TLP **5** (4 & 5): 419–440, 2005. © 2005 Cambridge University Press
doi:10.1017/S147106840500236X Printed in the United Kingdom

A CHR-based implementation of known arc-consistency

MARCO ALBERTI, MARCO GAVANELLI, EVELINA LAMMA
Dipartimento di Ingegneria, Universitá degli Studi di Ferrara, Ferrara, Italy
(*e-mail:* malberti@ing.unife.it)

PAOLA MELLO, MICHELA MILANO
Dipartimento di Elettronica, Informatica e Sistemistica, Universitá degli Studi di Bologna, Bologna, Italy

submitted 31 August 2002; revised 13 November 2003, 30 June 2004; accepted 9 August 2004

Abstract

In classical CLP(*FD*) systems, domains of variables are completely known at the beginning of the constraint propagation process. However, in systems interacting with an external environment, acquiring the whole domains of variables before the beginning of constraint propagation may cause waste of computation time, or even obsolescence of the acquired data at the time of use. For such cases, the Interactive Constraint Satisfaction Problem (ICSP) model has been proposed (Cucchiara *et al.* 1999a) as an extension of the CSP model, to make it possible to start constraint propagation even when domains are not fully known, performing acquisition of domain elements only when necessary, and without the need for restarting the propagation after every acquisition. In this paper, we show how a solver for the two sorted CLP language, defined in previous work (Gavanelli *et al.* 2005) to express ICSPs, has been implemented in the Constraint Handling Rules (CHR) language, a declarative language particularly suitable for high level implementation of constraint solvers.

KEYWORDS: constraints, known arc-consistency, constraint handling rules

1 Introduction

Constraint Logic Programming on Finite Domains (CLP(*FD*)) represents one of the most successful implementations of declarative languages. By means of constraints, the user can give the specifications of a combinatorial problem and possibly solve it, exploiting efficient propagation algorithms. CLP(*FD*) languages have been successfully used for solving a variety of industrial and academic problems. However, in some constraint problems, where domain elements need to be acquired, it may not be wise to perform the acquisition of the whole domains of variables before the beginning of the constraint propagation process. For instance, in configuration problems (Mailharro 1998; ILOG 1999) domain elements represent components, which have to be synthesized before being used. The set of components is not known beforehand, and sometimes even the size of the set cannot be estimated.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

Often, a minimization of the set of components is required, thus the constraint solver produces a new component only when it is strictly necessary.

In systems that need to interact with an external environment, domain elements can be produced by an acquisition system that retrieves information about the outer world. An example is given by Faltings and Macho-Gonzalez (2003) where Internet applications are faced and obviously not all the information can be computed before starting the constraint satisfaction process. As another example, consider a visual search system (Cucchiara *et al.* 1999b) where domain elements are basic visual features (like segments, points, or surface patches) extracted from the image. In a classical CLP(*FD*) computation, all domain values must be known when defining the variables, so all the possible visual features would have to be extracted before starting the visual search process, even if only a small subset of them will be actually used. The synthesis of visual features is usually very time consuming, because the information encoded with signals must be converted into symbolic form. Thus, the extraction of domain elements that will not be used can result in a significant waste of computation time. Also, in systems that interact with an evolving environment, full acquisition of all the domain elements is not wise (Barruffi *et al.* 1999). In fact, if all the possible information is acquired beforehand, some of the information might be obsolete at the end of the acquisition.

For all these reasons, a new model called Interactive Constraint Satisfaction Problem (ICSP) has been proposed (Cucchiara *et al.* 1999a) as an extension of the widely used Constraint Satisfaction Problem (CSP) model. In an ICSP, domains consist of a known part, containing the available elements, plus a variable that semantically represents a set of values that could be added to the domain in the future. In a sense, in an ICSP, domains can be considered as streams of information from one system to the constraint solver. Constraint propagation can be performed even when the domains are not completely known, and domain values can be requested from an acquisition system during constraint propagation; in other words, constraint propagation and value acquisition *interact* (thus *Interactive* in the name of the framework) and are interleaved, whereas in classical CSP frameworks domain elements are completely known before the beginning of the propagation. In this way, the acquisition system can possibly extract only elements consistent with the imposed constraints, thus focusing the attention only on significant data. Various propagation algorithms have been proposed (Cucchiara *et al.* 2001) for exploiting the available information and acquiring new domain values only when strictly necessary. Reducing the number of extracted elements can provide a notable speedup (Cucchiara *et al.* 1999a).

Gavanelli *et al.* (2005) describe a corresponding CLP language. Our language is two sorted. The first sort is the classical sort on Finite Domains (*FD*). The second sort, called $\mathscr{I}$-Set, is based on a structure similar to *streams*, and represents domains of the *FD* variables. From the constraints on the *FD* sort, the system can start propagation before having full knowledge of domain elements. Each element will be inserted *on demand* in the domain, without having to restart constraint propagation from scratch. Moreover, constraints can be imposed on domains, thus helping the

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

user defining I-Sets[1] declaratively. In this paper, we present an implementation of the two-sorted language in Constraint Handling Rules (CHR). CHR (Frühwirth 1998) is a declarative language for defining new constraint solvers at a very high level. CHR can be used for rapid prototyping, and has proven effective in various real life applications. The purpose of this paper is to show how Constraint Handling Rules can be effectively used to implement the solver for our two sorted ICSP-based language. In previous work (Cucchiara *et al.* 1999a), the propagation algorithms have been proposed and separately implemented, but we did not describe the full implementation of the solver for ICSP problems. The algorithms were implemented using non fully declarative constructs (e.g. metaterms with destructive assignment). The high level, declarative encoding in CHR consists of a solver for the $\mathscr{I}$-Set sort, a solver for the *FD* sort, and an interface between them designed to exploit the advantages of the ICSP model in systems interacting with external acquisition modules.

Of course, other aspects in the model of a problem, besides domain elements, could be unknown: in a CSP there could be unknown variables or unknown constraints. Typically, in all Constraint Programming systems new constraints can be easily added, while removal of constraints is more complex (Dechter and Dechter 1988). The addition of variables has been taken into account by Dynamic CSP models (Mittal and Falkenhainer 1990). Our work is focussed on unknown domain elements, and proposes an interaction based on the acquisition, from an external system, of domain elements.

The rest of the paper is organized as follows. The declarative and operational semantics of the language defined by Gavanelli *et al.* (2005) are briefly recalled in Section 2. In Section 3 we describe the architecture of the language from an implementation viewpoint, and in Section 4 we show how it is implemented in the CHR language. Discussion of related work (including a detailed comparison with Mailharro (1998), which is very related to our work from the operational viewpoint) and conclusions follow.

2 Syntax and semantics

The language defined by Gavanelli *et al.* (2005) is based on a two sorted CLP, where the first sort is the classical sort on Finite Domains (*FD*) and the second is the sort on $\mathscr{I}$-Set. I-Sets are used both as domains for *FD* variables and as communication channels with an external source providing elements.

