Computing and Informatics, Vol. 25, 2006, 1001–1021, V 2006-Mar-28

EXPLANATIONS AND PROOF TREES

Gérard Ferrand, Willy Lesaint, Alexandre Tessier

Laboratoire d'Informatique Fondamentale d'Orléans
Rue Léonard de Vinci
BP 6759
45067 Orléans Cedex 2
France
e-mail: {Gerard.Ferrand, Willy.Lesaint, Alexandre.Tessier}@univ-orleans.fr

Revised manuscript received 3 February 2006

Abstract. This paper proposes a model for explanations in a set theoretical framework using the notions of closure or fixpoint. In this approach, sets of rules associated with monotonic operators allow to define proof trees. The proof trees may be considered as a declarative view of the trace of a computation. We claim they are explanations of the results of a computation. This notion of explanation is applied to constraint logic programming, and it is used for declarative error diagnosis. It is also applied to constraint programming, and used for constraint retraction.

Keywords: Explanation, proof tree, rule, fixpoint, closure, constraint logic programming, constraint satisfaction problem, error diagnosis, constraint retraction

1 INTRODUCTION

The motivation for our notion of explanation is to explain a result of a computation or an intermediate result of a computation. Many applications need to understand how a result is obtained, that is to have an explanation of it. The explanations may be useful for applications such as:

- program debugging, in order to help error diagnosis as shown for example in Section 3.4;

- computation optimization, for example for intelligent backtracking in constraint programming [6];
- to reason on results, for example to prove some properties of the results of the computations;
- to justify a result, for example when the user wants to understand a result;
- as a theoretical tool, for example to prove the correctness of a system as shown in Section 4.4 for the correctness of constraint retraction algorithms.

It is important to note that to compute a result is exactly to find a proof of it. A trace of the computation can be considered as an explanation; but a trace may be very large and intricate. Furthermore, it may contain events that are not useful because they are not significant to explain the result (for example the order of application of the operators in constraint programming).

It is then necessary to have high level explanations without all the details of a particular operational semantics. We would like to define proofs of results which are *declarative*.

Declarative languages, such as logic programming and constraint programming, may be described in terms of fixpoint computations by monotonic operators.

This paper proposes a model for explanations in a set theoretical framework using the notions of closure or fixpoint. In this approach, sets of rules associated with monotonic operators allow to define proof trees [1]. We claim they are explanations of the result of a computation. These explanations may be considered as a declarative view of the trace of a computation.

First, the general scheme is given. This general scheme is applied to constraint logic programming. Two notions of explanations are given: positive explanations and negative explanations. Their use in declarative error diagnosis is proposed.

Next, the general scheme is applied to constraint programming. In this framework, two definitions of explanations are described as well as an application to constraint retraction.

2 PROOF TREES AND FIXPOINT

Our model for explanations is based on the notion of *proof tree*. To be more precise, from a formal point of view we see an explanation as a proof tree, which is built with *rules*. Here is an example: the following tree

```
   a
  /|\
 b c d
 |  / \
 e f   g
```

is built with 7 rules including the rule $a \leftarrow \{b, c, d\}$; the rule $b \leftarrow \{e\}$; the rule $e \leftarrow \emptyset$ and so on. From an intuitive point of view the rule $a \leftarrow \{b, c, d\}$ is an immediate

explanation of a by the set $\{b, c, d\}$, the rule $e \leftarrow \emptyset$ is a *fact* which means that e is given as an axiom. The whole tree is a complete explanation of a.

For legibility, we do not write braces in the body of rules: the rule $a \leftarrow \{b, c, d\}$ is written $a \leftarrow b, c, d$, the fact $e \leftarrow \emptyset$ is written $e \leftarrow$.

2.1 Rules and Proof Trees

Rules and proof trees [1] are abstract notions which are used in various domains in logic and computer science such as *proof theory* [31] or *operational semantics of programming languages* [30, 28, 13].

A rule $h \leftarrow B$ is merely a pair (h, B) where B is a set. If B is empty the rule is a *fact* denoted by $h \leftarrow$. In general h is called the *head* and B is called the *body* of the rule $h \leftarrow B$. In some contexts h is called the *conclusion* and B the *set of premises*.

A tree is *well founded* if it has no infinite branch. In any tree t, with each node ν is associated a rule $h \leftarrow B$ where h is the label of ν and B is the set of the labels of the children of ν. Obviously a fact is associated with a *leaf*.

A set of rules $\mathcal{R}$ defines a notion of *proof tree*: a tree t is a proof tree wrt $\mathcal{R}$ if it is well founded and the rules associated with its nodes belong to $\mathcal{R}$.

2.2 Least Fixpoints and Upward Closures

In logic and computer science, interesting sets are often defined as least fixpoints of monotonic operators. Our framework is set-theoretical, so here an *operator* is merely a mapping $T : \mathcal{P}(S) \to \mathcal{P}(S)$ where $\mathcal{P}(S)$ is the power set of a set S. T is *monotonic* if $X \subseteq Y \subseteq S \Rightarrow T(X) \subseteq T(Y)$. From now on T is supposed monotonic.

X is *upward closed* by T if $T(X) \subseteq X$ and X is a *fixpoint* of T if $T(X) = X$.

The least X which is upward closed by T exists (it is the intersection of all these X) and it is also the *least fixpoint* of T, denoted by $\mathrm{lfp}(T)$ (it is a particular case of the classical *Knaster-Tarski* theorem). Since $\mathrm{lfp}(T)$ is the least X such that $T(X) \subseteq X$, to prove $\mathrm{lfp}(T) \subseteq X$ it is sufficient to prove $T(X) \subseteq X$. It is the *principle of proof by induction* and $\mathrm{lfp}(T)$ is called the set *inductively defined* by T.

Sometimes an interesting set is not directly defined as the least fixpoint of a monotonic operator, it is defined as upward closure of a set by a monotonic operator, but it is basically the same machinery: the *upward closure* of X by T is the least Y such that $X \subseteq Y$ and $T(Y) \subseteq Y$, that is to say the least Y such that $X \cup T(Y) \subseteq Y$, which is the least fixpoint of the operator $Y \mapsto X \cup T(Y)$ (but it is not necessarily a fixpoint of T itself).

Now let $\mathcal{R}$ be a given set of rules. In practice a set S is supposed to be given such that $h \in S$ and $B \subseteq S$ for each rule $(h \leftarrow B) \in \mathcal{R}$. In this context the set of rules $\mathcal{R}$ defines the operator $T_{\mathcal{R}} : \mathcal{P}(S) \to \mathcal{P}(S)$ by

$$T_{\mathcal{R}}(X) = \{h \mid \exists B \subseteq X \, (h \leftarrow B) \in \mathcal{R}\}$$

which is obviously monotonic.

For example $h \in T_{\mathcal{R}}(\emptyset)$ if and only if $h \leftarrow$ is a rule (fact) of $\mathcal{R}$; $h \in T_{\mathcal{R}}(T_{\mathcal{R}}(\emptyset))$ if and only if there is a rule $h \leftarrow B$ in $\mathcal{R}$ such that $b \in T_{\mathcal{R}}(\emptyset)$ for each $b \in B$ ($b \leftarrow$ is a rule of $\mathcal{R}$); it is easy to see that members of $T^n_{\mathcal{R}}(\emptyset)$ are proof tree roots.

