

CONSTRUCTIVE NEGATION BY PRUNING

FRANÇOIS FAGES

- ▷ We show that a simple concurrent pruning mechanism over standard SLD derivation trees, called constructive negation by pruning, provides a complete operational semantics for normal constraint logic programs (CLP) w.r.t. Fitting-Kunen's 3-valued logic semantics. The principle of concurrent pruning is the only extra machinery needed to handle negation, in particular there is no need for considering complex subgoals with explicit quantifiers outside the constraint part.
- The main result of the paper is the definition of a fixpoint semantics for normal CLP programs which is fully abstract for the observation of computed answer constraints. This allows to generalize the s-semantics approach to normal CLP programs, and provides a fixpoint characterization of Kunen's semantics. The definition is based on a non-ground continuous finitary version of Fitting's operator.
- We relate also these results to an important aspect of CLP programming practice: optimization. We investigate various forms of goal optimization within CLP languages, and provide a declarative semantics for them via a translation to normal CLP programs. We show that constructive negation by pruning specializes for these classes of programs to a more efficient concurrent branch and bound like procedure. ◁

Contents

1	INTRODUCTION	2
2	PRELIMINARIES AND NOTATIONS	3
2.1	Constraint Languages with Negation	3
2.2	CLP(A) Programs	4
3	CONSTRUCTIVE NEGATION BY PRUNING	6
3.1	Procedural Interpretation on SLD Derivation Forests	6

Address correspondence to LIENS, URA 1327 CNRS, Ecole Normale Supérieure, 45 rue d'Ulm, 75230 Paris Cedex 05, France. E-mail: fages@dmi.ens.fr.

3.2	Operational Semantics	8
3.2.1	Uniform Derivations	8
3.2.2	Non-Uniform Derivations	12
4	FULLY ABSTRACT FIXPOINT SEMANTICS	15
5	THREE-VALUED LOGICAL SEMANTICS	21
6	COMPARISON WITH OTHER WORKS	23
7	VARIATIONS ON A SCHEME FOR OPTIMIZATION PREDICATES	25
7.1	Constraint Minimization	25
7.2	Query Optimization	26
7.3	Global Optimization Predicates	28
7.4	Local Optimization Predicates	31
8	CONCLUSION	32

1. INTRODUCTION

Constraint logic programming and concurrent constraint programming are simple and powerful models of computation that have been implemented in several systems over the last decade, and proved successful in a variety of applications ranging from combinatorial optimization problems to complex system modeling [16]. Extending these classes of languages with a negation operator is a major issue as it allows the user to express arbitrary logical combinations of relations.

Negation in logic programming has been extensively studied due to the problems of non-monotonicity and non recursive enumerability of the canonical model approach [1] [19]. On the theoretical side these difficulties have been satisfactorily solved by Kunen [18] and Fitting [12] who proposed to define the declarative semantics of a program by the set of the 3-valued logical consequences of its Clark's completion, and to construct fixpoint semantics in the semi-lattice of partial interpretations. On the implementation side, most constraint logic programming systems allow restricted forms of negation, but the operational mechanism based for instance on negation by failure is too weak w.r.t. Kunen's logical semantics, and the restriction to negative goals containing no variable doesn't fit well with constraint programming. Other ad hoc mechanisms are thus added in most CLP systems for dealing with optimization predicates for instance [28].

Constructive negation, as introduced by Chan [6] [7] for logic programs, and generalized to CLP programs by Stuckey [26], provides an operational mechanism that is correct and complete w.r.t. Kunen's three-valued logical semantics of programs with negation. However the schemes proposed by Chan and Stuckey are not easily amenable to a practical implementation as they necessitate dealing with explicitly quantified complex subgoals, and computing the disjunctive normal form of a complex formula at each resolution step with a negative subgoal. The compilative version proposed by Bruscoli et al. [5], named intensional negation, performs all disjunctive normal form transformations once and for all at compile time, but still

all quantifiers need be explicit at run time and derivation rules need be defined for complex goals.

In this paper we present a scheme for constructive negation based solely on a pruning mechanism over standard SLD-derivation trees, without the need for considering explicitly quantified complex subgoals outside the constraint part. The formalism we develop is based on a simple frontier calculus. The resulting execution model is essentially equivalent to the one proposed independently by Drabent for normal logic programs [8]. We argue that this scheme is simple enough to lead to practical implementations as the principle of concurrent pruning is the only extra machinery needed to handle negation.

The main result of the paper is the definition of a fixpoint semantics for normal CLP programs which is fully abstract for the observation of computed answer constraints. The definition is based on a non-ground continuous finitary version of Fitting's operator, that is similar to (yet different from) the operators studied in [26], [5] and [3]. This result allows to generalize the s-semantics approach [4] to normal CLP programs. It provides also a fixpoint characterization of Kunen's semantics.

In the last section we relate these results to an important aspect of CLP programming practice: optimization. We investigate various forms of goal optimization within CLP languages, and provide a declarative semantics for them via a translation to normal CLP programs. We study these special classes of normal CLP programs and derive from the general scheme of constructive negation by pruning a more efficient concurrent branch and bound like procedure, that is proved correct and complete without any restriction on the degree of nesting of, and on the degree of recursion through, optimization predicates in the program.

2. PRELIMINARIES AND NOTATIONS

We recall the basic concepts of constraint logic programming (CLP) as defined in [15], with some different emphasis due to our interest in negation. Concerning the declarative semantics of CLP programs we focus on the logical semantics instead of the algebraic semantics which is highly undecidable, doing so some conditions such as solution compactness [15] become irrelevant. We adopt also the point of view of [13] and [20] that for a programming language the observation of computed answer constraints is a more natural choice of observable than the success set considered in [15], and that the formal semantics of CLP programs should characterize the set of computed answer constraints. We shall thus present formal semantics accordingly with sets of constrained atoms [4]. Before that we fix notations and make precise the constraint languages and structures considered for CLP programs with negation.

2.1. Constraint Languages with Negation

The first-order language of constraints is defined on a countably infinite set of variables V and on a signature Σ composed of a set of predicate symbols containing *true* and $=$, and of sets of n -place function symbols for each arity n (constants are functions with arity 0). A *primitive constraint* is an atomic proposition of the form $p(t_1, \dots, t_n)$, where p is a predicate symbol in Σ and the t_i 's are Σ, V -terms. A *constraint* is a well-formed first-order Σ, V -formula. The set of free variables in an

expression e is denoted by $V(e)$. Sets of variables will be denoted by X, Y, \dots and we shall sometimes write $e(X)$ if $V(e) = X$. For a constraint c , we shall use the notation $\exists c$ (resp. $\forall c$) to represent the closed constraint $\exists X c$ (resp. $\forall X c$) where $X = V(c)$.

The intended interpretation of constraints is defined by fixing a Σ -structure \mathcal{A} . An \mathcal{A} -valuation for a Σ, V -expression is a mapping $\theta : V \rightarrow \mathcal{A}$ which extends by morphism to terms and primitive constraints. Logical connectives and quantifiers are interpreted as usual, a constraint c is \mathcal{A} -solvable iff $\mathcal{A} \models \exists c$.

It is not necessary for our purpose to suppose that \mathcal{A} is solution compact [15] [21], we suppose only that the constraints are decidable in \mathcal{A} , so that \mathcal{A} can be presented by a decidable first-order theory $th(\mathcal{A})$, i.e. satisfying:

1. (soundness) $\mathcal{A} \models th(\mathcal{A})$,
2. (satisfaction completeness) either $th(\mathcal{A}) \models \exists c$ or $th(\mathcal{A}) \models \neg \exists c$, for any constraint c .

As a constraint is any Σ, V -formula, these conditions are equivalent to say that $th(\mathcal{A})$ is a complete first-order theory, and thus that all models of $th(\mathcal{A})$ are elementary equivalent. For example, Clark's equational theory CET (augmented with the domain closure axiom DCA if the signature is finite) provides such a complete decidable theory for the Herbrand universe with first-order equality constraints [17].

In practice however, the language of constraints will often be a restricted class of Σ, V -formulae, assumed to be closed only by renaming, conjunction and existential quantification, not by negation. Stuckey [26] calls such a restriction a language of *admissible constraints*, which intuitively represents the constraints the solver can deal with. A structure \mathcal{A} is then said to be *admissible* if the negation of an admissible constraint is equivalent to a disjunction of admissible constraints:

$$\mathcal{A} \models \forall X (\neg \exists Y c(X, Y) \leftrightarrow \exists Z_1 d_1(X, Z_1) \vee \dots \vee \exists Z_n d_n(X, Z_n))$$

For the sake of simplicity, we shall assume in this paper that the language of constraints is closed by negation, but we shall indicate latter in section 6 how our scheme can be easily modified to deal with admissible constraints only, when the structure \mathcal{A} is admissible.

2.2. CLP(\mathcal{A}) Programs

$CLP(\mathcal{A})$ programs are defined using an extra finite set of program predicate symbols Π disjoint from constraint predicate symbols Σ . An *atom* has the form $p(t_1, \dots, t_n)$ where $p \in \Pi$ and the t_i 's are Σ, V -terms. A *literal* is either an atom (positive literal) or a negated atom $\neg A$ (negative literal).

A *definite (resp. normal) CLP(\mathcal{A}) program* is a finite set of clauses of the form $A \leftarrow c | L_1, \dots, L_n$ where $n \geq 0$, A is an atom, called the head, c is a constraint, and L_1, \dots, L_n are atoms (resp. literals). The *local variables* of a program clause is the set of free variables in the clause which do not occur in the head. A *definite (resp. normal) goal* is a formula $c | L_1, \dots, L_n$ where L_1, \dots, L_n are atoms (resp. literals). We will identify conjunction “ \wedge ” and multiset union, greek letters, α, β, \dots will be used to denote multisets of literals, so that a goal (resp. a clause) will be sometimes written $c | \alpha$ (resp. $A \leftarrow c | \alpha$). We shall denote by α^+ (resp. α^-) the multiset of positive (resp. negative) literals in α , and by \square the empty multiset. The set of

goals is denoted by \mathcal{G} . In the rest of this paper we shall assume that all atoms in programs and goals contain no constant, no function symbol and no multiple occurrences of a same variable. Of course this is not a restriction as any program or goal can be rewritten in such a *standard form* by introducing new variables and equality constraints with terms. For instance the clause $p(x+1) \leftarrow p(x)$ will be written as $p(y) \leftarrow y = x+1 | p(x)$.

Following the s-semantics approach of [5], the formal semantics of definite $CLP(\mathcal{A})$ programs will be defined by sets of constrained atoms. A *constrained atom* is a couple $c|A$ where c is an \mathcal{A} -solvable constraint such that $V(c) \subseteq V(A)$. The set of constrained atoms is denoted by \mathcal{B} . A *constrained interpretation* is a subset of \mathcal{B} . The set of *ground instances* of a constrained atom over \mathcal{A} is defined by:

$$[c|A]_{\mathcal{A}} = \{A\theta \mid \theta : V \rightarrow \mathcal{A}, \mathcal{A} \models c\theta\}$$

We denote also by $[I]_{\mathcal{A}}$ the set of ground instances of a constrained interpretation I . A ground atom $A\theta$ is true (resp. false) in I if $A\theta \in [I]_{\mathcal{A}}$ (resp. $A\theta \notin [I]_{\mathcal{A}}$).

Constraint entailment defines a natural preorder on constrained atoms, called the *covering preorder*: $c|A \sqsubseteq d|A$ iff $th(\mathcal{A}) \models c \rightarrow d$. Note that as $th(\mathcal{A})$ is a complete theory, $c|A \sqsubseteq d|A$ is equivalent to $[c|A]_{\mathcal{A}} \subseteq [d|A]_{\mathcal{A}}$. The covering preorder extends to sets of constrained atoms in two ways: *strong covering* (used for strong completeness results),

$$I \sqsubseteq J \text{ iff } \forall c|A \in I \exists d|A \in J \text{ } th(\mathcal{A}) \models c \rightarrow d$$

and *finite covering*,

$$I \sqsubseteq_f J \text{ iff } \forall c|A \in I \exists \{d_1|A, \dots, d_n|A\} \subseteq J \text{ } th(\mathcal{A}) \models c \rightarrow \bigvee_{i=1}^n d_i.$$

The operational semantics of *definite* $CLP(\mathcal{A})$ programs is based on a simple transition relation on definite goals, defined as the least relation satisfying the following SLD derivation rule:

$$SLD : c|\alpha, p(X), \alpha' \rightarrow c \wedge c_i|\alpha, \alpha_i, \alpha'$$

for each renamed clause $p(X) \leftarrow c_i|\alpha_i$ defining p in P such that $\mathcal{A} \models \exists(c \wedge c_i)$. We note \rightarrow^* the reflexive transitive closure of \rightarrow . A computed answer constraint (c.a.c.) for a definite goal $c|\alpha$ is a constraint of the form $\exists Y \ d$ such that $c|\alpha \rightarrow^* d|\square$ and $Y = V(d) \setminus V(c|\alpha)$. An and-compositionality lemma (3.7) states that a c.a.c. d for a composite goal $c|A_1, \dots, A_n$ is of the form $d = c \wedge \bigwedge_{i=1}^n c_i$ where the c_i 's are c.a.c. for atomic goals $true|A_i$. Thus the operational behavior of definite $CLP(\mathcal{A})$ programs w.r.t. answer constraints is fully characterized by the following set of constrained atoms:

$$\mathcal{O}(P) = \{\exists Y \ c|p(X) \in \mathcal{B} : true|p(X) \rightarrow^* c|\square, Y = V(c) \setminus X\}$$

Taking as logical semantics

$$\mathcal{L}(P) = \{c|p(X) \in \mathcal{B} : P, th(\mathcal{A}) \models c \rightarrow p(X)\}$$

we obtain the well-known soundness, $\mathcal{O}(P) \subseteq \mathcal{L}(P)$, and completeness, $\mathcal{L}(P) \sqsubseteq_f$

$\mathcal{O}(P)$, results of *SLD*-resolution for definite *CLP*(\mathcal{A}) programs w.r.t. answer constraints [20] [13].

The logical semantics of *normal CLP*(\mathcal{A}) programs is defined via the Clark's completion of the program. The *Clark's completion* of a *CLP*(\mathcal{A}) program P is the conjunction of $th(\mathcal{A})$ with a formula P^* obtained from P by putting in a conjunction the following formula:

$$\forall X \ p(X) \leftrightarrow \bigvee_{i=1}^n \exists Y_i \ c_i \wedge \alpha_i$$

for each predicate symbol p defined in P with a set of clauses $\{p(X) \leftarrow c_i | \alpha_i\}_{1 \leq i \leq n} \in P$, where $Y_i = V(c_i | \alpha_i) \setminus X$, and the formula $\forall X \neg p(X)$ for the other predicate symbols which don't appear in any head in P .