In this section, we briefly recall the syntax and semantics of the language.

In the following, we comply to the conventions in Jaffar *et al.* (1998). In particular, every constraint domain $\mathscr{C}$ (where $\mathscr{C}$ can be *FD*, $\mathscr{I}$-Set, or *FD*+$\mathscr{I}$-Set) contains: the constraint domain signature $\Sigma_{\mathscr{C}}$, the class of constraints $\mathscr{L}_{\mathscr{C}}$ (a set of first-order Σ-formulas), the domain of computation $\mathscr{D}_{\mathscr{C}}$ (a Σ-structure that is the intended

[1] In this article, the $\mathscr{I}$-Set notation (with calligraphic $\mathscr{I}$) will denote the sort, while the notation I-Set (with non-calligraphic I) will denote a particular I-Set.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

interpretation of constraints), the constraint theory $\mathscr{T}_{\mathscr{C}}$ (a Σ-theory that describes the logical semantics of the constraints), and the solver $solv_{\mathscr{C}}$.

2.1 The $\mathscr{I}$-Set sort

The $\mathscr{I}$-Set sort is meant to provide domains for variables in the *FD* sort. Domains are thus considered as first-class objects, and they can be declaratively defined by means of constraints. In a sense, they can be considered as *streams*, but they are intrinsically non-ordered, and do not contain repeated elements. Declaratively, I-Sets are sets; thus unification and constraints should consider $\mathscr{I}$-Set terms modulo the set theory (Dovier *et al.* 1996): $\{A|\{A|B\}\} = \{A|B\}$, $\{A|\{B|C\}\} = \{B|\{A|C\}\}$, which states that sets do not contain repeated elements and order is not important. Non-ground elements are forbidden in an I-Set; this restriction can be exploited for more efficient propagation algorithms. In fact, we represent an I-Set as the union of a set of ground elements (which we name the *known part* of the I-Set) and a variable representing its *unknown part*.

In CLP($\mathscr{I}$-Set), the constraint domain signature, $\Sigma_{\mathscr{I}\text{-Set}}$, contains the following constraints:

- $s\text{-}member(E, S) \Leftrightarrow E \in S$,
- $union(A, B, C) \Leftrightarrow A \cup B = C$,
- $intersection(A, B, C) \Leftrightarrow A \cap B = C$,
- $difference(A, B, C) \Leftrightarrow A \setminus B = C$,
- $inclusion(A, B) \Leftrightarrow A \subseteq B$,

where E represents a ground term, and A, B, and C represent I-Sets.

The operational semantics of the $\mathscr{I}$-Set sort is defined in terms of *state* of I-Sets, *primitives* used to check and modify the state, and *events* over I-Sets.

The state of an I-Set is defined by its *known part*, i.e. the *set* of the elements which are known to belong to the I-Set (as opposed to its *unknown part*, representing the elements which have not yet been acquired for the I-Set), and its *open-closed* condition, i.e. an I-Set is *open* if new elements can be inserted into it, *closed* otherwise.

A convenient notation to express the state of an I-Set is one based on Prolog-like lists. An I-Set is represented by a structure S defined by $S ::= \{\}$, or $S ::= \{T|S\}$, or $S ::= V$, where T is a ground term and V is a variable.

The known part of an I-Set is the set of all the ground elements in the list representing it; an I-Set is closed if its continuation (tail) is ground, open otherwise. For example, I-Set $\{1, 2, 3, 4|T\}$ is open, and its known part is the set $\{1, 2, 3, 4\}$; I-Set $\{1, 2, 3, 4\}$ has the same known part, but it is closed.

We have primitives to modify and check the state of an I-Set. Two primitives are used to modify the state, namely:

- *ensure_member(Element,Iset)*: enforces *Element* to be a member of the known part of *Iset*, possibly adding it if it is not already member, or failing if it is not already member and the I-Set is closed;

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

- *close(Iset)*: closes *Iset*; after execution of this primitive, no new elements can be added to the I-Set.

The following primitives are used to check the state of an I-Set:

- *known(Iset,KnownPart)*: *KnownPart* is the list of elements in the known part of *Iset*;
- *is_closed(Iset)*: checks if the *Iset* is closed.

An *event* is a notification of how the status of the computation has been modified, which may be relevant for the rest of the computation and is to be processed properly. The semantics of an event is defined by rules specifying its interactions with the constraints in the store. It is quite apparent that the concept of event finds a natural representation as a CHR constraint; nevertheless, we prefer, in defining the operational semantics of the language, to keep events and proper $\mathscr{I}$-Set constraints distinct. The concept of event makes it possible to express the semantics of $\mathscr{I}$-Set constraints with simple (CHR-like) rules. For example, this rule is all that is needed to define the fact that, in the *inclusion*/2 constraint, all elements in the first I-Set also appear in the second:

$$inserted(Element, Iset\,1), inclusion(Iset\,1, Iset\,2) \Longrightarrow \\ ensure_member(Element, Iset\,2)$$

This rule simply states that, if the *inserted(Element,Iset1)* event is raised (i.e., if *Element* has been inserted into *Iset1*) and *Iset1* is known to be included in *Iset2*, the propagation process must make sure that the element is also present in *Iset2*. This may imply the insertion of *Element* into *Iset2*, with a subsequent *inserted(Element,Iset2)* event, if *Iset2* is open and *Element* is not already a member, or a failure, if *Iset2* is closed and *Element* is not already a member.

Propagation of closure can also be managed with ease. For instance, the rule

$$\begin{array}{l} closed(Iset\,2), inclusion(Iset\,1, Iset\,2), \\ known(Iset\,1, K\,1), known(Iset\,2), K\,2 \Longrightarrow \\ permutation(K\,1, K\,2)| \\ close(Iset\,1) \end{array} \tag{1}$$

states that if $Iset\,1 \subseteq Iset\,2$, *Iset*2 is closed (as indicated by the *closed(Iset2)* event), and the known elements in *Iset*1 are all the known elements in *Iset*2, then also *Iset*1 is closed (by the *close(Iset1)* primitive).

2.2 *The FD sort*

The *FD* sort shares the same declarative semantics of the classical *FD* sort. Thus, the usual constraints in CLP(*FD*) are considered (arithmetic, relational constraints plus user-defined constraints). We suppose that the symbols $<, \leqslant, +, -, \times, \ldots$ belong to Σ_{FD} and are interpreted as usual.

Since we want cope with incompletely specified variable domains avoiding useless value acquisition, we use a constraint propagation based on the available knowledge, when domains are still partially specified. For this reason, we proposed (Gavanelli

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

et al. 2005) an extension, for the partially known case, of the concept of consistency, called *known consistency*. In this paper, we provide only the definition of node and arc-consistency; the extension to higher degrees of consistency is straightforward.

