digital

Residuation and Guarded Rules
for Constraint Logic Programming

June 1991   Gert Smolka

12

**Residuation and Guarded Rules
for Constraint Logic Programming**

Gert Smolka

June 1991

Publication Notes

Abstract

A major difficulty with logic programming is combinatorial explosion: since goals are solved with possibly indeterminate (*i.e.*, branching) reductions, the resulting search trees may grow wildly. Constraint logic programming systems try to avoid combinatorial explosion by building in strong determinate (*i.e.*, non-branching) reduction in the form of constraint simplification. In this paper, we present two concepts, residuation and guarded rules, for further strengthening determinate reduction. Both concepts apply to constraint logic programming in general and yield an operational semantics that coincides with the declarative semantics. Residuation is a control strategy giving priority to determinate reductions. Guarded rules are logical consequences of programs adding otherwise unavailable determinate reductions.

Résumé

Une difficulté majeure rencontrée en programmation logique est l'explosion combinatoire : puisque les buts sont résolus selon des réductions potentiellement indéterminées (*i.e.*, branchantes), les arbres de recherches qui en résultent peuvent grossir démesurément. Les systèmes de programmation par logique de contraintes tentent d'éviter l'explosion combinatoire en intégrant des primitives de réductions fortement déterminées (*i.e.*, non branchantes) en forme de simplification de contraintes. Dans ce papier, nous présentons deux concepts, la résiduation et les règles guardées, pour renforcer davantage les réductions déterminées. Ces deux concepts s'appliquent à la programmation logique par contrainte en général et fournissent une sémantique opérationnelle qui coïncide avec la sémantique déclarative. La résiduation est une stratégie de contrôle donnant priorité aux réductions déterminées. Les règles guardées sont des conséquences logiques de programmes ajoutant des réductions déterminées qui ne sont pas autrement disponibles.

Keywords

Acknowledgements

Contents

1   Introduction

A major difficulty with logic programming is combinatorial explosion: since goals are solved with possibly indeterminate (*i.e.*, branching) reductions, the resulting search trees may grow wildly. Constraint logic programming systems [7, 13, 9] try to avoid combinatorial explosion by building in strong determinate (*i.e.*, non-branching) reduction in the form of constraint simplification. In this paper we present two concepts, residuation and guarded rules, for further strengthening determinate reduction. Both concepts apply to constraint logic programming in general and yield an operational semantics that coincides with the declarative semantics.

1.1   Residuation

Residuation[1] is a control strategy for constraint logic programming meant to replace the rigid depth first strategy of Prolog, which amounts to eager generation of usually wrong assumptions. Roughly, residuation makes determinate reduction the rule, and indeterminate reduction the exception. The latter is left to the user's discretion by declaring relations as *generating*. Given a goal, an atom is called *determinate* if reduction with all but possibly one clause defining the atom immediately fails due to constraint simplification. *Residuation* is now the following control strategy:

- given a goal that contains determinate atoms, a determinate atom must be reduced;

- given a goal that contains no determinate atoms, an atom whose relation is declared as *generating* must be reduced.

Thus, the user controls which atoms can be reduced indeterminately by declaring some relations as generating. If no relation is declared generating, indeterminate reduction cannot occur. Even with generating relations, indeterminate reduction can only occur if determinate reduction is not possible. A relation is called *residuating* if it is not declared generating. Given a goal, an atom is called *residuated* if it is not determinate and its relation is residuating. An important feature of the residuation strategy is that goals whose atoms are all residuated are taken as answers. Often such complex answers are fine as they are. For instance, if length is a length predicate for lists, the goal

$$\exists N \left(\text{length}(L, N) \wedge N \leq 47 \right)$$

("L is a list with at most 47 elements") may be a perfectly satisfactory answer. If the user is not satisfied with such a complex answer, he can request indeterminate reduction of a residuated atom.

Residuation is similar to the control strategy of the Andorra model [10, 11], with the difference that residuation performs indeterminate reduction only on atoms whose relation is explicitly declared as generating. The philosophy behind residuation is that for most relations

---

[1]The term residuation was coined by Hassan Aït-Kaci [1, 3, 2] for delaying control schemes.

indeterminate reduction simply does not make sense, and that complex answers are often appropriate.

In the examples of this paper, we will assume a constraint system with trees and linear integer arithmetic.

A length relation for lists can be defined as follows (constraints are written in *italic font*):

$$\mathsf{length}\big(\mathsf{L},\mathsf{N}\big) \quad \leftrightarrow \quad L = nil \wedge N = 0$$
$$\vee \quad \exists \mathsf{H},\mathsf{R},\mathsf{M}\,\big(L = H.R \wedge N > 0 \wedge M = N - 1$$
$$\wedge\ \mathsf{length}\big(\mathsf{R},\mathsf{M}\big)\big).$$

Instead of the conventional definite clause syntax we use definite equivalences. Such an equivalence states explicitly that the relation on the left hand side is completely defined by the right-hand side (we are committed to least model semantics).[2]