The completion of a normal program can be inconsistent, e.g. with the program $P = \{p \rightarrow \neg p\}$, $P^* = (p \leftrightarrow \neg p)$, in that case any constraint should be a correct answer constraint for any goal. In order to define a faithful logical semantics for normal programs, such contradictions must be localized in the program, the solution proposed by Kunen is to define the logical semantics as the set of 3-valued logical consequences of $P^*, th(\mathcal{A})$. The usual strong 3-valued interpretations of the connectives and quantifiers are assumed, except for the connective $a \leftrightarrow b$ which is interpreted as t if a and b have the same truth value (f , t or u), and f otherwise (i.e. Lukasiewicz's 2-valued interpretation of \leftrightarrow). In the previous example we can assign the undefined truth value to predicate p so that $u \leftrightarrow \neg u$ is true, more generally Fitting [12] showed that any normal logic program has a three-valued model.

The formal semantics of normal *CLP*(\mathcal{A}) programs will be thus defined by partial interpretations. A *partial constrained interpretation* (partial interpretation for short here) is a couple of sets of constrained atoms, $I = \langle I^+, I^- \rangle$, satisfying the following consistency condition: $[I^+]_{\mathcal{A}} \cap [I^-]_{\mathcal{A}} = \emptyset$. The set of partial interpretations forms a semi-lattice for set inclusion on true and false constrained atoms, we denote it by $(\mathcal{I}, \subseteq_3)$. It is not a lattice as the union of two partial interpretations may not be a partial interpretation due to the consistency condition. The preorder \sqsubseteq extends to partial interpretation by $I \sqsubseteq J$ iff $I^+ \sqsubseteq J^+$ and $I^- \sqsubseteq J^-$. And similarly for \sqsubseteq_f .

The (Kunen's) *logical semantics* of a normal *CLP*(\mathcal{A}) program P is defined as the following partial interpretation:

$$\begin{aligned} \mathcal{L}(P) &= \langle \mathcal{L}^+(P), \mathcal{L}^-(P) \rangle \text{ where} \\ \mathcal{L}^+(P) &= \{c | p(X) \in \mathcal{B} : P^*, th(\mathcal{A}) \models_3 c \rightarrow p(X)\}, \\ \mathcal{L}^-(P) &= \{c | p(X) \in \mathcal{B} : P^*, th(\mathcal{A}) \models_3 c \rightarrow \neg p(X)\}. \end{aligned}$$

3. CONSTRUCTIVE NEGATION BY PRUNING

3.1. Procedural Interpretation on SLD Derivation Forests

Constructive negation by pruning can be presented informally as a simple pruning mechanism over standard *SLD*-derivation trees. The idea to resolve a goal $c | \alpha, \neg A$ where $\neg A$ is the selected literal is to develop concurrently two *SLD*-derivation trees, one Ψ_1 for $c | \alpha, (\neg A)$ in which $\neg A$ is not selected, and one Ψ_2 for $c | A$ (see figure 1).

Once a successful derivation is found in Ψ_2 , say with answer constraint d , then Ψ_1 is pruned by adding the constraint $\neg \exists Y d$ where $Y = V(d) \setminus V(c | A)$, to the nodes

in Ψ_1 where that constraint is satisfiable, and by removing the other nodes. This operation is called “pruning by success” (PBS).

Once a successful derivation is found in Ψ_1 , say with answer constraint e , we get a successful derivation for the main goal with answer constraint $f = e \wedge \bigwedge_{i=1}^n \neg \exists Y_i d_i$ where $Y_i = V(d_i) \setminus V(c|A)$, for each frontier¹ $\{d_i|\alpha_i\}_{1 \leq i \leq n}$ in Ψ_2 such that f is satisfiable (the deeper the frontier, the more general the computed answer). This operation is called “success by pruning” (SBP).

The main goal is finitely failed if Ψ_1 gets finitely failed after pruning (note the crucial role of the PBS rule with this respect).




FIGURE 1. Constructive negation by pruning.

Example 3.1. The nesting of negation, and the importance of the PBS rule, can be illustrated by the following program:

$p(X) :- X=0.$
 $p(X) :- p(X).$
 $q(X) :- \text{not } p(X).$


with the goal:

$? \text{ not } q(X)$

¹ A frontier in a SLD-derivation tree is a finite set of nodes in the tree such that every derivation in the tree is either finitely failed or passes through exactly one node of the frontier.

$X=0$

As the query contains no positive literal the first derivation tree is initially trivial. A second derivation tree is developed for $\text{true}|q(X)$, that tree contains one derivation to the goal $\text{true}|\neg p(X)$, thus a third derivation tree is developed for $\text{true}|p(X)$. As $X = 0$ is a success for $p(X)$, the second tree can be pruned with $X \neq 0$ by using the PBS rule (note that the SBP rule doesn't apply here as any frontier in the third tree contains the goal $\text{true}|p(X)$ whose constraint cannot be negated). Then by negating the frontier in the second tree after pruning and by applying the SBP rule we get a successful derivation for the query with answer constraint $X = 0$.



3.2. Operational Semantics

3.2.1. Uniform Derivations We shall first define the operational semantics of constructive negation by pruning with a simple calculus on frontiers of *uniform* SLD trees, i.e. SLD trees such that a tree for $c|\alpha, \alpha'$ is a combination of a tree for $c|\alpha$ and of a tree for $c|\alpha'$.

The set of frontiers is the set $\mathcal{P}_f(\mathcal{G})$ of finite sets of goals. The calculus is based on a binary operator: the cross product of frontiers, \times , and on a negation operator for frontiers w.r.t. a set of variables V , noted $\neg_V F$, which associates to a frontier F the constraint representing the negation of the projection on V of the constraints in F .

Definition 3.2. Given two frontiers $F = \{c_i|\alpha_i\}_{i \in I}$, $F' = \{d_j|\beta_j\}_{j \in J}$, let us define

$$F \times F' = \{(c_i \wedge d_j|\alpha_i, \beta_j) \mid i \in I, j \in J, \mathcal{A} \models \exists(c_i \wedge d_j)\}$$

$$c \times F = \{c|\Box\} \times F = \{(c \wedge c_i|\alpha_i) \mid i \in I, \mathcal{A} \models \exists(c \wedge c_i)\}$$

$$\neg_V F = \bigwedge_{i \in I} \neg \exists Y_i c_i, \text{ where } Y_i = V(c_i) \setminus V$$

$$\mathcal{S}(F) = \{c|\Box \in F \mid \mathcal{A} \models \exists c\}$$

$\mathcal{S}(F)$ is the set of successes in F . $c \times F$ is called the pruning of F by constraint c , that operation will be used to formalize the “pruning by success” rule (PBS) of the previous section.

One can easily check that $(\mathcal{P}_f(\mathcal{G}), \cup, \emptyset, \times, \{\text{true}|\Box\})$ is a commutative semi-ring:

$$\times \text{ is associative and commutative,} \tag{1}$$

$$F \times \emptyset = \emptyset, \tag{2}$$

$$\times \text{ distributes over } \cup, \quad (3)$$

$$\text{furthermore, } \neg_V \emptyset = \text{true}, \quad (4)$$

$$(\neg_V F) \times F = \emptyset, \quad (5)$$

$$\neg_V (F \cup F') = (\neg_V F) \wedge (\neg_V F'), \quad (6)$$

$$\neg_V (F \times F') = (\neg_V F) \vee (\neg_V F'), \quad (7)$$

$$\mathcal{S}(F \times F') = \mathcal{S}(F) \times \mathcal{S}(F'). \quad (8)$$

Now the relation $\triangleleft \in \mathcal{G} \times \mathcal{P}_f(\mathcal{G})$ which associates a frontier to a goal, can be defined inductively as the least relation satisfying the following axiom and rules:

TRIV:	$c \alpha \triangleleft \{c \alpha : \mathcal{A} \models \exists(c)\}$
RES:	$\frac{c \wedge c_1 \alpha_1 \triangleleft F_1 \quad \dots \quad c \wedge c_k \alpha_k \triangleleft F_k}{c p(X) \triangleleft F_1 \cup \dots \cup F_k}$ <p>where $\{(p(X) \leftarrow c_i \alpha_i)\}_{1 \leq i \leq k}$ is the set of renamed apart clauses defining $p(X)$ in P such that $\mathcal{A} \models \exists(c \wedge c_i)$, and we assume $(V(F_i) \setminus V(c \wedge c_i \alpha_i)) \cap X = \emptyset$ in order to avoid variable clashes.</p>
FRT:	$\frac{c \alpha_1 \triangleleft F_1 \quad c \alpha_2 \triangleleft F_2}{c \alpha_1, \alpha_2 \triangleleft F_1 \times F_2}$ <p>where $\alpha_1 \neq \square, \alpha_2 \neq \square$ and in order to avoid variable clashes we assume $V_1 \cap V(\alpha_2) = \emptyset, V_2 \cap V(\alpha_1) = \emptyset, V_1 \cap V_2 = \emptyset$ where $V_1 = V(F_1) \setminus V(c \alpha_1)$ and $V_2 = V(F_2) \setminus V(c \alpha_2)$.</p>
PRN:	$\frac{c A \triangleleft F}{c \neg A \triangleleft c \times \{\neg_V S \neg A, \neg_V F \square\}}$ <p>where $S \subseteq \mathcal{S}(F)$ and $V = V(c A)$.</p>

FIGURE 2. Inductive definition of the goal-frontier relation for uniform derivations.

Note that this presentation of the operational semantics is not in the SOS format of Plotkin insofar as we do not specify a transition relation over states, corresponding to elementary execution steps, but directly its transitive closure representing the possible results of a computation². Rule RES is the usual resolution rule for positive literals. Rule FRT expresses the formation of frontiers by cross products (a more standard operational semantics where frontiers are not formed by cross products but by elementary SLD resolution steps is studied in the next section). The last rule called “pruning” (PRN) is the new inference rule introduced for negative

²It is of course possible to give an incremental SOS presentation of our system but we did not find it elegant nor useful for our purpose. Similar difficulties have been noted for the definition of SLDNF resolution (see [1]). An inductive definition of SLDNF resolution is given in [18].

literals. The two elements of the inferred frontier formalize the pruning by success rule (PBS) and the success by pruning rule (SBP) of the procedural interpretation respectively³. Note that the negation as failure rule is the restriction of the pruning rule to the case $F = \emptyset$ (if $c|A \triangleleft \emptyset$ then $c|\neg A \triangleleft \{c|\neg A, c|\Box\}$ by equation 4 and PRN).

Definition 3.3. A computed answer constraint (c.a.c.) for a goal $c|\alpha$ is a constraint of the form $\exists Y d$ such that $c|\alpha \triangleleft \{d|\Box\} \cup F$ and $Y = V(d) \setminus V(c|\alpha)$. A goal $c|\alpha$ is finitely failed if $c|\alpha \triangleleft \emptyset$.

Example 3.4. Going back to example 3.1, the answer constraint $x = 0$ for the goal $\text{true}|\neg q(x)$ can be obtained by the following proof tree:

$$\begin{array}{c}
 \text{RES} \quad \frac{x = 0|\Box \triangleleft \{x = 0|\Box\} \quad \text{true}|p(x) \triangleleft \{\text{true}|p(x)\}}{\text{PRN} \quad \frac{\text{true}|p(x) \triangleleft \{x = 0|\Box, \text{true}|p(x)\}}{\text{RES} \quad \frac{\text{true}|\neg p(x) \triangleleft \{x \neq 0|\neg p(x)\}}{\text{PRN} \quad \frac{\text{true}|q(x) \triangleleft \{x \neq 0|\neg p(x)\}}{\text{true}|\neg q(x) \triangleleft \{\text{true}|\neg q(x), x = 0|\Box\}}}}
 \end{array}$$

By a simple inspection of the rules we can easily state several lemma on the goal-frontier relation \triangleleft . For some proofs we shall use the principle of structural induction on proof trees for \triangleleft , that is we shall show that a property holds for \triangleleft , simply by showing that it holds for the axiom TRIV, and for the conclusion of the rules RES, FRT and PRN assuming it holds for the premises of these rules.

Lemma 3.5. (instantiation lemma) If $c|\alpha \triangleleft F$ then for any constraint d there exists a frontier F' such that $c \wedge d|\alpha \triangleleft F'$ and $F' = d \times F$.

PROOF. The proof is by structural induction on a proof tree for $c|\alpha \triangleleft F$.

TRIV: We have $F = \{c|\alpha : \mathcal{A} \models \exists c\}$. By rule TRIV we have also $c \wedge d|\alpha \triangleleft F'$ with $F' = \{c \wedge d|\alpha : \mathcal{A} \models \exists(c \wedge d)\} = d \times F$.

RES: We have $\alpha = p(X)$ and $F = \bigcup_{i \in I} F_i$ where $\{p(X) \leftarrow c_i|\alpha_i\}_{i \in I}$ is the set of renamed rules defining $p(X)$ in P such that $\mathcal{A} \models \exists(c \wedge c_i)$, and $c \wedge c_i|\alpha_i \triangleleft F_i$. By the induction hypothesis we get $c \wedge c_i \wedge d|\alpha_i \triangleleft d \times F_i$. Let $J \subseteq I$ be the subset of indices such that $c \wedge c_i \wedge d$ is \mathcal{A} -satisfiable, then by the RES rule we get $c \wedge d|p(X) \triangleleft F'$ with $F' = \bigcup_{j \in J} d \times F_j = d \times \bigcup_{i \in I} F_i = d \times F$.

FRT: We have $\alpha = \alpha_1, \alpha_2$, $c|\alpha_1 \triangleleft F_1$, $c|\alpha_2 \triangleleft F_2$ and $F = F_1 \times F_2$. By induction we get $c \wedge d|\alpha_1 \triangleleft d \times F_1$ and $c \wedge d|\alpha_2 \triangleleft d \times F_2$, hence by rule FRT and equation 1 we have $c \wedge d|\alpha \triangleleft d \times F$.

PRN: We have $\alpha = \neg A$ and $F = c \times \{\neg_V S|\neg A, \neg_V F''|\Box\}$ with $c|A \triangleleft F''$, $S \subseteq \mathcal{S}(F'')$, and $V = V(c|A)$.

By the induction hypothesis we get $c \wedge d|A \triangleleft d \times F''$, so by the PRN rule we

³The fact that in the procedural interpretation the SBP rule need be applied only to successful derivations in the main tree is justified at the end of this section (cf. prop. 3.12).

have $c \wedge d \mid \neg A \triangleleft F'$ with $F' = (c \wedge d) \times \{\neg_{V'}(d \times S) \mid \neg A, \neg_{V'}(d \times F'') \mid \square\}$ and $V' = V \cup V(d)$. Now

$$\begin{aligned} F' &= c \times (d \times \{\neg_{V'} d \vee \neg_{V'} S \mid \neg A, \neg_{V'} d \vee \neg_{V'} F'' \mid \square\}) \text{ by eq. 1 and 7,} \\ &= c \times \{d \wedge \neg_{V'} S \mid \neg A, d \wedge \neg_{V'} F'' \mid \square\}, \\ &= d \times c \times \{\neg_{V'} S \mid \neg A, \neg_{V'} F'' \mid \square\}, \\ &= d \times F. \end{aligned}$$

□

Lemma 3.6. (lifting lemma) If $c \mid \alpha \triangleleft F$ then there exists F' such that $\text{true} \mid \alpha \triangleleft F'$ and $F = c \times F'$.

PROOF. By structural induction, similarly to the proof of lemma 3.5. □

Lemma 3.7. (And-compositionality of uniform derivations)

$c \mid \alpha_1, \alpha_2 \triangleleft F$ if and only if there exist F_1 and F_2 such that $\text{true} \mid \alpha_1 \triangleleft F_1$, $\text{true} \mid \alpha_2 \triangleleft F_2$, and $F = c \times F_1 \times F_2$.

