Logical Algorithms meets CHR

A Meta-Complexity Result for Constraint Handling Rules with Rule Priorities

Leslie De Koninck

Department of Computer Science, K.U.Leuven, Belgium
(e-mail: *FirstName.LastName@cs.kuleuven.be*)

submitted 20 December 2007; revised ; accepted

Abstract

This paper investigates the relationship between the Logical Algorithms language (LA) of Ganzinger and McAllester and Constraint Handling Rules (CHR). We present a translation schema from LA to CHR^{P} : CHR with rule priorities, and show that the meta-complexity theorem for LA can be applied to a subset of CHR^{P} via inverse translation. Inspired by the high-level implementation proposal for Logical Algorithm by Ganzinger and McAllester and based on a new scheduling algorithm, we propose an alternative implementation for CHR^{P} that gives strong complexity guarantees and results in a new and accurate meta-complexity theorem for CHR^{P} . It is furthermore shown that the translation from Logical Algorithms to CHR^{P} combined with the new CHR^{P} implementation, satisfies the required complexity for the Logical Algorithms meta-complexity result to hold.

1 Introduction

Constraint Handling Rules (CHR) (Frühwirth 1998) is a high-level rule based language, originally designed for the implementation of constraint solvers, but also increasingly used as a general purpose programming language. Recently, it was shown that all algorithms can be implemented in CHR while preserving both time and space complexity (Sneyers et al. 2005). We assume some familiarity with CHR and refer to (Frühwirth 1998) for more details.

In “Logical Algorithms” (LA) (Ganzinger and McAllester 2002) (and based on previous work in (Ganzinger and McAllester 2001; McAllester 1999)), Ganzinger and McAllester present a bottom-up logic programming language for the purpose of facilitating the derivation of complexity results of algorithms described by logical inference rules. This problem is far from trivial because the runtime is not necessarily proportional to the derivation length (i.e., the number of rule applications), but also includes the cost of pattern matching for multi-headed rules, as well as the costs related to high-level execution control which is specified using rule priorities in the Logical Algorithms language. The language of Ganzinger and McAllester resembles CHR in many ways and has often been referred to in the discussion of complexity results of CHR programs (Christiansen 2005; Frühwirth 2002b; Schrijvers and Frühwirth 2006; Sneyers et al. 2006a). In particular, in (Christiansen 2005), Christiansen uses the meta-complexity theorem that accompanies the Logical Algorithms language, and notes that the CHR system used (SICStus CHR by

Holzbaaur et al. (Holzbaaur and Frühwirth 1998)) does not always exhibit the right complexity because previously computed partial rule matches are not stored.

The aim of this paper is to investigate the relationship between both languages. More precisely, we look at how the meta-complexity theorem for Logical Algorithms can be applied to (a subset of) CHR, and how CHR can be used to implement Logical Algorithms with the correct complexity. First, we present a translation schema from Logical Algorithms to CHR^{P} : CHR extended with rule priorities (De Koninck et al. 2007b). Logical Algorithms derivations of the original program correspond to CHR^{P} derivations in the translation and vice versa. We also show how to translate a subclass of CHR^{P} programs into Logical Algorithms. This allows us to apply the meta-complexity theorem for Logical Algorithms to these CHR^{P} programs as well. Because the Logical Algorithms meta-complexity theorem is based on an optimized implementation, it gives more accurate results than the implementation independent meta-complexity theorem of (Frühwirth 2002a; Frühwirth 2002b) while being more general than the ad-hoc complexity derivations in (Schrijvers and Frühwirth 2006; Sneyers et al. 2006a).

Our current implementation of CHR^{P} as presented in (De Koninck et al. 2007) does not guarantee the complexity required for the meta-complexity theorem for Logical Algorithms to hold via translation to CHR^{P} . Another issue is that the translation from CHR^{P} to Logical Algorithms is restricted to a subset of CHR^{P} . Therefore, we propose a new implementation of CHR^{P} , designed such that it supports a new meta-complexity theorem for the complete CHR^{P} language, while also ensuring that Logical Algorithms programs translated into CHR^{P} are executed with the correct complexity.

The implementation is based on the high-level implementation proposal for Logical Algorithms as given in (Ganzinger and McAllester 2002), and on a new scheduling data structure proposed in (De Koninck 2007). By using a CHR system with advanced indexing support, such as the K.U.Leuven CHR system (Schrijvers and Demoen 2004), our implementation achieves the complexity required to enable a new and accurate meta-complexity result for the whole CHR^{P} language.

Overview The rest of this paper is organized as follows. In Section 2, the syntax and semantics of the Logical Algorithms language and CHR^{P} are reviewed and the known meta-complexity theorems for both languages are presented. In Section 3 a translation of LA programs to CHR^{P} programs is presented and in Section 4, the opposite is done for a subset of CHR^{P} . Section 5 proposes an alternative implementation for CHR^{P} which enables a new meta-complexity theorem for this language, given in Section 6. Some concluding remarks are given in Section 7.

2 Logical Algorithms and CHR^{P}

In this section, we give an overview of the syntax and semantics of Logical Algorithms (Section 2.1) and CHR^{P} (Section 2.2). In Section 2.3, we review the meta-complexity results that are known for both languages.

2.1 Logical Algorithms

This subsection gives an overview of the syntax and semantics of the Logical Algorithms language.

