

# Automatic Implication Checking for CHR Constraint Solvers

*Tom Schrijvers    Bart Demoen    Gregory Duck*  
*Peter Stuckey    Thom Frühwirth*

*Report CW402, January 2005*



Katholieke Universiteit Leuven  
Department of Computer Science  
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

# Automatic Implication Checking for CHR Constraint Solvers

*Tom Schrijvers    Bart Demoen    Gregory Duck*  
*Peter Stuckey    Thom Frühwirth*

*Report CW402, January 2005*

Department of Computer Science, K.U.Leuven

## **Abstract**

Constraint Handling Rules (CHRs) are a high-level rule-based programming language commonly used to define constraint solvers. We present a method for automatic implication checking between constraints of CHR solvers. Our method does not copy the entire constraint store, but performs the check in place using a trailing mechanism. The necessary code enhancements can be done by automatic program transformation based on the rules of the solver. We extend our method to work for hierarchically organized modular CHR solvers. We show the soundness of our method and its completeness for a restricted class of canonical solver as well as for specific existing non-canonical CHR solvers. We evaluate our trailing method experimentally by comparing with the copy approach: runtime is almost halved.

**Keywords :** Constraint Handling Rules, constraint implication, constraint solver hierarchy.

**CR Subject Classification :** D.3.2, D.3.3, F.4.1

# Automatic Implication Checking for CHR Constraints

Tom Schrijvers\*, Bart Demoen,  
Gregory Duck, Peter Stuckey,  
Thom Frühwirth

January 28, 2005

## Abstract

Constraint Handling Rules (CHRs) are a high-level rule-based programming language commonly used to define constraint solvers. We present a method for automatic implication checking between constraints of CHR solvers. Our method does not copy the entire constraint store, but performs the check in place using a trailing mechanism. The necessary code enhancements can be done by automatic program transformation based on the rules of the solver. We extend our method to work for hierarchically organized modular CHR solvers. We show the soundness of our method and its completeness for a restricted class of canonical solver as well as for specific existing non-canonical CHR solvers. We evaluate our trailing method experimentally by comparing with the copy approach: runtime is almost halved.

## 1 Introduction

Constraint Handling Rules (CHRs) are a high-level rule-based programming language commonly used to define constraint systems. The promise of user-defined constraint systems as stated in [2]: “For the theoretician meta-theorems can be proved and analysis techniques invented once and for all; for the implementor different constructs (backward and forward chaining, suspension, compiler optimization, debugging) can be implemented once and for all; for the user only one set of ideas need to be understood, though with rich (albeit disciplined) variations (constraint systems).”

Constraint handling rules [7] (CHRs) are a very flexible formalism for writing incremental constraint solvers and other reactive systems. In effect, the rules define transitions from one constraint set to an equivalent constraint set. Transitions serve to simplify constraints and detect satisfiability and unsatisfiability.

---

\*Research Assistant of the fund for Scientific Research - Flanders (Belgium)(F.W.O. - Vlaanderen)

CHRs have been used extensively (see e.g. [10]). Efficient implementations are already available for the languages SICStus Prolog and Eclipse Prolog, SWI Prolog, HAL and Java. CHRs provide the best mechanism to date for creating executable user-defined constraint solvers.

In this paper we investigate how to automatically extend a CHR constraint solver to not only answer questions of satisfiability, but also to answer questions about implication.

Previously, it was already shown how to extend built-in solvers with CHR solvers in [4]. In this paper we have added a means to extend CHR solvers with other CHR solvers.

In [3] a general technique is presented for extending basic ask constraints to implication checks of arbitrary logic formulas. The CHR implication checks presented in this paper can be extended to arbitrary formulas in that way.

The first related technique for CHR was already sketched in [6]. However, that technique does not represent an ask constraint, but rather a reified constraint. It also performs its operations in place, but spuriously added constraints are never removed explicitly. Instead, they are kept around until they are removed by the CHR solver. The technique is not safe in the general case: it either allows the removal of constraints present in the initial constraint store or can lead to non-termination if removal is suppressed.

This paper is structured as follows. In the next section, CHR constraint solvers and related notions are introduced. Section 3 presents the basic method of implication checking and soundness as well as completeness results. Section 4 extends this technique to hierarchical solvers. The completeness of implication checking with our method is studied for several concrete CHR solvers in Section 5. Section 6 presents some experimental results. Finally, Section 7 concludes.

## 2 CHR Solvers

### 2.1 CHR Syntax and Operational Semantics

In this subsection we briefly introduce the syntax and semantics of CHR programs or solvers. The other relevant aspects of CHR solvers are covered in the rest of this section. For a more extensive introduction and survey of CHR we refer the reader to [7].

### 2.2 Syntax

A CHR solver *CS* consists of a sequence of CHR rules. There are three kinds of CHR rules. We introduce them with an example.

**Example 2.1** The following four CHR rules define an equality solver *eq*, with *eq/2* the equality constraint:

```

reflexive @ eq(X,Y) <=> X == Y | true.
redundant @ eq(X1,Y1) \ eq(X2,Y2) <=> X1 == X2, Y1 == Y2 | true.

```

```

symmetric @ eq(X,Y) ==> eq(Y,X).
transitive @ eq(X1,Y1), eq(X2,Y2) ==> Y1 == X2 | eq(X1,Y2).

```

with the arguments of `eq/2` variable identifiers and `==/2` syntactic identity.

The `reflexive` rule is an example of a simplification rule, the `symmetric` and `transitive` rule are propagation rules, and the `redundant` rule is a simpagation rule.

The part before the `@` symbol is the name of the rule, it is optional. The part before the arrow is called the *head* of the rule. The part between the arrow and the pipe symbol (`|`) is called the *guard* of the rule (optional). The part after the pipe symbol is the *body* of the rule.

The head of a rule is a sequence of CHR constraints. The guard is a sequence of given built-in constraints. The body is a sequence of both CHR and built-in constraints.