Now, given a goal whose constraint is ϕ, an atom $\mathsf{length}\big(\mathsf{L},\mathsf{N}\big)$ in this goal is determinate if either the constraint $\phi \wedge L = nil \wedge N = 0$ simplifies to $\bot$, or the constraint $\phi \wedge \exists H, R\,\big(L = H.R \wedge N > 0\big)$ simplifies to $\bot$, where $\bot$ is the canonical unsatisfiable constraint. Assuming a sufficiently powerful constraint simplifier, the goal $\mathsf{length}\big(\mathsf{X},\mathsf{N}\big) \wedge \mathsf{N} \geq 2$ reduces in two steps determinately to the goal

$$\exists \mathsf{Y},\mathsf{Z},\mathsf{U},\mathsf{M}\,\big(\mathsf{X} = \mathsf{Y}.\mathsf{Z}.\mathsf{U} \wedge \mathsf{M} = \mathsf{N} - 2 \wedge \mathsf{M} \geq 0 \wedge \mathsf{length}\big(\mathsf{U},\mathsf{M}\big)\big),$$

which is an answer if the relation length is residuating. In any case, it would not make sense to reduce this goal further.

Residuation is a simple and powerful alternative to delay primitives such as the delay annotations of IC-Prolog [5], the freeze construct of Prolog II [8], or the wait declarations of MU-Prolog [17]. There are major advantages offered by residuation over these delay primitives.

- Residuation applies to every constraint system (rather than to tree systems only).

- No annotations in clauses are needed—the programmer only decides which relations should be generating.

- Residuation is much more flexible—even if all relations are declared generating, the search space is considerably pruned since determinate reductions are performed first.

An idealized method for solving problems with residuation splits the problem solver into a propagating part and a generating part:

- a predicate $\mathsf{propagate}\big(\mathsf{S}\big)$ that holds if and only if S is a solution of the problem, and that depends only on residuating relations;

---

[2]For the special case of Horn clause programming, the translation from the conventional definite clause syntax to definite equivalences is given by Clark's completion [4].

- a predicate $\mathsf{generate}(\mathsf{S})$ that defines candidates for (partial) solutions and depends on generating relations.

A problem instance is then given as a query:

$$\phi \wedge \mathsf{propagate}(\mathsf{S}) \wedge \mathsf{generate}(\mathsf{S}),$$

where the constraint ϕ describes the particular problem instance. With residuation, the formula $\phi \wedge \mathsf{propagate}(\mathsf{S})$ will be reduced determinately to a constraint propagation network consisting of residuated atoms and a shared constraint. In general, the constraint propagation network alone is too weak to exhibit solutions. Thus, $\mathsf{generate}(\mathsf{S})$ is needed to generate incrementally assumptions about the value of the variable S. As soon as an assumption is made, the constraint propagation network becomes active since atoms that were residuated before can now fire. Typically, most of the generated assumptions will be invalidated immediately by constraint propagation leading to failure. To obtain a feasible search space, two considerations are essential: (1) the careful design of the propagation and generation component, and (2) an expressive underlying constraint system.

1.2   Guarded rules

Guarded rules are logical consequences of the program introducing additional determinate reduction rules. We will see that guarded rules can significantly strengthen the propagation component of a problem solver.

Consider the following definition of list concatenation:

$$\mathsf{app}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \quad \leftrightarrow \quad \mathsf{X} = \mathsf{nil} \wedge \mathsf{Y} = \mathsf{Z}$$
$$| \quad \mathsf{X} = \mathsf{H}.\mathsf{R} \wedge \mathsf{Z} = \mathsf{H}.\mathsf{U} \wedge \mathsf{app}(\mathsf{R}, \mathsf{Y}, \mathsf{U}).$$

It is written using a sugared syntax (indicated by writing $|$ rather than $\vee$), which suppresses existential quantification of auxiliary variables and allows nesting of constraint terms.

With this definition, the goal $\mathsf{app}(\mathsf{X}, \mathsf{Y}, \mathsf{Y})$ is not reduced determinately although it is equivalent to $\mathsf{X} = \mathsf{nil}$. In fact, the relation app satisfies the formula

$$\mathsf{Y} = \mathsf{Z} \rightarrow (\mathsf{app}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \leftrightarrow \mathsf{X} = \mathsf{nil}),$$

which validates the determinate reduction of the atom $\mathsf{app}(\mathsf{X}, \mathsf{Y}, \mathsf{Z})$ to the constraint $\mathsf{X} = \mathsf{nil}$ if the constraint of the goal entails the "guard" $\mathsf{Y} = \mathsf{Z}$.

A *guarded rule* is a formula

$$\phi \rightarrow (A \leftrightarrow G),$$

for convenience written as

$$\phi \,\square\, A \,\triangleright\, G,$$

where ϕ is a constraint (called the *guard*), A is an atom, and G is a goal. A guarded rule is *admissible* if it is valid in every model of the declarative semantics (again, we are committed to

least model semantics). Thus, admissible guarded rules are redundant as far as the declarative semantics is concerned.