2.1.1 Syntax

A Logical Algorithms program $P = \{r_1, \dots, r_n\}$ is a set of rules. In (Ganzinger and McAllester 2002), a graphical notation is used to represent rules. We use a new textual representation that is closer to the syntax of CHR. A Logical Algorithms rule is an expression

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

where r is the rule *name*, the atoms A_i (for $1 \leq i \leq n$) are the *antecedents* and C is the *conclusion*, which is a conjunction of atoms whose variables appear in the antecedents. Rule r has *priority* p where p is an arithmetic expression whose variables (if any) occur in the first antecedent A_1 . If p contains variables, then r is called a dynamic priority rule. Otherwise, it is called a static priority rule. In the graphical notation of (Ganzinger and McAllester 2002), the above rule is represented as shown below.

$$\begin{array}{c} A_1 \\ \vdots \\ A_n \\ (r,p) \text{ ---} \\ C \end{array}$$

The arguments of an atom are either Herbrand terms or (integer) arithmetic expressions. There are two types of atoms: comparisons and user-defined atoms. A comparison has the form $x < y$, $x \leq y$, $x = y$ or $x \neq y$ with x and y arithmetic expressions or, in case of $(=)/2$ and $(\neq)/2$, Herbrand terms. Comparisons are only allowed in the antecedents of a rule and all variables in a comparison must appear in earlier antecedents. A user-defined atom can be positive or negative. A negative user-defined atom has the form $del(A)$ where A is a positive user-defined atom. A ground user-defined atom is called an assertion.

Example 1

An example rule (from Dijkstra's shortest path algorithm as presented in (Ganzinger and McAllester 2002)) with name `d2` and priority 1 is

```
d2 @ 1 : dist(V,D1), dist(V,D2), D2 < D1 => del(dist(V,D1)).
```

The antecedent $D_2 < D_1$ is a comparison, the atoms `dist(V,D1)` and `dist(V,D2)` are positive user-defined antecedents. The negative ground atom `del(dist(a,5))` is an example of a negative assertion.

2.1.2 Operational Semantics

A Logical Algorithms state σ consists of a set of (positive and negative) assertions. A state can simultaneously contain the positive assertion A and the negative assertion $del(A)$. Let $\mathcal{D}$ be the usual interpretation for the comparisons. Given a program P , the following transition converts one state into the next:

- 1. Apply** $\sigma \xrightarrow{LA}_P \sigma \cup \theta(C)$ if there exists a (renamed apart) rule r in P of priority p of the form

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

and a ground substitution θ such that for every antecedent A_i ,

- $\mathcal{D} \models \theta(A_i)$ if A_i is a comparison
- $\theta(A_i) \in \sigma$ and $del(\theta(A_i)) \notin \sigma$ if A_i is a positive user-defined atom
- $\theta(A_i) \in \sigma$ if A_i is a negative user-defined atom

Furthermore, $\theta(C) \not\subseteq \sigma$ and no rule of priority p' and substitution θ' exists with $\theta'(p') < \theta(p)$ for which the above conditions hold.

A state is called final if no more transitions apply to it. A non-final state has priority p if the next firing rule instance has priority p . The condition $\theta(C) \not\subseteq \sigma$ ensures that no rule instance fires more than once and prevents trivial non-termination. This condition, combined with the fact that each transition only creates new assertions, causes the consecutive states in a derivation to be monotone increasing. Although the priorities restrict the possible derivations, the choice of which rule instance to fire from those with equal priority is non-deterministic.

2.2 CHR^{rp}: CHR with Rule Priorities

CHR^{rp} is CHR extended with user-definable rule priorities. It is introduced in (De Koninck et al. 2007b) as a solution to the lack of high-level execution control in CHR. In this section, we review the syntax and semantics of CHR^{rp}.

2.2.1 Syntax

A constraint $c(t_1, \dots, t_n)$ is an atom of predicate c/n with t_i a host language value (e.g., a Herbrand term in Prolog) for $1 \leq i \leq n$. There are two types of constraints: built-in constraints and CHR constraints (also called user-defined constraints). The CHR constraints are solved by the CHR program whereas the built-in constraints are solved by an underlying constraint solver (e.g., the Prolog unification algorithm).

There are three types of Constraint Handling Rules: *simplification rules*, *propagation rules* and *simpagation rules*. They have the following form:

$$\begin{array}{lll} \textbf{Simplification} & p :: r @ & H^r \iff g \mid B \\ \textbf{Propagation} & p :: r @ H^k & \implies g \mid B \\ \textbf{Simpagation} & p :: r @ H^k \setminus H^r & \iff g \mid B \end{array}$$

where p is the rule priority, r is the rule *name*, H^k and H^r are non-empty sequences

1. **Solve** $\langle \{c\} \uplus G, S, B, T \rangle_n \xrightarrow{\omega_p}_P \langle G, S, c \wedge B, T \rangle_n$ where c is a built-in constraint.
2. **Introduce** $\langle \{c\} \uplus G, S, B, T \rangle_n \xrightarrow{\omega_p}_P \langle G, \{c\#n\} \cup S, B, T \rangle_{n+1}$ where c is a CHR constraint.
3. **Apply** $\langle \emptyset, H_1 \cup H_2 \cup S, B, T \rangle_n \xrightarrow{\omega_p}_P \langle \theta(C), H_1 \cup S, B, T \cup \{t\} \rangle_n$ where P contains a rule of priority p of the form

$$p :: r @ H'_1 \setminus H'_2 \iff g \mid C$$

and a matching substitution θ such that $\text{chr}(H_1) = \theta(H'_1)$, $\text{chr}(H_2) = \theta(H'_2)$, $\mathcal{D} \models B \rightarrow \exists_B(\theta \wedge g)$, $\theta(p)$ is a ground arithmetic expression and $t = \langle r, \text{id}(H_1) ++ \text{id}(H_2) \rangle \notin T$. Furthermore, no rule of priority p' and substitution θ' exists with $\theta'(p') < \theta(p)$ for which the above conditions hold.

 Table 1. Transitions of ω_p

of CHR constraints and are called the *heads* of the rule. The rule *guard* g is a sequence of built-in constraints and the rule *body* B is a sequence of both CHR and built-in constraints. The rule priority is either a number in which case the rule is called a *static* priority rule, or an arithmetic expression whose variables appear in the heads H^k and/or H^r in which case the rule is called a *dynamic* priority rule. We say that priority p is higher than priority p' if $p < p'$. For simplicity, we sometimes assume priorities are integers and the highest priority is 1. Finally, a program P is a set of CHR rules. Apart from the rule priorities, CHR^{P} is identical to CHR.