A rule partially defines the constraints in the head. Built-in constraints are undefined, they are implemented in an already existing solver.  $\square$

In general we will write a CHR rule as a simpagation rule  $H_k \setminus H_r \Leftarrow G \mid B$ , where  $H_k$  and  $H_r$  are the sequences of head constraints,  $G$  is the guard and  $B$  is the body of the rule. If  $H_k$  is empty, the rule is a simplification rule and if  $H_r$  is empty, the rule is a propagation rule.

### 2.3 Operational Semantics

The operational semantics of CHR can be described as a state transition system [1]. Execution starts from an initial state  $\sigma_0$  comprising the initial multi-set of constraints (also query or constraint store). A CHR rule application corresponds with a state transition  $\sigma_i \mapsto_{CS} \sigma_{i+1}$  and updates the multi-set of constraints. A rule is applicable if there are distinct constraints in the current state that match the head of the rule and if—under this matching—the guard is implied by the built-in constraints of the current state.

A terminating CHR derivation  $d$  is one that starts from an initial state  $\sigma_0$  and reaches, after a finite number of transition steps, a final state  $\sigma$ . A final state  $\sigma$  is one in which no more transition step is possible. In general for the same initial state many different final states may be reached through many different derivations. We denote a particular derivation  $d$  from  $\sigma_0$  to  $\sigma$  as  $\sigma_0 \mapsto_{CS}^d \sigma$ . Also, we denote the final state obtained through a derivation  $d$  from an initial state  $\sigma_0$  as  $solve_d(\sigma_0)$ . Note that CHR is a committed-choice language, i.e. rule applications cannot be undone. So a particular execution of a query will involve exactly one of the possible derivations.

In this paper we will use a particular instance of the semantics. The refined operational semantics [5] is the actual operational semantics implemented by most CHR systems, such as those in SICStus [9], HAL [11] and the K.U.Leuven CHR system [12]. These semantics describe a particular execution strategy that considerably reduces the number of different possible derivations for any initial state. CHR rules are applied in a top-to-bottom manner. The state is separated into two parts: a sequence of goals to be processed from left to right,

and a multiset of constraints used for matching. Rules are applied using the notion of an *active* constraint, the last constraint added, is exhaustively used in a matching then becoming inactive. For more details see [5].

**Example 2.2** The following example informally describes a derivation under the refined operational semantics.

Consider the execution of the goal  $\mathbf{eq(a,b)}, \mathbf{eq(b,c)}$  for the  $\mathbf{eq/2}$  solver defined in Example 2.1.

Adding  $eq(a,b)$  the *reflexive* rule is not applicable because  $a == b$  does not hold, the *redundant* rule misses a second constraint, the *symmetric* rule adds  $eq(b,a)$ . This constraint adds  $eq(a,b)$  using the *symmetric* rule, which is then deleted by the *redundant* rule. The *transitive* rule now can match  $(eq(a,b), eq(b,a) ==> b == b | eq(a,a))$  so adds  $eq(a,a)$  which is removed by the *reflexive* rule. Similarly the *transitive* rule adds  $eq(b,b)$  which is deleted by the *reflexive* rule. The store is currently  $\{eq(a,b), eq(b,a)\}$ .

The addition of  $eq(b,c)$  causes the addition of  $eq(c,b)$  using the *symmetric* rule, and the *transitive* rule adds  $eq(b,b)$  and  $eq(c,c)$  which are immediately deleted by the *reflexive* rule, as well as  $eq(c,a)$  ( $eq(c,b), eq(b,a) ==> b == b | eq(c,a)$ ) which adds  $eq(a,c)$  using the *symmetric* rule. Later the *transitive* rule adds another copy of  $eq(a,c)$  ( $eq(a,b), eq(b,c) ==> b == b | eq(a,c)$ ) which is deleted using the *redundant* rule. The final store is  $\{eq(a,b), eq(b,a), eq(b,c), eq(c,b), eq(a,c), eq(c,a)\}$

Part of the derivation is illustrated below, with the two parts of the state shown as  $\langle Goal, Store \rangle$  and the active constraint in bold, and matching constraints underlined.

$$\begin{array}{l}
\langle [eq(a,b), eq(b,c)], \emptyset \rangle \\
\mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(\mathbf{a}, \mathbf{b})}\} \rangle \\
\text{symmetric} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(\mathbf{b}, \mathbf{a})}\} \rangle \\
\text{symmetric} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(\mathbf{a}, \mathbf{b})}\} \rangle \\
\text{redundant} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(\mathbf{b}, \mathbf{a})}\} \rangle \\
\text{transitive} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(\mathbf{b}, \mathbf{b})}\} \rangle \\
\text{reflexive} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(\mathbf{b}, \mathbf{a})}\} \rangle \\
\text{transitive} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(\mathbf{a}, \mathbf{a})}\} \rangle \\
\text{reflexive} \quad \mapsto_{CS} \langle [eq(b,c)], \{\underline{eq(a,b)}, \underline{eq(b,a)}\} \rangle \\
\mapsto_{CS} \langle [], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(\mathbf{b}, \mathbf{c})}\} \rangle \\
\text{symmetric} \quad \mapsto_{CS} \langle [], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(b,c)}, \underline{eq(\mathbf{c}, \mathbf{b})}\} \rangle \\
\text{symmetric} \quad \mapsto_{CS} \langle [], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(b,c)}, \underline{eq(c,b)}, \underline{eq(\mathbf{b}, \mathbf{c})}\} \rangle \\
\text{redundant} \quad \mapsto_{CS} \langle [], \{\underline{eq(a,b)}, \underline{eq(b,a)}, \underline{eq(b,c)}, \underline{eq(c,b)}\} \rangle \\
\vdots \\
\mapsto_{CS} \langle [], \{eq(a,b), eq(b,a), eq(b,c), eq(c,b), eq(a,c), eq(c,a)\} \rangle
\end{array}$$

□

Note that propagation rules are applied at most once to the same sequence of constraints to avoid trivial non-termination.

## 2.4 Declarative Semantics

Constraints are special predicates of first order logic, their meaning is defined by a constraint theory. A CHR constraint solver program  $CS$  has a *logical reading* (meaning)  $\llbracket CS \rrbracket$ . This theory  $\llbracket CS \rrbracket$  contains the constraint theory for the built-in constraints and one formula for each rule in the program  $CS$ .