Definition 2.1
A unary constraint $c(X_i)$ is *known node-consistent* iff

$$\forall v_i \in K_i^c,\ v_i \in c(X_i),$$

where K_i^c is the known part of the domain of X_i. A binary constraint $c(X_i, X_j)$ is *known arc-consistent* iff

$$\forall v_i \in K_i^c,\ \exists v_j \in K_j^c\ s \cdot t \cdot\ (v_i, v_j) \in c(X_i, X_j),$$

where K_i^c and K_j^c are the known parts of the domains of X_i and X_j, respectively. A constraint network is *known arc-consistent (KAC)* iff all unary constraints are known node-consistent and all binary constraints are known arc-consistent.

The following proposition shows the link between *known arc-consistency* and *arc-consistency* (Mackworth 1977). The proof can be found in Gavanelli *et al.* (2005).

Proposition 1
Every algorithm achieving KAC (i.e. any algorithm that computes an equivalent problem that is KAC) and that ensures at least a known element in each variable domain is able to detect inconsistency in the same instances as an algorithm achieving AC.

In other words, if there exists an arc-consistent sub-domain, then there exists a maximal arc-consistent sub-domain; so if KAC does not detect inconsistency, AC will not detect inconsistency either.

KAC is equivalent to AC when domains are completely known. The advantage in using KAC is that the check for known arc-consistency can be performed lazily, without full knowledge of all the elements in every domain.

2.3 Linking the two sorts

Intuitively, we want to bind CLP(*FD*) with CLP($\mathscr{I}$-Set) with the intended semantics that I-Sets provide domains for *FD* variables.

Given the two CLP languages $\mathscr{L}_{FD}$ and $\mathscr{L}_{\mathscr{I}\text{-Set}}$, we define the CLP language $\mathscr{L}$ as the union of the two languages, with a further constraint, $::$, defined as follows:

- the signature $\Sigma = \Sigma_{FD} \cup \Sigma_{\text{I-Set}} \cup \{ :: \}$;
- the intended interpretation $\mathscr{D}$ keeps the original mappings in the *FD* and $\mathscr{I}$-Set sorts, i.e. $\mathscr{D}|_{\Sigma_{FD}} = \mathscr{D}_{FD}$ and $\mathscr{D}|_{\Sigma_{\text{I-Set}}} = \mathscr{D}_{\text{I-Set}}$.

The declarative semantics of the constraint is

$$X :: S \leftrightarrow X \in S$$

where X is a *FD* variable. The $:: /2$ constraint links a *FD* variable to its domain as in most *CLP*(*FD*) frameworks, with the difference that S, being and I-Set, may

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

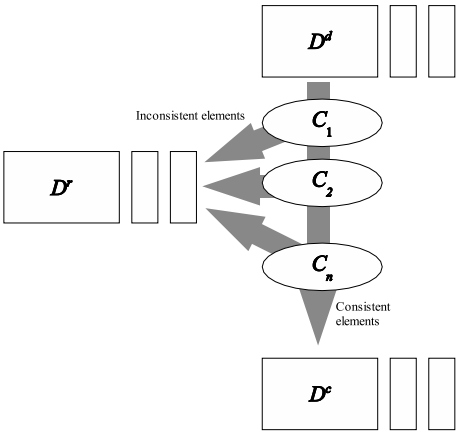


Fig. 1. FD Constraints as filters on variable domains.

be non-completely specified. The $::/2$ constraint should not be confused with the *s-member*/2 constraint of Section 2.1, which represents the constraint of set membership between a ground element and an I-Set.

In CLP(*FD*) systems, domains provide ancillary information about variables. Domains contain the possible values that a variable can take; if a value is not consistent with the imposed constraints, it is operationally removed from the domain. This helps many systems (Dincbas *et al.* 1988; Puget 1994; IC-Parc 2001; SICStus 2003) to obtain higher performance; in fact, domain wipe-outs are detected early and many alternatives are efficiently pruned. On the other hand, in the $\mathscr{I}$-Set sort, domains must be manipulated as logical entities: if an element declaratively belongs to a domain, it cannot be removed.

Suppose that we have constraint $\{1,2,3\} \subseteq D$ stating that the elements 1, 2 and 3 should belong to the set D, the constraint $X :: D$ that links variable X to the domain D, and $X \neq 1$ that states that X should be different from 1. Usual constraint propagation of the constraint $X \neq 1$ would remove element 1 from the domain D, but this is inconsistent with the constraint $\{1,2,3\} \subseteq D$ so the computation would fail. This behavior is not correct, because the set of constraints $\{1,2,3\} \subseteq D, X \in D$, and $X \neq 1$ is satisfiable.

For this reason, in our framework, the domain of a variable X is represented by two streams: a *Definition Domain*, D_X^d, that contains all the values synthesized for the variable, and a stream of *Removed Values*, D_X^r, that is the set of elements proven to be inconsistent with the imposed constraints. The set of available items for X, also called *Current Domain*, D_X^c, is given by the relation:

$$D_X^c = D_X^d \setminus D_X^r. \tag{2}$$

It should be noticed that D_X^c remains open until D_X^d is closed. D_X^r, instead, is always open, so to make it possible to move elements into it from the current domain even if the definition domain has been closed.

With the imposed constraints, it is possible to declaratively add a newly synthesized value to the definition domain (imposing the constraint $v \in D_X^d$) or remove an

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

inconsistent element from the current domain ($w \notin D_X^c$ or, equivalently, $w \in D_X^r$; see Figure 1). In this way, during search, an inconsistent element will not be tried if it belongs to D_X^r and a possible domain wipe-out will be detected when D_X^c is empty.

Note that the relation in (2) automatically propagates two types of information from the definition domain to the current domain of a variable. First, whenever a domain element is synthesized, it is considered in the current domain and can be exploited for propagation. Second, if no more values are available to the definition domain, also the current domain becomes closed.

Notice that more than one FD variable can range on the same definition domain or, equivalently, their definition domains can be linked by an equality constraint; however, each of them will have its own current domain and set of removed values. By definition, the user cannot close the current domain of a variable: the user should only access the definition domain directly. This is not a restriction, because if one wants to close independently the current domain of different variables (as in the previous example), he can define two different definition domains (and, possibly, impose some constraints among them, e.g. $D_X^d \subseteq D_Y^d$).

To achieve KAC, it is necessary to remove elements and to *promote* elements, i.e. to move ideally some elements from the unknown part to the known part. Elements can then be removed (i.e., prevented from entering the current domain) if they are shown to be inconsistent. An algorithm for achieving KAC is shown by Gavanelli *et al.* (2005).

The addition, besides the deletion, of elements to the domain might seem non monotonic. However, the current domain is kept open as long as new acquisitions are possible (i.e. until the definition domain is closed), the unknown part representing the set of future acquisitions. Thus, declaratively, when a new element enters the known part of the current domain, it is not properly added: it is just made explicit. A similar behavior is also achieved by Mailharro (1998) with the use of a *wildcard* representing values *entering* the domain.