2.2.2 Operational Semantics

Operationally, CHR constraints have a multi-set semantics. To distinguish between different occurrences of syntactically equal constraints, CHR constraints are extended with a unique identifier. An identified CHR constraint is denoted by $c\#i$ with c a CHR constraint and i the identifier. We write $\text{chr}(c\#i) = c$ and $\text{id}(c\#i) = i$.

The operational semantics of CHR^{P} , called the priority semantics and denoted by ω_p , is given in (De Koninck et al. 2007b) as a state transition system, similar to the approach of (Duck et al. 2004) for the theoretical and refined operational semantics of CHR. A CHR execution state σ is represented as a tuple $\langle G, S, B, T \rangle_n$ where G is the goal, a multi-set of constraints to be solved; S is the CHR constraint store, a set of identified CHR constraints; B is the built-in store, a conjunction of built-in constraints; T is the propagation history, a set of tuples denoting the rule instances that have already fired; and n is the next free identifier, used to identify new CHR constraints. The transitions of ω_p are shown in Table 1 where $\mathcal{D}$ denotes the built-in constraint theory and $\exists_X Y$ denotes the existential closure of Y apart from the variables in X . The transitions are exhaustively applied starting from the state $\langle G, \emptyset, \text{true}, \emptyset \rangle_1$ with G the initial goal.

The following theorem on the correspondence between the ω_p semantics of CHR^{P} and the ω_t semantics of CHR, is proven in (De Koninck et al. 2007b).

Theorem 1

Every derivation D under ω_p is also a derivation under ω_t . If a state σ is a final state under ω_p , then it is also a final state under ω_t .

In the refined operational semantics of CHR (Duck et al. 2004), the textual order of the program rules determines which rule is tried next for the current *active* constraint. However, only rule instances in which the active constraint takes part are considered, and so a higher priority fireable rule instance in which the active constraint does not participate, will not fire. The textual rule order also does not support dynamic rule priorities.

2.2.3 Differences compared to Logical Algorithms

CHR^{rp} differs from Logical Algorithms in the following ways:

- A Logical Algorithms state is a set of ground assertions, while the CHR constraint store is a multi-set and may also contain non-ground constraints.
- In Logical Algorithms, built-in constraints are restricted to ask constraints and only include comparisons; CHR^{rp} supports any kind of built-in constraints.
- A removed CHR constraint may be reasserted and can then participate again in rule firings whereas a removed LA assertion cannot be asserted again.
- A Logical Algorithms rule may contain negated heads. In contrast, CHR^{rp} requires all heads to be positive.¹
- In the Logical Algorithms language, the priority of a dynamic priority rule is determined by the variables in the left-most head, whereas in CHR^{rp} it may depend on multiple heads.

We note that rules for which the priority depends on more than one head, can easily be transformed into the correct form as follows. Given a Logical Algorithms rule of the form

$$r @ p : A_1, \dots, A_m, A_{m+1}, \dots, A_n \Rightarrow C$$

where the priority expression p is fully determined by the variables from the antecedents $A_1, \dots, A_m$. This rule can be transformed into the equivalent rules

$$\begin{aligned} r_1 @ 1 : A_1, \dots, A_m &\Rightarrow \text{priority}_r(p) \\ r_2 @ p : \text{priority}_r(p), A_1, \dots, A_m, A_{m+1}, \dots, A_n &\Rightarrow C \end{aligned}$$

where $\text{priority}_r/1$ is a new user-defined predicate. Now the first head of the dynamic priority rule determines the rule priority.

2.3 Meta-Complexity Results

The Logical Algorithms language was designed with a meta-complexity result in mind. Such a result has also been formulated for CHR. In this subsection, we review both results and give a first intuition on how they relate to each other.

¹ See (Van Weert et al. 2006) for an extension of CHR with negation as absence.

2.3.1 The Logical Algorithms Meta-Complexity Result

A *prefix instance* of a Logical Algorithms rule $r @ p : A_1, \dots, A_n \Rightarrow C$ is a tuple $\langle \theta(r), i \rangle$ with θ a ground substitution defined on the variables occurring in $A_1, \dots, A_i$ and $1 \leq i \leq n$. Its antecedents are $\theta(A_1), \dots, \theta(A_i)$. A *strong prefix firing* is a prefix instance whose antecedents hold in a state with priority lower or equal to the prefix' rule priority.² The time complexity for running Logical Algorithms programs is given in (Ganzinger and McAllester 2002) as $\mathcal{O}(|\sigma_0| + P_s + (P_d + A_d) \cdot \log N)$ where σ_0 is the initial state and $|\sigma_0|$ is its size. P_s is the number of strong prefix firings of static priority rules and P_d is the number of strong prefix firings of dynamic priority rules; A_d is the number of assertions that may participate in a dynamic priority rule instance; and N is the number of distinct priorities. The following example is adapted from (Ganzinger and McAllester 2002).

Example 2 (Dijkstra's Shortest Path)

The rules below implement Dijkstra's single source shortest path algorithm.

```

d1 @ 1    : source(V) => dist(V,0).
d2 @ 1    : dist(V,D1), dist(V,D2), D2 < D1 => del(dist(V,D1)).
d3 @ D+2  : dist(V,D), e(V,C,U) => dist(U,D+C).
    
```

A `source(V)` fact means that V is the (unique) source node for the algorithm. A `dist(V,D)` fact means that the shortest path distance from the source node to node V does not exceed D . Finally, an `e(V,C,U)` fact means that there is an edge from node V to node U with cost (weight) C . Given an initial state consisting of one `source/1` fact and e `e/3` facts, we can derive that the number of strong prefix firings is $\mathcal{O}(1)$ for rule `d1`, and $\mathcal{O}(e)$ for both rules `d2` and `d3`. This result is based on the fact that at priority 2 and lower, there is at most one (positive) `dist/2` fact for each node, and each of these facts represent the shortest path distance from the source node to this node. This means that at most e `dist/2` facts are ever created. Using the meta-complexity theorem, we find that the total complexity is $\mathcal{O}(e \log e)$.