Let  $vars(F)$  denote the set of free variables in formula  $F$ , let  $\forall F$  denote the universal closure of a formula  $\forall vars(F) F$ , and  $\exists_V F$  denote the formula  $\exists W F$  where  $W = vars(F) - V$ , called projection of  $F$  on  $V$ .

The logical reading of a simpagation rule  $H_k \setminus H_r \Leftarrow G \mid B$  is:

$$\forall (G \rightarrow (H_k \rightarrow (H_r \leftrightarrow \exists \bar{y} B))).$$

where  $\bar{y}$  are the variables that appear only in the body  $B$ , i.e.  $vars(B) - (vars(G) \cup vars(H_k) \cup vars(H_r))$ . If either of  $H_k$  or  $H_r$  is empty, it is considered *true* in the formula.

Typically the intention of the solver writer is for  $\llbracket CS \rrbracket$  to approximate some constraint theory  $\mathcal{D}$ , i.e.  $\llbracket CS \rrbracket$  covers only the part of  $\mathcal{D}$  that is relevant for a particular use. For correctness of the approximation, we require that  $\llbracket CS \rrbracket$  is a model of a constraint theory  $\mathcal{D}$ , i.e.  $\mathcal{D} \models \llbracket CS \rrbracket$  (but typically  $\llbracket CS \rrbracket \neq \mathcal{D}$ ).

Because of the logical reading  $\llbracket CS \rrbracket$ , logical equivalence between successive states is preserved.

**Example 2.3** The logical reading  $\llbracket eq \rrbracket$  of the *eq* solver of Example 2.1 is the following set of formulas (where we use  $\equiv$  for syntactic identity):

reflexive	$x \equiv y \rightarrow (eq(x, y) \leftrightarrow true)$
redundant	$x_1 \equiv x_2 \wedge y_1 \equiv y_2 \rightarrow (eq(x_1, y_1) \rightarrow (eq(x_2, y_2) \leftrightarrow true))$
symmetric	$eq(x, y) \rightarrow eq(y, x)$
transitive	$y_1 \equiv x_2 \rightarrow (eq(x_1, y_1) \wedge eq(x_2, y_2) \rightarrow eq(x_1, y_2))$

Clearly the above rules describe the classical properties of equality.  $\square$

## 2.5 Solver Program Properties

A highly useful property of constraint solvers is *confluence*, which ensures that each possible derivation for a goal leads to the same result.

**Definition 2.1 (Confluent Solver)** A CHR constraint solver  $CS$  is confluent if:

$$\forall C, d, d' : C_1 = solve_d(C) \wedge C_2 = solve_{d'}(C) \Rightarrow \models (\exists_{vars(C)} C_1) \leftrightarrow (\exists_{vars(C)} C_2)$$

See [1] for a decidable, necessary and sufficient condition for confluence of terminating CHR programs under the general operational semantics.

We typically do not mention the specific derivation  $d$  for confluent solvers, if we are not interested in intermediate states and write  $C \mapsto^* C_1$  or simply  $C_1 = solve(C)$ .

**Example 2.4** The equality solver  $eq$  is a confluent solver. For example, for the query  $eq(\mathbf{a}, \mathbf{a})$  a single application of the *reflexive* rule or one application of the *symmetric* and two of the *reflexive* rule both yield an empty constraint store.  $\square$

A property that combines confluence with the declarative semantics is *canonicity*. It requires that semantically equivalent goals yield the same result.

**Definition 2.2 (Canonical Solver)** A confluent CHR constraint solver  $CS$  is canonical if:

$$\forall C, C', d, d' : (\llbracket CS \rrbracket \models C \leftrightarrow C') \wedge (C_1 = solve_d(C)) \wedge (C_2 = solve_{d'}(C')) \Rightarrow \models (\exists_{vars(C \wedge C')} C_1) \leftrightarrow (\exists_{vars(C \wedge C')} C_2)$$

In [15] a sufficient, but not necessary, condition for canonical solvers is given. A CHR solver which is confluent, range-restricted and where all simplification rules are single-headed, gives a canonical solver. A CHR solver is *range-restricted* if all variables appearing in the guard and body of the rule appear in the head. In general, it may be non-trivial to show that a solver is canonical.

**Example 2.5** The equality solver  $eq$  is a canonical range-restricted solver. It is obvious that it is range-restricted. Showing it is canonical relies on showing that it returns a store  $\{eq(x, y), eq(y, x) \mid x \text{ and } y \text{ are connected in the graph created by all the } eq \text{ constraints in the goal}\}$ .  $\square$

In this paper, we are interested in CHR solvers with a syntactical restriction, *CHR-only solvers*, of which the rules do not contain any built-in constraints.

**Lemma 2.3 (Range-Restricted CHR-Only Solvers)**

If a CHR-only solver  $CS$  is range-restricted, then:

$$\forall C : C \mapsto_{CS} C' \Rightarrow vars(C') \subseteq vars(C)$$

**Proof** Obvious from the definitions.  $\square$

### 3 Basic Implication Checking

Based on the properties of logical implication and conjunction, we can use the following technique to verify whether a constraint  $c$  is implied by a conjunction of constraints.

$$D \models C \rightarrow c \Leftrightarrow D \models (C \wedge c) \leftrightarrow C$$

Namely we can use the equivalence of the conjunctions to conclude implication.

In principle, we will verify equivalence in the following way. The solved forms  $solve(C)$  and  $solve(C \wedge c)$  are computed along some derivations. These solved forms are then projected on the variables of  $C \wedge c$  and checked for syntactical equivalence.

This approach is not necessarily complete, but sound, i.e. if the projections of the solved forms are syntactically identical, then  $c$  is implied by  $C$ .