Clearly, the repeated acquisition of the same element would result in a loop; this can be avoided by having an acquisition module that does not provide twice the same element to the same I-Set (as hypothesized, for instance, by Mailharro (1998)).

Example: Numeric CSP. Various applications could be thought exploiting interactive constraint propagation, as many systems need to interact with an external environment during propagation: some real-world examples can be found in Gavanelli *et al.* (2005). In this section, we give a simple example, aimed at showing a typical ICSP computation, rather than at showing the power of the language.

With the given language, we can state in a natural way the following problem:

$$:\!-X :: D_X, Y :: D_Y, Z :: D_Z, intersection(D_X, D_Y, D_Z), Z > X \cdot$$

defining three variables, X, Y, and Z, with constraints on them and their domains. KAC propagation can start even with domains fully unknown, i.e. when D_X, D_Y and D_Z are variables. Let us suppose that the elements are acquired through interaction with a user and the first element retrieved for X is 1. This element is inserted in

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

the definition domain of X, i.e. $D_X^d = \{1|(D_X^d)'\}$; since it is consistent with FD constraints, it is not redirected to the D_X^r stream; value 1 is considered in the current domain, i.e. $D_X^c = \{1|(D_X^c)'\}$. Then KAC propagation tries to find a support for this element in each domain of those variables linked by *FD* constraints; in our instance D_Z. A value is requested for D_Z and the user gives a (possibly consistent) value: 2. This element is inserted into the definition domain of Z: $D_Z^d = \{2|(D_Z^d)'\}$. The constraint imposed on domains can propagate, so element 2 has to be inserted in the definition domain of Y and X, thus $D_Y^c = D_Y^d = \{2|D_Y'\}$ and $D_X^c = D_X^d = \{1, 2|D_X''\}$.

Since the acquired element is consistent with FD constraints involving Z, it enters the current domain of Z. KAC propagation must now find a known support for element 2 for X, so another request is performed. If the user replies that there is no other element in the domain of Z, then 2 is sent to the stream of removed values for X (thus, it is not considered in the current domain of X). It will remain in the definition domain because the element semantically belongs to the domain even if no consistent solution can exist containing it.

When KAC propagation reaches the quiescence, each element of the current domain of each variable has a support for each FD constraint involving that variable.

3 Implementation concepts

Since one of the aims of the ICSP model is to manage efficiently those problems in which value acquisition is costly, it is useful to exploit the link between the two sorts to avoid unnecessary acquisition of domain elements. For instance, as shown in the example in Section 2.3, it is possible to infer the presence of a value in the known part of an I-Set from $\mathscr{I}$-Set constraints, without acquiring it directly.

Thus, for the sake of efficiency, the implementation of the language must be aware of the link between the two sorts, given by the :: /2 constraint (see Section 2.3), and exploit it when useful. Nonetheless, the two sorts are clearly distinct, and need to be handled by different computational mechanisms.

In more detail, it is possible to devise two constraint solvers: one for the *FD* sort, meant to ensure KAC (see Definition 2.1) of the *FD* constraint network, and one for the $\mathscr{I}$-Set sort, meant to ensure satisfaction of $\mathscr{I}$-Set constraints.

These solvers only need to interact in two cases:

- when, during *FD* KAC check, a new element is acquired, $\mathscr{I}$-Set propagation must be activated in order to satisfy $\mathscr{I}$-Set constraints;
- conversely, when a new element enters the known part of an I-Set, *FD* KAC check must be activated over all the *FD* variables which have their definition domain in the I-Set.

The CHR language is a perfectly suitable tool for this purpose, letting us deal with the two sorts separately, and to manage the link between them with ease.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

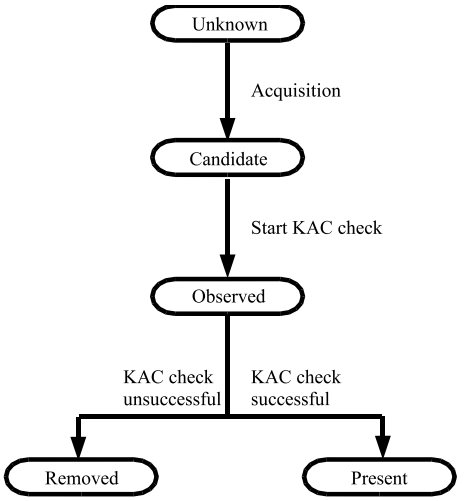


Fig. 2. State transitions for (*Variable*,*Element*) pairs.

3.1 FD sort concepts

The aim of the *FD* solver is to keep the constraint network known arc-consistent. This is achieved by preventing elements from entering current domains if no KAC support is found for the constraints involving them. In more detail, each time a new element E is promoted from the unknown to the known part of an I-Set C, it is inserted into the definition domain of each variable V such that $V :: C$ ($E \in D_V^d$ with the notation of Section 2.3). The algorithm then tries to find a KAC support for E: if support is found, then E is also inserted into the *Current Domain* of V ($E \in D_V^c$); otherwise, although it stays in the definition domain of V, it enters the stream of *Removed Values* for V ($E \in D_V^r$).

In this way, at any step of the propagation, all of the values in current domains are certainly supported.

3.1.1 Values' states

The process described above can be clearly formalized as a sequence of state transitions for (V, E) pairs. These states are:

- *unknown*: E has not yet been acquired as a value for V;
- *candidate*: E has been inserted into D_V^d, but (V, E) has not yet been chosen for supporting another value and no KAC control is being run on it;
- *observed*: (V, E) provides support for some other *observed* pair, but it has not yet been proven to be supported;
- *present*: (V, E) is supported, and has thus been inserted into D_V^c;
- *removed*: E is not supported, and has thus been inserted into D_V^r.

The possible state transitions are shown in Figure 2.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

3.1.2 Lazy value acquisition

The ICSP framework is more suitable for those applications in which the process of acquiring domain values is computationally costly. In this perspective, the KAC check procedure has been designed so as to acquire new values only when it is really necessary, i.e. when no support for a value can be found among already known values.

3.2 $\mathscr{I}$-Set sort concepts

The task performed by the $\mathscr{I}$-Set solver is to ensure the satisfaction of the $\mathscr{I}$-Set sort constraints, and to propagate the closure of definition domains, when possible. As shown in the example below, it is possible to infer that an element belongs to an I-Set from $\mathscr{I}$-Set sort constraints, thus reducing the number of necessary acquisitions and improving efficiency.

In particular, each time an element is added to an I-Set, an *inserted*/2 event (see Section 2.1) is raised, which starts a check for the satisfaction of the $\mathscr{I}$-Set constraints involving the I-Set itself.