2.3.2 The "As Time Goes By" Approach

In (Frühwirth 2002a; Frühwirth 2002b), an upper bound on the worst case time complexity of a CHR program P is given as

$$\mathcal{O} \left(D \sum_{r \in P} (c_{max}^{n_r} (O_{H_r} + O_{G_r}) + (O_{C_r} + O_{B_r})) \right) \quad (1)$$

where D is the maximal derivation length (i.e., the maximal number of rule firings), c_{max} is the maximal number of CHR constraints in the store, and for each rule $r \in P$:

- n_r is the number of heads in r

² In (Ganzinger and McAllester 2002), also the concept of a *weak* prefix firing is defined, but it is of no importance for our purposes.

- O_{H_r} is the cost of head matching, i.e. checking that a given sequence of n_r constraints match with the n_r heads of rule r
- O_{G_r} is the cost of checking the guard
- O_{C_r} is the cost of adding built-in constraints after firing
- O_{B_r} is the cost of adding and removing CHR constraints after firing

For programs with simplification and simpagation rules only, the maximal derivation length can be derived using an appropriate ranking on constraints that decreases after each rule firing (Frühwirth 2000a). We note that finding such a ranking is not trivial. The meta-complexity result is based on a very naive CHR implementation, and therefore on the one hand gives an upper bound on the time complexity for any reasonable implementation of CHR, but on the other hand often largely overestimates the worst case time complexity on optimized implementations.³ The following example is adapted from (Frühwirth 2002a).

Example 3 (Boolean)

The rules below implement the boolean $\text{and}(X, Y, X \wedge Y)$ constraint given that 1 represents *true* and 0 represents *false*.

$\text{and}(0, Y, Z) \iff Z = 0.$	$\text{and}(X, 0, Z) \iff Z = 0.$
$\text{and}(X, 1, Z) \iff X = Z.$	$\text{and}(1, Y, Z) \iff Y = Z.$
$\text{and}(X, X, Z) \iff X = Z.$	$\text{and}(X, Y, 1) \iff X = 1, Y = 1.$

Let the rank of an $\text{and}/3$ constraint be one, then the rank of the head of each rule equals one, and the rank of the body equals zero.⁴ For a goal consisting of n $\text{and}/3$ constraints, the derivation length is n , which is also the maximal number of CHR constraints in the store. The cost of head matching, (implicit) guard checking, removing CHR constraints and asserting built-in constraints can all be considered constant. Then using (1), we derive that the total runtime complexity is $\mathcal{O}(n^2)$.

2.3.3 A First Comparison

Although at this point we do not intend to make a complete comparison between both results, we can already show that the Logical Algorithms result in a sense is at least as accurate as Frühwirth's approach, at least as far as programs without built-in tell constraints are concerned. The reasoning is as follows. In each derivation step, a constant number of atoms (constraints) are asserted. Let c_{max} be the maximal number of (strictly) positive atoms in the database in any given state. Furthermore assume rules have positive heads only, then each of the asserted atoms can participate in at most $\sum_{r \in P} (n_r \cdot c_{max}^{n_r-1})$ strong prefix firings. Because only $\mathcal{O}(c + D)$ constraints are ever asserted where c is the number of CHR constraints

³ Built-in constraints may lead to a worse complexity in practical optimized implementations if many constraints are repeatedly reactivated without this resulting in new rule firings. We return to this issue in Section 6.3.

⁴ Built-in constraints have a rank of zero by definition.

in the initial goal and D is the derivation length, the total number of strong prefix firings $P_s + P_d$ is

$$\mathcal{O} \left((c + D) \cdot \sum_{r \in P} c_{max}^{n_r - 1} \right)$$

and because $c = \mathcal{O}(c_{max})$ we also have the following bound

$$\mathcal{O} \left(D \cdot \sum_{r \in P} c_{max}^{n_r} \right) \quad (2)$$

In absence of (dynamic) priorities, the total runtime complexity according to the Logical Algorithms meta-complexity result is bounded by the same formula (2) and hence is at least as accurate as the result of (Frühwirth 2002b) given that the cost of both head matching (O_{H_r}) and adding and removing CHR constraints (O_{B_r}) is constant for each rule r .

3 Translating Logical Algorithms into CHR^{rp}

In this section, we show how Logical Algorithms programs can be translated into CHR^{rp} programs. CHR states of the translated program can be mapped onto LA states of the original. With respect to this mapping, both programs have the same derivations.

3.1 The Translation Schema

The translation of a LA program P is denoted by $T(P) = T_{S/D}(P) \cup T_R(P)$. The definitions of $T_{S/D}(P)$ and $T_R(P)$ are given below.

3.1.1 Set and Deletion Semantics

We represent Logical Algorithms assertions as CHR constraints consisting of the assertion itself and an extra argument, called the *mode indicator*, denoting whether it is positively asserted (“p”), negatively asserted (“n”) or both (“b”). For every user-defined predicate a/n occurring in P , $T_{S/D}(P)$ contains the following rules to deal with a new positive or negative assertion:

$$\begin{aligned} 1 &:: a_r(\bar{X}, M) \setminus a(\bar{X}) \iff M \neq \mathbf{n} \mid true \\ &1 :: a_r(\bar{X}, \mathbf{n}), a(\bar{X}) \iff a_r(\bar{X}, \mathbf{b}) \\ &2 :: a(\bar{X}) \iff a_r(\bar{X}, \mathbf{p}) \\ 1 &:: a_r(\bar{X}, M) \setminus del(a(\bar{X})) \iff M \neq \mathbf{p} \mid true \\ &1 :: a_r(\bar{X}, \mathbf{p}), del(a(\bar{X})) \iff a_r(\bar{X}, \mathbf{b}) \\ &2 :: del(a(\bar{X})) \iff a_r(\bar{X}, \mathbf{n}) \end{aligned}$$

If a representation already exists, one of the priority 1 rules updates this representation. Otherwise, one of the priority 2 rules generates a new representation. At

lower priorities, it is guaranteed that every assertion, whether asserted positively, negatively or both, is represented by exactly one constraint in the store.