**Theorem 3.1 (Soundness)** *For any CHR solver CS:*

$$\forall C, c : \models \exists_{\text{vars}(C \wedge c)} \text{solve}(C) \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C \wedge c) \Rightarrow \llbracket CS \rrbracket \models (C \rightarrow c)$$

**Proof** We have that  $CS \models C \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C)$  and  $CS \models C \wedge c \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C)$  from Theorem 4.2 of [7]. Hence if  $\models \exists_{\text{vars}(C \wedge c)} \text{solve}(C) \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C \wedge c)$  we have that  $CS \models C \leftrightarrow (C \wedge c)$  and hence  $CS \models C \rightarrow c$ .  $\square$

In practical implementations, logical equivalence testing ( $\exists_{\text{vars}(C \wedge c)} \text{solve}(C) \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C \wedge c)$ ) of projected constraint is restricted to syntactic equivalence of multisets ( $\text{solve}(C) \equiv \text{solve}(C \wedge c)$ ). For range-restricted CHRs the two tests are equivalent since  $\text{vars}(\text{solve}(C)) \subseteq \text{vars}(C \wedge c)$  and  $\text{vars}(\text{solve}(C \wedge c)) \subseteq \text{vars}(C \wedge c)$ .

The straightforward implementation approach for implication checking is to make a copy  $C'$  of  $C$ , compute  $C' \wedge c \mapsto^* C''$  and then check equivalence of  $C$  with  $C''$ . We call this approach the *copy approach*.

However, the above copying approach may be quite expensive.  $C$  may consist of two parts  $C = C_1 \wedge C_2$  such that  $C_1$  is a minimal set of constraints that imply  $c$ . By copying  $C$  in its entirety  $C_2$  is copied unnecessarily and causes unnecessary overhead in the final equivalence test.

We propose the *trailing* approach for CHR-only solvers. It only looks at a minimal set of constraints: the conjunction  $C \wedge c$  is solved in place and a trail of changes is maintained. Analysis of the trail afterward tells us whether the resulting store is equivalent to the original. If that is the case, the updates to the store may remain. Otherwise, the trail is used to revert to the original situation.

The above soundness result is still valid for our restricted trailing approach.

**Theorem 3.2 (Specific Soundness)** *For a canonical range-restricted solver CS:*

$$\forall C, c : \models \text{solve}(C) \equiv \text{solve}(C \wedge c) \Rightarrow \llbracket CS \rrbracket \models (C \rightarrow c)$$

**Proof** The theorem is derived from Theorem 3.1. The two projections in Theorem 3.1 are omitted for range-restricted CHR-only solvers based on Lemma 2.3. The logical equivalence ( $\leftrightarrow$ ) is replaced with syntactical equivalence ( $\equiv$ ) based on the definition of canonical solvers 2.2.  $\square$

**Example 3.1** The following CHR rule conceptually represents the above strategy. Keep in mind the refined semantics with sequential left-to-right execution of the constraints and top to bottom trial of rules.

```
implication @ check_eq(X,Y) <=> eq(X,Y), analyse_trail.
```

The `check_eq/2` constraint represents the implication check. Calling this constraint will either succeed or fail, since `analyse_trail/0` succeeds if the resulting store is equivalent and fails if it is not.  $\square$

Our trail analysis has to look at the addition and removal of constraints to decide equivalence. Roughly, if any constraint is added or deleted during the implication checking, the resulting store will not be equivalent to the original. More precisely, stores are also equivalent if a constraint is only *temporarily* added or deleted, since addition and deletion of the same constraint cancel each other out.

The following set of CHR rules reflect this approach for `analyse_trail/0`:

```
temporary @ analyse_trail \ added(C), removed(C) <=> true.
addition  @ analyse_trail \ added(C) <=> fail.
removal   @ analyse_trail \ removed(C) <=> fail.
success   @ analyse_trail <=> true.
```

Here `added/1` and `removed/1` represent trail entries of added and deleted constraints.

The code above makes use of the refined semantics to work correctly. A call to `analyse_trail` looks for matching `added/1` and `removed/1` constraints, and removes them using the first rule. If any (unmatched) `added/1` and `removed/1` constraints remain, the second or third rule causes it to fail. Otherwise it reaches the fourth rule which simply succeeds.

In general the original solver is transformed as follows to maintain information about changes. For each constraint predicate  $p$  we add a rule (before other rules mentioning it)  $p(\bar{x}) \implies \text{added}(p(\bar{x}))$ . Each rule  $H_k \setminus H_r \iff G \mid B$  with  $H_r = [p_1(\bar{x}_1), \dots, p_n(\bar{x}_n)]$ , is replaced by the rule:

$$H_k \setminus H_r \iff G \mid \text{removed}(p_1(\bar{x}_1)), \dots, \text{removed}(p_n(\bar{x}_n)), B$$

**Example 3.2** The `eq/2` solver is transformed as follows to explicitly generate the necessary `added/1` and `removed/1` constraints.

```
new          @ eq(X,Y) ==> added(eq(X,Y)).
reflexive    @ eq(X,Y) <=> X == Y | removed(eq(X,Y)).
redundant    @ eq(X1,Y1) \ eq(X2,Y2) <=>
               X1 == X2, Y1 == Y2 | removed(eq(X2,Y2)).
symmetric    @ eq(X,Y) ==> eq(Y,X).
transitive   @ eq(X1,Y1), eq(X2,Y2) ==> Y1 == X2 | eq(X1,Y2).
```

We want to check whether  $c \equiv eq(a,c)$  given  $C \equiv eq(a,b) \wedge eq(b,c)$ . We call the goal `eq(a,b)`, `eq(b,c)`, `check_eq(a,c)`. The first two constraints lead to a store  $eq(a,b), eq(b,c), eq(a,c), eq(b,a), eq(c,b), eq(c,a)$  as shown in Example 2.2. The constraint `check_eq(a,c)` first adds `added(eq(a,c))` using the *new* rule, then the *redundant* rule succeeds adding `removed(eq(a,c))`. The call to `analyse_trail` removes both of these using the *temporary* rule, then succeeds using the *success* rule.  $\square$

The transformed solver program above has the disadvantage that it always trails. The following set of rules only enable explicit trailing during implication checking. These rules require `trail_off` to be in the constraint store initially.

```

implication  @ check_eq(X,Y) <=> enable_trail,
              eq(X,Y), analyse_trail,
              disable_trail.
enable       @ trail_off, enable_trail <=> trail_on.
disable     @ trail_on, disable_trail <=> trail_off.
filter_add  @ trail_off \ added(C) <=> true.
filter_remove @ trail_off \ removed(C) <=> true.