PROOF.

⇒ The proof is by cases on the root rule of a proof tree for $c \mid \alpha_1, \alpha_2 \triangleleft F$.

TRIV: we have $F = \{c \mid \alpha_1, \alpha_2 : \mathcal{A} \models \exists c\}$. By rule TRIV we can take $F_1 = \{\text{true} \mid \alpha_1\}$ and $F_2 = \{\text{true} \mid \alpha_2\}$, thus $F = c \times F_1 \times F_2$.

RES: we have $\alpha_1 = p(X)$ and $\alpha_2 = \square$, by lifting lemma 3.6 we get $\text{true} \mid p(X) \triangleleft F_1$ with $F = c \times F_1$, and by rule TRIV we can take $F_2 = \{\text{true} \mid \square\}$ so that $F = c \times F_1 \times F_2$.

FRT: By lifting lemma 3.6 we immediately get $F = c \times F_1 \times F_2$.

PRN: same proof as for the RES case.

⇐ By instantiation lemma 3.5, we get $c \mid \alpha_1 \triangleleft c \times F_1$ and $c \mid \alpha_2 \triangleleft c \times F_2$, hence by rule FRT we have $c \mid \alpha_1, \alpha_2 \triangleleft F$ with $F = (c \times F_1) \times (c \times F_2) = c \times F_1 \times F_2$.

□

Corollary 3.8. (Canonical proof trees)⁴

Any derivation admits a canonical proof tree in which in each application of the FRT rule, α_1 is a literal.

PROOF. By taking the first literal of the goal for α_1 in lemma 3.7 we can build recursively a canonical proof tree for any derivation. □

Corollary 3.9. (And-compositionality of computed answer constraints)

d is a computed answer constraint for the goal $c \mid A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$, if and only if there exists computed answer constraints c_1, \dots, c_n for the goals $\text{true} \mid A_1, \dots, \text{true} \mid A_m, \text{true} \mid \neg A_{m+1}, \dots, \text{true} \mid \neg A_n$ respectively, such that $d = c \wedge \bigwedge_{i=1}^n c_i$ is \mathcal{A} -satisfiable.

PROOF. By n applications of the lemma. □

Uniform derivations can thus be decomposed into elementary derivations, one for each literal in the query. That fundamental property does not hold for arbitrary SLD derivations, but we shall show in the next subsection that any finite SLD derivation can be extended to a uniform derivation (theorem 3.19).

⁴Canonical proof trees will be used only in the proof of lemma 3.18.

Lemma 3.10. (finite failure lemma) If c is a computed answer constraint for $\text{true}|\neg p(X)$ then $c|p(X) \triangleleft \emptyset$. Conversely, if $c|p(X) \triangleleft \emptyset$ then there exists a computed answer constraint d for $\text{true}|\neg p(X)$ such that $\mathcal{A} \models c \rightarrow d$.

PROOF. First let us suppose $\text{true}|\neg p(X) \triangleleft F$ with $d|\Box \in \mathcal{S}(F)$, $c = \exists Y d$, $Y = V(d) \setminus X$. Necessarily the PRN rule is applied at the root of a proof tree for $\text{true}|\neg p(X) \triangleleft F$, hence we have $\text{true}|p(X) \triangleleft F'$ with $F' = \text{true} \times \{\neg_X S | \neg p(X), \neg_X F' | \Box\}$ and $S \subseteq \mathcal{S}(F')$. Thus $d = \neg_X F' = c$. Hence by instantiation lemma 3.5, we have $c|p(X) \triangleleft c \times F'$, and $c \times F' = \neg_X F' \times F' = \emptyset$ by eq. 5.

Conversely, let us suppose $c|p(X) \triangleleft \emptyset$. Then by applying the PRN rule we get $c|\neg p(X) \triangleleft c \times \{\text{true}|\neg p(X), \text{true}|\Box\}$, hence c is a c.a.c. for $c|\neg p(X)$. Therefore by corollary 3.9, there exists a c.a.c. d for $\text{true}|\neg p(X)$ such that $\mathcal{A} \models c \rightarrow d$. \square

In view of these lemmas, the observation of finite failure on an atom is equivalent to the observation of a success on the negation of the atom (lemma 3.10), and the computed answer constraints for a goal can be retrieved from the computed answer constraints for the unconstrained literals that appear in the goal (lifting lemma 3.6, and corollary 3.9). Therefore we can define the operational semantics of the program as the set of computed answer constraints for unconstrained literals solely.

Definition 3.11. $\mathcal{O}(P) = \langle \mathcal{O}^+(P), \mathcal{O}^-(P) \rangle$

$$\mathcal{O}^+(P) = \{c|p(X) \in \mathcal{B} : c \text{ is a c.a.c. for the goal } \text{true}|p(X)\}$$

$$\mathcal{O}^-(P) = \{c|p(X) \in \mathcal{B} : c \text{ is a c.a.c. for the goal } \text{true}|\neg p(X)\}$$

Note that in the procedural interpretation of the previous section the SBP rule need be applied only to the success nodes in the main tree, not to all nodes as in the PRN rule. This difference obviously does not affect successful derivations in the main tree, nor does it affect the negation of a frontier in that tree:

Proposition 3.12. (negation of frontiers obtained by the PRN rule) Let U, V be two sets of variables and S, F be two frontiers s.t. $S \subseteq F$. Then $\neg_U \{\neg_V S | \alpha, \neg_V F | \beta\} = \neg_U \{\neg_V S | \alpha\}$.

PROOF. Let $F = \{c_i | \alpha_i\}_{i \in I}$, and $S = \{c_j | \alpha_j\}_{j \in J}$ where $J \subseteq I$. For all $i \in I$ let $Y_i = V(c_i) \setminus V$ and $Z_i = (V(c_i) \setminus V) \setminus U$. We have

$$\begin{aligned} \neg_U \{\neg_V S | \alpha, \neg_V F | \beta\} &= \neg_U \{\bigwedge_{j \in J} \neg \exists Y_j c_j | \alpha, \bigwedge_{i \in I} \neg \exists Y_i c_i | \beta\} \\ &= \bigvee_{j \in J} \exists Z_j \exists Y_j c_j \wedge \bigvee_{i \in I} \exists Z_i \exists Y_i c_i \\ &= \bigvee_{j \in J} \exists Z_j \exists Y_j c_j \\ &= \neg_U \{\neg_V S | \alpha\}. \end{aligned}$$

\square

The successful derivations are thus the same in the procedural interpretation and in the \triangleleft relation. This shows that the operational semantics and the procedural interpretation are indeed equivalent w.r.t. computed answer constraints.

3.2.2. Non-Uniform Derivations Standard SLD trees are formed by elementary SLD derivation steps instead of cross products. Although not necessary for the rest of this paper, it is thus interesting to study the goal-frontier relation $\triangleleft \in \mathcal{G} \times \mathcal{P}_f(\mathcal{G})$ defined as the \triangleleft relation except that the RES and FRT rules are replaced by the standard SLD resolution rule, and the PRN rule is generalized to conjunctive goals. The inductive definition of the \triangleleft relation is given in figure 3.

TRIV:	$c \alpha \ll \{c \alpha : \mathcal{A} \models \exists c\}$
SLD:	$\frac{c \wedge c_1 \alpha, \alpha_1, \alpha' \ll F_1 \quad \dots \quad c \wedge c_k \alpha, \alpha_k, \alpha' \ll F_k}{c \alpha, p(X), \alpha' \ll F_1 \cup \dots \cup F_k}$ <p>where $\{(p(X) \leftarrow c_i \alpha_i)\}_{1 \leq i \leq k}$ is the set of renamed clauses defining p in P such that $\mathcal{A} \models \exists(c \wedge c_i)$, and $(V(F_i) \setminus V(c \wedge c_i \alpha, \alpha_i, \alpha')) \cap X = \emptyset$.</p>
PRN:	$\frac{c A \ll F_1 \quad c \alpha, \alpha' \ll F_2}{c \alpha, \neg A, \alpha' \ll c \times \{\neg_V S \neg A, \neg_V F_1 \Box\} \times F_2}$ <p>where $S \subseteq \mathcal{S}(F_1)$, $V = V(c A)$ and $(V(F_2) \setminus V(c \alpha, \alpha')) \cap V = \emptyset$.</p>

FIGURE 3. Inductive definition of the goal-frontier relation for non-uniform derivations.

Example 3.13. Let $P = \{p(x) \leftarrow x = 0, p(x) \leftarrow x = 1, q(x, y) \leftarrow p(x), p(y)\}$. We have the following proof tree for the goal $\text{true}|q(x, y)$:

$$\begin{array}{c}
 \frac{x = 1 \wedge y = 0|\Box \ll \{x = 1 \wedge y = 0|\Box\} \quad x = 1 \wedge y = 1|\Box \ll \{x = 1 \wedge y = 1|\Box\}}{x = 0|p(y) \ll \{x = 0|p(y)\} \quad x = 1|p(y) \ll \{x = 1 \wedge y = 0|\Box, x = 1 \wedge y = 1|\Box\}} \\
 \frac{\text{true}|p(x), p(y) \ll \{x = 0|p(y), x = 1 \wedge y = 0|\Box, x = 1 \wedge y = 1|\Box\}}{\text{true}|q(x, y) \ll \{x = 0|p(y), x = 1 \wedge y = 0|\Box, x = 1 \wedge y = 1|\Box\}}
 \end{array}$$

hence $c = (x \neq 0 \wedge (x \neq 1 \vee (y \neq 0 \wedge y \neq 1)))$ is now a computed answer constraint for the query $\text{true}|\neg q(x, y)$. On the other hand we have

$$\text{true}|q(x, y) \triangleleft \{x = 0|p(y), x = 1|p(y)\}, \text{ or}$$

$$\text{true}|q(x, y) \triangleleft \{x = 0 \wedge y = 0|\Box, x = 0 \wedge y = 1|\Box, x = 1 \wedge y = 0|\Box, x = 1 \wedge y = 1|\Box\}$$

but the answer constraint c for $\text{true}|\neg q(x, y)$ cannot be computed by a uniform derivation, as $p(y)$ cannot be developed in one branch and not in another as in a non-uniform derivation.

Definition 3.14. Let us define the operational semantics for non-uniform derivations as $\tilde{\mathcal{O}}(P) = \langle \tilde{\mathcal{O}}^+(P), \tilde{\mathcal{O}}^-(P) \rangle$ where

$$\tilde{\mathcal{O}}^+(P) = \{c|p(X) \in \mathcal{B} : c \text{ is a } \ll\text{-c.a.c. for the goal } \text{true}|p(X)\}$$

$$\tilde{\mathcal{O}}^-(P) = \{c|p(X) \in \mathcal{B} : c \text{ is a } \ll\text{-c.a.c. for the goal } \text{true}|\neg p(X)\}$$

Not surprisingly, one can easily check that uniform derivations can be simulated by non-uniform derivations:

Proposition 3.15. If $c|\alpha \triangleleft F$ then $c|\alpha \ll F$.

PROOF. By structural induction on a proof tree for $c|\alpha \triangleleft F$. \square

Corollary 3.16. $\mathcal{O}(P) \subseteq \tilde{\mathcal{O}}(P)$.

Of course the previous example shows that the converse of that proposition doesn't hold but we can show that any non-uniform derivation can be extended to a uniform derivation, and thus that computed answers obtained by non-uniform derivations are covered by computed answers obtained by uniform derivations. For this result an extra technical lemma is needed on uniform derivations.

Definition 3.17. Let V be a set of variables, and S, T be two sets of success goals. The (strong) covering preorder w.r.t. V is defined by $S \sqsubseteq_V T$ iff for all $c|\square \in S$ there exists $d|\square \in T$ s.t. $\mathcal{A} \models \exists Y c \rightarrow \exists Z d$ where $Y = V(c) \setminus V$ and $Z = V(d) \setminus V$.

Lemma 3.18. Let $c|\alpha$ be a goal and $V = V(c|\alpha)$. If $c|\alpha \triangleleft F$ and $c|\alpha \triangleleft F'$ then there exists F'' such that $c|\alpha \triangleleft F''$ with $\mathcal{S}(F) \sqsubseteq_V \mathcal{S}(F'')$, $\mathcal{S}(F') \sqsubseteq_V \mathcal{S}(F'')$, $\mathcal{A} \models \neg_V F \rightarrow \neg_V F''$ and $\mathcal{A} \models \neg_V F' \rightarrow \neg_V F''$.

PROOF. The proof is by structural induction on the cartesian product of canonical proof trees (cf. proposition 3.8) for $c|\alpha \triangleleft F$ and $c|\alpha \triangleleft F'$. As the rules RES, FRT and PRN are mutually exclusive there are only 5 cases.

TRIV — We just have to take $F'' = F'$.

— TRIV We take $F'' = F$.

RES-RES Then $\alpha = p(X)$, let $\{p(X) \leftarrow c_k|\alpha_k\}_{k \in K}$ be the set of clauses defining p in P s.t. $c \wedge c_k$ is \mathcal{A} -satisfiable. We have $F = \bigcup_{k \in K} F_k$ with $c \wedge c_k|\alpha_k \triangleleft F_k$ for all $k \in K$, and $F' = \bigcup_{k \in K} F'_k$ with $c \wedge c_k|\alpha_k \triangleleft F'_k$ for all $k \in K$. By the induction hypothesis for all $k \in K$ there exist F''_k such that $c \wedge c_k|\alpha_k \triangleleft F''_k$, $\mathcal{S}(F_k) \sqsubseteq_V \mathcal{S}(F''_k)$, $\mathcal{S}(F'_k) \sqsubseteq_V \mathcal{S}(F''_k)$, $\mathcal{A} \models \neg_V F_k \rightarrow \neg_V F''_k$ and $\mathcal{A} \models \neg_V F'_k \rightarrow \neg_V F''_k$. Hence by the RES rule we get $c|p(X) \triangleleft F''$ with $F'' = \bigcup_{k \in K} F''_k$. Furthermore $\mathcal{S}(F) = \bigcup_{k \in K} \mathcal{S}(F_k) \sqsubseteq_V \mathcal{S}(F'')$, and similarly $\mathcal{S}(F') \sqsubseteq_V \mathcal{S}(F'')$. We have also $\neg_V F = \bigwedge_{k \in K} \neg_V F_k$ and $\neg_V F'' = \bigwedge_{k \in K} \neg_V F''_k$ by eq. 6, thus $\mathcal{A} \models \neg_V F \rightarrow \neg_V F''$ and similarly $\mathcal{A} \models \neg_V F' \rightarrow \neg_V F''$.

FRT-FRT As the proof trees are canonical, we have $\alpha = L, \alpha_2$, $c|L \triangleleft F_1$, $c|\alpha_2 \triangleleft F_2$, $F = F_1 \times F_2$, $c|L \triangleleft F'_1$, $c|\alpha_2 \triangleleft F'_2$, and $F' = F'_1 \times F'_2$. Now let $F'' = F''_1 \times F''_2$ where $c|L \triangleleft F''_1$ and $c|\alpha_2 \triangleleft F''_2$ are given by the induction hypothesis, by the FRT rule we have $c|\alpha \triangleleft F''$ and we easily check that the rest of the induction hypothesis is satisfied.