Example Let us suppose the constraint *intersection*(D_X,D_Y,D_Z) is given among the D_X, D_Y and D_Z I-Sets, and that the current values of the domains are

$$D_X = \{2, 4|D'_X\};\ D_Y = \{3, 4|D'_Y\};\ D_Z = \{4|D'_Z\}.$$

If 5 is added to D_Z, the system immediately ensures 5 to be a member of both D_X and D_Y, without the need for acquiring the element. If 3 is added to D_Z, the system only needs to add it to D_X. Conversely, if 3 is added to D_X, it is also added to D_Z; if 1 is added to D_X, instead, the system cannot infer anything, and thus makes no additions.

$\mathscr{I}$-Set constraints also have to be taken into account to propagate closure of I-Sets, when possible (as in the example shown in Section 2.1). For this purpose, we use a *closed*/1 event, which notifies that an I-Set has been closed.

3.3 Interaction between sorts

The link between the two sorts, represented by the :: /2 constraint (see Section 2.3), is implemented as an interaction mechanism between the solvers activated when, on the one hand, a new element is acquired during a KAC check and, on the other hand, an element is inserted into an I-Set.

Semantically, we use the *inserted*/2 event for both of these cases (see Section 2.1).

At the implementation level, the event causes the control to pass from one solver to another.

3.4 Algorithms

3.4.1 Priority among supporting values

A *(Variable,Element)* pair is supported if, for each constraint involving *Variable*, support is found in the known part of other variables' domains. But supporting

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

pairs may not be supported in their turn: this implies that, if a pair is found to be unsupported, new support needs to be found for all of the pairs that it was supporting. For efficiency, we keep track of which pairs support which other pairs for which constraints, so that, in case a pair is found unsupported, it is possible to revise only the affected constraints.

However, because of states (see Section 3.1.1), it is possible to reduce the number of supports to keep track of. If an element is supported by an element whose state is *present*, there is no need for keeping track of the support, because a present element is already known to be supported, as explained in Section 3.1.1. Thus, when seeking support for a constraint, present elements should be tried first.

Likewise, an element whose state is *observed* is more convenient than an element whose state is *candidate*, because the algorithm is already seeking support for it, and can thus avoid seeking support for a new element (note that all candidates will eventually be checked; but in this way, when their turn comes, there will be, in general, more present elements).

So, a priority can be established among eligible supporting elements: first present, then observed, then candidates. Only when no support is found among these elements, a new element needs to be acquired: if the definition domain of the corresponding variable is open, an element can be requested, otherwise the support seek procedure fails.

3.4.2 The support graph

Support dependencies can be represented by a (directed) support graph, in which nodes represent observed *(Variable,Element)* pairs, and arcs represent support dependencies.

When a candidate pair becomes observed, the corresponding node is added to the graph. This may happen at the beginning of graph construction, when the graph is empty and the candidate is chosen to be the first node, or later, during the support seeking procedure, if the candidate provides support for an observed.

A new node needs support for the *FD* constraints involving it. For each *FD* constraint, the following attempts, in this order (see Section 3.4.1) are made, until one succeeds:

1. support is found in present elements only: the graph needs no modifications;
2. one of the supporting elements is observed: an arc is added from the supporting node to the supported, marked with the constraint;
3. one of the supporting elements is candidate: a new node is added for the candidate, an arc is added from the supporting node to the supported, marked with the constraint.

If none of these attempts succeeds, a variable (other than *Variable*) is chosen for acquisition among those involved by the constraint, and the search for support is restarted from the support-in-candidates attempt (step 3). The search for support needs to be restarted because, in order to satisfy $\mathcal{I}$-Set constraints, new candidates may have been added, which may provide support for the considered *FD* constraint.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

A node is deleted from the graph if it is found to have no support; in this case, new support is sought for all of the nodes that it was supporting.

Once a graph is fully built, all of its nodes can be inserted into the current domains (observed-to-present transition), whereas the unsupported nodes will be inserted into the stream of removed values.

The procedure reaches quiescence only when there are no more candidates to search support for.

4 CHR implementation

In this section, we show how the concepts explained in Section 3 have been implemented in a constraint solver using the Constraint Handling Rules library of SICStus Prolog.

4.1 $\mathcal{I}$-Set variables and constraints

I-Sets are represented by CHR variables: these variables (also "link variables" hereafter) act as a link for the information stored in constraints, and are never instantiated. To keep memory of the state of I-Sets we use CHR constraints, meant to interact to ensure satisfaction of $\mathcal{I}$-Set constraints.

The state of an I-Set is represented by means of three CHR constraints: *iset_known(Iset,Known)* links an I-Set to its known part, *iset_open(Iset)* indicates that the I-Set is open (i.e. other elements can be inserted), and *closed(Iset)* indicates that the I-Set is closed.

An I-Set is created by the user either explicitly (*new_iset_object*/3 predicate) or implicitly, imposing constraint *icsp_def_domain*/2 (which implements the :: /2 constraint of Sect. 2.3) between an *FD* variable and the link variable of an I-Set. It is possible to define the domain as empty or as having a starting known part.

Each constraint among I-Sets (see Section 2.1 for the formal definitions) is represented by a corresponding CHR constraint among the link variables of the I-Sets that it involves, namely the *iset_member*/2 (in this case, the first argument is a ground *FD* value), *iset_union*/3, *iset_intersection*/3, *iset_inclusion*/2, and *iset_difference*/3 constraints.

The satisfaction of $\mathcal{I}$-Set constraints is ensured incrementally (see Section 3.2). For instance, the following rule is used for implementing the *iset_intersection*/3 constraint:

```
intersection_right_to_left @
        iset_inserted(Iset3,Element),
        iset_intersection(Iset1,Iset2,Iset3)
                        # _iset_intersection
        ==> ensure_membership(Iset1,Element),
        ensure_membership(Iset2,Element)
        pragma passive(_iset_intersection).
```

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

The rule is activated when a new element is added to the I-Set represented by the third argument of an *iset_intersection*/3 constraint: this is notified by the *iset_inserted*/2 constraint, which implements the *inserted* event of Section 2.1. The rule imposes constraint *ensure_membership*/2 on each of the other two I-Sets and the element; if necessary, the element will be inserted into the I-Sets. Notice that *ensure_membership*/2 is defined so as to fail if the I-Set is closed and the element is not already in its known part. Constraint *iset_intersection*/3 is declared passive for efficiency, because it is not necessary to generate code for it ($\mathcal{I}$-Set constraints as *iset_intersection*/3 are imposed only at the beginning of the computation).

Two more rules (not shown here for lack of space) manage the case of an element having been inserted into the first or second argument of the *iset_intersection*/3 constraint.