```

Our method of implication checking is complete for canonical range-restricted CHR-only solvers.

**Theorem 3.3 (Canonical Completeness)** *If the CHR solver  $CS$  is canonical, then implication checking is complete. That is*

$$\forall C, c : (\llbracket CS \rrbracket \models C \Rightarrow c) \Rightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C) \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C \wedge c)$$

**Proof** Direct from Definition 2.2  $\square$

**Theorem 3.4 (Completeness)** *If  $CS$  is a canonical range-restricted solver, then implication checking is complete. That is*

$$\forall C, c : (\llbracket CS \rrbracket \models C \Rightarrow c) \Rightarrow \text{solve}(C) \equiv \text{solve}(C \wedge c)$$

**Proof** From Definition 2.2 we have that

$$\forall C, c : (\llbracket CS \rrbracket \models C \leftrightarrow (C \wedge c) \Rightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C) \leftrightarrow \exists_{\text{vars}(C \wedge c)} \text{solve}(C \wedge c)$$

Since  $CS$  is range-restricted  $\text{vars}(\text{solve}(C)) \subseteq \text{vars}(C \wedge c)$  and  $\text{vars}(\text{solve}(C \wedge c)) \subseteq \text{vars}(C \wedge c)$ . Hence  $\text{solve}(C) \equiv \text{solve}(C \wedge c)$ .  $\square$  Note that in the special that  $C \wedge c$  fails,  $c$  is not implied by  $C$  as  $C \wedge c$  is not satisfiable. This case is correctly covered by our approach.

A constraint solver need not be canonical for our implication checking to be complete. In Section 5 we will show that our method is also complete for several non-canonical, even non-confluent, CHR solvers.

## 4 Implication Checking for Modular Solver Hierarchies

In this section we extend the implication checking technique of the previous section to modular CHR solver hierarchies. Constraint solvers may be nested hierarchically: one solver depends on some solvers that on their turn depend on other solvers. We say a solver *depends* on another solver if it uses constraints that are defined in the other solver. In particular a CHR solver can use constraints of other solvers in guards and bodies of CHR rules.

Previously it was only possible to use constraints of builtin solvers in guards [4]. Here we show how to construct CHR solver hierarchies where CHR constraints can be used in guards. In a CHR solver hierarchy the dependency graph is

acyclic. We will use parent solver and child solver to refer to one solver that depends on the other.

A *modular* CHR solver is a CHR solver that can be compiled using the interface of its dependencies only. In particular, no knowledge of dependents is required.

**Example 4.1** The following less-than-or-equal-to solver *leq* depends on the *eq* solver:

```

leq_new          @ leq(X,Y) <=> check_eq(X,Y) | true.
leq_antisymmetric @ leq(X1,Y1), leq(X2,Y2) <=>
                  check_eq(X1,Y2), check_eq(X2,Y1) | eq(X1,Y1).
leq_redundant    @ leq(X1,Y1) \ leq(X2,Y2) <=>
                  check_eq(X1,X2), check_eq(Y1,Y2) | true.
leq_transitive   @ leq(X1,Y1), leq(X2,Y2) ==>
                  check_eq(Y1,X2) | leq(X1,Y1).

```

This *leq* solver depends on the *eq* solver in two ways. Firstly, it calls *eq/2* constraints in the body of the *leq\_antisymmetric* rule. Secondly, it also uses the *check\_eq/2* implication check in the guard of all its rules. As both the constraint and the implication check can easily be exported from the *eq* solver this does not violate modularity.

There is a part of the operational semantics of the guard that we have not addressed yet. We call an *event* the addition of a constraint of the child solver or one of the solvers it depends on, then the guard may be satisfied. A CHR rule application may not succeed immediately because a guard is not satisfied, but an event may cause it to be satisfied at a later point.

The semantics of CHR require that CHR constraints of the parent solver are re-activated in case of an event that now satisfies a previously unsatisfied guard. In practice, CHR implementations overestimate the impact of events, i.e. re-activate more than necessary. Typically for builtin solvers, relevant events are provided in the solver interface by the solver programmer together with a mechanism to notify interested parties.

The following rules describe the necessary operations for such a mechanism of events and notifications for the *eq* solver:

```

new_event @ eq(X,Y) ==> touched(X), touched(Y).
trigger   @ touched(X), delayed(X,Goal,ID) ==> call(Goal).
end_event @ touched(X) <=> true.
kill_goal @ kill(ID) \ delayed(X,Goal,ID) <=> true.
kill_end  @ kill(ID) <=> true.

```

The *eq* solver provides a *touched(X)* event in its interface, without knowing anything about particular uses. The *new\_event* rule generates the *touched* event for every variable involved in a new *eq/2* constraint. Users of the interface, such as the *leq* solver will be notified of these events by calling the *delayed/3* constraint. This constraint supplies a callback goal, that is called when the

appropriate `touched/1` event fires and allows the notified party to take due action. The `kill/1` constraint allows for the removal of one or more delayed callbacks, based on an identifier and allows the notified party to no longer receive any events.

The following pseudo-code shows how the *leq* solver subscribes itself to `touched` events. It is pseudo-code because it accesses some internals of the CHR implementation.

```
listen @ leq(X,Y) # CID ==> new_delay_id(ID),
                               delay(X,reactivate(CID),ID),
                               delay(Y,reactivate(CID),ID),
                               listening(CID,ID).
```

Here `CID` is an internal identifier of the CHR constraint. This pseudo-rule is executed when the `leq(X,Y)` constraint is first activated. The call to `new_delay_id/1` generates a new notification identifier. With the two calls to `delay` the  $\leq$  solver will be notified of the relevant events. Upon notification the internal goal `reactivate(CID)` is called which reactivates the corresponding constraint. The call to `listening` internally associates the notification identifier with the corresponding `leq/2` constraint. When the `leq/2` constraint with identifier `CID` is removed, internally the `kill/1` constraint is called on all associated notification identifiers. This avoids reactivation of removed constraints.  $\square$

However, several modifications to the implication checking are now necessary to the original scheme, to accommodate both the hierarchy and the modularity. In the following we explain how to do trailing for multiple CHR solvers, how to distinguish between trails of recursively called implication checks and how implication checking should interact with the event mechanism.