PRN-PRN Here $\alpha = \neg A$, $F = c \times \{\neg_V S_1 | \neg A, \neg_V F_1 | \square\}$ with $c|A \triangleleft F_1$, $S_1 \subseteq \mathcal{S}(F_1)$, and $F' = c \times \{\neg_V S_2 | \neg A, \neg_V F_2 | \square\}$ with $c|A \triangleleft F_2$, $S_2 \subseteq \mathcal{S}(F_2)$. By the induction hypothesis there exists F''_1 such that $c|A \triangleleft F''_1$, $\mathcal{S}(F_1) \sqsubseteq_V \mathcal{S}(F''_1)$, $\mathcal{S}(F_2) \sqsubseteq_V \mathcal{S}(F''_1)$, $\mathcal{A} \models \neg_V F_1 \rightarrow \neg_V F''_1$ and $\mathcal{A} \models \neg_V F_2 \rightarrow \neg_V F''_1$. Now let $F'' = c \times \{\neg_V S | \neg A, \neg_V F''_1 | \square\}$ where $S = \mathcal{S}(F''_1)$. By the PRN rule we have $c|\neg A \triangleleft F''$, furthermore $\mathcal{S}(F'') = \{c \wedge \neg_V F''_1 | \square\} \sqsubseteq_V \mathcal{S}(F)$, similarly $\mathcal{S}(F') \sqsubseteq_V \mathcal{S}(F'')$. Finally by proposition 3.12 we have $\neg_V F = \neg c \vee \bigvee_{s|\square \in S_1} \exists Y_s s$ and $\neg_V F'' = \neg c \vee \bigvee_{s|\square \in S} \exists Y_s s$ where $Y_s = V(s) \setminus V$, thus $\mathcal{A} \models \neg_V F \rightarrow \neg_V F''$ and similarly $\mathcal{A} \models \neg_V F' \rightarrow \neg_V F''$.

\square

Theorem 3.19. (Extension to uniform derivations) Let G be a goal and $V = V(G)$. If $G \triangleleft F$ then there exists F' such that $G \triangleleft F'$, $\mathcal{S}(F) \sqsubseteq_V \mathcal{S}(F')$ and $\mathcal{A} \models \neg_V F \rightarrow \neg_V F'$.

PROOF. The proof is by structural induction on a proof tree for $G \triangleleft F$.

TRIV: We take $F' = F$.

SLD: We have $G = \alpha, p(X), \alpha', F = \bigcup_{k \in K} F_k$ where $\{p(X) \leftarrow c_k | \alpha_k\}_{k \in K}$ is the set of clauses defining p in P such that $\mathcal{A} \models \exists(c \wedge c_k)$, and $c \wedge c_k | \alpha, \alpha_k, \alpha' \triangleleft F_k$. By induction for all $k \in K$ there exist F'_k such that $c \wedge c_k | \alpha, \alpha_k, \alpha' \triangleleft F'_k$, $\mathcal{S}(F_k) \sqsubseteq \mathcal{S}(F'_k)$ and $\mathcal{A} \models \neg F_k \rightarrow \neg F'_k$.

By lemmas 3.7 and 3.5, for all $k \in K$ there exist F''_k and F'''_k such that $c \wedge c_k | \alpha_k \triangleleft F''_k$, $c | \alpha, \alpha' \triangleleft F'''_k$ and $F'_k = F''_k \times F'''_k$.

Hence by the RES rule we get

$$c | p(X) \triangleleft \bigcup_{k \in K} F''_k.$$

Furthermore by lemma 3.18 there exists F''' such that

$$c | \alpha, \alpha' \triangleleft F'''$$

and for all $k \in K$ $\mathcal{S}(F''_k) \sqsubseteq_V \mathcal{S}(F''')$ and $\mathcal{A} \models \neg_V F''_k \rightarrow \neg_V F'''$. Let $F' = \bigcup_{k \in K} F''_k \times F'''$, by the FRT rule we get

$$c | \alpha, p(X), \alpha' \triangleleft F'$$

$$\begin{aligned} \text{Now } \mathcal{S}(F) &= \bigcup_{k \in K} \mathcal{S}(F_k) \\ &\sqsubseteq_V \bigcup_{k \in K} \mathcal{S}(F'_k) \\ &\sqsubseteq_V \bigcup_{k \in K} \mathcal{S}(F''_k) \times \mathcal{S}(F''') \text{ by eq. 8} \\ &\sqsubseteq_V \bigcup_{k \in K} \mathcal{S}(F''_k) \times \mathcal{S}(F''') \\ &\sqsubseteq_V \mathcal{S}(F') \text{ by eq. 8.} \end{aligned}$$

Furthermore $\neg_V F = \bigwedge_{k \in K} \neg_V F_k$ by eq. 6,

thus $\mathcal{A} \models \neg_V F \rightarrow \bigwedge_{k \in K} \neg_V F''_k \vee \neg_V F'''$ by eq. 7

$$\mathcal{A} \models \neg_V F \rightarrow \bigwedge_{k \in K} \neg_V F''_k \vee \neg_V F'''$$

$$\mathcal{A} \models \neg_V F \rightarrow \neg_V F' \text{ by eq. 6 and 7.}$$

PRN: We have $G = c | \alpha, \neg A, \alpha', F = c \times \{\neg_V S | \neg A, \neg_V F_1 | \square\} \times F_2$ with $S \subseteq \mathcal{S}(F_1)$, $c | A \triangleleft F_1$ and $c | \alpha, \alpha' \triangleleft F_2$. By induction there exist F'_1 and F'_2 such that $c | A \triangleleft F'_1$, $c | \alpha, \alpha' \triangleleft F'_2$, $\mathcal{S}(F_1) \sqsubseteq_V \mathcal{S}(F'_1)$, $\mathcal{S}(F_2) \sqsubseteq_V \mathcal{S}(F'_2)$, $\mathcal{A} \models \neg_V F_1 \rightarrow \neg_V F'_1$ and $\mathcal{A} \models \neg_V F_2 \rightarrow \neg_V F'_2$. Let $F' = c \times \{\neg_V \mathcal{S}(F'_1) | \neg A, \neg_V F'_1 | \square\} \times F'_2$, by the PRN rule we have $G \triangleleft F'$ and we easily check that $\mathcal{S}(F) \sqsubseteq_V \mathcal{S}(F')$ and $\mathcal{A} \models \neg_V F \rightarrow \neg_V F'$ by proposition 3.12.

□

Corollary 3.20. $\tilde{\mathcal{O}}(P) \sqsubseteq \mathcal{O}(P)$.

4. FULLY ABSTRACT FIXPOINT SEMANTICS

In this section we define a continuous non-ground variant of Fitting's operator for constraint logic programs. We show that the least fixed point of that operator is in fact *equal* to the operational semantics of constructive negation by pruning. Such

a full abstraction result allows to generalize the s-semantics approach [4] to normal CLP programs. We first recall the definition of Fitting's operator Φ_P^A .

Given a $\Pi - \Sigma$ -algebra \mathcal{A} , the \mathcal{A} -ground base $\mathcal{B}_A = [\mathcal{B}]_A$ is the set of ground instances of the base \mathcal{B} of constrained atoms. A *partial ground interpretation* is a couple $I = \langle I^+, I^- \rangle$ such that $I^+, I^- \subseteq \mathcal{B}_A$ and $I^+ \cap I^- = \emptyset$. A ground atom A is true (resp. false) in I iff $A \in I^+$ (resp. $A \in I^-$). A constraint is true in I if it is true in \mathcal{A} . A first-order $\Sigma - \Pi$ -formula ψ is true in I , noted $I \models_3 \psi$, if it is true under the usual strong three-valued interpretation of the logical symbols. The set of partial ground interpretations forms a semi-lattice for set inclusion on true and false atoms, we denote it by $(\mathcal{GI}, \subseteq_3)$.

Definition 4.1 [26][18][12]. Let P be a normal CLP(\mathcal{A}) program, the immediate consequence operator $\Phi_P^A : \mathcal{GI} \rightarrow \mathcal{GI}$ is defined by:

$$\begin{aligned} \Phi_P^A(I) &= \{A \in \mathcal{B}_A \mid \text{there exist a clause in } P, p(X) \leftarrow c \mid \alpha, \\ &\quad \text{and a valuation } \theta \text{ such that } A = p(X)\theta \text{ and } I \models_3 (c \wedge \alpha)\theta\} \\ \Phi_P^A(I) &= \{A \in \mathcal{B}_A \mid \text{for any clause in } P, p(X) \leftarrow c \mid \alpha, \\ &\quad \text{and any valuation } \theta \text{ such that } A = p(X)\theta \text{ then } I \models_3 \neg(c \wedge \alpha)\theta\}. \end{aligned}$$

Φ_P^A is a monotonic operator in the semi-lattice of partial ground interpretations. It thus admits a least fixpoint which is the least three-valued \mathcal{A} -model of the program's completion [12]. It is not continuous however, so its power at ordinal ω is generally not a fixpoint (cf. example 4.3).

In order to abstract from a given algebra \mathcal{A} and to prove completeness results, Stuckey [26] defined a non-ground version of Fitting's operator based on partial constrained interpretations. In his definition the downward closure of constrained atoms by their instances prevents however a characterization of the operational behavior of the program w.r.t. answer constraints. Furthermore the operator of Stuckey is not continuous either, so it doesn't provide CLP programs with a fixpoint semantics.

The idea of the operator T_P for obtaining a fully abstract fixpoint semantics is simply to take the finitary, hence continuous, non-downward closed constraint based version of Fitting's operator. So a constrained atom will be true (resp. false) in $T_P(I)$ if the constraint in the constrained atom is a combination of constraints in a finite part of I which validates the body of a program clause for the atom (resp. invalidates the body of all program clauses for the atom).

Definition 4.2. Let P be a CLP(\mathcal{A}) program. T_P is an operator over $2^{\mathcal{B}} \times 2^{\mathcal{B}}$ defined by $T_P(I) = \langle T_P^+(I), T_P^-(I) \rangle$ where:

$$\begin{aligned} T_P^+(I) &= \{c \mid p(X) \in \mathcal{B} : \text{there exist a clause in } P \text{ with local variables } Y, \\ &\quad p(X) \leftarrow d \mid A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n. \\ &\quad c_1 \mid A_1, \dots, c_m \mid A_m \in I^+, c_{m+1} \mid A_{m+1}, \dots, c_n \mid A_n \in I^- \\ &\quad \text{such that } c = \exists Y \, d \wedge \bigwedge_{i=1}^n c_i \text{ is } \mathcal{A}\text{-satisfiable}\} \\ T_P^-(I) &= \{c \mid p(X) \in \mathcal{B} : \text{for each clause defining } p \text{ in } P \text{ with local variables } Y_k, \\ &\quad p(X) \leftarrow d_k \mid A_{k,1}, \dots, A_{k,m_k}, \alpha_k^+, \alpha_k^-. \\ &\quad \text{there exist } e_{k,1} \mid A_{k,1}, \dots, e_{k,m_k} \mid A_{k,m_k} \in I^-, \\ &\quad e_{k,m_k+1} \mid A_{k,m_k+1}, \dots, e_{k,n_k} \mid A_{k,n_k} \in I^+, \\ &\quad \text{where for } m_k + 1 \leq j \leq n_k, \neg A_{k,j} \text{ occurs in } \alpha_k^-, \\ &\quad \text{such that } c = \bigwedge_k \forall Y_k (\neg d_k \vee \bigvee_{i=1}^{n_k} e_{k,i}) \text{ is } \mathcal{A}\text{-satisfiable}\}. \end{aligned}$$

Note that in the definition of T_P^+ , for each literal in the body of a program clause

defining p , exactly one constrained atom is taken in I . In the definition of T_P^- , if p is not defined in P then we have $c = \text{true}$, otherwise for each clause defining p , a finite number of constrained atoms are taken in I to invalidate the body of the clause. Note that for each positive literal in the body at most one constrained atom is taken in I^- , whereas for each negative literal a finite number of constrained atoms can be taken in I^+ . This is crucial for the completeness w.r.t. the logical semantics. For instance, with the program P :

$$\begin{aligned} p &\leftarrow \neg q(X). \\ q(X) &\leftarrow X \geq 0. \\ q(X) &\leftarrow X < 0. \end{aligned}$$

we have $T_P^+(\emptyset) = \{X \geq 0 | q(X), X < 0 | q(X)\}$ and $T_P^-(T_P(\emptyset)) = \{\text{true} | p\}$. If in the definition of T_P^- only one constrained atom was taken in I^+ for a negative literal in the clause, then p would not be false in the iteration of T_P . Allowing to take similarly a finite number of constrained atoms in I^- for a same positive literal, instead of at most one, would not change the definition of T^- as we shall see that the finite powers of T_P are closed by disjunction on false atoms (proposition 5.5).

Example 4.3. Let us consider an example over the Herbrand domain formed with a constant 0 and an unary function symbols s . Clark's equational theory CET augmented with the domain closure axiom DCA is a complete theory for that structure [21], in particular we have $CET + DCA \models (\forall y \ x \neq s(y)) \leftrightarrow x = 0$. The following program is a classical example that shows that Fitting's operator Φ_P^A is not continuous:

$$\begin{aligned} p(x) &\leftarrow x = s(y) | p(y). \\ q &\leftarrow p(x). \end{aligned}$$

No atom is true in the powers of Φ_P^A and T_P . At ordinal ω , all ground instances of $p(x)$ are false both in $\Phi_P^A \uparrow \omega$ and $[T_P \uparrow \omega]$, whereas the atom q becomes false in $\Phi_P^A \uparrow \omega + 1$ and stays undefined in $[T_P \uparrow \omega + 1]$:

α	$(\Phi_P^A \uparrow \alpha)^-$	$(T_P \uparrow \alpha)^-$
0	\emptyset	\emptyset
1	$\{p(0)\}$	$\{x = 0 p(x)\}$
2	$\{p(0), p(s(0))\}$	$\{x = 0 p(x), x = 0 \vee x = s(0) p(x)\}$
...
ω	$\{p(s^i(0) \mid i \geq 0)\}$	$\{x = 0 \vee \dots \vee x = s^i(0) p(X) \mid i \geq 0\}$
$\omega + 1$	$\{q\} \cup \{p(s^i(0) \mid i \geq 0)\}$	$\{x = 0 \vee \dots \vee x = s^i(0) p(X) \mid i \geq 0\}$

The definition of $\Phi_P^A(I)$ based on valuations allows to infer that q is false in $\Phi_P^A \uparrow \omega + 1$ whilst the definition of T_P based on finite subsets of I does not.

Proposition 4.4. T_P is an operator over partial interpretations.

PROOF. We just have to prove that if I is a partial interpretation, then $[T_P^+(I)] \cap [T_P^-(I)] = \emptyset$.

Let $c | p(X) \in T_P^+(I)$ and θ be any valuation of the variables in X such that $c\theta$ is true. There exists a clause in P , $p(X) \leftarrow d | A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$, such that for all $1 \leq i \leq m$ there exists $c_i | A_i \in I^+$, for all $m+1 \leq j \leq n$ there exists $c_j | A_j \in I^-$, such that $c = \exists Y (d \wedge \bigwedge_{i=1}^n c_i)$. As $c\theta$ is true in \mathcal{A} let ρ be a valuation extending θ to the variables in Y such that $(d \wedge \bigwedge_{i=1}^n c_i)\rho$ is true.

