 3rd International Workshop on Metaprogramming in Logic*, A. Pettorossi, Ed. Number 649 in Lecture Notes in Computer Science. Springer-Verlag, Uppsala, Sweden, 294–307.

Bueno, F., de la Banda, M. G., and Hermenegildo, M. 1994. Effectiveness of Global Analysis in Strict Independence-Based Automatic Program Parallelization. In *Proceedings of the 1994 International Symposium on Logic Programming.* MIT Press, 320–336.

Codish, M., Mulkers, A., Bruynooghe, M., García de la Banda, M., and Hermenegildo, M. 1995. Improving abstract interpretations by combining domains. *ACM Trans. Program. Lang. Syst. 17,* 1 (Jan.), 28–44.

Codognet, P. and Filé, G. 1992. Computations, Abstractions and Constraints in Logic Programs. In *Proceedings of the 4th International Conference on Computer Languages*, J. Cordy, Ed. IEEE Computer Society Press, 155–164.

Colmerauer, A. 1990. An introduction to PROLOGIII. *Commun. ACM 30,* 7, 69–96.

Cousot, P. and Cousot, R. 1977. Abstract Interpretation: a Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proceedings of the 4th ACM Symposium on Principles of Programming Languages.* Los Angeles, 238–252.

Cousot, P. and Cousot, R. 1979. Systematic Design of Program Analysis Frameworks. In *Proceedings of the 6th ACM Symposium on Principles of Programming Languages.* San Antonio, Texas, 269–282.

Cousot, P. and Cousot, R. 1992a. Abstract Interpretation and Application to Logic Programs. *J. Logic Program. 13,* 2 & 3, 103–179.

Cousot, P. and Cousot, R. 1992b. Comparing the Galois Connection with Widening/Narrowing Approaches to Abstract Interpretation. In *Proceedings of the 4th International Symposium on*

Programming Language Implementation and Logic Programming, M. Bruynooghe and M. Wirsing, Eds. Number 631 in Lecture Notes in Computer Science. Springer-Verlag, Leuven, Belgium, 269–295.

Dart, P. 1988. Dependency analysis and query interfaces for deductive databases. Ph.D. thesis, U. of Melbourne, Australia.

Debray, S. K. 1989. Static Inference of Modes and Data Dependencies in Logic Programs. *ACM Trans. Program. Lang. Syst. 11,* 3, 418–450.

Debray, S. K. 1992a. Efficient Dataflow Analysis of Logic Programs. *Journal of the ACM 39,* 4 (Oct.), 949–984.

Debray, S. K., Ed. 1992b. *J.Logic Program., Special Issue: Abstract Interpretation.* Vol. 13(2& 3). North-Holland.

Debray, S. K. 1995. On the Complexity of Dataflow Analysis of Logic Programs. *ACM Trans. Program. Lang. Syst. 17,* 2 (Mar.), 331–365.

Dumortier, V. 1994. Freeness and Related Analyses of Constraint Logic Programs Using Abstract Interpretation. Ph.D. thesis, K.U.Leuven, Dept. of Computer Science.

Dumortier, V. and Janssens, G. 1994. Towards A Practical Full Mode Inference System for CLP(H,N). In *Proceedings of the 11th International Conference on Logic Programming*, P. Van Hentenryck, Ed. MIT Press, Santa Margherita Ligure, Italy, 569–583.

Dumortier, V., Janssens, G., Bruynooghe, M., and Codish, M. 1993. Freeness Analysis in the Presence of Numerical Constraints. In *Proceedings of the 10th International Conference on Logic Programming*, D. S. Warren, Ed. MIT Press, Budapest, Hungary, 100–115.

Englebert, V., Le Charlier, B., Roland, D., and Van Hentenryck, P. 1992. Generic Abstract Interpretation Algorithms for Prolog : Two Optimization Techniques and Their Experimental Evaluation. In *Proceedings of the 4th International Symposium on Programming Language Implementation and Logic Programming (PLILP 92)*, M. Bruynooghe and M. Wirsing, Eds. Number 631 in Lecture Notes in Computer Science. Springer-Verlag, Leuven, Belgium, 311–325. Also in *Software Practice and Experience*, 23(4):419–460, 1993.

García de la Banda, M. 1994. Independence, Global Analysis, and Parallelism in Dynamically Scheduled Constraint Logic Programming. Ph.D. thesis, Universidad Politécnica de Madrid (UPM).

García de la Banda, M. and Hermenegildo, M. 1993. A Practical Approach to the Global Analysis of CLP Programs. In *Proceedings of the 1993 International Logic Programming Symposium*, D. Miller, Ed. MIT Press, Vancouver, Canada, 437–455.

García de la Banda, M., Hermenegildo, M., and Marriott, K. 1993. Independence in Constraint Logic Programs. In *Proceedings of the 1993 International Logic Programming Symposium*, D. Miller, Ed. MIT Press, Vancouver, Canada, 130–146.