Conversely, in this set-theoretical framework, each monotonic operator T is defined by a set of rules, that is to say $T = T_{\mathcal{R}}$ for some $\mathcal{R}$ (for example take the trivial rules $h \leftarrow B$ such that $h \in T(B)$).

Now $\mathrm{lfp}(T_{\mathcal{R}})$ is called the set inductively defined by $\mathcal{R}$ and to prove $\mathrm{lfp}(T_{\mathcal{R}}) \subseteq X$ by induction, that is to say to prove merely $T_{\mathcal{R}}(X) \subseteq X$ is exactly to prove $B \subseteq X \Rightarrow h \in X$ for each rule $h \leftarrow B$ in $\mathcal{R}$.

A significant property is that the members of the least fixpoint of $T_{\mathcal{R}}$ are exactly the *proof tree roots* wrt $\mathcal{R}$. Let R the set of the proof tree roots wrt $\mathcal{R}$. It is easy to prove $\mathrm{lfp}(T_{\mathcal{R}}) \subseteq R$ by induction. $R \subseteq \mathrm{lfp}(T_{\mathcal{R}})$ is also easy to prove: if t is a proof tree, by well-founded induction all the labels of the nodes of t are in $\mathrm{lfp}(T_{\mathcal{R}})$.

Note that for each monotonic operator T there are possibly many $\mathcal{R}$ such that $T = T_{\mathcal{R}}$. In each particular context there is often one $\mathcal{R}$ that is natural, which can provide a natural notion of explanation for the membership of the least fixpoint of T. Here, the operator T is associated to a program and there exists a set of rules that can be naturally deduced from the program and that gives an interesting notion of explanation for the members of the least fixpoint of T.

2.3 Upward Iterations

The least fixpoint of a monotonic operator $T : \mathcal{P}(S) \to \mathcal{P}(S)$ can be computed by *iterating* T from the empty set: let $X_0 = \emptyset, X_{n+1} = T(X_n)$. It is easy to see that $X_0 \subseteq X_1 \subseteq \cdots \subseteq X_n$ and that $X_n \subseteq \mathrm{lfp}(T)$. If S is finite, obviously there exists a natural number n such that $X_n = X_{n+1}$. It is easy to see that $X_n = \mathrm{lfp}(T)$.

In the general case the iteration must be *transfinite*: n may be any *ordinal*. $X_n = T(X_{n-1})$ if n is a *successor ordinal*, and $X_n = \bigcup_{\nu<n} X_\nu$ if n is a *limit ordinal*. Then for some ordinal α we have $X_\alpha = X_{\alpha+1}$ which is $\mathrm{lfp}(T)$. The first such α is the *(upward) closure ordinal* of T.

In practice S is not necessarily finite but often $T = T_{\mathcal{R}}$ for a set $\mathcal{R}$ of rules which are *finitary*, that is to say that, in each rule $h \leftarrow B$, B is finite. In that case the closure ordinal of T is $\leq \omega$ that is to say $\mathrm{lfp}(T) = X_\omega = \bigcup_{n<w} X_n = \bigcup_{n\in\mathbb{N}} X_n$ (intuitively, the natural numbers are sufficient because each proof tree is a finite tree).

More generally the upward closure of X by T can be computed by iterating T from X by defining: $X_0 = X, X_{n+1} = X_n \cup T(X_n), \ldots$

Sometimes the operator T is defined by $T(X) = \bigcup_{i\in I} T_i(X)$ where several operators $T_i : \mathcal{P}(S) \to \mathcal{P}(S)$ $(i \in I)$ are given. It is the case when $T = T_{\mathcal{R}}$ with $\mathcal{R} = \bigcup_{i\in I} \mathcal{R}_i$ and $T_i = T_{\mathcal{R}_i}$. The upward closure of X by T is also called the upward closure of X by the T_i $(i \in I)$ and it is the least Y such that $X \subseteq Y$ and $T_i(Y) \subseteq Y$ for each $i \in I$. Instead of computing this closure by using $T(X) = \bigcup_{i\in I} T_i(X)$, it is in practice more efficient to use an *upward chaotic iteration* [9, 14, 2] of the T_i $(i \in I)$, where at each step only one T_i is chosen and applied: $X_0 = X, X_{n+1} =$

$X_n \cup T_{i_{n+1}}(X_n), \ldots$ with $i_{n+1} \in I$. The sequence $i_1, i_2, \ldots$ is called a *run* and is a formalization of the choices of the T_i. If S is finite, obviously for some natural number n, we have $X_n = X_{n+1}$ that is to say $T_{i_{n+1}}(X_n) \subseteq X_n$ but X_n is the closure only if $T_i(X_n) \subseteq X_n$ for all $i \in I$. If also I is finite it is easy to see that *finite runs* $i_1, i_2, \ldots, i_n$ exist such that X_n is the closure, for example by choosing each i in turn.

In general, from a theoretical point of view a *fairness* condition on the (infinite) run is presupposed to ensure that the closure is reached. Such a run is called a *fair run* but the details are beyond the scope of the paper. For the application below to constraint satisfaction problems, I and S may be supposed to be finite.

2.4 Duality and Negative Information

Sometimes the interesting sets are greatest fixpoint or downward closures of some monotonic operators.

Each monotonic operator $T : \mathcal{P}(S) \to \mathcal{P}(S)$ has a *greatest fixpoint*, denoted by $\mathrm{gfp}(T)$, that is to say the greatest X such that $T(X) = X$. In fact $\mathrm{gfp}(T)$ is also the greatest X which is *downward closed* by T, that is $X \subseteq T(X)$. It is the reason why to prove $X \subseteq \mathrm{gfp}(T)$ it is sufficient to prove $X \subseteq T(X)$ (*principle of proof by co-induction*, $\mathrm{gfp}(T)$ is called the set *coinductively defined* by T).

The *downward closure* of X by T is the greatest Y such that $Y \subseteq X$ and $Y \subseteq T(Y)$, that is to say the greatest Y such that $Y \subseteq X \cap T(Y)$, which is the greatest fixpoint of the operator $Y \mapsto X \cap T(Y)$ (but it is not necessarily a fixpoint of T itself).

Greatest fixpoint and downward closure can be computed by *downward iterations*, similar to the upward iterations of the previous section, reversing $\subseteq$, replacing $\cup$ by $\cap$ and $\emptyset$ by S. Each monotonic operator has a (downward) closure ordinal which is obviously finite (natural number) if S is a finite set. If S is infinite, the downward closure ordinal may be $> \omega$ even if the upward closure ordinal is $\leq \omega$. For example, it is the case for the application to constraint logic programming (but it is outside the scope of this paper).

However, this apparent symmetry between least fixpoint and greatest fixpoint is misleading because we are mainly interested in the notion of proof tree, as a model for explanations, so we are interested in a set of rules $\mathcal{R}$, which defines an operator $T = T_{\mathcal{R}}$. It is only the least fixpoint of $T_{\mathcal{R}}$ which has the significant property that its members are exactly the proof tree roots wrt $\mathcal{R}$. The greatest fixpoint can also be described in terms of trees, but these trees are not necessarily well founded and they are not in the scope of this paper. In this paper a tree must be well founded in order to be an explanation.

However, concerning greatest fixpoint and downward closure, we are going to see that a proof tree can be an explanation for the non-membership, that is to say to deal with *negative information*.