Let us suppose that there exists $e | p(X) \in T_P^-(I)$ such that $e\theta$ is true. Then for

the previous clause defining p there exists $p \leq m$ and $\{e_1|A_1, \dots, e_p|A_p\} \subseteq I^-$, there exists $\{e_{p+1}|B_{p+1}, \dots, e_q|B_q\} \subseteq I^+$, where for all $p+1 \leq j \leq q$ $B_j \in \{A_{m+1}, \dots, A_n\}$, such that $e = \forall Y(\neg d \vee \bigvee_{j=1}^q e_j)$ is satisfied by θ . Hence $(\neg d \vee \bigvee_{j=1}^q e_j)\rho$ is true.

Now as $d\rho$ is true, $e_j\rho$ must be true for some $j \in [1, q]$, with $e_j|A_i \in I^-$ for some $i \in [1, m]$ (or $e_j|A_i \in I^+$ for some $i \in [m+1, n]$). Hence we have $c_i|A_i \in I^+$ (or $c_i|A_i \in I^-$) with $c_i\rho$ true, so we get a contradiction: $c_i \wedge e_j$ is satisfied by ρ and we have $c_i|A_i \in I^+$ and $e_j|A_i \in I^-$ (or $c_i|A_i \in I^-$ and $e_j|A_i \in I^+$), i.e. I is not a partial interpretation. \square

Proposition 4.5. T_P is monotonic in the semi-lattice $(\mathcal{I}, \subseteq_3)$.

PROOF. If $I \subseteq_3 J$ then $I^+ \subseteq J^+$ and $I^- \subseteq J^-$, so it is straightforward to verify that by definition of T_P we have both $T_P^+(I) \subseteq T_P^+(J)$ and $T_P^-(I) \subseteq T_P^-(J)$, thus $T_P(I) \subseteq_3 T_P(J)$. \square

Proposition 4.6. T_P is continuous in the semi-lattice $(\mathcal{I}, \subseteq_3)$.

PROOF. The result follows from the fact that an operator f over a powerset, monotonic w.r.t. set inclusion, is continuous if it is finitary, i.e. $\forall x, y \ x \in f(y) \Rightarrow \exists y' \subseteq y$ finite s.t. $x \in f(y')$. From its definition T_P is clearly finitary. \square

As T_P is continuous we can take the least fixpoint of T_P as the fixpoint semantics of the program. We then show a strong equivalence theorem with the operational semantics which shows that the fixpoint semantics fully characterizes the operational behavior of normal CLP programs w.r.t. answer constraints.

Definition 4.7. (Fixpoint semantics) $\mathcal{F}(P) = \text{lfp}(T_P) = T_P \uparrow \omega$.

Main theorem 4.8. (Full abstraction for answer constraints computed by uniform derivations) $\mathcal{O}(P) = \mathcal{F}(P)$.

PROOF.

\subseteq_3 : We show more generally that if $c|\alpha \triangleleft F$ where $\alpha \neq \square$ then, let $V = V(c|\alpha)$,

- 1) if $d|\square \in F$ then for each occurrence of an atom A_i in α , $1 \leq i \leq m$, there exists $c_i|A_i \in \mathcal{F}(P)^+$, for each occurrence of a negative literal $\neg A_j$ in α , $m+1 \leq j \leq n$, there exists $c_j|A_j \in \mathcal{F}(P)^-$, such that $\exists Y d = c \wedge \bigwedge_{i=1}^n c_i$ where $Y = V(d) \setminus V$.
- 2) if $\neg_V F$ is \mathcal{A} -satisfiable then there exist occurrences of atoms in α , let A_1, \dots, A_m , $m \geq 0$, such that for all $1 \leq i \leq m$ there exists $c_i|A_i \in \mathcal{F}^-$, and there exists $n \geq m$ such that for all $m+1 \leq j \leq n$ there exists $c_j|A_j \in \mathcal{F}^+$ where $\neg A_j$ is a negative literal in α , such that $\neg_V F = \neg c \vee \bigvee_{i=1}^n c_i$.

Therefore taking $c|\alpha = \text{true}|p(x)$ in 1) we get $\mathcal{O}(P)^+ \subseteq \mathcal{F}(P)^+$, and taking $c|\alpha = \text{true}|\neg p(x)$ in 1) we get $\mathcal{O}(P)^- \subseteq \mathcal{F}(P)^-$. The proof is by structural induction on a proof tree for $c|\alpha \triangleleft F$.

TRIV: 1) $\{d|\square\} \notin F$ as $F = \{c|\alpha : \mathcal{A} \models \exists c\}$ and $\alpha \neq \square$.

2) We have $F = \{c|\alpha : \mathcal{A} \models \exists c\}$ and $\neg_V F = \neg c$.

RES: We have $\alpha = p(X)$ and $F = \bigcup_{i=1}^k F_i$ where $\{p(X) \leftarrow c_i|\alpha_i\}_{1 \leq i \leq k}$ is the set of renamed apart clauses defining $p(X)$ in P with local variables Y_i such that $c \wedge c_i$ is \mathcal{A} -satisfiable, and $c \wedge c_i|\alpha_i \triangleleft F_i$ for all $1 \leq i \leq k$.

1) Let us suppose $d|\square \in F$, then $d|\square \in F_i$ for some i . By the induction

hypothesis applied to $c \wedge c_i | \alpha_i \triangleleft F_i$ we get that for each atom's occurrence A_j (resp. negative literal's occurrence $\neg A_j$) in α_i , $1 \leq j \leq n$, there exists $e_j | A_j \in \mathcal{F}(P)^+$ (resp. $e_j | A_j \in \mathcal{F}(P)^-$) such that $\exists Z d = c \wedge c_i \wedge \bigwedge_{j=1}^n e_j$ where $Z = V(d) \setminus V(c \wedge c_i | \alpha_i)$.

Now let $e = \exists Y_i (c_i \wedge \bigwedge_{j=1}^n e_j)$, then by the definition of T_P^+ we get that $e | p(X) \in \mathcal{F}(P)^+$, thus $\exists Y d = \exists Y_i \exists Z d = c \wedge e$.

- 2) Let us suppose $\neg_V F$ is \mathcal{A} -satisfiable. As $\neg_V F = \bigwedge_{i=1}^k \neg_V F_i$ by eq. 6, $\neg_V F_i$ is \mathcal{A} -satisfiable for all $1 \leq i \leq k$. Let $V_i = V(c \wedge c_i | \alpha_i)$, we have $\neg_V F = \bigwedge_{i=1}^k \forall Y_i \neg_{V_i} F_i$, hence for all $1 \leq i \leq k$, $\neg_{V_i} F_i$ is \mathcal{A} -satisfiable as well. By applying the induction hypothesis to $c \wedge c_i | \alpha_i \triangleleft F_i$, we get that for all $1 \leq i \leq k$,

$$\neg_{V_i} F_i = \neg c \vee \neg c_i \vee \bigvee_{j=1}^{n_i} c_j^i$$

where $n_i \geq 0$ and for all $1 \leq j \leq n_i$, $c_j^i | A_j^i \in \mathcal{F}^-$ (resp. $c_j^i | A_j^i \in \mathcal{F}^+$) where A_j^i is an atom's occurrence in α_i (resp. $\neg A_j^i$ is a literal in α_i).

Now let $d = \bigwedge_{i=1}^k \forall Y_i (\neg c_i \vee \bigvee_{j=1}^{n_i} c_j^i)$, we get from the definition of T_P^- that $d | p(X) \in \mathcal{F}(P)^-$. Thus $\neg_V F = \bigwedge_{i=1}^k \forall Y_i \neg_{V_i} F_i = \neg c \vee d$.

FRT: We have $\alpha = \alpha_1, \alpha_2$, $c | \alpha_1 \triangleleft F_1$, $c | \alpha_2 \triangleleft F_2$ and $F = F_1 \times F_2$. Let $V_1 = V(c | \alpha_1)$ and $V_2 = V(c | \alpha_2)$.

- 1) If $d | \square \in F$ then there exist $d_1 | \square \in F_1$ and $d_2 | \square \in F_2$ such that $d = d_1 \wedge d_2$. Let $Y = V(d) \setminus V(c | \alpha)$, $Y_1 = V(d_1) \setminus V(c | \alpha_1)$ and $Y_2 = V(d_2) \setminus V(c | \alpha_2)$. By the hypothesis on variable clashes in the FRT rule we have $\exists Y d_1 = \exists Y_1 d_1$ and $\exists Y d_2 = \exists Y_2 d_2$, therefore $\exists Y d = \exists Y_1 d_1 \wedge \exists Y_2 d_2$. Now the induction hypothesis 1) applied to $c | \alpha_1 \triangleleft F_1$ and $c | \alpha_2 \triangleleft F_2$ immediately concludes the proof.
- 2) By eq. 7 we have $\neg_V F = \neg_V F_1 \vee \neg_V F_2$. Furthermore by the hypothesis on variable clashes in the FRT rule we have $\neg_{V_1} F_1 = \neg_V F_1$ and $\neg_{V_2} F_2 = \neg_V F_2$, therefore $\neg_V F = \neg_{V_1} F_1 \vee \neg_{V_2} F_2$. Now the induction hypothesis 2) applied to $c | \alpha_1 \triangleleft F_1$ and $c | \alpha_2 \triangleleft F_2$ also immediately concludes the proof.

PRN: We have $\alpha = \neg p(X)$ and $F = c \times \{\neg_V S | \neg p(X), \neg_V F' | \square\}$ with $c | p(X) \triangleleft F'$ and $S \subseteq \mathcal{S}(F')$.

- 1) If $d = c \wedge \neg_V F'$ is \mathcal{A} -satisfiable, then $\neg_V F'$ is \mathcal{A} -satisfiable, hence by the induction hypothesis 2) applied to $c | p(X) \triangleleft F'$ we get that there exists $c_1 | p(X) \in \mathcal{F}(P)^-$ such that $\neg_V F' = \neg c \vee c_1$. Thus $\exists Y d = c \wedge (\neg c \vee c_1) = c \wedge c_1$.
- 2) Let us suppose that $\neg_V F$ is \mathcal{A} -satisfiable. Let $S = \{s_1 | \square, \dots, s_m | \square\}$, proposition 3.12 gives $\neg_V F = \neg c \vee \bigvee_{i=1}^m \exists Y_i s_i$ where $Y_i = V(s_i) \setminus V$. Now by the induction hypothesis 1) applied to $c | p(X) \triangleleft F'$ we get that there exists $\{c_1 | p(X), \dots, c_m | p(X)\} \subseteq \mathcal{F}(P)^+$ such that for all $1 \leq i \leq m$, $\exists Y_i s_i = c_i$. Thus $\neg_V F = \neg c \vee \bigvee_{i=1}^m c_i$ with $c_i | p(X) \in \mathcal{F}(P)^+$.

\supseteq_3 : We prove by induction on n that $\mathcal{O}^+(P) \supseteq T_P \upharpoonright n^+$ and $\mathcal{O}^-(P) \supseteq T_P \upharpoonright n^-$ for all $n \geq 0$. The base case $n = 0$ is trivial. Let us consider the induction

step.

Let $c|p(X) \in (T_P \uparrow n)^+$. There exists a clause with local variables Y

$$p(X) \leftarrow d|A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$$

for all $1 \leq i \leq m$ there exists $c_i|A_i \in (T_P \uparrow n-1)^+$ for all $m+1 \leq j \leq n$ there exists $c_j|A_j \in (T_P \uparrow n-1)^-$ such that $c = \exists Y d \wedge \bigwedge_{i=1}^n c_i$.

By induction we get that c_i is a computed answer to the goal $true|A_i$ for all $1 \leq i \leq m$ and to the goal $true|\neg A_i$ for all $m+1 \leq i \leq n$.

Hence by corollary 3.9, $\bigwedge_{i=1}^n c_i$ is a computed answer to

$$true|A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n.$$

By the instantiation lemma 3.5, we get that c is a computed answer to the goal $d|A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$, hence by the RES rule we get $c|p(X) \in \mathcal{O}^+(P)$.

Let $c|p(X) \in (T_P \uparrow n)^-$. For any clause defining p in P , with local variables Y_k ,

$$p(X) \leftarrow d_k|A_{k,1}, \dots, A_{k,m_k}, \alpha_k$$

there exist $e_{k,1}|A_{k,1}, \dots, e_{k,m_k}|A_{k,m_k} \in (T_P \uparrow n-1)^-$, $e_{k,m_k+1}|A_{k,m_k+1}, \dots, e_{k,n_k}|A_{k,n_k} \in (T_P \uparrow n-1)^+$, where for all $m_k+1 \leq j \leq n_k$, $\neg A_{k,j}$ is a negative literal in α_k , such that $c_k = \forall Y_k (\neg d_k \vee \bigvee_{i=1}^{n_k} e_{k,i})$ is \mathcal{A} -satisfiable, and $c = \bigwedge_k c_k$ is \mathcal{A} -satisfiable.

By induction we have that for all $1 \leq i \leq m_k$, $e_{k,i}$ is a c.a.c. for the goal $true|\neg A_{k,i}$. As the PRN rule is necessarily applied at the root, we have

$$true|A_{k,i} \triangleleft F_{k,i}$$

with $e_{k,i} = \neg_{V(A_{k,i})} F_{k,i}$.

Similarly by induction we have that for all $m_k+1 \leq j \leq n_k$, $e_{k,j}$ is c.a.c. for the goal $true|A_{k,j}$. Hence by the PRN rule, taking the singleton $\{e_{k,j}|\Box\}$ as success set, we have

$$true|\neg A_{k,j} \triangleleft F_{k,j}$$

with $\neg e_{k,j}|\neg A_{k,j} \in F_{k,j}$. By proposition 3.12 we get $\neg_{V(A_{k,j})} F_{k,j} = e_{k,j}$.

Let $\beta_k = \alpha_k \setminus \bigcup_{j=m_k+1}^{n_k} \neg A_{k,j}$, by lemma 3.7 and rule TRIV we have

$$d_k|A_{k,1}, \dots, A_{k,m_k}, \neg A_{k,m_k+1}, \dots, \neg A_{k,n_k}, \beta_k \triangleleft F_k$$

where $F_k = d_k \times \times_{i=1}^{n_k} F_{k,i} \times \{d_k|\beta_k\}$. Note that by equation 7 we have $\neg_X F_k = \forall Y_k (\neg d_k \vee \bigvee_{i=1}^{n_k} e_{k,i}) = c_k$.

Now by applying the RES rule we get $true|p(X) \triangleleft F$ where $F = \bigcup_{k \in K} F_k$. By equation 6 we have $\neg_X F = \bigwedge_k c_k = c$, hence by the PRN rule we get $c|p(X) \in \mathcal{O}^-(P)$.

□

Full abstraction does not hold for non uniform derivations, however the full abstraction theorem 4.8, together with corollaries 3.16 and 3.20 show that non-uniform derivations are nevertheless sound and complete w.r.t. the fixpoint semantics.