The operational semantics of guarded rules is defined as follows. Given a goal G,

$$\exists X\big(\phi \wedge A \wedge R\big),$$

and a guarded rule,

$$\psi \,\square\, A \,\triangleright\, G',$$

the goal G can reduce determinately to:

$$\exists X\big(\phi \wedge G' \wedge R\big)$$

if the constraint ϕ entails the constraint ψ—that is, if the implication $\phi \to \psi$ is valid in every model of the constraint system. Note that $\exists X\big(\phi \wedge G' \wedge R\big)$ is logically equivalent to G in all models of the declarative semantics if the guarded rule is admissible. Moreover, $\exists X\big(\phi \wedge G' \wedge R\big)$ is a goal up to constraint simplification and minor syntactic rearrangement.

Two further admissible guarded rules for app are:

$$Y = \mathsf{nil} \,\square\, \mathsf{app}(X, Y, Z) \,\triangleright\, X = Z \wedge \mathsf{list}(X)$$
$$X = Z \,\square\, \mathsf{app}(X, Y, Z) \,\triangleright\, Y = \mathsf{nil} \wedge \mathsf{list}(X),$$

where the relation list is defined as follows:

$$\mathsf{list}(L) \;\leftrightarrow\; L = \mathsf{nil} \mid L = \mathsf{H.R} \wedge \mathsf{list}(R).$$

Our notion of admissible guarded rules is a new concept that must not be confused with the guarded clauses of committed-choice languages such as Concurrent Prolog [18] or Parlog [6]. In these languages guarded clauses are used to define agents, while in our framework, relations are defined by definite equivalences and admissible guarded rules are logical consequences of the definitions. Moreover, committed-choice languages usually do not have a declarative semantics. Maher [16] has given a declarative semantics for a strongly restricted class of committed-choice languages, where guards must be mutually exclusive. This is usually not the case for guarded rules, as can be seen in the list concatenation example.

Guarded rules have some similarity with the demon predicates of CHIP [9], but are much more general. First, demon predicates in CHIP are defined by guarded rules only, while in our approach the relation is defined independently by clauses. Second, guards in CHIP are restricted to positive tree patterns. Third, in our approach guarded rules can be given for generating relations, while demon predicates in CHIP are residuating by definition. Last but not least, to our knowledge, no declarative semantics for CHIP's demon predicates has yet been even outlined.

In the presence of guarded rules, an atom in a goal is called determinate if it is either determinate as defined before, or reducible using a guarded rule. Residuation is defined as before, except that it now relies on the stronger notion of determinate atoms.

Residuation with guarded rules yields a surprisingly strong constraint propagation mechanism, which we will illustrate with two further examples. Consider the following relational definition of the Boolean "and" function:

$$\mathsf{and}(X, Y, Z) \quad \leftrightarrow \quad X = 1 \wedge Y = Z \wedge \mathsf{bool}(Y)$$
$$\mid \quad X = 0 \wedge Z = 0 \wedge \mathsf{bool}(Y)$$

$$\mathsf{bool}(X) \quad \leftrightarrow \quad X = 1 \mid X = 0.$$

First, note that the definition of and in the presence of residuation already realizes four implicit guarded rules:

$$X \neq 1 \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = 0 \wedge Z = 0 \wedge \mathsf{bool}(Y)$$
$$Y \neq Z \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = 0 \wedge Z = 0 \wedge \mathsf{bool}(Y)$$
$$X \neq 0 \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = 1 \wedge Y = Z \wedge \mathsf{bool}(Y)$$
$$Z \neq 0 \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = 1 \wedge Y = Z \wedge \mathsf{bool}(Y).$$

The second and fourth rule could be optimized since under their guards we have $Y = 1$, but residuation will reduce $\mathsf{bool}(Y)$ anyway to $Y = 1$. By exploiting the symmetry of and with respect to its first two arguments, we obtain the admissible guarded rules:

$$Y \neq 1 \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, Y = 0 \wedge Z = 0 \wedge \mathsf{bool}(X)$$
$$X \neq Z \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = 1 \wedge Y = 0 \wedge Z = 0$$
$$Y \neq 0 \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = Z \wedge Y = 1 \wedge \mathsf{bool}(X).$$

By adding two further admissible guarded rules:

$$X = Y \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, X = Z \wedge \mathsf{bool}(X)$$
$$X \neq Y \,\square \quad \mathsf{and}(X, Y, Z) \,\triangleright\, Z = 0 \wedge \mathsf{bool}(X) \wedge \mathsf{bool}(Y),$$

we obtain an optimal constraint propagation.