3.1.2 Rules

Given a LA rule $r \in P$ of the form

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

We first split up the antecedents into user-defined antecedents and comparison antecedents by using the `split` function defined below.

$$\begin{aligned} \text{split}([A|T]) &= \begin{cases} \langle [A|A^u], A^c \rangle & \text{if } A \text{ is a user-defined atom} \\ \langle A^u, [A|A^c] \rangle & \text{if } A \text{ is a comparison} \end{cases} \\ &\quad \text{where } \text{split}(T) = \langle A^u, A^c \rangle \\ \text{split}([]) &= \langle [], [] \rangle \end{aligned}$$

In the Logical Algorithms language, a given assertion may participate multiple times in the same rule instance, whereas in CHR all constraints in a single rule instance must be different. To overcome this semantic difference, a single LA rule is translated as a set of CHR rules such that every CHR rule covers a case of syntactically equal head constraints. Let $\langle A^u, A^c \rangle = \text{split}([A_1, \dots, A_n])$ with $A^u = [A_1^u, \dots, A_m^u]$ and $A^c = [A_1^c, \dots, A_l^c]$. Let $\mathcal{P}$ be the set of all partitions of $\{1, \dots, m\}$.⁵ For a given partition $\rho \in \mathcal{P}$, the following function returns the most general unifier that unifies all antecedents $\{A_i \mid i \in S\}$ for every $S \in \rho$ where $\text{mgu}(S)$ is the most general unifier of all elements in S .

$$\text{partition_to_mgu}(\rho, [A_1^u, \dots, A_m^u]) = \bigcirc_{S \in \rho} \text{mgu}(\{A_i^u \mid i \in S\})$$

Let $\mathcal{PU} = \{\langle \rho, \theta \rangle \mid \rho \in \mathcal{P} \wedge \theta = \text{partition_to_mgu}(\rho, A^u) \wedge \mathcal{D} \models \exists_{\theta} \theta(A^c)\}$. $\mathcal{PU}$ contains all partitions for which `partition_to_mgu` is defined and for which the comparison antecedents A^c are still satisfiable after applying the unifier. The next step is to filter out antecedents so that every set in the partition has only one representative. This is done by computing $\text{filter}(\theta(A^u), \rho)$ for each $\langle \rho, \theta \rangle \in \mathcal{PU}$ where the `filter` function is as follows:

$$\begin{aligned} \text{filter}([\theta(A_i^u)|T], \rho) &= \begin{cases} [\theta(A_i^u)|\text{filter}(T, \rho)] & \text{if } \exists S \in \rho : i = \min(S) \\ \text{filter}(T, \rho) & \text{otherwise} \end{cases} \\ \text{filter}([], _) &= [] \end{aligned}$$

⁵ $\mathcal{P}$ contains B_m elements in the worst case with B_m the m^{th} Bell number.

Finally, we add mode indicators to all remaining user-defined antecedents:

$$\begin{aligned} \text{modes}([A^{u'}|T]) &= \begin{cases} \langle [a_r(\bar{X}, \mathbf{p})|A^m], N \rangle & \text{if } A^{u'} = a(\bar{X}) \\ \langle [a_r(\bar{X}, N')|A^m], [N' \neq \mathbf{p}|N] \rangle & \text{if } A^{u'} = \text{del}(a(\bar{X})) \end{cases} \\ &\quad \text{where } \langle A^m, N \rangle = \text{modes}(T) \\ \text{modes}([]) &= \langle [], [] \rangle \end{aligned}$$

The `modes` function returns both the resulting antecedents and the necessary conditions on the mode indicators of these antecedents. For every $\langle \rho, \theta \rangle \in \mathcal{PU}$, the CHR translation $T_R(P)$ contains a rule

$$p + 2 :: r_\rho @ H \implies g_1, g_2 \mid C'$$

where $\langle H, g_1 \rangle = \text{modes}(\text{filter}(\theta(A^u), \rho))$, $g_2 = \theta(A^c)$ and $C' = \theta(C)$.

3.1.3 Examples

We illustrate the translation schema on some examples.

Example 4

A LA implementation of Dijkstra's shortest path algorithm is

```
d1 @ 1 : source(V) => dist(V,0).
d2 @ 1 : dist(V,D1), dist(V,D2), D2 < D1 => del(dist(V,D1)).
d3 @ D+2 : dist(V,D), e(V,C,U) => dist(U,D+C).
```

Its translation is

```
1 :: e_r(V,C,U,M) \ e(V,C,U) <=> M \= n | true.
1 :: e_r(V,C,U,n) , e(V,C,U) <=> e_r(V,C,U,b).
2 :: e(V,C,U) <=> e_r(V,C,U,p).

1 :: e_r(V,C,U,M) \ del(e(V,C,U)) <=> M \= p | true.
1 :: e_r(V,C,U,p) , del(e(V,C,U)) <=> e_r(V,C,U,b).
2 :: del(e(V,C,U)) <=> e_r(V,C,U,n).

... % (similar rules for source/1 and dist/2)

3 :: d1_1 @ source_r(V,p) ==> dist(V,0).
3 :: d2_{1/2} @ dist_r(V,D1,p), dist_r(V,D2,p) ==> D2 < D1 | del(dist(V,D1)).
D+4 :: d3_{1/2} @ dist_r(V,D,p), e_r(V,C,U,p) ==> dist(U,D+C).
```

Example 5

A rule from the union-find implementation of (Ganzinger and McAllester 2002) is the following:

```
uf4 @ 1 : union(X,Y), find(X,Z), find(Y,Z) => del(union(X,Y)).
```

Because antecedents `find(X,Z)` and `find(Y,Z)` are unifiable, this leads to the following two CHR rules:

```

3 :: uf41/2/3 @ unionr(X,Y,p), findr(X,Z,p), findr(Y,Z,p) ==> del(union(X,Y)).
3 :: uf41/23 @ unionr(X,X,p), findr(X,Z,p) ==> del(union(X,X)).