In this set-theoretical framework we can use complementation: for $X \subseteq S$, the *complementary* $S \setminus X$ is denoted by $\overline{X}$. The *dual* of $T : \mathcal{P}(S) \to \mathcal{P}(S)$ is the

operator $T' : \mathcal{P}(S) \to \mathcal{P}(S)$ defined by $T'(X) = \overline{T(\overline{X})}$. T' is obviously monotonic if T is monotonic. X is a fixpoint of T' if and only if $\overline{X}$ is a fixpoint of T. Since $X \subseteq Y$ if and only if $\overline{Y} \subseteq \overline{X}$, $\mathrm{gfp}(T) = \overline{\mathrm{lfp}(T')}$ and $\overline{\mathrm{gfp}(T)} = \mathrm{lfp}(T')$. So if $\mathcal{R}'$ is a natural set of rules defining T', a proof tree wrt $\mathcal{R}'$ can provide a natural notion of explanation for the membership of the complementary of the greatest fixpoint of T since it is the membership of the least fixpoint of T'.

Concerning iterations it is easy to see that downward iterations which compute greatest fixpoint of T and downward closures can be uniformly converted by complementation into upward iterations which compute least fixpoint of T' and upward closures.

Sometimes the operator T is defined by $T(X) = \bigcap_{i \in I} T_i(X)$ where several operators $T_i : \mathcal{P}(S) \to \mathcal{P}(S)$ $(i \in I)$ are given. The downward closure of X by the T_i $(i \in I)$ is the greatest Y such that $Y \subseteq X$ and $Y \subseteq T_i(Y)$ for each $i \in I$. Instead of computing this closure by using $T(X) = \bigcap_{i \in I} T_i(X)$, it is more efficient in practice to use a *downward chaotic iteration* of the T_i $(i \in I)$, where at each step only one T_i is chosen and applied. It is easy to see that downward chaotic iterations which compute downward closures can be uniformly converted by complementation into upward chaotic iterations which compute upward closures.

2.5 Explanations for Diagnosis

Intuitively, let us consider that a set of rules $\mathcal{R}$ is an abstract formalization of a computational mechanism so that a proof tree is an abstract view of a *trace*. The results of the possible computations are proof tree roots wrt $\mathcal{R}$, that is to say members of the least fixpoint of a monotonic operator $T = T_{\mathcal{R}}$. Infinite computations related to non-well founded trees and greatest fixpoint are outside the scope of this paper. For example, in the application below to constraint satisfaction problems the formalization uses a greatest fixpoint but in fact by the previous duality we consider proof trees related to a least fixpoint.

Now let us consider that the set of rules $\mathcal{R}$ may be *erroneous*, producing non expected results: some $r \in \mathrm{lfp}(T)$ are *non expected* and some others $r \in \mathrm{lfp}(T)$ are *expected*. From a formal viewpoint this is represented by a set $E \subseteq S$ such that, for each $r \in S$, r is expected if and only if $r \in E$. A $r \in \mathrm{lfp}(T) \setminus E$ (a non expected result) is called a *symptom* wrt E. If there exists a symptom, $\mathrm{lfp}(T) \not\subseteq E$ so $T(E) \not\subseteq E$ (otherwise $\mathrm{lfp}(T) \subseteq E$ by the principle of proof by induction). $T(E) \not\subseteq E$ means that there exists a rule $h \leftarrow B$ in $\mathcal{R}$ such that $B \subseteq E$ but $h \notin E$. Such a rule is called an *error* wrt E. Intuitively it is the existence of errors which explains the existence of symptoms. Diagnosis consists in locating errors in $\mathcal{R}$ from symptoms.

Now the notion of proof tree can *explain* how an error can be a *cause* of a symptom: if r is a symptom it is the root of a proof tree t wrt $\mathcal{R}$. In t we call *symptom node* a node whose label is a symptom (there is at least a symptom node which is the root). Since t is well founded, the relation parent-child is well founded, so there is at least a *minimal symptom node* wrt this relation. The rule $h \leftarrow B$ associated with a minimal symptom node is obviously an error since h is a symptom but no $b \in B$ is

a symptom. The proof tree t is an abstract view of a *trace* of a computation which has produced the symptom r. It explains how erroneous information is propagated to the root. Moreover, by inspecting some nodes in t it is possible to locate an error.

3 CONSTRAINT LOGIC PROGRAMMING

We consider the general scheme of constraint logic programming [25] called CLP(X), where X is the underlying constraint domain. For example, X may be the Herbrand domain, infinite trees, finite domains, $\mathbb{N}$, $\mathbb{R}$...

Two kinds of atomic formula are considered in this scheme: *constraints* (with built-in predicates) and *atoms* (with program predicates, i.e. predicates defined by the program).

A *clause* is a formula

$$a_0 \leftarrow c \wedge a_1 \wedge \cdots \wedge a_n \quad (n \geq 0)$$

where the a_i are atoms and c is a (possibly empty) conjunction of constraints. In order to simplify, we assume that each a_i is an atom $p_i(x_1^i, \ldots, x_{k_i}^i)$ and all the variables x_j^i ($i = 0, \ldots, n$, $j = 1, \ldots, k_i$) are different. This is always possible by adding equalities to the conjunction of constraints c.

Each program predicate p is defined by a set of clauses: the clauses that have an atom with the predicate symbol p in the left part (the head of the clause). This set of clauses is called *the packet* of p.

A *constraint logic program* is a set of clauses.

3.1 Positive Answer

In constraint logic programming, a *logical answer* to a goal $\leftarrow a$ (a is an atom) is a formula $c \rightarrow a$ where c is a conjunction of constraints and $c \rightarrow a$ is a *logical consequence* of the program. So $c \rightarrow a$ is true in each model $\mathcal{M}$ of the program expanding the underlying constraint domain $\mathcal{D}$. In other words, in $\mathcal{M}$, a is satisfied by each solution of c. A solution of c is a valuation v in $\mathcal{D}$ for which c is true. If a is the atom $p(x_1, \ldots, x_n)$ and v is a valuation assigning the value d_i to the variable x_i, then the expression $p(d_1, \ldots, d_n)$ is denoted by $v(a)$. A model $\mathcal{M}$ can be identified with a set of such expressions $v(a) = p(d_1, \ldots, d_n)$, $v(a) \in \mathcal{M}$ meaning that, in $\mathcal{M}$, a is true for the valuation v. The program has a least model, which is a formalisation of its declarative semantics, since $c \rightarrow a$ is a logical answer to the goal $\leftarrow a$ if and only if $c \rightarrow a$ is true in this least model.

If no valuation satisfies c then the answer is not interesting (because c is false and false $\rightarrow a$ is always true). *Computed answers* are logical answers which are defined by the operational semantics. A *reject criterion* tests the satisfiability of the conjunction of constraints built during the computation in order to end the computation when it detects that the conjunction is unsatisfiable. The reject criterion is often incomplete and it just ensures that rejected conjunctions of con-

straints are unsatisfiable in the underlying constraint domain. From an operational viewpoint, the reject criterion may be seen as an optimization of the computation (needless to continue the computation when the conjunction of constraints has no solution).

A monotonic operator may be defined such that its least fixpoint provides the declarative semantics of the program. A candidate is an operator similar to the well known *immediate consequence operator* (often denoted by T_P) in the framework of pure logic programming (logic programming is a particular case of constraint logic programming where unification is seen as equality constraint over terms).