In our next example, we consider the problem of solving a crossword puzzle. For this purpose, it will be useful to use a predicate $\mathsf{s}(I, U, J, V)$ to characterize the situation where the I-th letter of the word U is identical to the J-th letter of the word V. Using a predicate *suchthatat*(I, U, X) capturing the fact that the letter X is in position I in the word U, this predicate can be defined as:

$$\mathsf{s}(I, U, J, V) \quad \leftrightarrow \quad I = 1 \wedge U = H.R \wedge \mathsf{at}(J, V, H)$$
$$\mid \quad I > 1 \wedge U = H.R \wedge \mathsf{s}(I - 1, R, J, V)$$

$$\mathsf{at}(I, U, X) \quad \leftrightarrow \quad I = 1 \wedge U = X.R$$
$$\mid \quad I > 1 \wedge U = H.R \wedge \mathsf{at}(I - 1, R, X).$$

Now the goal $\mathsf{s}(2, U, J, V)$ is reduced to:

$$\exists X, Y, W \left(U = X.Y.W \wedge \mathsf{at}(J, V, Y) \right),$$

which makes explicit the fact that the word U consists of at least two characters. However, the symmetric goal $\mathsf{s}(\mathsf{I}, \mathsf{U}, 2, \mathsf{V})$ does not reduce determinately. This can be arranged by providing the following additional admissible guarded rules:

$$\mathsf{J} \leq 1 \,\square\; \mathsf{s}(\mathsf{I}, \mathsf{U}, \mathsf{J}, \mathsf{V}) \;\triangleright\; \exists \mathsf{H}, \mathsf{R}\, \big(\mathsf{J} = 1 \wedge \mathsf{V} = \mathsf{H.R} \wedge \mathsf{at}(\mathsf{I}, \mathsf{U}, \mathsf{H})\big)$$
$$\mathsf{J} \neq 1 \,\square\; \mathsf{s}(\mathsf{I}, \mathsf{U}, \mathsf{J}, \mathsf{V}) \;\triangleright\; \exists \mathsf{H}, \mathsf{R}\, \big(\mathsf{J} > 1 \wedge \mathsf{V} = \mathsf{H.R} \wedge \mathsf{s}(\mathsf{I}, \mathsf{U}, \mathsf{J} - 1, \mathsf{R})\big)$$
$$\neg \exists \mathsf{H}, \mathsf{R}\, \big(\mathsf{V} = \mathsf{H.R}\big) \,\square\; \mathsf{s}(\mathsf{I}, \mathsf{U}, \mathsf{J}, \mathsf{V}) \;\triangleright\; \bot.$$

1.3 Nondeclarative use of guarded rules

So far, only admissible guarded rules have been considered. Namely, those guarded rules that are logical consequences of, and are operationally compatible with, the declarative semantics of a relation defined as an equivalence. However, the operational semantics obtained by residuation and nonadmissible guarded rules is significantly stronger than what can be captured by a conventional declarative semantics. In fact, the object-oriented programming techniques developed for Concurrent Prolog [18] become available if determinate atoms are selected for reduction with a fair strategy.

For instance, an agent that reads two input streams X, Y and merges them into one output stream Z can be defined by four nonadmissible guarded rules:

$$\mathsf{X} = \mathsf{nil} \,\square\; \mathsf{merge}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \;\triangleright\; \mathsf{Y} = \mathsf{Z}$$
$$\mathsf{X} = \mathsf{H.R} \,\square\; \mathsf{merge}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \;\triangleright\; \exists \mathsf{U}\, \big(\mathsf{Z} = \mathsf{H.U} \wedge \mathsf{merge}(\mathsf{R}, \mathsf{Y}, \mathsf{U})\big)$$
$$\mathsf{Y} = \mathsf{nil} \,\square\; \mathsf{merge}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \;\triangleright\; \mathsf{X} = \mathsf{Z}$$
$$\mathsf{Y} = \mathsf{H.R} \,\square\; \mathsf{merge}(\mathsf{X}, \mathsf{Y}, \mathsf{Z}) \;\triangleright\; \exists \mathsf{U}\, \big(\mathsf{Z} = \mathsf{H.U} \wedge \mathsf{merge}(\mathsf{X}, \mathsf{R}, \mathsf{U})\big).$$

Operationally, this merging agent will behave exactly as expected: as soon as a message appears on one of the two input streams, it can fire and put the message on the output stream.

However, it is easy to see that no **merge** relation may be defined with an equivalence such that the given guarded rules be admissible. In this particular **merge** example, one could attempt modeling streams as bags (*i.e.*, lists whose order does not matter) rather than lists, but in general, this would destroy the declarative semantics of stream consumers.

The rest of the paper presents a simple and general framework for declarative constraint logic programming with residuation and admissible guarded rules. The complications of Jaffar and Lassez's framework [14] are avoided by not providing negation as failure.

2 Reduction Systems

The abstract notion of a well-founded reduction system captures important properties of logic programming. It builds on predicate logic in that it takes for granted first-order structures and formulae with the usual connectives and quantifiers. We assume that $\bot$ ("falsity") is a variable-free formula that is invalid in every structure.

A **reduction system** consists of the following:

- a set of formulae called **goals** containing **the trivial goal** $\bot$;

- a set of structures called **models** in which the goals are interpreted;

- a set of equivalences $G \leftrightarrow G_1 \vee \ldots \vee G_n$ called **reductions** such that:

    - G and $G_1, \ldots, G_n$ are goals, and $G \neq \bot$;
    - $G \leftrightarrow G_1 \vee \ldots \vee G_n$ is valid in every model.