## 4.1 Trailing Interface

Firstly, because of the hierarchy, during an implication check on a parent solver, constraints in the child solver may be added and deleted. Hence, the parent solver trail mechanism has to recursively rely on the child solver trail mechanism. The child solver has to export the necessary trail operations for this.

**Example 4.2** The following set of rules encode the trailing dependency of the `leq/2` solver on the `eq/2` solver:

```
rec_analysis @ leq_analyse_trail ==> eq_analyse_trail.
rec_enable   @ leq_enable_trail  ==> eq_enable_trail.
rec_disable  @ leq_disable_trail ==> eq_disable_trail.
```

$\square$

## 4.2 Implication Strata

Secondly, because of the hierarchy, implication checking may be recursive as well. For example, an implication check of a `leq/2` constraint may require

the implication check of a `eq/2` constraint. Our *shallow* trailing approach does not cover this any more. Indeed, it does not distinguish between those `eq/2` constraints added and deleted during the recursive `eq/2` implication check and those during the top level `leq/2` implication check. A more involved trailing mechanism is needed.

Our solution is to associate with each implication stratum (i.e. level of implication check nesting) a stratum identifier. The top level which is not inside any implication check has stratum identifier 0, an implication check called from top level has identifier  $-1$ , etc.

Every constraint is labeled with the stratum it is called in. For example, `eq(X,Y)` becomes `eq(X,Y,S)` if it is called in stratum `S`.

Constraints called in the top-level query are assigned stratum 0. Constraints called in the body of a rule inherit the lowest stratum of any constraints in the head and the implication checking lowers the stratum by one.

**Example 4.3** For example, the `leq/2` solver is transformed as follows:

```
leq_new          @ leq(X,Y,S) <=> check_eq(X,Y,S-1) | true.
leq_antisymmetric @ leq(X1,Y1,S1), leq(X2,Y2,S2) <=>
    check_eq(X1,Y2,min(S1,S2)-1), check_eq(X2,Y1,min(S1,S2)-1)
    | eq(X1,Y1,min(S1,S2)).
leq_redundant    @ leq(X1,Y1,S1) \ leq(X2,Y2,S2) <=>
    check_eq(X1,X2,min(S1,S2)-1), check_eq(Y1,Y2,min(S1,S2)-1)
    | true.
leq_transitive   @ leq(X1,Y1,S1), leq(X2,Y2,S2) ==>
    check_eq(Y1,X2,min(S1,S2)-1) | leq(X1,Y1,min(S1,S2)).
```

□

Now it is possible for the implication trailing operations to work on a single stratum by looking at the stratum identifiers: all the related constraints are extended with their stratum's identifier.

However, the implication checking is weakened, if the trailing operations are confined within a stratum. The reason is the `temporary` rule:

```
temporary @ analyse_trail(S) \ added(C,S), removed(C,S) <=> true.
```

This rule only cancels out additions and deletions in the same stratum. What is no longer canceled out, is a constraint added in a higher stratum that is removed in a lower stratum and re-added in that lower stratum.

It is possible to recapture this possibility as follows. With every deletion both the stratum of the deleted constraint and the lowest stratum of any of the head constraints is recorded. The latter is the *cause* of the removal. For example, the `leq_antisymmetric` rule then looks like:

```
leq_antisymmetric @ leq(X1,Y1,S1), leq(X2,Y2,S2) <=>
    check_eq(X1,Y2,min(S1,S2)-1),
    check_eq(X2,Y1,min(S1,S2)-1)
```

```

| removed(leq(X1,Y1),S1,min(S1,S2)),
  removed(leq(X2,Y2),S2,min(S1,S2)),
  eq(X1,Y1,min(S1,S2)).

```

The following rules deal with this new `removed/3` constraint.

```

temporary @ analyse_trail(S) \ added(leq(X,Y),S),
  removed(leq(X,Y),S,S)
  <=> true.
promotion @ analyse_trail(S) \ added(leq(X,Y),S), leq(X,Y,S),
  removed(leq(X,Y),Sr,S)
  <=> S < Sr | leq(X,Y,Sr).

```

The `temporary` rule still cancels out addition and deletion within the same stratum, but the `promotion` rule promotes a new constraint to the stratum of the previously deleted constraint. In this way the full power of the basic implication checking is retained for solver hierarchies.

The definition of `check_eq/3` becomes the following. We can get rid of explicit trail enabling and disabling now that we have the implication strata: stratum 0 never requires trailing and the other strata always do.

```

toplevel_add @ added(_,0)      <=> true.
toplevel_rem @ removed(_,_,0)  <=> true.

implication @ check_eq(X,Y,S) <=> eq(X,Y,S), eq_analyse_trail(S).

```

### 4.3 Inter-stratum Events

During an implication check which takes place in the child solver an event may be fired waking some parent solver constraints that cause some parent solver constraints to be added or deleted in a higher stratum.

However, it is not necessary for these events to travel across strata. An implication check can safely be resolved without propagating any information to the parent solver in the higher stratum: as a child solver does not depend on the parent solver the outcome of an implication check on the child solver should not require interaction with the parent solver.

The other way around, a higher stratum will never generate any event in the presence of a lower stratum, since it is temporarily suspended while execution goes on in the lower stratum and only disappears after the implication check in the lower stratum has finished and thus the lower stratum is gone altogether.

Hence, is safe and cheaper for events to only trigger callbacks within the same stratum. The following modified event code reflects this.

```

new_event @ eq(X,Y,S) ==> touched(X,S), touched(Y,S).
trigger   @ touched(X,S), delayed(X,Goal,S,ID) ==> call(Goal).
end_event @ touched(X,S) <=> true.