A *set of rules* may be associated with this monotonic operator. For example, a convenient set of rules is the set of all the $v(a_0) \leftarrow v(a_1), \ldots, v(a_n)$ such that $a_0 \leftarrow c \wedge a_1 \wedge \cdots \wedge a_n$ is a clause of the program and v is a valuation solution of c. This set of rules basically provides a notion of explanation. Because of the clause, if $v(a_1), \ldots, v(a_n)$ belong to the semantics of the program, then $v(a_0)$ belongs to the semantics of the program. The point is that the explanations defined by this set of rules are theoretical because they cannot always be expressed in the language of the program (for example, if the constraint domain is $\mathbb{R}$, each value of the domain does not correspond to a constant of the programming language). Moreover, it is better to use the same language for the program answers and their explanations. Indeed, if the user is able to understand the program answers then s/he should be able to understand the explanations (no need to understand the computational behaviour).

Another *monotonic operator* may be defined such that its least fixpoint is the set of computed answers $(c \rightarrow a)$.

Again, we can give a *set of rules* which inductively defines the operator. The rules come directly from the clauses of the program and the reject criterion [21]. The rules may be defined as follows:

- for all renamed clause $a_0 \leftarrow c \wedge a_1 \wedge \cdots \wedge a_n$
- for all conjunction of constraints $c_1, \ldots, c_n$

we have the rule:

$$(c_0 \rightarrow a_0) \leftarrow (c_1 \rightarrow a_1), \ldots, (c_n \rightarrow a_n)$$

where c_0 is not rejected by the reject criterion and c_0 is defined by $c_0 = \exists_{-a_0}(c \wedge c_1 \wedge \cdots \wedge c_n)$, $\exists_{-a_0}$ denotes the existential quantification except on the variables of a_0.

Other sets of rules can be given and are discussed in [36]. Each set of rules provides another notion of explanation. For each answer $c \rightarrow a$, there exists an explanation rooted by $c \rightarrow a$. Moreover, each node of an explanation is also an answer: a formula $c \rightarrow a$. An answer is explained as a consequence of other answers using a rule deduced from a clause of the program. This notion of explanation has been successfully used for declarative error diagnosis [36] in the framework of algorithmic debugging [34] as shown later.

3.2 Negative Answer

Because of the non-determinism of constraint logic programs, another level of answer may be considered. It is built from the answers of the first level. If $c_1 \rightarrow a, \ldots, c_n \rightarrow a$ are the answers of the first level to a goal $\leftarrow a$, we have $c_1 \vee \cdots \vee c_n \rightarrow a$ in the program semantics. For the second level of answer we now consider $c_1 \vee \cdots \vee c_n \leftarrow a$.

The answers of the first level (the $c_i \rightarrow a$) are called *positive answers* because they provide positive information on the goals (each solution of a c_i is a solution of a) whereas the answers of the second level (the $c_1 \vee \cdots \vee c_n \leftarrow a$) are called *negative answers* because they provide negative information on the goals (there is no other solution of a than the solutions of the c_i).

Again, the set of negative answers is the least fixpoint of a monotonic operator. A *set of rules* may be naturally associated with the operator, each rule is defined using the packet of clauses of a program predicate. The set of rules provides a notion of *negative explanation*.

It is not possible to give in few lines the set of (negative) rules because it requires several preliminary definitions (in particular, it needs to define very rigorously the CSLD[1]-search tree with the notion of skeleton of partial explanations). The reader may find details about some systems of negative rules and the explanations of negative answers in [21, 22, 35].

The nodes of a negative explanation are negative answers: formula $C \leftarrow a$, where C is a disjunction of conjunctions of constraints.

3.3 Links Between Explanations and Computation

In this article, the notion of answer is defined when the computation is finite, that is to say when the computation ends and provides a result.

The notion of positive computation corresponds to the notion of CSLD-derivation [29, 25]. It corresponds to the computation of a branch of the CSLD-search tree. With each finite branch of the CSLD-search tree a positive answer is associated (even when the CSLD-search tree is not finite).

The notion of negative computation corresponds to the notion of CSLD-resolution [29, 25]. It corresponds to the computation of the whole CSLD-search tree. Thus a negative answer is associated only with a finite CSLD-search tree.

A positive explanation explains an answer computed by a finite CSLD-derivation (a positive answer) while a negative explanation explains an answer computed by a finite CSLD-search tree (a negative answer).

The interesting point is that the nodes of the explanations are answers, that is, an answer is explained as a consequence of other answers.

[1] The *CSLD-resolution* is an adaptation of the well-known *SLD-resolution* (Linear resolution with a Selection rule for Definite programs) to constraint logic programs. The main adaptation is to use the *reject criterion* instead of the *unification*.

The explanations defined here may be seen as a declarative view of the trace: it contains all the declarative information of the trace without the operational details. This is important because in constraint logic programming, the programmer may write his/her program using only a declarative knowledge of the problem to solve. Thus it would be such a great pity that the explanations of answers used operational aspects of the computation.

3.4 Declarative Error Diagnosis

An unexpected answer of a constraint logic program is the *symptom* of an *error* in the program. Because we have an (unexpected) answer, the computation is finite. If we have a *positive symptom*, that is an unexpected positive answer, the finite computation corresponds to a finite branch of the CSLD-search tree. If we have a *negative symptom*, that is an unexpected negative answer, then the CSLD-search tree is finite.

Given some expected properties of a constraint logic program, given a (positive or negative) symptom, using the previous notions of explanations (positive explanations or negative explanations), using the general scheme for diagnosis given in Section 2.5, we can locate an error (or several errors) in the constraint logic program. The diagnoser asks an oracle (in practice, the user or a specification of the program) in order to know if a node of the explanation is a symptom. The diagnoser searches for a minimal symptom in the explanation. A minimal symptom exists because the root of the explanation is a symptom and the explanation is well founded (it is finite). The rule that links the minimal symptom to its children is erroneous in some sense:

- If the symptom is a positive symptom, then it is a positive rule and the clause used to define the rule is a *positive error*: the clause is *incorrect* according to the expected properties of the program. Moreover, the constraint in the minimal symptom provides a context in which the clause is not correct.
- If the symptom is a negative symptom, then it is a negative rule and the packet of clauses used to define the rule is a *negative error*: the packet of clauses is *incomplete* according to the expected properties of the program.

Different strategies may be used in order to locate a minimal symptom [36].

Thanks to the diagnosis, the programmer knows a clause, or a packet of clauses, that is not correct and can fix his/her program.

A positive symptom is a wrong positive answer. A negative symptom is a wrong negative answer, but a negative symptom is also the symptom of a missing positive answer. Another kind of negative error diagnosis has been developed for pure logic programs [16, 17]. It needs the definition of infinite (positive) explanations. The set of roots of infinite positive explanations is the greatest fixpoint of the operator defined by the set of positive rules. Note that if the programmer can notice that a positive answer is missing then the CSLD-search tree is finite (there is a negative

answer). Thus, if a positive answer is missing, then it is not in the greatest fixpoint of the operator defined by the positive rules (in that case, the missing positive answer is not also in the least fixpoint of the operator). Note however that, in this context, the good notion refers to the greatest fixpoint and infinite positive explanations. The principle of this other error diagnosis for missing positive answer consists in trying to build an infinite positive explanation rooted by the missing positive answer. Because it is not in the greatest fixpoint, the building of the infinite positive explanation fails. When it fails, it provides an error: a packet of clauses *insufficient* according to the expected properties of the program.