A reduction $G \leftrightarrow G_1 \vee \ldots \vee G_n$ **applies** to the goal G and no other goal. Typically, a reduction system contains many reductions with the same left hand side; that is, more than one reduction applies to a goal. A reduction system can be seen as a rewrite system which allows to rewrite a disjunction of goals into an equivalent disjunction of goals by replacing a goal according to a reduction. The idea is to rewrite until no further reduction applies. The reduction systems corresponding to logic programs are nonterminating in general; that is, there are goals from which infinite rewrite derivations issue.

A reduction system can be separated into a **declarative component** given by its goals and models on one hand, and an **operational component** given by its goals and reductions on the other hand.

We say that a goal G **reduces in one step to** G' and write $G \Rightarrow G'$ if there exists a reduction $G \leftrightarrow G_1 \vee \ldots \vee G_n$ such that $G' = G_i$ for some i. We say that a goal G **reduces to** G' if $G \Rightarrow^* G'$, where $\Rightarrow^*$ is the reflexive and transitive closure of $\Rightarrow$.

An **interpretation** is a pair consisting of a model $\mathcal{A}$ and a variable valuation α into $\mathcal{A}$. A **solution of a goal** G is an interpretation $(\mathcal{A}, \alpha)$ such that G is valid in $\mathcal{A}$ under α. A goal is **satisfiable** if it has at least one solution.

An **answer** is a goal to which no reduction applies. Note that $\bot$ is always an answer (the **trivial answer**). An **answer for a goal** G is an answer G' such that $G \Rightarrow^* G'$. A set of answers for a goal G is **complete** if it contains for every solution σ of G an answer G' such that σ is a solution of G'.

The **computational service** provided by a reduction system is that of **solving of goals**; that is, the enumeration of a complete set of answers for a given goal. The declarative component of a reduction system specifies a class of problems where every goal corresponds to a particular problem instance and the solutions of the goal are the solutions of the problem instance. The operational component of a reduction system specifies a method for solving problem instances, where solving means to enumerate a complete set of answers.

A reduction system is **well-founded** if there exists a well-founded ordering on pairs of goals and interpretations such that for every reduction $G \leftrightarrow G_1 \vee \ldots \vee G_n$ and every solution σ of G there exist an $i = 1, \ldots, n$ such that $(G, \sigma) > (G_i, \sigma)$ and σ is a solution of G_i. A well-founded reduction system has two important properties:

- every goal has a complete set of answers;

- a complete set of answers for a goal G can be enumerated as follows: if no reduction applies to G, then $\{G\}$ is a complete set of answers; otherwise, choose arbitrarily any reduction $G \leftrightarrow G_1 \vee \ldots \vee G_n$ and solve the goals $G_1, \ldots, G_n$ in parallel.

We will see that every Horn clause program yields a well-founded reduction system.

A reduction is **determinate** if its right hand side is a single goal. We say that G **reduces determinately to** G' if G reduces to G' using only determinate reductions. If G reduces determinately to G', then G and G' have exactly the same solutions. A reduction system is **determinate** if it has only determinate reductions. Note that, in well-founded and determinate reduction systems, there exist no infinite reduction chains $G \Rightarrow G_1 \Rightarrow G_2 \Rightarrow \cdots$ issuing from a satisfiable goal G.

A reduction system is **terminating** if there exists no infinite chain $G \Rightarrow G_1 \Rightarrow G_2 \Rightarrow \cdots$ of reduction steps. Note that a terminating reduction system is always well-founded, but not vice versa. Even a well-founded and determinate reduction system may not terminate on unsatisfiable goals.

3   Constraint Systems

A **constraint system** is a terminating and determinate reduction system whose goals are closed under conjunction, existential quantification, and variable renaming. In a constraint system we call the goals **constraints**, the answers **simplified constraints**, and the process of reducing a constraint to a simplified constraint **constraint simplification**. Note that in a constraint system, it is possible to compute a simplified constraint for every constraint. Moreover, if a constraint simplifies to the trivial constraint $\perp$, it must be unsatisfiable. A constraint system is called **complete** if a constraint is unsatisfiable if and only if it simplifies to $\perp$. Thus, constraint simplification in a complete constraint system is a decision algorithm for satisfiability of constraints.

The operational component of a constraint system is called a **constraint simplifier**, and the operational component of a complete constraint system is called a **constraint solver**. Our framework does not require that the underlying constraint system be complete. In practice, given a set of constraints with the corresponding models, one may prefer an incomplete constraint simplifier since a (tractable) constraint solver may not exist.

Our notion of a constraint system is deliberately very general: every set of formulae with a corresponding class of models can be seen as a constraint system if we provide no reductions and close the formulae under conjunction, existential quantification and variable renaming. Of course, such trivial constraint systems provide no computational service and are uninteresting in practice.

4 Definite Construction

We now introduce definite construction, a paradigm which is the principle underlying constraint logic programming. We obtain a very simple framework for constraint logic programming with residuation. The two theorems given in this section are consequences of the results in [12].

We assume that a constraint system and a set of **definite relation symbols** are given, where the definite relation symbols take a fixed number of arguments and do not occur in the constraint system.