García de la Banda, M., Marriott, K., and Stuckey, P. 1995. Efficient Analysis of Logic Programs with Dynamic Scheduling. In *Logic Programming, Proceedings of the 1995 International Symposium (ILPS'95)*, J. LLoyd, Ed. MIT Press, Portland, Oregon, 417–431.

Giacobazzi, R., Debray, S., and Levi, G. 1993. Generalized Semantics and Abstract Interpretation for Constraint Logic Programs. Draft, University of Pisa. Apr. Preliminary version in Proc. FGCS92.

Hanus, M. 1993. Analysis of Nonlinear Constraints in CLP(R). In *Proceedings of the 10th International Conference on Logic Programming*, D. S. Warren, Ed. MIT Press, Budapest, Hungary, 83–99.

Hanus, M. 1995. Analysis of Residuation in Logic Programs. *J. Logic Program. 24,* 3 (Sept.), 161–199.

Hermenegildo, M. and Greene, K. J. 1990. &-Prolog and its performance: Exploiting independent and-parallellism. In *Proceedings of the 7th International Conference on Logic Programming*, D. H. D. Warren and P. Szeredi, Eds. MIT Press, Jerusalem, Israel, 253–268.

Hermenegildo, M., Marriott, K., Puebla, G., and Stuckey, P. 1995. Incremental Analysis of Logic Programs. In *Proceedings of the 12th International Conference on Logic Programming*, L. Sterling, Ed. MIT Press, Kanagawa, Japan. 797–811.

Jacobs, D. and Langen, A. 1992. Static Analysis of Logic Programs for Independent And-Parallelism. *J. Logic Program. 13,* 2 & 3 (July), 291–314.

Jaffar, J. and Lassez, J.-L. 1987. Constraint Logic Programming. In *Proceedings of the 14th ACM Symp. Principles of Programming Languages.* ACM, Munich, Germany, 111–119.

Jaffar, J. and Maher, M. 1994. Constraint Logic Programming: A Survey. *J. Logic Program. 19 & 20*, 503–581.

Jaffar, J., Michaylov, S., Stuckey, P., and Yap, R. 1992. An Abstract Machine for CLP($\mathcal{R}$). In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation.* San Francisco, 128–139.

Janssens, G. and Bruynooghe, M. 1992. Deriving Descriptions of Possible Values of Program Variables by means of Abstract Interpretation. *J. Logic Program. 13,* 2 & 3 (July), 205–258.

Janssens, G., Bruynooghe, M., and Dumortier, V. 1995. A blueprint for an abstract machine for abstract interpretation of (constraint) logic programs. In *Logic Programming, Proceedings of the 1995 International Symposium (ILPS'95)*, J. LLoyd, Ed. MIT Press, Portland, Oregon, 336–350.

Jørgensen, N., Marriott, K., and Michaylov, S. 1991. Some Global Compile-Time Optimizations for CLP($\mathcal{R}$). In *Proceedings of the 1991 International Symposium on Logic Programming*, V. Saraswat and K. Ueda, Eds. MIT Press, San Diego, USA, 420–434.

Lassez, J.-L. and McAloon, K. 1992. A Canonical Form for Generalised Linear Constraints. *Journal of Symbolic Computation 13,* 1 (Jan.), 1–24.

Le Charlier, B., Musumbu, K., and Van Hentenryck, P. 1991. A Generic Abstract Interpretation Algorithm and its Complexity Analysis (extended abstract). In *Proceedings of the 8th International Conference on Logic Programming*, K. Furukawa, Ed. MIT Press, Paris, France, 64–78.

Le Charlier, B., Rossi, S., and Hentenryck, P. V. 1994. An abstract interpretation framework for almost full prolog. In *Proceedings of the 1994 International Logic Programming Symposium*, M. Bruynooghe, Ed. MIT Press, Ithaca, N.Y.

Le Charlier, B. and Van Hentenryck, P. 1994. Experimental Evaluation of a Generic Abstract Interpretation Algorithm for Prolog. *ACM Trans. Program. Lang. Syst. 16,* 1 (Jan.), 35–101.

Lloyd, J. W. 1987. *Foundations of Logic Programming*, Second, Extended ed. Springer Series : Symbolic Computation - Artificial Intelligence. Springer-Verlag.

Macdonald, A. D., Stuckey, P. J., and Yap, R. H. C. 1993. Redundancy of Variables in CLP($\mathcal{R}$). In *Proceedings of the 1993 International Logic Programming Symposium*, D. Miller, Ed. MIT Press, Vancouver, Canada, 75–93.

Marriott, K. 1993. Frameworks for Abstract Interpretation. *Acta Inf. 30*, 103–129.

Marriott, K., García de la Banda, M., and Hermenegildo, M. 1994. Analyzing Logic Programs with Dynamic Scheduling. In *Proceedings of the 20th Annual ACM Conference on Principles of Programming Languages.* ACM, 240–253.

Marriott, K. and Søndergaard, H. 1989. Semantics-Based Dataflow Analysis of Logic Programs. *Information Processing*, 601–606.