4 CONSTRAINT SATISFACTION PROBLEMS

Constraint Satisfaction Problems (CSP) [37, 3, 12] have proved to be efficient to model many complex problems. A lot of modern constraint solvers (e.g. chip, gnuProlog, Ilog solver, choco) are based on domain reduction to find the solutions of a CSP. But these solvers are often black-boxes whereas the need to understand the computations is crucial in many applications. Explanations have already proved their efficiency for such applications. Furthermore, they are useful for dynamic constraint satisfaction problems [5, 33, 7], over-constrained problems [27], search methods [32, 24, 6], declarative error diagnosis [19]...

Here, two notions of explanations are described: explanation-tree and explanation-set. The first one corresponds to the notion of proof tree. The second one, which can be deduced from explanation-tree, is less precise but sufficient for a lot of practical applications. A more detailed model of these explanations for constraint programming over finite domains is proposed in [18] and a more precise presentation of their application to constraint retraction can be found in [11].

4.1 CSP and Solutions

Following [37], a constraint satisfaction problem is made of two parts: a syntactic part and a semantic part. The syntactic part is a finite set V of variables, a finite set C of constraints and a function var : $C \rightarrow \mathcal{P}(V)$, which associates a set of related variables to each constraint. Indeed, a constraint may involve only a subset of V. For the semantic part, we need to consider various families $f = (f_i)_{i \in I}$. Such a family is referred to by the function $i \mapsto f_i$ or by the set $\{(i, f_i) \mid i \in I\}$.

$(D_x)_{x \in V}$ is a family where each D_x is a finite non empty set of possible values for x. We define the *domain of computation* by $\mathbb{D} = \bigcup_{x \in V}(\{x\} \times D_x)$. This domain allows simple and uniform definitions of (local consistency) operators on a power-set. For reduction, we consider subsets d of $\mathbb{D}$. Such a subset is called an *environment*. Let $d \subseteq \mathbb{D}$, $W \subseteq V$, we denote by $d|_W$ the set $\{(x, e) \in d \mid x \in W\}$. d is actually a family $(d_x)_{x \in V}$ with $d_x \subseteq D_x$: for $x \in V$, we define $d_x = \{e \in D_x \mid (x, e) \in d\}$. d_x is the *domain* of the variable x.

Constraints are defined by their set of allowed tuples. A *tuple* t on $W \subseteq V$ is a particular environment such that each variable of W appears only once: $t \subseteq \mathbb{D}|_W$

and $\forall x \in W, \exists e \in D_x, t|_{\{x\}} = \{(x, e)\}$. For each $c \in C$, T_c is a set of tuples on $\mathrm{var}(c)$, called the solutions of c. Note that a tuple $t \in T_c$ is equivalent to a family $(e_x)_{x \in \mathrm{var}(c)}$ and t is identified with $\{(x, e_x) \mid x \in \mathrm{var}(c)\}$.

We can now formally define a CSP and a solution:

> A *Constraint Satisfaction Problem* (CSP) is defined by: a finite set V of variables, a finite set C of constraints, a function $\mathrm{var} : C \to \mathcal{P}(V)$, a family $(D_x)_{x \in V}$ (the domains) and a family $(T_c)_{c \in C}$ (the constraints semantics). A *solution* for a CSP $(V, C, \mathrm{var}, (D_x)_{x \in V}, (T_c)_{c \in C})$ is a tuple t on V such that $\forall c \in C, t|_{\mathrm{var}(c)} \in T_c$.

4.2 Domain Reduction

To find the possibly existing solutions, the solvers are often based on *domain reduction*. In this framework, monotonic operators are associated with the constraints of the problem with respect to a notion of local consistency (in general, the more accurate is the consistency, the more expensive is the computation). These operators are called *local consistency operators*. In GNU-Prolog for example, these operators correspond to the indexicals (the *x in r*) [8].

For the sake of clarity, we will consider in our presentation that each operator is applied to the whole environment, but in practice it only removes from the environments of one variable some values which are inconsistent with respect to the environments of a subset of V.

> A *local consistency operator* is a monotonic function $r : \mathcal{P}(\mathbb{D}) \to \mathcal{P}(\mathbb{D})$.

Classically [4, 2], reduction operators are considered as monotonic, contracting and idempotent functions. However, on the one hand, *contractance* is not mandatory because environment reduction after applying a given operator r can be forced by intersecting its result with the current environment, that is $d \cap r(d)$. On the other hand, *idempotence* is useless from a theoretical point of view (it is only useful in practice for managing the propagation queue). This is generally not mandatory to design effective constraint solvers. We can therefore use only *monotonic* functions to define the local consistency operators.

The solver semantics is completely described by the set of such operators associated with the handled constraints. More or less accurate local consistency operators may be selected for each constraint. Moreover, this framework is not limited to arc-consistency but may handle *any local consistency* which boils down to domain reduction as shown in [18].

Of course local consistency operators should be *correct* with respect to the constraints. In practice, to each constraint $c \in C$ is associated a set of local consistency operators $R(c)$. The set $R(c)$ is such that for each $r \in R(c)$, $d \subseteq \mathbb{D}$ and $t \in T_c$: $t \subseteq d \Rightarrow t \subseteq r(d)$.

From a general point of view, domain reduction consists in applying these local consistency operators according to a *chaotic iteration* until their *common greatest*

fixpoint is reached. Note that finite domains and chaotic iteration ensure to reach this fixpoint. Obviously, the common greatest fixpoint is an environment which contains all the solutions of the CSP. It is the most accurate set which can be computed using a set of local consistency operators.

In practice, constraint propagation is handled through a *propagation queue*. The propagation queue contains local consistency operators that may reduce the environment (in other words, the operators which are not in the propagation queue cannot reduce the environment). Informally, starting from the given *initial environment* of the problem, a local consistency operator is selected from the propagation queue (initialized with all the operators) and applied to the environment resulting in a new one. If a domain reduction occurs, new operators may be added to the propagation queue. Note that the operators' selection corresponds to a fair run.

Of course, in practice the computations needs to be finite. *Termination* is reached when:

- a domain of a variable is emptied: there is no solution to the associated problem;
- the propagation queue is emptied: a common fixpoint (or a desired consistency state) is reached ensuring that further propagation will not modify the result.

Note that to obtain a solution, domain reduction steps are interlaced with splitting steps. Since splitting may be considered as additional constraints, it could be easily included in our model. This leads to no conceptual difficulties but this is not really necessary here.

4.3 Explanations

Now, we detail two notions of explanations for CSP: *explanation-set* and *explanation-tree*. These two notions explain why a value is removed from the environment. Note that explanation-trees are both more precise and general than explanation-sets, but explanation-sets may be sufficient for some applications. For example, the algorithm of constraint retraction [26] only uses explanation-sets, while its proof of correctness [11] needs explanation-trees.

> Let R be the set of all local consistency operators. Let $h \in \mathbb{D}$ and $d \subseteq \mathbb{D}$. We call *explanation-set* for h wrt d a set of local consistency operators $E \subseteq R$ such that $h \notin CL \downarrow (d, E)$ where $CL \downarrow (d, E)$ denotes the downward closure of d by E.