An **atom** takes the form $r(x_1, \ldots, x_n)$, where r is a definite relation symbol taking n arguments and $x_1, \ldots, x_n$ are pairwise distinct variables. A **definite goal** takes the form

$$\exists X \left(\phi \wedge R \right),$$

where X is a possibly empty set of existentially quantified variables, ϕ is a constraint, and R is a possibly empty conjunction of atoms. Note that the definite goals containing no atoms are exactly the constraints. A **definite equivalence** takes the form

$$A \leftrightarrow G_1 \vee \ldots \vee G_n,$$

where A is an atom and $G_1, \ldots, G_n$ are definite goals called the **clauses** of A. A **definite specification** is a set of definite equivalences containing for every definite relation symbol r exactly one equivalence with r appearing at the left hand side.

In the following, we assume that a definite specification is given. Moreover, we assume that ϕ and ψ range over constraints, A over atoms, R over possibly empty conjunctions of atoms, and G over definite goals. We will construct a reduction system for definite goals by defining definite models (the declarative semantics) and definite reductions (the operational semantics).

For convenience, we will often refer to definite goals simply as goals.

A **definite structure** is a structure that can be obtained from a model of the constraint system by adding interpretations for the definite relation symbols. Definite structures are partially ordered as follows: $\mathcal{A} \leq \mathcal{B}$ iff $\mathcal{A}$ and $\mathcal{B}$ extend the same constraint model and $r^{\mathcal{A}} \subseteq r^{\mathcal{B}}$ for every definite relation symbol r. A **definite quasi-model** is a definite structure that is a model of the definite specification. A **definite model** is a minimal definite quasi-model. The following theorem validates our declarative semantics.

Theorem 1 *For every model of the constraint system there exists exactly one definite model extending it.*

Next, we define the operational semantics. We assume that the order in which atoms are written in a definite goal does not matter.

An equivalence $G \leftrightarrow D$ is a **definite reduction** iff the following conditions are satisfied:

- $G = \exists X (\phi \wedge A \wedge R)$ is a definite goal;

- $A \leftrightarrow \bigvee_{i=1}^{n} \exists Y_i (\phi_i \wedge R_i)$ is obtained from a definite equivalence of the definite specification by variable renaming such that only the variables in A are shared with G;

- obtain for every clause $\exists Y_i (\phi_i \wedge R_i)$ of the definite equivalence the goal

$$G_i := \begin{cases} \bot & \text{if } \phi \wedge \phi_i \text{ simplifies to } \bot, \\ \exists X \cup Y_i (\psi_i \wedge R_i \wedge R) & \text{if } \phi \wedge \phi_i \text{ simplifies to } \psi_i \neq \bot; \end{cases}$$

- D is the disjunction of all $G_i \neq \bot$; if all G_i's are $\bot$, then $D = \bot$.

Note that our definition of definite reductions corresponds exactly to SLD-resolution [15] for the special case of Horn clauses.

Given a constraint system $\mathcal{C}$ and a definite specification $\mathcal{D}$ over $\mathcal{C}$, we define $\mathcal{R}(\mathcal{C}, \mathcal{D})$ as the reduction system whose goals are the corresponding definite goals, whose models are the corresponding definite models, and whose reductions are the corresponding definite reductions *together* with the reductions of the constraint system $\mathcal{C}$. It is easy to verify that $\mathcal{R}(\mathcal{C}, \mathcal{D})$ is in fact a reduction system.

Theorem 2 $\mathcal{R}(\mathcal{C}, \mathcal{D})$ *is a well-founded reduction system whose answers are exactly the simplified constraints.*

It is now straightforward to build in **residuation**. We only have to discard unnecessary indeterminate reductions:

- discard all indeterminate reductions for goals that do have determinate reductions;

- discard all indeterminate reductions obtained by reducing a residuating atom (an atom whose relation is not declared generating).

Let us call the obtained reduction system $\mathcal{R}^*(\mathcal{C}, \mathcal{D}, \mathcal{G})$, where $\mathcal{G}$ is the set of generating relation symbols. Clearly, $\mathcal{R}^*(\mathcal{C}, \mathcal{D}, \mathcal{G})$ is still a well-founded reduction system. Moreover, let $\mathcal{R}^*(\mathcal{C}, \mathcal{D})$ be the reduction system $\mathcal{R}^*(\mathcal{C}, \mathcal{D}, \mathcal{G})$ where all definite relations are declared generating. Then, $\mathcal{R}^*(\mathcal{C}, \mathcal{D})$ is well-founded and has again exactly the simplified constraints as answers (this follows immediately from the above theorem). The important difference between $\mathcal{R}(\mathcal{C}, \mathcal{D})$ and $\mathcal{R}^*(\mathcal{C}, \mathcal{D})$ is that $\mathcal{R}^*(\mathcal{C}, \mathcal{D})$ has significantly smaller search spaces (even for the case of Horn clauses), a fact that has only been realized recently in the Andorra model [10, 11].