The *close*/1 primitive of Section 2.1 is implemented by the following CHR:

```
close_iset @
        close(Iset), iset_open(Iset) # _iset_open
        <=> closed(Iset)
        pragma passive(_iset_open).
```

CHR constraint *close*/1 is imposed to close the I-Set; the effect of this rule is to remove constraint *iset_open*/1 for the I-Set, and to impose CHR constraint *closed*/1 for the I-Set. Constraint *closed*/1 also implements the event notifying that the argument I-Set has been closed, which can interact with $\mathcal{I}$-Set constraints so as to propagate closure when possible. For instance, the rule of Section 2.1 is implemented as follows:

```
closure_propagation_inclusion @
        closed(Iset2),
        iset_inclusion(Iset1,Iset2) # _iset_inclusion,
        iset_known(Iset1,K1) # _iset_known1,
        iset_known(Iset2,K2) # _iset_known2,
        ==> permutation(K1,K2) |
        close(Iset1) pragma passive(_iset_inclusion),
        passive(_iset_known1), passive(_iset_known2).
```

Predicate *permutation*/2 in the guard checks that the two I-Sets have the same elements.

4.2 FD variables and constraints

4.2.1 FD variables and domains

FD variables are represented as CHR constrained variables. Their domain is an I-Set; constraint :: /2 (see Section 2.3) is implemented as the *icsp_def_domain(Variable,Iset)* CHR constraint, whereas the current domain of the variable and its set of removed values are represented by the *icsp_curr_domain(Variable,(Present,Removed))* CHR constraint, where *Present* is the list of present elements and *Removed* is the list of removed elements.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

4.2.2 FD constraints

FD constraints are represented by the *fd_constraint(ConstraintName, ListOfArguments)* CHR constraint. For instance, *FD* constraint $X < Z$ is represented by the `fd_constraint(lt,[X,Z])` CHR constraint.

FD constraints are defined by the user simply as one or more clauses for the Prolog predicate *fd_verify*/2 needed to verify whether the constraint is satisfied by a list of ground arguments. This way of defining *FD* constraints supports non-binary constraints, and easy extensibility of the system.

For instance, a possible definition of the *FD* $</2$ constraint could be

```
fd_verify(lt,[A,B]):- A<B.
```

4.2.3 KAC procedure

Candidate (see Section 3.1.1) elements for each variable are stored in the *current_candidates(Variable,ListOfElements)* constraint.

KAC check over candidates is performed by building a support graph (see Section 3.4.2). Nodes of the graph are *(Variable, Element)* pairs. A pair is a node of the graph if and only if it is member of the list argument of the *observed_candidates*/1 CHR constraint; an arc (indicating support dependency) is represented by the *relies((Var1,Val1), (Var2,Val2), FDConstraint)* CHR constraint, where *FDConstraint* is the *FD* constraint for which *(Var2,Val2)* supports *(Var1,Val1)*.

Graph construction is achieved by the following CHR rule:

```
kac_check_start @
        kac_unlock,
        current_candidates(Var,[Candidate|MoreCandidates])
                # _current_candidates
        <=> current_candidates(Var,MoreCandidates),
        observed_candidate((Var,Candidate)),
        check_candidate((Var,Candidate)),
        !, flush_candidates, end_flush_candidates,
        kac_unlock pragma passive(_current_candidates).
```

Graph construction begins when there is at least one non-empty list of candidates for a variable: an element is chosen (in the current implementation, it is simply the first of the list), pair *(Variable, Element)* becomes the first node of the graph (constraint *observed_candidate*/1 adds its argument to the list argument of constraint *observed_candidates*/1), and KAC check starts from it (*check_candidate*/2 constraint); during check, new nodes may be added to the graph, which will be checked for support in their turn. When the construction is done, graph nodes are inserted into the current domains (*flush_candidates*/0 and *end_flush_candidates*/0 constraints), and the procedure can start again over the rest of candidates.

kac_unlock/0 is a constraint meant simply to start the KAC procedure. If there are no candidates when it is imposed, then if there is a variable with empty current domain an element is acquired for it which will become a candidate; otherwise, the procedure reaches quiescence.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

4.2.4 Support seek

Support seek for a *(Variable, Element)* pair is started by constraint *check_candidate((Variable,Element))*. This constraint collects from the store all the *FD* constraints that involve *Variable*, and for each of them tries to find an assignment of all the involved variables that satisfies it (i.e., that makes goal *fd_verify*/2 (see Section 4.2.2) succeed). Elements for the other variables are chosen following the priority described in Section 3.4.1.

4.2.5 Support seek through element acquisition

If no acceptable assignment is found among known, observed, and candidate elements, then:

- if the I-Set domain of at least one of the other variables is open, a new element is acquired for one of the other variables;
- if the I-Set domains of all the other variables are closed, failure is reported, and possibly backtracking is applied.

Which variable should be chosen for acquisition among those with open domain is not obvious for non-binary constraints, and application specific heuristics could be useful; in this implementation the first variable is simply chosen.

In the current system, element acquisition is obtained by asking the user for an element, and the I-Set is closed in case of a predefined input from the user. Obviously, in practical applications, elements would be provided by an acquisition system, as in Mailharro (1998), where the acquisition is done automatically by generation of new component instances, or in Cucchiara *et al.* (1999b), where elements are provided by a low-level segmentation system.

4.3 Constraints and rules for interaction between sorts

Only one CHR constraint needs to be added to the solver to link the two sorts: *iset_inserted(Iset,Element)*, imposed when *Element* is inserted into the known part of *Iset*. This constraint is the implementation of the *inserted*/2 event described in Section 2.1.

4.3.1 FD to $\mathscr{I}$-Set link

When a new element *Element* is acquired for a variable V whose definition domain is *Iset*, *iset_inserted(Iset,Element)* is imposed. This CHR constraint implements an event that triggers the $\mathscr{I}$-Set constraint check. From a procedural point of view, this makes control pass from the *FD* solver to the $\mathscr{I}$-Set solver.

4.3.2 $\mathscr{I}$-Set to FD link

When imposed, *iset_inserted(Iset,Element)*:

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

- first, interacts with the $\mathscr{I}$-Set constraints involving *Iset*, as shown in Section 4.1;
- then, inserts *Element* in the list of candidates of all the variables having *Iset* as definition domain, for later support seek.

The interaction with the $\mathscr{I}$-Set constraints may involve new insertions, which will, in their turn, impose further *iset_inserted*/2 constraints.

New candidates are added only when $\mathscr{I}$-Set propagation has been completed, for the reasons explained in Section 3.2.

5 Related work

I-Sets can be considered as streams with a set semantics (i.e. in an I-Set there are no repeated elements and elements are not sorted). Streams are widely used for communication purposes in concurrent logic programming (Shapiro 1987). Various communication protocols can be implemented using this simple yet powerful data structure (Shapiro 1989). An I-Set can only contain ground elements; this restriction prevents (open) I-Sets from being passed in an I-Set, but lets us achieve higher efficiency.

The first results of our research in the ICSP framework are reported by Cucchiara *et al.* (1999a), where the ICSP model is proposed as an extension of the CSP model. In this model, variables range on partially known domains which have a known part and an unknown part represented as a variable. Domain values are provided by an extraction module and the acquisition process is (possibly) driven by constraints. The model has been proven effective in a vision system (Cucchiara *et al.* 1999b), in randomly-generated problems (Cucchiara *et al.* 1999a), and in planning (Barruffi *et al.* 1999). This work can be considered as the language extension and CHR implementation of the ICSP framework, maintaining it as the core of the propagation engine on the *FD* side.