```

```
listen    @ leq(X,Y,S) # CID ==> new_delay_id(ID),
          delay(X,reactivate(CID),S,ID),
          delay(Y,reactivate(CID),S,ID),
          listening(CID,ID).
```

## 5 Case Studies: Non-Canonical Solvers

We have shown in Section 3 that our CHR implication checking is complete for canonical solvers. In this section we investigate the completeness for some classical, non-canonical solvers.

It will turn out that the implication checking is still complete in many cases, or can be made complete with a little customization in particular cases.

### 5.1 Naive Union-Find Equality Solver

In [14] a CHR implementation of the naive union-find algorithm is presented (see Appendix A for the source code). The `union/2` constraint in that implementation may serve as an equality constraint.

The naive union-find represents equal variables as nodes in the same tree. Any tree with the same variables in it represents the equality of its elements. There is not one preferred, canonical form. For this reason it is even non-confluent: the order of the `union/2` constraints, decides the shape of the tree.

If two variables are unioned that are already equal, their common tree is not modified, nor are any other constraints deleted or added. However, if two variables are not yet equal, a union will merge their trees into one.

Hence, our implication check is complete for this union-find equality solver.

### 5.2 Optimal Union-Find Equality Solver

Next to the naive algorithm also a CHR implementation of an optimal union-find algorithm is given in [13] (see Appendix B for the source code). This algorithm combines path compression with union-by-rank.

Again, when two variables are not equal, their respective trees are merged (by-rank) and this is detected by our implication method. However, in case the variables are already equal, path compression may still modify the tree by shortening paths from nodes to the root. Because the compressed tree is not syntactically identical to the initial tree, our implication method will reject it.

Nevertheless it is possible to customize the `trail_analysis/0` rules to overcome this problem and safely allow path compression, while rejecting truly new equalities. Namely, instead of these general rules:

```
addition  @ analyse_trail(S) \ added(C,S) <=> fail.
removal   @ analyse_trail(S) \ removed(C,_,S) <=> fail.
```

only the detection of the removal of a `root` constraint is required to detect the linking of two trees:

```

removal   @ analyse_trail(S) \ removed(root(_,_),_,S) <=> fail.
cleanup1  @ analyse_trail(S) \ added(_,S) <=> true.
cleanup2  @ analyse_trail(S) \ removed(_,_),S) <=> true.
cleanup3  @ analyse_trail(S) \ '~>'(X,Y,S) <=> '~>'(X,Y,S+1).

```

Since these rules do not consider path compression as possible non-equivalence of trees. Indeed, they even will not undo the path compression after the implication check, but promote newly created edges to the higher stratum. Hence the compacter tree representation is retained after a succeeding implication check, making future operations cheaper.

### 5.3 Finite Domain Solver

The following CHR solver is a typical bounds propagation based finite domain solver (see [8]). It maintains bounds consistency for variables. With every variable  $X$  a `domain(X,L,U)` constraint is associated that maintains the lower and upper bounds,  $L$  and  $U$  respectively, of  $X$ 's domain. Rule `consistency` ensures that the domain is non-empty and rule `intersect` retains the intersection, if two domains exist for  $X$ .

Constraint propagators, like `dom_eq/2` and `dom_leq/2`, propagate new domains for the involved variables, each time the domain of any of these variables changes. Propagators for other finite domain constraints can be defined analogously.

```

consistent @ domain(X,L,U) ==> L =< U.
redundant  @ domain(X,L1,U1) \ domain(X,L2,U2) <=>
             L1 >= L2, U1 =< U2 | true.
intersect  @ domain(X,L1,U1), domain(X,L2,U2) <=>
             domain(X,max(L1,L2),min(U1,U2)).
eq         @ dom_eq(X,Y), domain(X,LX,UX), domain(Y,LY,UY) ==>
             domain(X,LY,UY), domain(Y,LX,UX).
leq        @ dom_leq(X,Y), domain(X,LX,UX), domain(Y,LY,UY) ==>
             domain(X,LX,UY), domain(Y,LX,UY).

```

With these rules and our implication checking it is possible to ask whether `domain(X,L,U)` is implied. Either the checked domain is empty and the check fails correctly or the `redundant` rule holds and the check succeeds, or `intersect` rule will combine the check with the domain already present in the store creating a new, smaller domain. Then the trail analysis discovers the change and fails the check. Otherwise the check correctly succeeds. Hence the implication check is optimal for `domain/3` checks.

As the `dom_eq/2` and `dom_leq/3` propagators are never removed, checking for their implication always fails.

	Approach	$solve(C)$	$solve(C \wedge c)$	$C \implies c$	
$c = eq(V_1, V_n)$	copy	4.34	4.42	9.16	100.0%
	trailing	5.02	5.07	5.07	55.2%
$c = eq(V_1, V_{n+1})$	copy	4.34	5.61	9.95	100.0%
	trailing	5.02	6.53	6.53	65.6%
$c = union(V_1, V_n)$	copy	2.78	2.80	5.58	100.0%
	trailing	3.75	3.77	3.77	67.6%
$c = union(V_1, V_{n+1})$	copy	2.78	2.79	5.57	100.0%
	trailing	3.75	3.78	3.78	67.9%

Table 1: Experimental Results

## 6 Experimental Evaluation

We compare our trailing approach with the naive copy approach to validate it. For this reason we consider a particular benchmark for the *eq* solver.  $n - 1$  constraints  $eq(V_i, V_{i+1})$  are imposed for  $1 \leq i < n$ . This conjunction of equality constraints we call  $C$ . The constraint we test for,  $c$ , is  $eq(V_1, V_n)$  in one case and  $eq(V_1, V_{n+1})$  in the other. The former test succeeds and the latter fails.