5   Guarded Rules

Let a constraint system $\mathcal{C}$ and a definite specification $\mathcal{D}$ over $\mathcal{C}$ be given. A *guarded rule* is a formula:

$$\phi \to (A \leftrightarrow G),$$

where ϕ is a constraint (called the *guard*), A is an atom, and G is a definite goal. A guarded rule is *admissible* if it is valid in every definite model.

Let $\mathcal{F}$ be a set of admissible guarded rules. Then $G \leftrightarrow G'$ is called a *forward reduction* iff the following conditions are satisfied:

- $G = \exists X\big(\phi \wedge A \wedge R\big)$ is a definite goal;

- $\psi \rightarrow \big(A \leftrightarrow \exists Y(\phi' \wedge R')\big)$ is obtained from a guarded rule in $\mathcal{F}$ by variable renaming such that only the variables in the atom A are shared with G;

- $\phi \wedge \neg\psi$ is a constraint that simplifies to $\perp$;

- $G' = \begin{cases} \perp & \text{if } \phi \wedge \phi' \text{ simplifies to } \perp, \\ \exists X \cup Y\big(\phi'' \wedge R' \wedge R\big) & \text{if } \phi \wedge \phi' \text{ simplifies to } \phi'' \neq \perp. \end{cases}$

If $\phi \wedge \neg\psi$ simplifies to $\perp$, then ϕ entails ψ; that is, the implication $\phi \rightarrow \psi$ is valid in every model of the constraint system. Moreover, if the constraint system is complete, then $\phi \wedge \neg\psi$ simplifies to $\perp$ if and only if ϕ entails ψ.

The reduction system $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ is obtained from $\mathcal{R}\big(\mathcal{C}, \mathcal{D}\big)$ by adding the forward reductions defined by the admissible guarded rules in $\mathcal{F}$. It is easy to verify that $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ is in fact a reduction system, and that every goal of $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ has a complete set of answers.

In general, $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ is not well-founded; consider, for instance, the admissible guarded rule $\neg\perp \rightarrow \big(A \leftrightarrow A\big)$. It is the responsibility of the programmer to design the guarded rules in $\mathcal{F}$ such that $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ be well-founded. Further research is necessary to find good sufficient conditions for the well-foundedness of $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$.

Residuation for $\mathcal{R}\big(\mathcal{C}, \mathcal{D}, \mathcal{F}\big)$ is defined as before.

6 Conclusions

Residuation is a control strategy for CLP meant to replace the rigid depth first strategy of Prolog, which amounts to eager generation of usually wrong assumptions. Residuation makes determinate reduction the rule, and indeterminate reduction the exception. Consequently, residuation may produce complex answers containing residuated atoms.

Guarded rules are logical consequences of programs adding otherwise unavailable determinate reductions. Together with residuation guarded rules yield a general and powerful constraint propagation mechanism resulting in drastically smaller search spaces.

Residuation overcomes the strictly sequential computation strategy of Prolog. With residuation every determinate atom can be reduced next, which amounts to multiple threads of computation if a fair selection strategy is used.

The operational semantics of residuation and nonadmissible guarded rules is more expressive

than what can be captured by classical declarative semantics. In fact, the object-oriented programming techniques developed for Concurrent Prolog [18] can be expressed.

Topics for further research include: investigation of abstract incrementality properties ensuring efficient implementation if satisfied by constraint simplifiers; design of an abstract machine separating control from constraint simplification; and investigation of parallel reduction strategies.

References

1. Hassan Aït-Kaci and Roger Nasr. Residuation: A new paradigm for integrating logic and functional programming. MCC Technical Report Number AI-359-86, Microelectronics and Computer Technology Corporation, Austin, TX (October 1986).

2. Hassan Aït-Kaci and Roger Nasr. Integrating logic and functional programming. *Lisp and Symbolic Computation*, 2:51–89 (1989).

3. Hassan Aït-Kaci, Roger Nasr, and Patrick Lincoln. Le Fun: Logic, equations, and Functions. In *Proceedings of the Symposium on Logic Programming (San Francisco, CA)*, pages 17–23, Washington, DC (1987). IEEE, Computer Society Press.

4. K. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Databases*, pages 293–322. Plenum Press, New York, NY (1978).

5. K. L. Clark and F. G. McCabe. The control facilities of IC-PROLOG. In D. Mitchie, editor, *Expert Systems in the Micro-Electronic Age*. Edinburgh University Press, Edinburgh, Scotland (1979).

6. K.L. Clark and S. Gregory. PARLOG: Parallel programming in logic. *ACM Transactions on Programming Languages and Systems*, 8(1):1–49 (1986).

7. A. Colmerauer. An introduction to PROLOG III. *Communications of the ACM*, pages 70–90 (July 1990).

8. A. Colmerauer, H. Kanoui, and M. Van Caneghem. Prolog, theoretical principles and current trends. *Technology and Science of Informatics*, 2(4):255–292 (1983).

9. M. Dincbas, P. Van Hentenryck, H. Simonis, A. Aggoun, T. Graf, and F. Berthier. The constraint logic programming language CHIP. In *Proceedings of the International Conference on Fifth Generation Computer Systems FGCS-88*, pages 693–702, Tokyo, Japan (December 1988).