Operationally, achieving KAC has some similarities with achieving *Lazy Arc Consistency* (LAC) (Schiex *et al.* 1996). LAC is an algorithm that finds an arc-consistent sub-domain (not necessarily a maximal one) and tries to avoid the check for consistency of all the elements in every domain. KAC looks for an arc-consistent sub-domain as well, but it is aimed at avoiding unnecessary information retrieval, rather than unnecessary constraint checks.

Codognet and Diaz (1996) describe a method for compiling constraints in CLP(*FD*). There is only one primitive constraint (X `in` R), used to implement all the other constraints. R represents a collection of objects and can also be a user function. Thus, in CLP(*FD*) domains are managed as first-class objects; our framework can be fruitfully implemented in systems exploiting this idea.

Sergot (1983) proposes a framework to deal with interaction with the user in a logic programming environment. Our work can be used for interaction in a CLP framework; it lets the user interactively provide domain values.

Zweben and Eskey (1989) propose an algorithm that evaluates domain elements only when necessary; domains are streams and constraints are filters on domains. Dent and Mercer (1994) show the effectiveness of such an approach when constraint

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

checks are expensive operations. In our proposal, implementation of delayed evaluation is quite natural and simple, even if our work is aimed at reducing domain values extractions, not constraint checks.

Dynamic Constraint Satisfaction (DCS) (Dechter and Dechter 1988) is a promising field of AI taking into account dynamic changes of the constraint store such as the addition and deletion of values and constraints. The difference between DCS and our approach concerns the way of handling these changes. DCS approaches propagate constraints as if they worked in a *closed world*. Basically, in a DCS one can add or remove a constraint; thus, one can also add and delete domain elements provided that they are all known from the beginning. In an ICSP, instead, domain elements that are unknown can be requested and inserted in the domain. DCS solvers record dependencies between constraints and the corresponding propagation in proper data structures (Schiex and Verfaillie 1993) so as to tackle modifications of the constraint store as soon as data change. In this perspective, we also cope with changes since the acquisition of new values can be seen as a modification of the constraint store. However, we work in an *open world* where domains are left *open* thanks to their unknown part. Unknown domain parts represent intensionally future acquisitions, i.e. future changes.

Another DCS framework was given for configuration problems (Mittal and Falkenhainer 1990). This framework considers dynamicity in the set of variables; variables are introduced or removed during search by means of constraints. The aim is to find a solution where only some of the variables are assigned a value, while the others are *inactive*. However, differently from our framework, the set of domain values is given at the specification of the problem.

Many systems implement sets, because sets have powerful description capabilities. In particular, some have been described as instances of the general CLP($\mathcal{X}$) framework (Jaffar *et al.* 1998), like $\{\log\}$ (Dovier and Rossi 1993; Dovier *et al.* 1996; Dovier *et al.* 2001), CLPS (Legeard and Legros 1991), or Conjunto (Gervet 1997). Others apply an object-oriented approach (Puget 1992; Puget 1994).

In $\{log\}$(Dovier and Rossi 1993; Dovier *et al.* 1996; Dovier *et al.* 2001), a set can be either the empty set $\emptyset$, or defined by a constructor `with` which, given a set S and an element e, returns the set composed of $S \cup \{e\}$. This language is very powerful, allowing sets and variables to belong to sets. However, set unification and the propagation of other constraints has an exponential time complexity in the worst case. If we allow non-ground elements in an I-Set, we obtain a more expressive (although less efficient) framework, like $\{\log\}$; we are currently studying this extension.

Set variables (Puget 1992; Puget 1994) can range on set domains. Each domain is represented by its greatest lower bound and its least upper bound. Each element in a set must be ground, sets are finite, and they cannot contain sets. These restrictions avoid the non-deterministic unification algorithm, and give good performance results; they are implemented in most Constraint Programming systems (Puget 1994; Smolka 1995; SICStus 2003; Gervet 1997; IC-Parc 2001). However, these systems do not deal with problems in which some domain elements are not known;

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

Table 1. *Comparison between two systems*

ICSP	ILOG-Configurator
FD Variables	Port Variables
I-Set	Component Types
Known Part	Set of generated instances
Unknown Part	Set of not yet generated instances (Wildcard)
::/2	Link port-type
Definition Domain D_X^d	Set of instances of the target type
Current Domain D_X^c	Possible set of a port
Removed Elements D_X^r	Instances set of the target type\possible set
Inserted Event	Extension delta domain of ports

in fact, the user must provide the least upper bound of each set, thus giving the universe set at the beginning of the computation.

ILOG-Solver (Puget 1994) does not deal directly with those problems, but has an extension module, called ILOG-Configurator (Mailharro 1998; ILOG 1999) whose main added value is to implement open domains in a way similar to our system. This system is very related with ours, thus we give a more detailed comparison.

Comparison with ILOG-Configurator. In Ilog-Configurator, there are variables called "ports" whose domains are defined by the set of instances of a given component type. The set of instances is not known in advance and instances are generated on demand during constraint propagation. In this way, domains of port variables are dynamically extended; a domain that can be extended contains an element called "wildcard".

Our system has many points in common with the one described by Mailharro (1998). To compare the two, the reader can refer to Table 1.

The differences that we envisage between the two systems are the following. ILOG-Configurator is based on an object-oriented technology, while ours is based on logic programming; this is reflected in some specific choices made in the two systems. For example, both the systems reason about the possible closure of a domain. This information is carried by a special element, called "wildcard" by Mailharro (1998), while in our system it is the unknown part of the domain. In our system, if the continuation of the domain is a variable, other elements can be added, otherwise, if it is the empty set, the insertion of other elements will be forbidden by unification. This has a declarative meaning from a set viewpoint; while a set containing a wildcard must be updated through some destructive assignment, our formalization lets the user specify the value of a logical variable: the continuation of the domain will be extended, as in a stream. In other words, in the ICSP formulation we keep set membership and set inclusion distinct. In fact, in Mailharro (1998), the wildcard element represents a set of elements, but is also an element of the domain. Thanks to our formalization, some future extensions are possible: for example, it is possible to extend the type of sets in a domain, and to have domains that can contain sets

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

themselves, like in $\{log\}$ (Dovier *et al.* 1996). The propagation of FD constraints is operationally performed in a similar way (Known Arc Consistency propagation); however, the property of Known Arc Consistency (KAC), which was not defined by Mailharro (1998), enabled us to prove properties of the algorithms achieving KAC (Gavanelli *et al.* 2005).

Finally, in our framework it is possible to define I-Sets as the combination of other I-Sets with set operators.