Table 1 lists the experimental results in seconds obtained for this benchmark with  $n = 20$  using the K.U.Leuven CHR system in SWI-Prolog on an Intel Pentium 4 2.0GHz with 512MB of RAM. The total time to perform implication checking for the copy approach is equal to the sum of the times for  $solve(C)$  and  $solve(C \wedge c)$ <sup>1</sup>. With our trailing approach, the time to compute  $C \implies c$  corresponds with the time for  $solve(C \wedge c)$ . While there is about 16% overhead for ordinary use ( $solve(C)$ ), the trailing approach is clearly superior to the copy approach for implication testing: it is almost twice as fast for our benchmark. The succeeding test performs a little better than the failing one.

The table also lists the results for a similar benchmark using the naive union-find program. The trailing version has been specialized to not trails additions and removals of constraint that are never stored. The results are similar as for the *eq* solver.

## 7 Conclusion

In this paper we have presented a new approach for automatic implication checking in CHR solvers. We have established the soundness of our trailing approach as well as its completeness a for the class of canonical CHR-only solvers. In addition we have studied the completeness for several existing CHR solvers. Experimental evaluation supports our claim that the trailing approach is more efficient than a naive copying approach. In addition we have extended our trailing approach to CHR solver hierarchies. We can now use CHR constraints of

<sup>1</sup>The time to compare the constraint stores is negligible for this benchmark.

one solver in the guards of rules of another solver.

In future work we intend to extend implication checking techniques for larger classes of CHR solvers. In particular, we would like to be able to project away local variables in built-in and CHR constraints. An important challenge there is automatic inference of constraint projection for CHR constraint solvers, as opposed to simply have the user supply appropriate projection operations.

In addition, we would like to explore the possibility of automatically inferring specific conditions or events that cause constraints to be implied. This would replace user-supplied events. The combination of automatic implication checks with automatic events allows for automatic reified constraints.

## Acknowledgments

We would like to thank Marc Meister for commenting on a preliminary version of this paper. Part of this research was conducted while Tom Schrijvers was visiting the University of Melbourne in July 2004 and the Universität Ulm in November 2004.

## References

- [1] Abdennadher, S., *Operational semantics and confluence of constraint propagation rules*, in: G. Smolka, editor, *Proceedings of the Third International Conference on Principles and Practice of Constraint Programming*, 1997, pp. 252–266.
- [2] ACM, *The constraint programming working group*, Technical report, ACM-MIT SDRC Workshop, Report Outline (1996).
- [3] Duck, G., M. Garcia de la Banda and P. Stuckey, *Compiling ask constraints*, in: B. Demoen and V. Lifschitz, editors, *Proceedings of the 20th International Conference on Logic Programming*, LNCS (2004), pp. 105–119.
- [4] Duck, G. J., P. J. Stuckey, M. G. de la Banda and C. Holzbaur, *Extending arbitrary solvers with constraint handling rules*, in: *Proceedings of the 5th ACM SIGPLAN international conference on Principles and practice of declarative programming* (2003), pp. 79–90.
- [5] Duck, G. J., P. J. Stuckey, M. García de la Banda and C. Holzbaur, *The refined operational semantics of constraint handling rules*, in: *20th International Conference on Logic Programming (ICLP'04)*, Saint-Malo, France, 2004, pp. 90–104.
- [6] Frühwirth, T., *Entailment Simplification and Constraint Constructors for User-Defined Constraints*, in: *3rd Workshop on Constraint Logic Programming (WCLP 93)*, 1993.

- [7] Frühwirth, T., *Theory and practice of constraint handling rules*, Journal of Logic Programming **37** (1998), pp. 95–138.
- [8] Frühwirth, T. and S. Abdennadher, “Essentials of Constraint Programming,” Cognitive Technologies, Springer, 2003.
- [9] Holzbaur, C. and T. Frühwirth, *Compiling constraint handling rules into Prolog with attributed variables*, in: G. Nadathur, editor, *Proceedings of the International Conference on Principles and Practice of Declarative Programming*, number 1702 in LNCS (1999), pp. 117–133.
- [10] Holzbaur, C. and T. Frühwirth, *Constraint handling rules, special issue*, Journal of Applied Artificial Intelligence **14** (2000).
- [11] Holzbaur, C., M. García de la Banda, P. J. Stuckey and G. J. Duck, *Optimizing compilation of constraint handling rules in hal*, Special Issue of Theory and Practice of Logic Programming on Constraint Handling Rules (2005), to appear.
- [12] Schrijvers, T. and B. Demoen, *The K.U.Leuven CHR system: Implementation and application*, in: *First workshop on constraint handling rules: selected contributions*, 2004, pp. 1–5.
- [13] Schrijvers, T. and T. Frühwirth, *Implementing and Analysing Union-Find in CHR*, Report CW 389, K.U.Leuven, Department of Computer Science (2004).
- [14] Schrijvers, T. and T. Frühwirth, *Optimal Union-Find in Constraint Handling Rules*, Theory and Practice of Logic Programming (2005), to appear.
- [15] Stuckey, P. J. and M. Sulzmann, *A Theory of Overloading*, in: *Proceedings of the seventh ACM SIGPLAN international conference on Functional programming* (2002), pp. 167–178.

## A Source Code: Naive Union-Find

```

Naive Union-Find
make      @ make(X) <=> root(X).

union     @ union(X,Y) <=> find(X,A), find(Y,B), link(A,B).

findNode  @ X ~> PX \ find(X,R) <=> find(PX,R).
findRoot  @ root(X) \ find(X,R) <=> R=X.

linkEq    @ link(X,X) <=> true.
link      @ link(X,Y), root(X), root(Y) <=> Y ~> X, root(X).

```

## B Source Code: Optimal Union-Find

```

make      @ make(X) <=> root(X,0).

findNode  @ X ~> PX , find(X,R) <=> find(PX,R), X ~> R.
findRoot  @ root(X,_) \ find(X,R) <=> R=X.

linkEq    @ link(X,X) <=> true.
linkLeft  @ link(X,Y), root(X,RX) root(Y,RX) <=> RX >= RY |
           Y ~> X, NRX is max(RX,RX+1), root(X,NRX).
linkRight @ link(X,Y), root(Y,RX) root(X,RX) <=> RX >= RY |
           X ~> Y, NRY is max(RY,RX+1), root(Y,NRY).

```