```

3.2 The Correspondence between LA and CHR^{TP} Derivations

In this subsection, we show that every derivation of the original program under the Logical Algorithms semantics, corresponds to a derivation of the translation under the ω_p semantics of CHR^{TP}. In order to do so, we introduce a mapping function `chr_to_la` between *reachable* CHR execution states and Logical Algorithms states.⁶ Reachability is considered with respect to initial states of the form $\langle G, \emptyset, true, \emptyset \rangle_n$ where the user-defined constraints in G are of the form $a(\bar{X})$ and $del(a(\bar{X}))$ and do not include constraints of the form $a_r(\bar{X}, M)$.

$$\begin{aligned} \text{chr_to_la}(\sigma) = \{ & a(\bar{X}) \mid a(\bar{X}) \in A \vee (a_r(\bar{X}, M) \in A \wedge M \neq \mathbf{n}) \} \\ & \cup \{ del(a(\bar{X})) \mid del(a(\bar{X})) \in A \vee (a_r(\bar{X}, M) \in A \wedge M \neq \mathbf{p}) \} \end{aligned}$$

where $\sigma = \langle G, S, B, T \rangle_n$ and $A = G \cup \text{chr}(S)$. The mapping function also takes into account the constraints that are still in the goal and those for which the set and deletion semantics rules have not yet fired. In the rest of this section, we first show how CHR execution states are normalized and then show that in a Logical Algorithms state and its corresponding normalized CHR execution state, corresponding rule instances can fire. We start by defining a *pre-normal form*.

Definition 1 (Pre-normal Form)

A (reachable) state σ is in pre-normal form if and only if $\sigma = \langle \emptyset, S, true, T \rangle_n$, all constraints in S are of the form $a_r(\bar{X}, M) \# i$, and if $a_r(\bar{X}, M_1) \# i_1 \in S$ and $a_r(\bar{X}, M_2) \# i_2 \in S$ then $i_1 = i_2$ (and consequently $M_1 = M_2$).

The following lemma shows that every reachable state is *pre-normalized* before rules are tried with priority > 2 .

Lemma 1 (Pre-normalization)

For every reachable state σ , there exists a finite derivation $D = \sigma \xrightarrow{\omega_p^*}_{T(P)} \sigma^*$ such that σ^* is in pre-normal form, $\text{chr_to_la}(\sigma) = \text{chr_to_la}(\sigma^*)$, and all rules fired in D have priority 1 or 2. Every state has a unique pre-normal form with respect to the `chr_to_la` mapping function.

Proof

We introduce the following ranking function on CHR states:

$$\|\sigma\| = 2 \cdot |\{a(\bar{X}) \mid a(\bar{X}) \in A\} \uplus \{del(a(\bar{X})) \mid del(a(\bar{X})) \in A\}| + |G|$$

where $\sigma = \langle G, S, true, T \rangle_n$, $A = G \uplus \text{chr}(S)$ and if X is a (multi-)set, $|X|$ is its cardinality. Clearly, the rank of any state is positive, and if $\|\sigma\| = 0$, state σ is in pre-normal form. If σ is not in pre-normal form, then there exists at least one transition

⁶ See (Duck et al. 2007) for a formal definition of reachability.

$\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$. We show that for all such transitions $\text{chr_to_la}(\sigma) = \text{chr_to_la}(\sigma')$ and $\|\sigma'\| < \|\sigma\|$, which ensures termination.

If the goal G is not empty, then only the **Introduce** transition is applicable. Every application of this transition moves a CHR constraint from the goal to the CHR constraint store, so $\|\sigma'\| = \|\sigma\| - 1$. By definition, $\text{chr_to_la}(\sigma') = \text{chr_to_la}(\sigma)$ (because the chr_to_la function does not distinguish between the goal and the CHR constraint store).

If the goal G is empty then given that σ is not in pre-normal form, $\text{chr}(S)$ contains a constraint of the form $a(\bar{X})$ or $\text{del}(a(\bar{X}))$. We look into detail to the case of $a(\bar{X}) \in \text{chr}(S)$; the case of $\text{del}(a(\bar{X})) \in \text{chr}(S)$ is similar. We start by showing that at least one rule of priority 1 or 2 is applicable. Next, we show that each rule application decreases the norm and maintains the invariance with respect to the chr_to_la function.

Assume $a(\bar{X}) \in \text{chr}(S)$. If $a_r(\bar{X}, \mathbf{p}) \in \text{chr}(S)$ or $a_r(\bar{X}, \mathbf{b}) \in \text{chr}(S)$ then the following rule of $T(P)$ is applicable:

$$1 :: a_r(\bar{X}, M) \setminus a(\bar{X}) \iff M \neq \mathbf{n} \mid \text{true}$$