10. S. Haridi. A logic programming language based on the Andorra model. *New Generation Computing*, 7:109–125 (1990).

11. S. Haridi and S. Janson. Kernel Andorra Prolog and its computation model. In D.H.D. Warren and P. Szeredi, editors, *Logic Programming, Proceedings of the 7th International Conference*, pages 31–48, Cambridge, MA (June 1990). The MIT Press.

12. Markus Höhfeld and Gert Smolka. Definite relations over constraint languages. LILOG Report 53, IWBS, IBM Deutschland, Postfach 80 08 80, 7000 Stuttgart 80, Germany (October 1988). To appear in the Journal of Logic Programming.

13. J. Jaffar and S. Michaylov. Methodology and implementation of a CLP system. In J.-L. Lassez, editor, *Proceedings of the 4th International Conference on Logic Programming*, Cambridge, MA (1987). The MIT Press.

14. Joxan Jaffar and Jean-Louis Lassez. Constraint logic programming. In *Proceedings of the 14th ACM Symposium on Principles of Programming Languages*, pages 111–119, Munich, Germany (January 1987).

15. J. W. Lloyd. *Foundations of Logic Programming*. Symbolic Computation. Springer-Verlag, Berlin, Germany, second, extended edition (1987).

16. Michael J. Maher. Logic semantics for a class of committed-choice programs. In Jean-Louis Lassez, editor, *Logic Programming, Proceedings of the Fourth International Conference*, pages 858–876, Cambridge, MA (1987). The MIT Press.

17. L. Naish. Automating control for logic programs. *Journal of Logic Programming*, 3:167–183 (1985).

18. E. Shapiro and A. Takeuchi. Object oriented programming in Concurrent Prolog. *New Generation Computing*, 1:24–48 (1983).

PRL Research Reports

The following documents may be ordered by regular mail from:

Research Report 1: *Incremental Computation of Planar Maps*. Michel Gangnet, Jean-Claude Hervé, Thierry Pudet, and Jean-Manuel Van Thong. May 1989.

Research Report 2: *BigNum: A Portable and Efficient Package for Arbitrary-Precision Arithmetic*. Bernard Serpette, Jean Vuillemin, and Jean-Claude Hervé. May 1989.

Research Report 3: *Introduction to Programmable Active Memories*. Patrice Bertin, Didier Roncin, and Jean Vuillemin. June 1989.

Research Report 4: *Compiling Pattern Matching by Term Decomposition*. Laurence Puel and Ascánder Suárez. January 1990.

Research Report 5: *The WAM: A (Real) Tutorial*. Hassan Aït-Kaci. January 1990.

Research Report 6: *Binary Periodic Synchronizing Sequences*. Marcin Skubiszewski. May 1991.

Research Report 7: *The Siphon: Managing Distant Replicated Repositories*. Francis J. Prusker and Edward P. Wobber. May 1991.

Research Report 8: *Constructive Logics. Part I: A Tutorial on Proof Systems and Typed λ-Calculi*. Jean Gallier. May 1991.

Research Report 9: *Constructive Logics. Part II: Linear Logic and Proof Nets*. Jean Gallier. May 1991.

Research Report 10: *Pattern Matching in Order-Sorted Languages*. Delia Kesner. May 1991.

Research Report 11: *Towards a Meaning of LIFE*. Hassan Aït-Kaci and Andreas Podelski. June 1991 (Revised, October 1992).

Research Report 12: *Residuation and Guarded Rules for Constraint Logic Programming*. Gert Smolka. June 1991.

Research Report 13: *Functions as Passive Constraints in LIFE*. Hassan Aït-Kaci and Andreas Podelski. June 1991 (Revised, November 1992).

Research Report 14: *Automatic Motion Planning for Complex Articulated Bodies*. Jérôme Barraquand. June 1991.

Research Report 15: *A Hardware Implementation of Pure Esterel*. Gérard Berry. July 1991.

Research Report 16: *Contribution à la Résolution Numérique des Équations de Laplace et de la Chaleur*. Jean Vuillemin. February 1992.

Research Report 17: *Inferring Graphical Constraints with Rockit*. Solange Karsenty, James A. Landay, and Chris Weikart. March 1992.

Research Report 18: *Abstract Interpretation by Dynamic Partitioning*. François Bourdoncle. March 1992.

Research Report 19: *Measuring System Performance with Reprogrammable Hardware*. Mark Shand. August 1992.

Research Report 20: *A Feature Constraint System for Logic Programming with Entailment*. Hassan Aït-Kaci, Andreas Podelski, and Gert Smolka. November 1992.

Research Report 21: *The Genericity Theorem and the notion of Parametricity in the Polymorphic λ-calculus*. Giuseppe Longo, Kathleen Milsted, and Sergei Soloviev. January 1993.

Research Report 22: *Sémantiques des langages impératifs d'ordre supérieur et interprétation abstraite*. François Bourdoncle. January 1993.

Research Report 23: *Dessin à main levée et courbes de Bézier : comparaison des algorithmes de subdivision, modélisation des épaisseurs variables*. Thierry Pudet. January 1993.

digital