Explanation-sets allow a direct access to direct and indirect consequences of a given constraint c. For each $h \notin CL \downarrow (d, R)$, expl($h$) represents any explanation-set for h. Notice that for any $h \in CL \downarrow (d, R)$, expl($h$) does not exist.

Several explanations generally exist for the removal of a given value. [26] shows that a good compromise between precision (small explanation-sets) and ease of computation of explanation-sets is to use the solver-embedded knowledge. Indeed, constraint solvers always know, although it is scarcely explicit, *why* they remove values

from the environments of the variables. By making that knowledge explicit and therefore kind of *tracing* the behavior of the solver, quite precise explanation-sets can be computed. Indeed, explanation-sets are a compact representation of the necessary constraints to achieve a given domain reduction.

A more complete description of the interaction of the constraints responsible for this domain reduction can be introduced through *explanation-trees* which are closely related to actual computation.

According to the solver mechanism, domain reduction must be considered from a dual point of view. Indeed, we are interested in the values which may belong to the solutions, but the solver keeps in the domains values for which it cannot prove that they do not belong to a solution. In other words, it only computes proofs for removed values.

With each local consistency operator considered above, its dual operator (the one removing values) can be associated. Then, these dual operators can be defined by sets of rules. Note that for each operator many such systems of rules can exist, but in general one is more natural to express the notion of local consistency used. Examples for classical notions of consistency are developed in [18]. At first, we need to introduce the notion of deduction rule related to dual of local consistency operators.

A *deduction rule* is a rule $h \leftarrow B$ such that $h \in \mathbb{D}$ and $B \subseteq \mathbb{D}$.

The intended semantics of a deduction rule $h \leftarrow B$ can be presented as follows: if all the elements of B are removed from the environment, then h does not appear in any solution of the CSP and may be removed harmlessly (in the arc-consistency case, the elements of B represent the support set of h).

A set of deduction rules $\mathcal{R}_r$ may be associated with each dual of local consistency operator r. It is intuitively obvious that this is true for arc-consistency enforcement but it has been proved [18] that for any dual of local consistency which boils down to domain reduction it is possible to associate such a set of rules (moreover, it shows that there exists a natural set of rules for classical local consistencies). Note that, in the general case, there may exist several rules with the same head but different bodies.

We consider the set $\cup_{r \in R} \mathcal{R}_r$ of all the deduction rules for all the local consistency operators of R. But the initial environment must also be taken into account in the set of deduction rules: the iteration starts from an environment $d \subseteq \mathbb{D}$; it is therefore necessary to add facts (deduction rules with an empty body) in order to directly deduce the elements of $\overline{d}$: let $\mathcal{R}^d = \{h \leftarrow \emptyset \mid h \in \overline{d}\}$ be this set. We denote by $\mathcal{R}$ the set $(\cup_{r \in R} \mathcal{R}_r) \cup \mathcal{R}^d$.

A *proof tree* with respect to a set of rules $\mathcal{R}$ is a finite tree such that for each node labelled by h, let B be the set of labels of its children, $h \leftarrow B \in \mathcal{R}$.

Proof trees are closely related to the computation of domain reduction. Let $d = d^0, d^1, \ldots, d^i, \ldots$ be an iteration. For each i, if $h \notin d^i$ then h is the root of

a proof tree with respect to $\mathcal{R}$. More generally, $\overline{CL \downarrow (d, R)}$ is the set of the roots of proof trees with respect to $\mathcal{R}$.

Each deduction rule used in a proof tree comes from a packet of deduction rules, either from a packet $\mathcal{R}_r$ defining a local consistency operator r, or from $\mathcal{R}^d$. A set of local consistency operators can be associated with a proof tree as follows.

> Let t be a proof tree. A *set X of local consistency operators associated with t* is such that for each node of t: let h be the label of the node and B the set of labels of its children: either $h \notin d$ (and $B = \emptyset$); or there exists $r \in X, h \leftarrow B \in \mathcal{R}_r$.

Note that there may exist several sets associated with a proof tree. Moreover, each super-set of a set associated with a proof tree is also convenient (R is associated with all proof trees). It is important to recall that the root of a proof tree does not belong to the closure of the initial environment d by the set of local consistency operators R. So there exists an explanation-set for this value.

> If t is a proof tree, then each set of local consistency operators associated with t is an explanation-set for the root of t.

From now on, a proof tree with respect to $\mathcal{R}$ is therefore called an *explanation-tree*. As we just saw, explanation-sets can be computed from explanation-trees.

Let us consider a fixed iteration $d = d^0, d^1, \ldots, d^i, \ldots$ of R with respect to $r^1, r^2, \ldots, r^{i+1}, \ldots$ In order to incrementally define explanation-trees during an iteration, let $(S^i)_{i \in \mathbb{N}}$ be the family recursively defined as (where $\mathrm{cons}(h, T)$ is the tree defined by h is the label of its root and T is the set of its subtrees, and where $\mathrm{root}(\mathrm{cons}(h, T)) = h$):

- $S^0 = \{\mathrm{cons}(h, \emptyset) \mid h \notin d\}$;
- $S^{i+1} = S^i \cup \{\mathrm{cons}(h, T) \mid h \in d^i, T \subseteq S^i, h \leftarrow \{\mathrm{root}(t) \mid t \in T\} \in \mathcal{R}_{r^{i+1}}\}$.

It is important to note that some explanation-trees do not correspond to any iteration, but when a value is removed there always exists an explanation-tree in $\bigcup_i S^i$ for this value removal.

Among the explanation-sets associated with an explanation-tree $t \in S^i$, one is preferred. This explanation-set is denoted by $\mathrm{expl}(t)$ and defined as follows (where $t = \mathrm{cons}(h, T)$):

- if $t \in S^0$ then $\mathrm{expl}(t) = \emptyset$;
- else there exists $i > 0$ such that $t \in S^i \setminus S^{i-1}$, then $\mathrm{expl}(t) = \{r^i\} \cup \bigcup_{t' \in T} \mathrm{expl}(t')$.

In fact, $\mathrm{expl}(t)$ is a $\mathrm{expl}(h)$ previously defined where t is rooted by h.

Obviously explanation-trees are more precise than explanation-sets. An explanation-tree describes the value removal thanks to deduction rules. Each deduction rule comes from a set of deduction rules that defines an operator. The explanation-set just provides these operators.

Note also that in practice explanation-trees can easily be extracted from a trace following the process described in [20].

In the following, we will associate a single explanation-tree, and therefore a single explanation-set, to each element h removed during the computation. This set will be denoted by $\text{expl}(h)$.

4.4 Constraint Retraction

We detail an application of explanations to constraint retraction algorithms [11]. Thanks to explanations, sufficient conditions to ensure the correctness of any incremental constraint retraction algorithms are given.

Dynamic constraint retraction is performed through the three following steps [23, 26]: *disconnecting* (i.e. removing the retracted constraint), *setting back values* (i.e. reintroducing the values removed by the retracted constraint) and *repropagating* (i.e. some of the reintroduced values may be removed by other constraints).

4.4.1 Disconnecting

The first step is to cut the retracted constraints C' from the constraint network. C' needs to be completely disconnected (and therefore will never get propagated again in the future).