If $a_r(\bar{X}, \mathbf{n}) \in \text{chr}(S)$ then the rule below applies:

$$1 :: a_r(\bar{X}, \mathbf{n}), a(\bar{X}) \iff a_r(\bar{X}, \mathbf{b})$$

Finally, if no rule of priority 1 can be applied, which implies that no constraint of the form $a_r(\bar{X}, M) \in \text{chr}(S)$, then the following $T(P)$ rule can fire:

$$2 :: a(\bar{X}) \iff a_r(\bar{X}, \mathbf{p})$$

This covers all possibilities. Now we look at what happens after firing one of the priority 1 or 2 rules. The rule

$$1 :: a_r(\bar{X}, M) \setminus a(\bar{X}) \iff M \neq \mathbf{n} \mid \text{true}$$

removes a constraint $a(\bar{X})\#i$ from S and has an empty body, so $\|\sigma'\| = \|\sigma\| - 2$. Since $M \neq \mathbf{n}$ the removed constraint was already represented by the $a_r(\bar{X}, M)$ constraint and so $\text{chr_to_la}(\sigma') = \text{chr_to_la}(\sigma)$. Firing

$$1 :: a_r(\bar{X}, \mathbf{n}), a(\bar{X}) \iff a_r(\bar{X}, \mathbf{b})$$

causes the removal of two constraints from S , namely $a_r(\bar{X}, \mathbf{n})\#i$ and $a(\bar{X})\#j$. Furthermore, it adds a new constraint $a_r(\bar{X}, \mathbf{b})$ to G . This results in $\|\sigma'\| = \|\sigma\| - 1$. The new constraint represents the combined mode of both removed constraints and hence $\text{chr_to_la}(\sigma') = \text{chr_to_la}(\sigma)$. Finally, the rule

$$2 :: a(\bar{X}) \iff a_r(\bar{X}, \mathbf{p})$$

is only applicable if $\text{chr}(S)$ does not contain a constraint of the form $a_r(\bar{X}, M)$. It removes a constraint $a(\bar{X})\#i$ from S and adds a new constraint $a_r(\bar{X}, \mathbf{p})$ to G , resulting in $\|\sigma'\| = \|\sigma\| - 1$. The new representation covers the positive assertion and so $\text{chr_to_la}(\sigma') = \text{chr_to_la}(\sigma)$.

In summary, if the goal is empty and σ is not in pre-normal form, a rule of

priority 1 or 2 can fire and so no rule with lower priority is applicable. All applicable transitions strictly decrease the value of the ranking function and so the pre-normalization terminates. Finally, none of the possible transitions changes the value of `chr_to_la`. $\square$

The state σ^* is called a pre-normalization of σ .

Definition 2 (Implied Rule Instance)

A rule instance $\theta(r)$ is implied in a state σ if $\theta(C) \subseteq \text{chr_to_la}(\sigma)$ with $\theta(\sigma)$ the conclusion of $\theta(r)$.

Lemma 2 (Normalization)

Let there be given a pre-normalized state $\sigma = \langle \emptyset, S, \text{true}, T \rangle_n$. If there exists a transition $\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$ in which an implied rule instance fires, then the pre-normalization of σ' has the form $\langle \emptyset, S, \text{true}, T' \rangle_{n'}$ with $T' \supsetneq T$. In other words $\text{chr_to_la}(\sigma) = \text{chr_to_la}(\sigma')$ and the CHR constraint store after pre-normalization is unchanged from the one before the implied rule instance fired while the propagation history is increased.

Proof

Let $\theta(r)$ be the implied rule instance with conclusion $\theta(C)$. Since $\theta(C) \subseteq \text{chr_to_la}(\sigma)$ with $\sigma = \langle \emptyset, S, \text{true}, T \rangle_n$, we have $\sigma' = \langle \theta(C), S, \text{true}, T \cup \{t\} \rangle_n$ and $\text{chr_to_la}(\sigma) = \text{chr_to_la}(\sigma')$ with t the propagation history tuple corresponding to $\theta(r)$. The goal G of σ' equals $\theta(C)$ and so it holds that if $a(\bar{X}) \in G$ then $a(\bar{X}, \mathbf{p}) \in \text{chr}(S)$ or $a(\bar{X}, \mathbf{b}) \in \text{chr}(S)$ and if $\text{del}(a(\bar{X})) \in G$ then $a(\bar{X}, \mathbf{n}) \in \text{chr}(S)$ or $a(\bar{X}, \mathbf{b}) \in \text{chr}(S)$. Now all constraints in the goal are first introduced in the CHR constraint store. Next, the newly introduced CHR constraints are removed one by one using one of the following rules:

$$\begin{aligned} 1 &:: a_r(\bar{X}, M) \setminus a(\bar{X}) \iff M \neq \mathbf{n} \mid \text{true} \\ 1 &:: a_r(\bar{X}, M) \setminus \text{del}(a(\bar{X})) \iff M \neq \mathbf{p} \mid \text{true} \end{aligned}$$

These rules remove all the constraints that were introduced from the goal and do not change the rest of the CHR constraint store, hence after pre-normalization, the CHR constraint store equals that of state σ again. $\square$

Because the CHR constraint store remains unchanged after firing an implied rule instance and pre-normalizing the resulting state, only finitely many such rule instances can fire before either reaching a final execution state, or a state in which a non-implied rule instance can fire. We call such a state *normalized*.

Definition 3 (Normal Form)

A pre-normalized CHR execution state σ is in normal form if it is a final state ($\sigma \not\xrightarrow{\omega_p}_{T(P)}$) or there exists a transition $\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$ such that $\text{chr_to_la}(\sigma') \not\subseteq \text{chr_to_la}(\sigma)$, i.e., in which a non-implied rule instance is fired.

Lemma 3

For every Logical Algorithms state σ_{LA} and every normalized CHR execution state $\sigma = \langle \emptyset, S, true, T \rangle_n$ such that $\sigma_{LA} = \text{chr_to_la}(\sigma)$, there exists a transition $\sigma_{LA} \xrightarrow{LA}_P \sigma'_{LA}$ if and only if there exists a transition $\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$ firing a non-implied rule instance such that $\sigma'_{LA} = \text{chr_to_la}(\sigma')$.