Marriott, K. and Søndergaard, H. 1990. Analysis of Constraint Logic Programs. In *Proceedings of the 1990 North American Conference on Logic Programming*, S. Debray and M. Hermenegildo, Eds. MIT Press, Austin, 531–547.

Marriott, K., Søndergaard, H., Stuckey, P. J., and Yap, R. H. 1994. Optimizing Compilation for CLP($\mathcal{R}$). In *Proceedings of the 17th Annual Computer Science Conference.* Christchurch, New Zealand.

Marriott, K. and Stuckey, P. 1993. The 3 R's of Optimizing Constraint Logic Programs : Refinement, Removal and Reordering. In *Proceedings of the 20th ACM Symposium on Principles of Programming Languages.* Charleston, South Carolina, 334–344.

Marriott, K. and Stuckey, P. 1994. Approximating Interaction Between Linear Arithmetic Constraints. In *Proceedings of the 1994 International Symposium on Logic Programming*, M. Bruynooghe, Ed. MIT Press, 571–585.

Martelli, A. and Montanari, U. 1982. An Efficient Unification Algorithm. *ACM Trans. Program. Lang. Syst. 4,* 3, 258–282.

Mellish, C. 1986. Abstract Interpretation of Prolog Programs. In *Proceedings of the 3rd International Conference on Logic Programming*, E. Shapiro, Ed. Number 225 in Lecture Notes in Computer Science. Springer-Verlag, London, UK, 463–475.

Mulkers, A. 1993. *Live Data Structures in Logic Programs, Derivation by Means of Abstract Interpretation.* Number 675 in Lecture Notes in Computer Science. Springer-Verlag.

Mulkers, A., Simoens, W., Janssens, G., and Bruynooghe, M. 1994. On the Practicality of Abstract Equation Systems. Tech. Rep. CW198, Department of Computer Science, Katholieke Universiteit Leuven. Nov.

Mulkers, A., Simoens, W., Janssens, G., and Bruynooghe, M. 1995. On the practicality of abstract equation systems. In *Proceedings of the 12th International Conference on Logic Programming*, L. Sterling, Ed. MIT Press, Kanagawa, Japan, 781–795.

Mulkers, A., Winsborough, W., and Bruynooghe, M. 1990. Analysis of Shared Data Structures for Compile-Time Garbage Collection in Logic Programs. In *Proceedings of the 7th International Conference on Logic Programming*, D. H. D. Warren and P. Szeredi, Eds. MIT Press, Jerusalem, Israel, 747–762.

Mulkers, A., Winsborough, W., and Bruynooghe, M. 1994. Live-structure dataflow analysis for Prolog. *ACM Trans. Program. Lang. Syst. 16,* 2 (Mar.), 205–258.

Muthukumar, K. and Hermenegildo, M. 1989. Determination of Variable Dependence Information at Compile-Time Through Abstract Interpretation. In *Proceedings of the 1989 North American Conference on Logic Programming*, E. Lusk and R. Overbeek, Eds. MIT Press, Cleveland, Ohio, 166–189.

Muthukumar, K. and Hermenegildo, M. 1990. Deriving A Fixpoint Computation Algorithm for Top-Down Abstract Interpretation of Logic Programs. Tech. Rep. ACT-DC-153-90, Microelectronics and Computer Technology Corporation (MCC), Austin, TX 78759. Apr.

Muthukumar, K. and Hermenegildo, M. 1992. Compile-time Derivation of Variable Dependency Using Abstract Interpretation. *J. Logic Program. 13,* 2 & 3 (July), 315–347.

Nielson, F. 1988. Strictness Analysis and Denotational Abstract Interpretation. *Information and Computation 76,* 1, 29–92.

Plaisted, D. A. 1984. The Occur-Check Problem in Prolog. *New Gen. Comput. 2,* 4, 309–322. Also in: *Proceedings of the 1984 International Symposium on Logic Programming*, pages 272–280, Atlantic City. IEEE Computer Society Press.

Ramachandran, V. and Van Hentenryck, P. 1995. LSign reordered. In *International Static Analysis Symposium (SAS'95).* Number 983 in Lecture Notes in Computer Science. Springer-Verlag, Glasgow, 330–347.

Sterling, L. and Shapiro, E. 1994. *The Art of Prolog: Advanced Programming Techniques*, second ed. Logic Programming Series. MIT Press.

Van Hentenryck, P. 1989. *Constraint Satisfaction in Logic Programming.* MIT Press.

Van Hentenryck, P. and Ramachandran, V. 1994. Backtracking without trailing in CLP($\mathcal{R}_{\mathbf{Lin}}$). In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation.* 349–360.

Van Roy, P. and Despain, A. M. 1992. High-Performance Logic Programming with the Aquarius Prolog Compiler. *IEEE Computer*, 54–67.

Warren, R., Hermenegildo, M., and Debray, S. 1988. On the Practicality of Global Flow Analysis of Logic Programs. In *Proceedings of the 5th International Conference and Symposium on Logic Programming*, R. Kowalski and K.A.Bowen, Eds. MIT Press, Seattle, Washington, 684–699.

Received January 1995; revised October 1995; February 1996; accepted April 1996