Disconnecting a set of constraints C' amounts to remove all their related operators from the current set of active operators. The resulting set of operators is $R^{\text{new}} \subseteq R$, where $R^{\text{new}} = \bigcup_{c \in C \setminus C'} R(c)$. Constraint retraction amounts to computing the closure of d by R^{new}.

4.4.2 Setting Back Values

The second step is to undo the past effects of the retracted constraints: both direct (each time the constraint operators have been applied) and indirect (further consequences of the constraints through operators of other constraints) effects of those constraints. This step results in the enlargement of the environment: values are put back.

Here, we want to benefit from the previous computation of d^i instead of starting a new iteration from d. Thanks to explanation-sets, we know the values of $d \setminus d^i$ which have been removed because of a retracted operator (that is an operator of $R \setminus R^{\text{new}}$). This set of values is defined by $d' = \{h \in d \mid \exists r \in R \setminus R^{\text{new}}, r \in \text{expl}(h)\}$ and must be re-introduced in the domain. Notice that all incremental algorithms for constraint retraction amount to compute an (often strict) super-set of this set. The next result (proof in [11]) ensures that we obtain the same closure if the computation starts from d or from $d^i \cup d'$ (the correctness of all the algorithms which re-introduce a super-set of d'):

$$CL \downarrow (d, R^{\text{new}}) = CL \downarrow (d^i \cup d', R^{\text{new}})$$

4.4.3 Repropagating

Some of the put back values can be removed applying other active operators (i.e. operators associated with non retracted constraints). Those domain reductions need to be performed and propagated as usual. At the end of this process, the system will be in a consistent state. It is exactly the state (of the domains) that would have been obtained if the retracted constraint would not have been introduced into the system.

In practice the iteration is done with respect to a sequence of operators which is dynamically computed thanks to the propagation queue. At the ith step, before setting values back, the set of operators which are in the propagation queue is denoted by R^i. Obviously, the operators of $R^i \cap R^{\text{new}}$ must stay in the propagation queue. The other operators $(R^{\text{new}} \setminus R^i)$ cannot remove any element of d^i, but they may remove an element of d' (the set of re-introduced values). So we have to put some of them back in the propagation queue: the operators of the set $R' = \{r \in R^{\text{new}} \mid \exists h \leftarrow B \in \mathcal{R}_r, h \in d'\}$. The next result (proof in [11]) ensures that the operators which are not in $R^i \cup R'$ do not modify the environment $d^i \cup d'$, so it is useless to put them back into the propagation queue (the correctness of all algorithms which re-introduce a super-set of R' in the propagation queue):

$$\forall r \in R^{\text{new}} \setminus (R^i \cup R'), d^i \cup d' \subseteq r(d^i \cup d').$$

Therefore, by the two previous results, any algorithm which restarts with a propagation queue including $R^i \cup R'$ and an environment including $d^i \cup d'$ is proved correct.

Note that the presented constraint retraction process encompasses both information recording methods and recomputation-based methods. The only difference relies on the way values to set back are determined. The first kind of methods record information to allow an easy computation of values to set back into the environment upon a constraint retraction. [5] and [10] use *justifications*: for each value removal the applied responsible constraint (or operator) is recorded. [15] uses a dependency graph to determine the portion of past computation to be reset upon constraint retraction. More generally, those methods amount to record some dependency information about past computation. A generalization [6] of both previous techniques relies upon the use of explanation-sets.

Note that constraint retraction is useful for dynamic constraint satisfaction problems but also for over-constrained problems. Indeed, users often prefer to have a solution to a relaxed problem than no solution for their problem. In this case, explanation does not only allow to compute a solution to the relaxed problem but it may also help the user choose the constraint to retract [7].

5 CONCLUSION

The paper recalls the notions of closure and fixpoint. When program semantics can be described by some notions of closure or fixpoint, proof trees are suitable

to provide explanations: the computation has proved a result and a proof tree is a declarative explanation of this result.

The paper shows two different domains where these notions apply: Constraint Logic Programming (CLP) and Constraint Satisfaction Problems (CSP).

Obviously, these proof trees are explanations because they can be considered as a declarative view of the trace of a computation and so, they may help understand how the results are obtained. Consequently, debugging is a natural application for explanations. Considering the explanation of an unexpected result it is possible to locate an error in the program (in fact an incorrect rule used to build the explanation, and this rule can be associated to an incorrect piece of program). As an example, the paper presents the declarative error diagnosis of constraint logic programs. The same method has also been investigated for constraint programming in [19]. In this framework, a symptom is a removed value which was expected to belong to a solution and the error is a rule associated with a local consistency operator.

It is interesting to note the difference between the application to CLP and CSP. In CLP, it is easier to understand a wrong (positive) answer because a wrong answer is a logical consequence of the program then there exists a proof of it (which should not exist). In CSP, it is easier to understand a missing answer because explanations are proofs of value removals. A finite domain constraint solver just tries to prove that some values cannot belong to a solution, but it does not prove that remaining values belong to a solution.

In constraint programming, when a constraint is removed from the set of constraints, a first possibility is to restart the computation of the new solutions from the initial domain. But it may be more efficient to benefit from the past computations. This is achieved by a constraint retraction algorithm. The paper has shown how explanations can be used to prove the correctness of a large class of constraint retraction algorithm [11]. In practice, such algorithms use explanations for dynamic problems, for intelligent backtracking during the search, for failure analysis...

The notion of explanation presented in this article has been applied to CLP and CSP. We claim that this general framework should easily apply to every semantics based on rules.

REFERENCES

[1] Aczel, P.: An Introduction to Inductive Definitions. In: J. Barwise (Ed.): Handbook of Mathematical Logic, Vol. 90 of Studies in Logic and the Foundations of Mathematics, chapter C. 7, pp. 739–782. North-Holland Publishing Company, 1977.

[2] Apt, K. R.: The Essence of Constraint Propagation. Theoretical Computer Science, Vol. 221, 1999, Nos. 1–2, pp. 179–210.

[3] Apt, K. R.: Principles of Constraint Programming. Cambridge University Press, 2003.

[4] Benhamou, F.: Heterogeneous Constraint Solving. In: M. Hanus and M. Rodríguez-Artalejo (Eds.): Proceedings of the 5th International Conference on Algebraic and

Logic Programming, ALP 96, Vol. 1139 of Lecture Notes in Computer Science, pp. 62–76. Springer-Verlag, 1996.

[5] Bessière, C.: Arc Consistency in Dynamic Constraint Satisfaction Problems. In: Proceedings of the 9th National Conference on Artificial Intelligence, AAAI 91, Vol. 1, pp. 221–226. AAAI Press, 1991.

[6] Boizumault, P.—Debruyne, R.—Jussien, N.: Maintaining Arc-Consistency within Dynamic Backtracking. In: Proceedings of the Sixth International Conference on Principles and Practice of Constraint Programming, No. 1894 in Lecture Notes in Computer Science, pp. 249–261. Springer-Verlag, 2000.

[7] Boizumault, P.—Jussien, N.: Best-first search for property maintenance in reactive constraints systems. In: J. Maluszynski (Ed.): Proceedings of the 1997 International Symposium on Logic Programming, pp. 339–353. MIT Press, 1997.

[8] Codognet, P.—Diaz, D.: Compiling Constraints in `clp(fd)`. Journal of Logic Programming, Vol. 27, 1996, No. 3, pp. 185–226.