Theorem 4.9. (Soundness and completeness of non-uniform derivations) $\tilde{\mathcal{O}}(P) \subseteq \mathcal{F}(P)$ and $\mathcal{F}(P) \subseteq \tilde{\mathcal{O}}(P)$.

5. THREE-VALUED LOGICAL SEMANTICS

The main theorem of [17], extended to CLP programs in [26], characterizes the three-valued logical consequences of the Clark's completion with the finite powers of Fitting's operator Φ_P^A :

Theorem 5.1 [17][26]. Let P be a normal CLP(\mathcal{A}) program and ψ be a Π, Σ, V -formula, then $P^, th(\mathcal{A}) \models_3 \psi$ iff ψ is true in $\Phi_P^A \uparrow n$ for some integer n .*

In this section we show that the finite powers of T_P coincide with those of Fitting's operator Φ_P^A as in [26], and thus, by the previous theorem, that the fixpoint and operational semantics are correct and complete w.r.t. the three-valued logical consequences of the program's completion.

Proposition 5.2. If I is a finite partial interpretation then $T_P(I)$ is finite.

PROOF. Obvious from the definition of T_P . \square

Corollary 5.3. For all $n \geq 0$, $T_P \uparrow n$ is finite.

Definition 5.4. A constrained interpretation I is closed by disjunction if whenever $c|p(X) \in I$, $c'|p(X) \in I$ then there exists $d|p(X) \in I$ such that $\mathcal{A} \models (c \vee c') \rightarrow d$.

Proposition 5.5. Let I be a partial constrained interpretation. If I^- is closed by disjunction then so is $T_P^-(I)$.

PROOF. Let $c|p(X), c'|p(X) \in T_P^-(I)$. For any clause defining p in P , with local variable Y_k , $p(X) \leftarrow A_{k,1}, \dots, A_{k,m_k}, \alpha_k$, where $m_k \geq 0$, there exist $\{e_{k,r}|A_{k,r}\}_{r \in R} \subseteq I^-$, and $\{e'_{k,r'}|A_{k,r'}\}_{r' \in R'} \subseteq I^-$, where R and R' are subsets of $\{1, \dots, m_k\}$, there exist finite sets $\{e_{k,s}|A_{k,s}\}_{s \in S} \subseteq I^+$ and $\{e'_{k,s'}|A_{k,s'}\}_{s' \in S'} \subseteq I^+$ where $R \cap S = \emptyset$, $R' \cap S' = \emptyset$ and for all $j \in S \cup S'$, $\neg A_{k,j}$ is a negative literal in α_k , such that $c_k = \forall Y_k (\neg d_k \vee \bigvee_{i \in R \cup S} e_{k,i})$, $c'_k = \forall Y_k (\neg d_k \vee \bigvee_{j \in R' \cup S'} e'_{k,j})$, $c = \bigwedge_k c_k$, and $c' = \bigwedge_k c'_k$ are \mathcal{A} -satisfiable ($c = c' = \text{true}$ if p is not defined in P).

Now for any clause defining p in P , let

$$\{f_{k,l}\}_{l \in L} = \{e_{k,s}\}_{s \in S} \cup \{e'_{k,s'}\}_{s' \in S'} \cup \{e_{k,r}\}_{r \in R \setminus R'} \cup \{e'_{k,r'}\}_{r' \in R' \setminus R} \cup \{g_{k,r}\}_{r \in R \cap R'}$$

where, as I is a partial interpretation closed by disjunction on false atoms, we can define $g_{k,r}$ for all $r \in R \cap R'$ by choosing $g_{k,r}|A_{k,r} \in I^-$ such that $\mathcal{A} \models e_{k,r} \vee e'_{k,r} \rightarrow g_{k,r}$.

Let $f_k = \forall Y_k (\neg d_k \vee \bigvee_{l \in L} f_{k,l})$, we have $\mathcal{A} \models (c_k \vee c'_k) \rightarrow f_k$, hence f_k is \mathcal{A} -satisfiable. Let $f = \bigwedge_k f_k$, we have $\mathcal{A} \models c \vee c' \rightarrow f$, hence f is \mathcal{A} -satisfiable. Therefore by definition of T_P^- we conclude that $f|p(X) \in T_P^-(I)$ with $\mathcal{A} \models c \vee c' \rightarrow f$.

\square

Corollary 5.6. For all $n \geq 0$, $T_P \uparrow n$ is closed by disjunction on false atoms.

Corollary 5.7. $\mathcal{F}(P)$ (and $\mathcal{O}(P)$) are closed by disjunction on false constrained atoms.

Lemma 5.8. $[T_P(I)]_{\mathcal{A}} = \Phi_P^{\mathcal{A}}([I]_{\mathcal{A}})$ for all finite partial interpretation I closed by disjunction on false atoms.

PROOF. We consider both inclusions on positive and negative parts separately.

\subseteq^+ : Let $c|p(X) \in T_P^+(I)$, and θ be any \mathcal{A} -valuation of X such that $c\theta$ is true. From the definition of T_P^+ there exists a clause in P with local variables Y ,

$$p(X) \leftarrow d|A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$$

such that for all $1 \leq i \leq m$ there exists $c_i|A_i \in I^+$, for all $m+1 \leq j \leq n$ there exists $d_j|A_j \in I^-$ such that $c = \exists Y(d \wedge \bigwedge_i c_i \wedge \bigwedge_j d_j)$ is \mathcal{A} -satisfiable. Therefore there exists an \mathcal{A} -valuation ρ which extends θ to an \mathcal{A} -valuation of the variables in Y such that $(d \wedge \bigwedge_i c_i \wedge \bigwedge_j d_j)\rho$ is true. Hence $d\rho$ is true, $A_i\rho \in [I^+]_{\mathcal{A}}$ for all i , $1 \leq i \leq m$, and $A_j\rho \in [I^-]_{\mathcal{A}}$ for all j , $m+1 \leq j \leq n$. Hence by definition of $(\Phi_P^{\mathcal{A}})^+$ we have $p(X)\theta \in \Phi_P^{\mathcal{A}^+}([I]_{\mathcal{A}})$ for all θ such that $c\theta$ is true.

\subseteq^- : Let $c|p(X) \in T_P^-(I)$, and θ be any \mathcal{A} -valuation of X such that $c\theta$ is true.

For any clause in P defining p , with local variable Y_k ,

$$p(X) \leftarrow d_k|A_{k,1}, \dots, A_{k,m}, \alpha_k,$$

there exists $\{e_{k,1}|A_{k,1}, \dots, e_{k,m}|A_{k,m}\} \subseteq I^-$,

there exists $\{e_{k,m+1}|A_{k,m+1}, \dots, e_{k,n}|A_{k,n}\} \subseteq I^+$, where for all $m+1 \leq j \leq n$, $\neg A_{k,j}$ is a negative literal in α_k ,

such that $c_k = \forall Y_k(\neg d_k \vee \bigvee_{i=1}^n e_{k,i})$ is \mathcal{A} -satisfiable, and $c = \bigwedge_k c_k$ is \mathcal{A} -satisfiable.

Therefore for any clause defining p in P , and any \mathcal{A} -valuation ρ_k extending θ to an \mathcal{A} -valuation for the variables in Y_k , we have

either $d_k\rho_k$ false,

or $e_{k,i}\rho_k$ true for some $1 \leq i \leq m$, in which case $A_i\rho_k \in [I^-]_{\mathcal{A}}$,

or $e_{k,j}\rho_k$ true for some $m+1 \leq j \leq n$, in which case $A_j\rho_k \in [I^+]_{\mathcal{A}}$.

Hence by definition of $(\Phi_P^{\mathcal{A}})^-$ we have $p(X)\theta \in \Phi_P^{\mathcal{A}^-}([I]_{\mathcal{A}})$ for all θ such that $c\theta$ is true.

\supseteq^+ : Let $p(X)\theta \in \Phi_P^{\mathcal{A}^+}([I]_{\mathcal{A}})$. There exists a clause in P with local variables Y , $p(X) \leftarrow d|A_1, \dots, A_m, \neg A_{m+1}, \dots, \neg A_n$ such that $d\theta$ is true and for all $1 \leq i \leq m$, $m+1 \leq j \leq n$, $A_i\theta \in [I^+]_{\mathcal{A}}$ and $A_j\theta \in [I^-]_{\mathcal{A}}$.

Hence for all $1 \leq i \leq m$, $m+1 \leq j \leq n$, there exist $c_i|A_i \in I^+$, $d_j|A_j \in I^-$, such that $c_i\theta$ and $d_j\theta$ are true. Hence $c = \exists Y(d \wedge \bigwedge_i c_i \wedge \bigwedge_j d_j)$ is \mathcal{A} -satisfiable (by θ), and from the definition of T_P^+ we get $c|p(X) \in T_P^+(I)$.

\supseteq^- : Let $p(X)\theta \in \Phi_P^{\mathcal{A}^-}([I]_{\mathcal{A}})$. From the definition of $\Phi_P^{\mathcal{A}}$, for any clause in P defining p , with local variable Y_k , $p(X) \leftarrow d_k|\alpha_k$, and for any \mathcal{A} -valuation θ_k extending θ to the variables in Y_k , we have

either $d_k\theta_k$ false,

or $A\theta_k \in [I^+]_{\mathcal{A}}$ for some negative literal $\neg A$ in α_k , in which case there exists $c|A \in I^+$ with $c\theta_k$ true,

or $A'\theta_k \in [I^-]_{\mathcal{A}}$ for some positive literal A' in α_k , in which case there exists $c'|A' \in I^-$ with $c'\theta_k$ true.

Let us consider the classes of all constrained atoms taken in I^- and I^+ , for all \mathcal{A} -valuations extending θ to the variables in Y_k . By hypothesis I is finite, thus these classes are finite sets, say $\{c_{k,1}|A_{k,1}, \dots, c_{k,n}|A_{k,n_k}\} \subseteq I^+$ where

$\neg A_{k,1}, \dots, \neg A_{k,n_k}$ are negative literals in α_k , and $\{c'_{k,1}|A'_{k,1}, \dots, c'_{k,n'}|A'_{k,n'}\} \subseteq I^-$ where $A'_{k,1}, \dots, A'_{k,n'}$ are positive literals in α_k .

Let $\{A_{k,n_k+1}, \dots, A_{k,m_k}\} = \{A'_{k,i} \mid 1 \leq i \leq n'\}$. As I is closed by disjunction on false atoms, there exists $\{c_{k,n_k+1}|A_{k,n_k+1}, \dots, c_{m_k}|A_{k,m_k}\} \subseteq I^-$ such that for all $c'_{k,i}|A'_{k,i}$, $1 \leq i \leq n'$ there exists a j , $n_k + 1 \leq j \leq m_k$ such that $A'_{k,i} = A_{k,j}$ and $\mathcal{A} \models c'_{k,i} \rightarrow c_{k,j}$.

Now let $c_k = \forall Y_k (\neg d_k \vee \bigvee_{i=1}^{n_k} c_{k,i})$. For all k , $c_k \theta$ is true, hence $c = \bigwedge c_k$ is \mathcal{A} -satisfiable and from the definition of T_P^- we get $c|p(X) \in T_P^-(I)$.

□

Theorem 5.9. For all $n \geq 0$, $[T_P \uparrow n]_{\mathcal{A}} = \Phi_P^{\mathcal{A}} \uparrow n$.

PROOF. By induction on n . The base case $n = 0$ is trivial. For the induction step, we have $[T_P \uparrow n]_{\mathcal{A}} = [T_P(T_P \uparrow n - 1)]_{\mathcal{A}}$. By corollary 5.3 and 5.6, $T_P \uparrow n - 1$ is finite and closed by disjunction on false atoms, hence we get by lemma 5.8, $[T_P \uparrow n]_{\mathcal{A}} = \Phi_P^{\mathcal{A}}([T_P \uparrow n - 1]_{\mathcal{A}})$. Therefore by the induction hypothesis we conclude $[T_P \uparrow n]_{\mathcal{A}} = \Phi_P^{\mathcal{A}}(\Phi_P^{\mathcal{A}} \uparrow n - 1) = \Phi_P^{\mathcal{A}} \uparrow n$. □

Corollary 5.10. For all $n \geq 0$, $\Phi_P^{\mathcal{A}} \uparrow n$ has a finite cover.

Theorem 5.11. (Correctness and completeness of the fixpoint semantics w.r.t. the logical semantics) $\mathcal{F}(P) \subseteq \mathcal{L}(P)$, $\mathcal{L}^+(P) \sqsubseteq_f \mathcal{F}^+(P)$ and $\mathcal{L}^-(P) \sqsubseteq \mathcal{F}^-(P)$.

PROOF. Let $c|A \in \mathcal{F}^+(P)$, then $c|A \in T_P \uparrow n^+$ for some integer n , by theorem 5.9 for all \mathcal{A} -valuation ρ s.t. $\mathcal{A} \models c\rho$ we have $A\rho \in \Phi_P^{\mathcal{A}} \uparrow n^+$, so $c \rightarrow A$ is true in $\Phi_P^{\mathcal{A}} \uparrow n$, hence by theorem 5.1 we get $P^*, th(\mathcal{A}) \models_3 c \rightarrow A$, thus $c|A \in \mathcal{L}^+(P)$. The proof that $\mathcal{F}^-(P) \subseteq \mathcal{L}^-(P)$ is similar.

Conversely, let $c|p(X) \in \mathcal{L}^+(P)$, by theorem 5.1, $\forall X(c \rightarrow p(X))$ is true in $\Phi_P^{\mathcal{A}} \uparrow n$ for some n , thus by theorem 5.9, it is true in $T_P \uparrow n$ for some n . Now as $T_P \uparrow n$ is finite (corollary 5.3), there exists $\{d_1|p(X), \dots, d_k|p(X)\} \subseteq T_P \uparrow n$ such that $\mathcal{A} \models c \rightarrow \bigvee_{i=1}^k d_i$, so $\mathcal{L}^+(P) \sqsubseteq_f \mathcal{F}^+(P)$. We prove similarly that $\mathcal{L}^-(P) \sqsubseteq_f \mathcal{F}^-(P)$, yet by corollary 5.7 we get $\mathcal{L}^-(P) \sqsubseteq \mathcal{F}^-(P)$. □

Corollary 5.12. $\mathcal{F}(P)$ provides a fixpoint characterization of Kunen's logical semantics.

PROOF. Given a (partial) constrained interpretation I let us denote by \bar{I} its closure by finite disjunction ($c \vee d|A \in \bar{I}$ whenever $c|A \in \bar{I}$ and $d|A \in \bar{I}$) and by entailment ($d|A \in \bar{I}$ whenever $c|A \in \bar{I}$ and $\mathcal{A} \models d \rightarrow c$). We have $\mathcal{L}(P) = \overline{\mathcal{F}(P)}$.

Note alternatively that the least fixed point of the operator $T_P'(I) = \overline{T_P(I)}$ is equal to $\mathcal{L}(P)$. □

Theorem 5.13. (Correctness and completeness of the operational semantics w.r.t. the logical semantics) $\mathcal{O}(P) \subseteq \mathcal{L}(P)$, $\mathcal{L}^+(P) \sqsubseteq_f \mathcal{O}^+(P)$ and $\mathcal{L}^-(P) \sqsubseteq \mathcal{O}^-(P)$.