6 Conclusions

In this work, we presented CHR the implementation of a language that performs constraint propagation on variables with finite domains when information about domains is not fully known. Domains are channels of information, and are considered as first-class objects that can be themselves defined by means of constraints. The obtained language belongs to the CLP class and deals with two sorts: the *FD* sort on finite domains and the $\mathscr{I}$-Set sort for domains. We provide a propagation engine for the *FD* sort exploiting known arc-consistency, and one for the $\mathscr{I}$-Set sort, as well as a mechanism for their interaction.

The source code of the system is available on request.

Acknowledgements

This work is partially funded by the Information Society Technologies programme of the European Commission under the IST-2001-32530 project (SOCS) within the Global Computing initiative.

We wish to thank the anonymous reviewers for their useful comments.

References

Barruffi, R., Lamma, E., Mello, P. and Milano, M. 1999. Least commitment on variable binding in presence of incomplete knowledge. In *Proceedings of the European Conference on Planning (ECP99)*, S. Biundo and M. Fox, Eds. Lecture Notes in Computer Science, vol. 1809. Springer, Durham, UK, 159–171.

Codognet, P. and Diaz, D. 1996. Compiling constraints in `clp(FD)`. *Journal of Logic Programming 27,* 3 (June), 185–226.

Cucchiara, R., Gavanelli, M., Lamma, E., Mello, P., Milano, M. and Piccardi, M. 1999a. Constraint propagation and value acquisition: why we should do it interactively. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, T. Dean, Ed. Morgan Kaufmann, Stockholm, Sweden, 468–477.

Cucchiara, R., Gavanelli, M., Lamma, E., Mello, P., Milano, M. and Piccardi, M. 1999b. Extending CLP(FD) with interactive data acquisition for 3D visual object recognition. In *Proceedings of the First International Conference on the Practical Application of Constraint Technologies and Logic Programming*. Practical Application Company, London, 137–155.

Cucchiara, R., Gavanelli, M., Lamma, E., Mello, P., Milano, M. and Piccardi, M. 2001. From eager to lazy constrained data acquisition: A general framework. *New Generation Computing 19,* 4 (Aug), 339–367.

Dechter, R. and Dechter, A. 1988. Belief maintenance in dynamic constraint networks. In *Proceedings of the 7th National Conference on Artificial Intelligence*, T. M. Smith, Reid G.; Mitchell, Ed. Morgan Kaufmann, St. Paul, MN, 37–42.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

Dent, M. and Mercer, R. 1994. Minimal forward checking. In *Proceedings of the Sixth International Conference on Tools with Artificial Intelligence*. 432–438.

Dincbas, M., Van Hentenryck, P., Simonis, H., Aggoun, A., Graf, T. and Berthier, F. 1988. The constraint logic programming language CHIP. In *Proceedings of the International Conference on Fifth Generation Computer System*. OHMSHA Ltd. Tokyo and Springer-Verlag, Tokyo, Japan, 693–702.

Dovier, A., Omodeo, E., Pontelli, E. and Rossi, G. 1996. $\{log\}$: A language for programming in logic with finite sets. *Journal of Logic Programming 28(1)*, 1–44.

Dovier, A., Pontelli, E. and Rossi, G. 2001. Constructive negation and constraint logic programming with sets. *New Generation Computing 19,* 3 (May), 209–256.

Dovier, A. and Rossi, G. 1993. Embedding extensional finite sets in CLP. In *Proceedings of the 1993 International Symposium on Logic Programming*. MIT Press, British Columbia, Canada, 540–556.

Faltings, B. and Macho-Gonzalez, S. 2003. Open constraint optimization. In *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP-2003)*. Lecture Notes in Computer Science. Springer, 303–317.

Frühwirth, T. 1998. Theory and practice of constraint handling rules. *Journal of Logic Programming 37,* 1-3 (Oct.), 95–138.

Gavanelli, M., Lamma, E., Mello, P. and Milano, M. 2004. Dealing with incomplete knowledge on CLP(*FD*) variable domains. *ACM Transactions on Programming Languages and Systems 27(02)*, March 2005.

Gervet, C. 1997. Propagation to reason about sets: Definition and implementation of a practical language. *Constraints 1*, 191–244.

IC-Parc. 2001. *ECL^{i}PS^e User Manual, Release 5.2*. IC-Parc, Imperial College, London, UK.

ILOG 1999. *ILOG-Configurator, user manual*, 1.0 ed. ILOG.

Jaffar, J., Maher, M., Marriott, K. and Stuckey, P. 1998. The semantics of constraint logic programs. *Journal of Logic Programming 37(1–3)*, 1–46.

Legeard, B. and Legros, E. 1991. Short overview of the CLPS system. In *Proceedings of the 3rd Int. Symposium on Programming Language Implementation and Logic Programming, PLILP91*, J. Małuszyński and M. Wirsing, Eds. Lecture Notes in Computer Science. Springer-Verlag, Passau, Germany, 431–433.

Mackworth, A. 1977. Consistency in networks of relations. *Artificial Intelligence 8*, 99–118.

Mailharro, D. 1998. A classification and constraint-based framework for configuration. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing 12*, 383–397.

Mittal, S. and Falkenhainer, B. 1990. Dynamic constraint satisfaction problems. In *Proc. of AAAI-90*. Boston, MA, 25–32.

Puget, J. 1992. PECOS: A high level constraint programming language. In *Proceedings of the First Singapore International Conference on Intelligent Systems (SPICIS)*. Singapore, 137–142.

Puget, J. 1994. A C++ implementation of CLP. Tech. Rep. 94-01, ILOG Headquarters.

Schiex, T., Régin, J., Gaspin, C. and Verfaillie, G. 1996. Lazy arc consistency. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence and the Eighth Innovative Applications of Artificial Intelligence Conference*. AAAI Press/MIT Press, Menlo Park, 216–221.

Schiex, T. and Verfaillie, G. 1993. Nogood recording for static and dynamic constraint satisfaction problems. In *Proceedings of the 5th International Conference on Tools with Artificial Intelligence*. IEEE Computer Society Press, Los Alamitos, CA, USA, 48–55.

Sergot, M. 1983. A query-the-user facility for logic programming. In *Integrated Interactive Computing Systems*, P. Degano and E. Sandewall, Eds. North-Holland, 27–41.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press

Shapiro, E., Ed. 1987. *Concurrent Prolog – Vol. I*. MIT Press.

Shapiro, E. 1989. The family of concurrent logic programming languages. *ACM Computing Surveys 21,* 4 (Mar.), 413–510.

SICStus 2003. SICStus Prolog user manual, release 3.11.0. `http://www.sics.se/isl/sicstus/`.

Smolka, G. 1995. The Oz programming model. In *Computer Science Today*, J. van Leeuwen, Ed. Lecture Notes in Computer Science, vol. 1000. Springer-Verlag, Berlin, 324–343.

Zweben, M. and Eskey, M. 1989. Constraint satisfaction with delayed evaluation. In *IJCAI 89*. Morgan Kaufmann, Detroit, 875–880.

https://doi.org/10.1017/S147106840500236X Published online by Cambridge University Press