[9] Cousot, P.—Cousot, R.: Automatic Synthesis of Optimal Invariant Assertions Mathematical Foundation. In: Symposium on Artificial Intelligence and Programming Languages, Vol. 12, 1977, No. 8, ACM SIGPLAN Not., pp. 1–12.

[10] Debruyne, R.: Arc-Consistency in Dynamic CSPs Is No More Prohibitive. In 8th Conference on Tools with Artificial Intelligence, pp. 299–306, 1996.

[11] Debruyne, R.—Ferrand, G.—Jussien, N.—Lesaint, W.—Ouis, S.—Tessier, A.: Correctness of Constraint Retraction Algorithms. In: I. Russell and S. Haller (Eds.): Sixteenth International Florida Artificial Intelligence Research Society Conference, pp. 172–176. AAAI Press, 2003.

[12] Dechter, R.: Constraint Processing. Morgan Kaufmann, 2003.

[13] Despeyroux, J.: Proof of Translation in Natural Semantics. In: Symposium on Logic in Computer Science, pp. 193–205. IEEE Computer Society, 1986.

[14] Fages, F.—Fowler, J.—Sola, T.: A Reactive Constraint Logic Programming Scheme. In: L. Sterling (Ed.): Proceedings of the Twelfth International Conference on Logic Programming, pp. 149–163. MIT Press, 1995.

[15] Fages, F.—Fowler, J.—Sola, T.: Experiments in Reactive Constraint Logic Programming. Journal of Logic Programming, Vol. 37, 1998, Nos. 1–3, pp. 185–212.

[16] Ferrand, G.: Error Diagnosis in Logic Programming: An Adaptation of E. Y. Shapiro's method. Journal of Logic Programming, Vol. 4, 1987, pp. 177–198.

[17] Ferrand, G.: The Notions of Symptom and Error in Declarative Diagnosis of Logic Programs. In: P. A. Fritzson (Ed.): Proceedings of the First International Workshop on Automated and Algorithmic Debugging, Vol. 749 of Lecture Notes in Computer Science, pp. 40–57. Springer-Verlag, 1993.

[18] Ferrand, G.—Lesaint, W.—Tessier, A.: Theoretical Foundations of Value Withdrawal Explanations for Domain Reduction. Electronic Notes in Theoretical Computer Science, Vol. 76, 2002.

[19] Ferrand, G.—Lesaint, W.—Tessier, A.: Towards Declarative Diagnosis of Constraint Programs Over Finite Domains. In: M. Ronsse (Ed.): Proceedings of the Fifth International Workshop on Automated Debugging, pp. 159–170, 2003.

[20] Ferrand, G.—Lesaint, W.—Tessier, A.: Explanations to Understand the Trace of a Finite Domain Constraint Solver. In: S. Muñoz-Hernández, J. M. Gómez-Pérez, and P. Hofstedt (Eds.): 14th Workshop on Logic Programming Environments, pp. 19–33, 2004.

[21] Ferrand, G.—Tessier, A.: Positive and Negative Diagnosis for Constraint Logic Programs in Terms of Proof Skeletons. In: M. Kamkar (Ed.): Third International Workshop on Automatic Debugging, Vol. 2, No. 009-12 of Linköping Electronic Articles in Computer and Information Science, ISSN 1401-9841, pp. 141–154, 1997.

[22] Ferrand, G.—Tessier, A.: Correctness and completeness of CLP semantics revisited by (co)-induction. In: O. Ridoux (Ed.): Journées Francophones de Programmation Logique et Programmation par Contraintes, pp. 19–38. HERMES, 1998 (in French).

[23] Georget, Y.—Codognet, P.—Rossi, F.: Constraint Retraction in CLP(FD): Formal Framework and Performance results. Constraints, Vol. 4, 1999, No. 1, pp. 5–42.

[24] Ginsberg, M.: Dynamic Backtracking. Journal of Artificial Intelligence Research. Vol. 1, 1993, pp. 25–46.

[25] Jaffar, J.—Maher, M. J.—Marriott, K.—Stuckey, P. J.: Semantics of Constraint Logic Programs. Journal of Logic Programming, Vol. 37, 1998, Nos. 1–3, pp. 1–46.

[26] Jussien, N.: E-Constraints: Explanation-Based Constraint Programming. In: Workshop on User-Interaction in Constraint Satisfaction, 2001.

[27] Jussien, N.—Ouis, S.: User-Friendly Explanations for Constraint Programming. In: Proceedings of the 11th Workshop on Logic Programming Environments, 2001.

[28] Kahn, G.: Natural semantics. In: F.-J. Brandenburg, G. Vidal-Naquet, and M. Wirsing (Eds.): Annual Symposium on Theoretical Aspects of Computer Science, Vol. 247 of Lecture Notes in Computer Science, pp. 22–39. Springer-Verlag, 1987.

[29] Lloyd, J. W.: Foundations of Logic Programming. Springer-Verlag, 1987.

[30] Plotkin, G. D.: A Structural Approach to Operational Semantics. Technical Report DAIMI FN-19, University of Aarhus, 1981.

[31] Prawitz, D.: Natural Deduction: A Proof Theoretical Study. Almquist & Wiksell, 1965.

[32] Prosser, P.: Hybrid Algorithms for the Constraint Satisfaction Problem. Computational Intelligence, Vol. 9, 1993, No. 3, pp. 268–299.

[33] Schiex, T.—Verfaillie, G.: Nogood Recording for Static and Dynamic Constraint Satisfaction Problems. In: Proceeding of the Fifth IEEE International Conference on Tools with Artificial Intelligence, pp. 48–55, 1993.

[34] Shapiro, E.: Algorithmic Program Debugging. ACM Distinguished Dissertation. MIT Press, 1982.

[35] Tessier, A.: Approach, in Term of Proof Skeletons, of the Semantics and Declarative Error Diagnosis of Constraint Logic Programs. Ph.D. thesis, Université d'Orléans, 1997 (in French).

[36] Tessier, A.—Ferrand, G.: Declarative Diagnosis in the Clp Scheme. In: P. Deransart, M. Hermenegildo, and J. Małuszyński (Eds.): Analysis and Visualization

Tools for Constraint Programming, Vol. 1870 of Lecture Notes in Computer Science, Chapter 5, pp. 151–174. Springer, 2000.

[37] Tsang, E.: Foundations of Constraint Satisfaction. Academic Press, 1993.

Gérard Ferrand, Ph. D. in Mathematics (1982) and Habilitation in Computer Science (1988) was professor at the University of Orléans (France) and is retired since 2005. He remains researcher in the Constraints and Machine Learning project of the LIFO. His research interests are mainly concerned with semantics of programming languages, in particular logic programming and constraint programming, including epistemological aspects.

Willy Lesaint submitted his thesis in computer science in 2003. He is now assistant professor at the University of Orléans (France) in the Constraints and Machine Learning project. He is interested in the concept of explanation for constraint programs and its applications, more particularly constraint retraction and declarative diagnosis.

Alexandre Tessier obtained his Ph. D. in computer science in 1997. He is now in a position of assistant professor. He is a researcher in the Constraints and Machine Learning project of the LIFO (University of Orléans, France). Its domains of research include constraint logic programming and constraint satisfaction problems and its topics of interests are mainly the semantics of programming languages and declarative debugging of programs.