Similarly $\tilde{\mathcal{O}}(P) \subseteq \mathcal{L}(P)$, $\mathcal{L}^+(P) \sqsubseteq_f \tilde{\mathcal{O}}^+(P)$ and $\mathcal{L}^-(P) \sqsubseteq \tilde{\mathcal{O}}^-(P)$.

PROOF. By theorems 5.11 and 4.8 (resp. 4.9). □

6. COMPARISON WITH OTHER WORKS

The constructive negation scheme of Chan [7] for logic programs, and Stuckey [26] for constraint logic programs, relies on a transition relation over explicitly quantified

complex goals. The transition relation is defined by the usual SLD resolution rule for positive subgoals and by the following constructive negation rule CN for complex subgoals:

$$CN : (c|\alpha, (\neg\exists Y \beta), \alpha') \rightarrow (c \wedge c_j|\alpha, \beta'_j, \alpha')$$

for each $j \in J$ where $\bigvee_{j \in J} c_j \wedge \beta'_j$ is a disjunctive normal form of $\bigwedge_{k \in K} \neg\exists Z_k (c \wedge d_k \wedge \beta_k)$ and where $\{c \wedge d_k|\alpha_k\}_{k \in K}$ is a frontier in a SLDCN derivation tree for $c|\beta$.

Not only the constraints but also the goals in the frontier of an auxiliary derivation tree are thus transformed into disjunctive normal form and reinjected in the resolvent at each resolution step with a negative subgoal. This makes the scheme hardly amenable to a practical implementation for normal CLP programs in all generality.

The compilative version proposed by Bruscoli et al. [5], named intensional negation, performs all disjunctive normal form transformations once and for all at compile time, but still all quantifiers need be explicit at run time and derivation rules need be defined for complex goals. The practical advantage of constructive negation by pruning is that it relies on standard SLD derivation trees for definite goals only. The only extra machinery to handle negation is a concurrent pruning mechanism over standard SLD derivation trees. It is remarkable that the exploitation of concurrency in the development of SLD derivation trees is sufficient to build a complete scheme for negation. This is the case also for the fail answers approach proposed recently by Drabent in [8] for normal logic programs, building on earlier work by Maluszyński and Näslund [22]. Drabent's execution model is essentially equivalent to constructive negation by pruning in that case, the success by pruning rule is a special case of the fail answer approach, we believe that both schemes define in fact the same set of computed answer substitutions for normal logic programs.

If we look at the nesting of negation, we can see that the effect of doubly negating a goal is to collect in a single answer constraint all the successes found for the positive goal. Corollary 5.7 shows that the computed answer constraints for negative goals are closed by disjunction, thus a simple way to obtain a strong completeness result w.r.t. the logical semantics (i.e. $\mathcal{L}(P) \sqsubseteq \mathcal{O}(P)$ instead of $\mathcal{L}(P) \sqsubseteq_f \mathcal{O}(P)$) is to put double negations on positive goals. On the other hand, in the intensional negation scheme double negations are eliminated by simplification. In this respect our scheme is nearer to the one of Chan and Stuckey.

The closure by disjunction property for negative literals can be seen also as a drawback as at some point in the execution all the current information on a negative literal need to be handled by the constraint solver. A general solution to this problem is to exploit the trade-off there is between the constraint solver and the non-deterministic derivation system. This is possible if the structure \mathcal{A} is admissible [26] (cf. section 2.1), in that case the language of constraints need not even be closed by negation. Constructive negation by pruning can be adapted mainly by changing the definition of $\neg_V F$. The negation of the projection of the constraints in a frontier $F = \{c_1|\alpha_1, \dots, c_n|\alpha_n\}$ over a set of variables V is then no longer a constraint but a frontier defined as:

$$\neg_V F = \{d_{1,1}|\square, \dots, d_{1,l_1}|\square\} \times \dots \times \{d_{n,1}|\square, \dots, d_{n,l_n}|\square\}$$

where $\mathcal{A} \models \forall V (\neg\exists Y_i c_i(V, Y_i) \leftrightarrow \exists Z_i (d_{i,1} \vee \dots \vee d_{i,l_i}))$ for all $1 \leq i \leq n$. This change amounts to replace in the procedural interpretation the pruning by success rule by a check of satisfiability with at least one of the disjunct, and the success by

pruning rule by the creation of a success for each satisfiable disjunct.

Another possible drawback of constructive negation by pruning is that once a derivation tree is developed for a negative literal it receives no more information from the resolution of the positive part of the goal. This is the price to pay for having a single derivation tree for a negative literal instead of duplicating resolution steps at all its occurrences. Many optimizations can nevertheless be imagined, such as sending back for pruning in the auxiliary tree the constraints which are entailed by the frontiers in the main tree.

On the theoretical side constructive negation was proved correct and complete w.r.t. Kunen's logical semantics by Stuckey for consistent fair computation rules [26]. Similar results were obtained for intensional negation [5] and by Drabent [8]. None of these schemes however were provided with a fixpoint semantics. The full abstraction theorem given for constructive negation by pruning allows to analyze and transform normal $CLP(\mathcal{A})$ programs by reasoning at the fixpoint semantics level of abstraction while preserving the program equivalence based on the observation of computed answer constraints. Note that a similar result is conjectured in [3]. Note also that the full abstraction result has been obtained without fixing a resolution strategy, it holds w.r.t. the computation of all frontiers in uniform derivation trees. This left open the problem of giving a fully abstract fixpoint semantics under specific strategies, such as breadth-first [26].

7. VARIATIONS ON A SCHEME FOR OPTIMIZATION PREDICATES

Because of their importance in CLP programming practice, most constraint logic programming systems, such as CHIP, CLP(R), Prolog III or Ilog Solver, include various constructs for searching not all solutions but only best solutions, i.e. optimal w.r.t. an objective function. The basic optimization procedure currently used in CLP systems is a variant of the branch and bound procedure, where constraints are used to prune the search space [28] [23]. That procedure can be used to find optimal solutions of the top-level query w.r.t. an objective function, but it becomes unsound when applied to subgoals of the program. The extension of CLP languages with optimization predicates is an important issue to solve multi-criteria optimization problems and modelize multi-component systems for which several optimization goals have to be combined in the query and/or the program. The problem is to reconcile the evaluation procedures for optimization goals with the declarative semantics of CLP, and its properties of compositionality.

In this section we relate several forms of optimization within CLP languages to special classes of CLP programs with negation, and we derive from the general scheme of constructive negation by pruning different execution models which are complete w.r.t. the Kunen's three-valued logical semantics of the program's completion.

First we define the constraint minimization problem that the constraint solver is assumed to solve, and the basic branch and bound procedure used for query optimization w.r.t. an objective function.

7.1. Constraint Minimization

We consider the case where the constraint solvers can treat minimization problems. A *constraint minimization problem* has the form

minimize t
 subject to c

where t is a Σ, V -term and c is a constraint. Our most general assumption is that (\mathcal{A}, \leq) is a *partial order* and that the constraint minimization problem in \mathcal{A} is decidable and is in fact expressible in the constraint language. More precisely we no longer assume that the language of constraints is closed by negation but by the following more restricted *optimization expressions*:

for any constraint $c(Y)$ and any term $f(Y)$, the formula $\neg(\exists Y c(Y) \wedge f(Y) < f(X))$ where $X \cap Y = \emptyset$ is also a constraint of the language.

In this way the constraint minimization problem can be expressed by the constraint $c(X) \wedge \neg(\exists Y c(Y) \wedge f(Y) < f(X))$. Note that an optimization expression is satisfiable if the set of values of $f(Y)$ under constraint $c(Y)$ admits a lower bound, not necessarily a minimum. If that set admits a greatest lower bound, noted $v = glb_{c(Y)} f(Y) = glb\{f(Y)\rho \mid \mathcal{A} \models c\rho\}$, then one can easily check that the optimization expression is \mathcal{A} -equivalent to the simple constraint $v \not< f(X)$, i.e. $f(X) \leq v$ if \mathcal{A} is a total order.

The execution models we describe in the following sections for goal optimization in CLP languages can be easily adapted to the general setting of partial orders, however for sake of clarity we shall present these execution models in the simpler setting of total orders where constraints define closed sets so that *constraint minimization problems are either unbounded or admit a minimum* (not only a greatest lower bound). Under these assumptions we shall note

$$min_c t = \min\{t\rho \mid \mathcal{A} \models c\rho\}$$

the minimum value of term t under constraint c , if it exists, if it doesn't exist then the term t is unbounded in c .

A typical example of these assumptions is provided by $CLP(\mathcal{R})$ systems where the Simplex algorithm used for linear constraint satisfaction can be used as well to solve the linear constraint minimization problem [16]. For constraints over finite domains, constraint minimization in all generality involves enumeration. Note however that the domains of the variables in Y can be used to bound the minimum value of $f(Y)$. Furthermore for some fragments of the constraint language for which arc-consistency algorithms are complete, e.g. conjunctions of linear inequalities over two variables, and for some objective functions, e.g. monotone functions, the constraint minimization problem can be solved without enumeration [29] [14].

7.2. Query Optimization

The optimization of a top-level query $G(X)$ w.r.t. an objective function $f(X)$ can be achieved with a simple pruning mechanism on the derivation tree of the query [23] [28]. The main procedure is a variant of the branch-and-bound procedure, pictured in figure 4. Once a successful derivation is found for the query $G(X)$, say with answer constraint c , then the optimal value v of the objective function f is computed for that derivation, $v = min_c f(X)$, and the tree is pruned by adding the constraint $f(X) < v$. If v doesn't exist it is a failure, otherwise whenever the tree gets finite after pruning, the last computed cost v gives the optimal cost, then the (possibly infinite) set of successful derivations leading to optimal solutions can be




FIGURE 4. Query optimization by pruning.

enumerated by SLD resolution with the goal $f(X) = v|G(X)$.

There are some problems however to use the branch and bound procedure recursively for optimization goals, noted $\text{minimize}(G(X), f(X))$ where $G(X)$ is a goal and $f(X)$ is a term to be minimized (i.e. the objective function). The difficulty can be illustrated by the following non-logical behavior well-known in current CLP systems:

```
p(X):- X>=0.
q(X):- X>=1.
? q(X), minimize(p(X),X).
X=1
? minimize(p(X),X), q(X).
fail
```

The problem comes from the interaction between the constraints imposed on X by the optimization subgoal and those imposed on X by the other subgoals. Whether the later should affect or not the optimization process is a choice of semantics for the optimization predicate, in the example the operational behavior corresponds to a different choice made according to the order of evaluation. One can thus define two kinds of optimization higher-order predicates: *global optimization predicates* (cf. def. 7.1), which define the optimal solutions to the goal passed as argument, and *local optimization predicates* (cf. def. 7.6), which define the optimal solutions to the goal passed as argument relatively to a set of protected variables. When the protected variables are further constrained by the context the optimal solutions w.r.t. this space are retained, whereas in the global optimization approach the globally optimal solutions are simply filtered by the added constraints. Under the global optimization semantics, the correct answer in the example is fail, and $X = 1$ is a correct answer to the goal $\text{minimize}((p(X), q(X)), X)$.

In [9] and [24], it is shown that both kinds of optimization higher-order predicates can be provided with a faithful logical semantics based on constructive negation. In the next section we show how constructive negation by pruning specializes to

a *complete* concurrent branch and bound like procedure for global optimization predicates.

7.3. Global Optimization Predicates

Definition 7.1. Let (\mathcal{A}, \leq) be a total order. The (global) minimization higher-order predicate

$$\text{minimize}(G(X), f(X))$$

where $G(X)$ is a goal and $f(X)$ is a term, is defined as an abbreviation for the formula:

$$G(X) \wedge \neg \exists Y (f(Y) < f(X) \wedge G(Y))$$

A $\mu\text{CLP}(\mathcal{A})$ program is a definite $\text{CLP}(\mathcal{A})$ program which can contain minimization predicates in clause bodies.

μCLP programs allow for the arbitrary composition of optimization subgoals in a program and for the recursive definition of predicates through their optimal solutions. μCLP programs can be transformed into normal CLP programs with the standard transformation rules [19], i.e. by writing $\text{minimize}(G(X), f(X))$ as:

$$G(X), \neg gf(X)$$

where gf is a new predicate symbol, and by adding the following clause to the program:

$$gf(X) \leftarrow f(Y) < f(X) | G(Y).$$

The transformation shows that μCLP programs can be executed, in principle, with a complete scheme for negation. We shall show here the completeness of a much simpler concurrent branch and bound like procedure.

To resolve a goal of the form

$$c | \text{minimize}(G(X), f(X)), \alpha$$

the idea is to develop two SLD derivation trees, one Ψ_1 for $c | G(X), \alpha$, and one Ψ_2 for $c \wedge f(Y) < f(X) | G(Y)$. Once a successful derivation is found in Ψ_2 , say with answer constraint d , then Ψ_1 is pruned by adding the constraint $f(X) \leq v$ if $v = \min_d f(Y)$ exists, false otherwise. Once a successful derivation is found in Ψ_1 , say with answer constraint e , then Ψ_1 and Ψ_2 are pruned by adding the constraint $f(X) \leq w$ if $w = \min_e f(X)$ exists, false otherwise. We get a successful derivation for the minimization goal when we get a successful derivation in Ψ_1 and Ψ_2 is finitely failed. The minimization goal gets finitely failed if Ψ_1 gets finitely failed after pruning. Figure 5 illustrates the mutual pruning mechanism.

Note that in the previous case of query optimization the context is empty, $c = \text{true}$, $\alpha = \square$, hence both SLD trees Ψ_1 and Ψ_2 can be taken identical up to variable renaming. The mutual pruning mechanism of the optimization scheme can thus be simplified into a single pruning operation as described in the previous section. This is no longer possible if the goal contains a constraint or an atom outside the minimization predicate.

Example 7.2. Consider the $\mu\text{CLP}(\mathcal{R})$ program




FIGURE 5. Subgoal optimization by pruning.

$p(X) :- X=0.$

$p(X) :- X \geq 1, p(X).$

and the goal $X \geq 1 / \text{minimize}(p(X), X)$. The first SLD tree for $X \geq 1 / p(X)$ is infinite. The second SLD tree for $X \geq 1, Y < X / p(Y)$ contains a success with answer constraint $Y = 0$. The first tree is thus pruned with the constraint $X \leq 0$, hence it gets finitely failed and the answer to the minimization goal is no, in accordance to the logical semantics.

Note that the optimization procedures described in [28] and [9] loop forever on this example. This shows the difficulty to define a complete scheme for optimization w.r.t. logical failures, and w.r.t. successes as well when minimization predicates are nested.

Now the completeness of that procedure can be proved by specializing the principle of constructive by pruning to optimization goals in μCLP programs. As the negative literals in the translation of a μCLP program all come from optimization predicates we can equivalently replace the PRN rule by the following OPT rule for μCLP programs:

$\text{OPT: } \frac{c G(X) \triangleleft F_1 \quad c, f(Y) < f(X) G(Y) \triangleleft F_2}{c \text{min}(G(X), f(X)) \triangleleft c \times F_1 \times \{\neg_V S \text{min}(G(X), f(X)), \neg_V F_2 \square\}}$ <p style="text-align: center;">where $S \subseteq \mathcal{S}(F_2)$, $V = V(c) \cup X$ and $Y \cap V = \emptyset$.</p>
--

It is easy to see that as the variables X and Y are related by the constraint

$f(Y) < f(X)$ solely, the negation of constraints involved in the OPT rule amount to a simple form of constraint minimization:

Proposition 7.3. (negation of constraints as constraint minimization) In the OPT rule for all $d|\alpha \in F_2$, let $Z = V(d) \setminus V$, if $v = \min_d f(Y)$ exists then $\mathcal{A} \models (c \wedge \neg \exists Z d) \leftrightarrow (c \wedge f(X) \leq v)$, otherwise $c \wedge \neg \exists Z d$ is \mathcal{A} -unsatisfiable.

PROOF.

As $Y \cap V = \emptyset$ we have $Y \subseteq Z$ and $d = c \wedge f(Y) < f(X) \wedge d'(Z)$. If $v = \min_d f(Y)$ exists then we have $v = \min_{d'(Z)} f(Y)$ thus $\mathcal{A} \models (c \wedge \neg \exists Z d) \leftrightarrow (c \wedge f(X) \leq v)$. If $\min_{d(X,Y,Z)} f(Y)$ doesn't exist then, under our assumptions on the constraint language, $f(Y)$ is unbounded over $d(X,Y,Z)$, thus for any value v , the constraint $f(Y) < v \wedge d'(Z)$ is \mathcal{A} -satisfiable, therefore $c \wedge \neg \exists Z d$ is \mathcal{A} -unsatisfiable. \square

The OPT rule can thus be interpreted procedurally with both a pruning by success rule (PBS) that prunes the main tree with the constraint $f(X) \leq v$ where v is the optimal value of the objective function for a success in the auxiliary tree (prune with false if v doesn't exist), and with a success by pruning rule (SBP) that negates frontiers in the auxiliary tree once a successful derivation is found in the main tree. It is worth noting however that the computation of frontiers is not necessary in this context, the following proposition shows that the SBP rule can be replaced by a reversed pruning operation and by a check for finite failure in the auxiliary tree.

Proposition 7.4. In the OPT rule, suppose $d|\square \in \mathcal{S}(F_1)$, if $w = \min_d f(X)$ exists and $(f(X) \leq w) \times F_2 = \emptyset$ then $\mathcal{A} \models (d \wedge \neg_X F_2) \leftrightarrow (d \wedge f(X) \leq w)$ otherwise $d \wedge \neg_X F_2$ is \mathcal{A} -unsatisfiable.

PROOF. Remark first that due to the similarity of the goals in the premises of the OPT rule, if $f(X)$ can take two values $v < v'$ under constraint d , then $f(Y)$ can take value v under constraint e for some $e|\alpha \in F_2$.

If $w = \min_d f(X)$ exists and $(f(X) \leq w) \times F_2 = \emptyset$ then for all $e|\alpha \in F_2$ $\mathcal{A} \models e \rightarrow f(Y) \geq w$, thus $\mathcal{A} \models (d \wedge f(X) \leq w) \rightarrow (d \wedge \neg_X F_2)$. Furthermore by (the contrapositive of) the previous remark we have $\mathcal{A} \models (d \wedge \neg_X F_2) \rightarrow (d \wedge f(X) = w)$.

Otherwise, either $\min_d f(X)$ doesn't exist, in which case by the previous remark $d \wedge \neg_X F_2$ is \mathcal{A} -unsatisfiable, or $(f(X) \leq w) \times F_2 \neq \emptyset$ in which case there exists $e|\alpha \in F_2$ such that $e \wedge f(X) \leq w$ is \mathcal{A} satisfiable, thus $d \wedge \neg \exists Z e$ where $Z = V(e) \setminus X$ is \mathcal{A} -unsatisfiable, and so is $d \wedge \neg_X F_2$. \square

Note finally that given a successful derivation with constraint d in the main tree such that $w = \min_d f(X)$ exists, even if the auxiliary tree doesn't get finitely failed by pruning, both the main tree and the auxiliary tree can be pruned with the constraint $f(X) \leq w$ as we already know there will be a similar successful derivation in the auxiliary tree with $f(Y) = w$. This provides evidence that the procedure given in the introduction of this section is a sound procedural interpretation of the principle of constructive negation by pruning. As the transformations preserve the equivalence with the general scheme, the completeness results of the previous sections continue to hold:

Theorem 7.5. Let P be a μ CLP program. The fixpoint semantics $\mathcal{F}(P)$ is fully abstract w.r.t. the answer constraints computed by rules TRIV, RES, FRT and OPT. The operational semantics based on rules TRIV, SLD and OPT is sound and complete w.r.t. the logical semantics $\mathcal{L}(P)$.

7.4. Local Optimization Predicates

The optimization predicates defined in [9] or [24] are more general than those considered in the previous section as they allow to protect a set of variables in the goal subject to optimization. The effect is to localize the optimization to the remaining variables, and relativize the result to the set of protected variables.

Definition 7.6. The local minimization predicate

$$\text{minimize}(G(X, Y), [X], f(X, Y))$$

where $[X]$ is the set of protected variables is defined as an abbreviation for the formula

$$G(X, Y) \wedge \neg \exists Z (f(X, Z) < f(X, Y) \wedge G(X, Z)).$$

The local maximization predicate is defined similarly.

Example 7.7. Local optimization predicates can be used to express the min-max method of game theory with the following goal:

$$\text{maximize}(\text{minimize}(\text{move}(X, Y), \text{move}(Y, Z)), [X, Y], \text{val}(Z)), [X], \text{val}(Z)).$$

Note that protected variables are necessary in this example to conform to the intended semantics.

The previous execution model for optimization predicates is not correct for local optimization predicates. This is not surprising as it is easy to see that any normal logic program can be encoded as a *CLP* program with local optimization predicates in place of negations. Therefore there is no hope to fundamentally improve a general scheme for negation in the context of local optimization predicates.

Proposition 7.8. Any normal logic program is equivalent to a CLP program containing local optimization predicates in place of negative literals.

PROOF. Given a normal logic program P and a normal goal G let us consider the *CLP* goal \overline{G} and the *CLP* program \overline{P} over the Herbrand domain and the natural numbers (Presburger arithmetic) obtained by replacing each negative literal $\neg p(X)$ by

$$\text{maximize}(q(X, y), [X], y)$$

where q is a new predicate symbol and y a new variable, and by adding the clauses

$$q(X, 0).$$

$$q(X, y) \leftarrow p(X).$$

One easily checks (by unfolding) that $\overline{P}^*, CET \models_3 \text{maximize}(q(X, y), [X], y) \leftrightarrow y = 0 \wedge \neg p(X)$, therefore we get $P^* \models_3 \exists G$ iff $\overline{P}^*, \mathcal{N} \models_3 \exists \overline{G}$. \square

The general principle of constructive negation by pruning can be used to interpret local optimization predicates, but now constraints need be negated in addition to be minimized (see figure 1). In its general form the procedure based on constructive negation by pruning may be very inefficient, however it can be simplified under some assumptions and can be used as a reference frame to prove the correctness of other procedures. For instance if the optimization goals are delayed until the protected variables get instantiated, then we can clearly rely on the procedure of the previous

section, thereby eliminating the need of negating constraints. This strategy is used in [24] with an incomplete scheme for global optimization predicates based on the basic branch and bound procedure, but if we replace it by the procedure of the previous section then we get a completeness result for local optimization predicates under the assumption of non-floundering.

It is also possible to design an alternative bottom-up evaluation procedure based on the immediate consequence operator T_P . Here again the direct implementation of the T_P operator is certainly very inefficient, but it can be optimized in many ways and can be a useful tool in some applications where at least a part of the optimization process should be processed bottom-up.

Note finally that constructive negation can be used as well to interpret directly preference predicates over solutions defined by CLP programs, that is to evaluate goals of the form

$$G(X) \wedge \neg \exists Y (G(Y) \wedge \text{better}(Y, X)).$$

where *better* is a user-defined preference predicate. This form of optimization, called relational optimization, doesn't need to encode preferences by objective functions, it is discussed and illustrated by one application in [11].

8. CONCLUSION

The principle of constructive negation by pruning (CNP) provides a correct and complete operational semantics for normal CLP programs w.r.t. Kunen's three-valued logical semantics. CNP is the first scheme to receive a fixpoint semantics which fully characterizes the operational behavior of normal CLP programs w.r.t. computed answer constraints. This generalizes the s-semantics approach [4] to normal CLP programs, and allows to model finite failure for definite CLP programs. Furthermore, the fixpoint semantics is based on a continuous finitary version of Fitting's operator which is interesting to study in its own right. We have shown that it provides a fixpoint characterization of Kunen's semantics.

Under this interpretation the operational semantics of normal logic programs is similar to the one of Drabent [8]. The operational semantics has been defined here with a frontier calculus which reflects the simplicity of the scheme: there are no complex subgoals with explicit quantifiers, no formula transformation at run-time or compile-time, only pruning over concurrent standard SLD derivation trees. It is remarkable that exploiting concurrency in the formation of standard SLD trees is sufficient to build a complete scheme for negation. This is an example of the potential power of concurrency in proof theory. We believe that CNP can lead to a practical scheme for handling negation in CLP systems. Of prime importance is the study of fair computation techniques and of efficient constraint solvers for constraint systems with negation over finite and infinite trees, linear arithmetic, finite domains, order-sorted domains, etc. In the context of global optimization higher-order predicates we have shown that constraint minimization is the only required extension of the solvers, and that CNP specializes to a concurrent branch and bound like procedure, without frontier computation in SLD trees.

Ongoing work concerns on the one hand the natural extension of the class *cc* of concurrent constraint programming languages [25] with constructive negation by pruning and optimization higher-order agents [10], and on the other hand the bottom-up abstract interpretation of normal CLP programs based on operator T_P .

Acknowledgements

The author is grateful to Annalisa Bossi, Alberto Bottoni, Patrick Cousot, Roberto Giacobazzi and Francesco Ranzato for valuable discussions of this material, as well as to the anonymous referees for many helpful comments and suggestions on an earlier version of this paper. This work has been partially supported by MESR contract 93S0051 PRC AMN.

REFERENCES

1. Apt, K., Bol, R., Logic Programming and Negation: a Survey, *Journal of Logic Programming*, 19-20, pp.9-71, 1994.
2. Apt, K., Blair, H.A., Walker, A., Towards a Theory of Declarative Knowledge, in: *Foundations of Deductive Databases and Logic Programming*, J. Minker Ed., Morgan Kaufmann Pub., pp.89-148, 1987.
3. Bossi, A., Fabris, M., Meo, M.C., A Bottom-up Semantics for Constructive Negation, *Proc of the 11th Int. Conf. on Logic Programming*, pp.520-534, MIT Press, 1994.
4. Bossi, A., Gabbrielli, M., Levi, G., Martelli, M., The s-Semantics Approach: Theory and Applications, *Journal of Logic Programming*, 19-20, pp.149-197, 1994.
5. Bruscoli, P., Levi, F., Levi, G., Meo, M.C., Compilative Constructive Negation in Constraint Logic Programs, in: S.Tsion (ed.), *Proceedings of CAAP'94*, LNCS 787, pp.52-67, Springer Verlag, 1994.
6. Chan, D., Constructive Negation Based on the Completed Database, in: R.A. Kowalski and K.A. Bowen (eds), *Proc. of the fifth International Conference on Logic Programming*, MIT Press, Cambridge, MA, pp.11-125, 1988.
7. Chan, D., An Extension of Constructive Negation and its Applications in Coroutining, *Proc. North American Conference on Logic Programming'89*, Cleveland, MIT Press, pp.477-493, 1989.
8. Drabent, W., What is Failure? An Approach to Constructive Negation, *Acta Informatica*, 32:1, pp.27-59, 1995.
9. Fages, F., On the Semantics of Optimization Predicates in CLP Languages, *Proc. 13th FSTTCS conference*, Bombay, LNCS 761, Springer-Verlag, pp. 193-204, 1993.
10. Fages, F., Constructive Negation by Pruning and Optimization Higher-order Predicates for CLP and CC languages, in: *Constraint Programming: Basics and Trends*, Ed. A. Podelski, Springer-Verlag LNCS 910, pp.68-89, 1995.
11. Fages, F., Fowler, J., Sola, T., Handling Preferences in Constraint Logic Programming with relational optimization, *Proc of PLILP'94*, Madrid, Springer Verlag, LNCS 844, pp.261-276, 1994.
12. Fitting, M., A Kripke/Kleene Semantics for Logic Programs, *Journal of Logic Programming*, 2(4), pp.295-312, 1985.
13. Falaschi, M., Levi, G., Martelli, M., Palamidessi, C., Declarative Modeling of the Operational Behavior of Logic Programs, *Theoretical Computer Science*, 69(3), pp.289-318, 1989.
14. Girodias, P., Cerny, E., Older, W.J., Solving Linear, Min and Max Constraint Systems using CLP based on Relational Interval Arithmetic, *Proc. 1st International Conference on Constraint Programming CP'95*, Cassis, France, Sprincger Verlag, LNCS 976, pp.186-203, 1995.
15. Jaffar, J., Lassez, J.L., Constraint Logic Programming, *Proc. of POPL'87*, Munich, pp.111-119, 1987.
16. Jaffar, J., Maher, M.J., Constraint Logic Programming: a Survey, *Journal of Logic Programming*, 19-20, 1994.

17. Kunen, K., Negation in Logic Programming, *Journal of Logic Programming*, 4(3), pp.289-308, 1987.
18. Kunen, K., Signed Data Dependencies in Logic Programming, *Journal of Logic Programming*, 7(3), pp.231-245, 1989.
19. Lloyd, J.W., Foundations of Logic Programming, Springer Verlag, 1987.
20. Maher, M.J., Logic Semantics for a Class of Committed-choice Languages, *Proc. 4th International Conference on Logic Programming*, pp.858-876, MIT Press, 1987.
21. Maher, M.J., A Logic Programming View of CLP, *Proc. 10th International Conference on Logic Programming*, pp.737-753, MIT Press, 1993.
22. Maluszyński, J., Näslund, T., Fail Substitutions for Negation as Failure, *Proc. North American Conference on Logic Programming'89*, Cleveland, MIT Press, pp.461-476, 1989.
23. Maher, M.J., Stuckey, P.J., Expanding Query Power in Constraint Logic Programming Languages, *Proc. North American Conference on Logic Programming'89*, Cleveland, MIT Press, pp.20-36, 1989.
24. Marriott, K., Stuckey, P.J., Semantics of Constraint Logic Programs with Optimization, *ACM letters on Programming Languages and Systems*, 2(1-4), 197-212, 1993.
25. Saraswat, V., Concurrent Constraint Programming, MIT Press, 1993.
26. Stuckey, P., Constructive Negation for Constraint Logic Programming, *Proc. LICS'91*, pp.328-339, 1991.
27. Stuckey, P., Negation and Constraint Logic Programming, *Information and Computation*, 118(1), pp.12-33, 1995.
28. Van Hentenryck, P., Constraint Satisfaction in Logic Programming, MIT Press 1989.
29. Van Hentenryck, P., Deville, Y., Teng, C.M., A Generic Arc-consistency Algorithm and its Specialization, *Artificial Intelligence*, 57, pp.291-321, 1992.