Proof

A transition of σ_{LA} to σ'_{LA} implies there exists a fireable rule instance $\theta(r)$ of a rule r in P with priority p of the form

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

Let $\langle A^u, A^c \rangle = \langle [A_1^u, \dots, A_m^u], [A_1^c, \dots, A_l^c] \rangle = \text{split}([A_1, \dots, A_n])$ where we use the **split** function defined in Section 3.1. The user-defined antecedents can be partitioned into sets of syntactically equal antecedents with respect to the matching substitution θ . The following function returns this partition:

$$\text{substitution_to_partition}(\theta, [A_1^u, \dots, A_m^u]) = \{S_1, \dots, S_m\}$$

where $S_i = \{j \mid \theta(A_j^u) = \theta(A_i^u)\}$. Let $\rho = \text{substitution_to_partition}(\theta, A^u)$. From the partition, we find the most general unifier θ' that unifies all antecedents $\{A_i^u \mid i \in S\}$ for every $S \in \rho$: $\theta' = \text{partition_to_mgu}(\rho, A^u)$ with **partition_to_mgu** as defined in Section 3.1. Clearly, θ' exists and is more general than θ . The applicability of the **Apply** transition means that for all comparison antecedents A_i^c with $1 \leq i \leq l$, $\mathcal{D} \models \theta(A_i^c)$ and so it holds that $\mathcal{D} \models \exists_0 \theta'(A_1^b \wedge \dots \wedge A_l^b)$ and consequently a rule r_ρ exists. This rule looks as follows:

$$p + 2 :: r_\rho @ H_1, \dots, H_k \Longrightarrow g_1, g_2 \mid C'$$

with $\langle [H_1, \dots, H_k], g_1 \rangle = \text{modes}(A^f)$, $A^f = [A_1^f, \dots, A_k^f] = \text{filter}(\theta'(A^u), \rho)$, $g_2 = \theta'(A^c)$ and $C' = \theta'(C)$. The **modes** and **filter** functions are as defined in Section 3.1.

Let θ'' be a ground matching substitution such that $\theta = \theta''|_{\text{vars}(\theta)} \circ \theta'$ where $\theta''|_{\text{vars}(\theta)}$ is the projection of θ'' on the variables in θ . Since θ' is more general than θ , θ'' exists. For all $i \in \{1, \dots, k\}$, if $A_i^f = a(\bar{X})$ then $H_i = a_r(\bar{X}, \mathbf{p})$. Because of the applicability of Logical Algorithms rule r in state σ_{LA} , $\theta''(a(\bar{X})) \in \sigma_{LA}$ and $\theta''(\text{del}(a(\bar{X}))) \notin \sigma_{LA}$, so $H'_i = \theta''(a_r(\bar{X}, \mathbf{p})) \# id_i \in S$ and $\theta''(H_i) = \text{chr}(H'_i)$. Similarly, if $A_i^f = \text{del}(a(\bar{X}))$ then $H_i = a_r(\bar{X}, N)$ and g_1 contains $N \neq \mathbf{p}$; $\theta''(\text{del}(a(\bar{X}))) \in \sigma_{LA}$ and as a result $H'_i = \theta''(a_r(\bar{X}, N')) \# id_i \in S$ with $N' = \mathbf{n}$ or $N' = \mathbf{b}$. Since N only appears in H_i and the guard $N \neq \mathbf{p}$, we can further impose that $\theta''(N) = N'$ and then $\theta''(H_i) = \text{chr}(H'_i)$.

All $\theta''(A_i^f)$ are different for $1 \leq i \leq k$, and therefore, all id_i must be different. From $\mathcal{D} \models \exists_0 \theta(A_i^c)$ for $1 \leq i \leq l$ and because $\theta''(g_1) = [N_1 \neq \mathbf{p}, \dots, N_o \neq \mathbf{p}]$ with $N_j = \mathbf{n}$ or $N_j = \mathbf{b}$ for $1 \leq j \leq o$, $\mathcal{D} \models true \rightarrow \exists_0 \theta''(g_1 \wedge g_2)$. We conclude that θ'' is a ground matching substitution that matches the head with constraints from S and for which the guard is entailed.

It is not possible that $\langle r_\rho, id(H) \rangle \in T$ because **chr_to_la** grows monotonically, which implies that $\theta(C) = \theta''(C') \in \text{chr_to_la}(\sigma) = \sigma_{LA}$ which contradicts with the applicability of $\theta(r)$ in σ_{LA} .

If we ignore rule priorities, all conditions are satisfied so that rule instance $\theta(r_\rho)$ can fire. The resulting state σ' has the form $\langle \theta(C), S, true, T \cup \{\langle r_\rho, id(H) \rangle\} \rangle_n$. Clearly, if $\sigma_{LA} = \text{chr_to_la}(\langle \emptyset, S, true, T \rangle_n)$ and $\sigma'_{LA} = \sigma_{LA} \cup \theta(C)$ then $\sigma'_{LA} = \text{chr_to_la}(\sigma')$. We now prove that every CHR transition firing a non-implied rule instance corresponds to a Logical Algorithms transition, also ignoring rule priorities. Both results combined give us that the priority of the highest priority rule instance is equal in both σ and σ_{LA} .

A transition of $\sigma = \langle \emptyset, S, true, T \rangle_n$ to σ' implies that $T(P)$ contains a rule

$$p + 2 :: r_\rho @ H \implies g_1, g_2 \mid C'$$

and so the Logical Algorithms program P contains a rule

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

Let $\langle A^u, A^c \rangle = \text{split}([A_1, \dots, A_n])$ and $\theta = \text{partition_to_mgu}(\rho, A^u)$. If $A_i = a(\bar{X}) \in A^u$ then $\theta(a_r(\bar{X}, \mathbf{p})) \in H$. If $A_i = \text{del}(a(\bar{X})) \in A^u$ then $\theta(a_r(\bar{X}, N)) \in H$ and $(N \neq \mathbf{p}) \in g_1$. Finally, if $A_i \in A^c$ then $\theta(A_i) \in g_2$. There exists a (ground) matching substitution θ' such that $\theta'(H) \in \text{chr}(S)$ and $\mathcal{D} \models \exists_\theta \theta'(g_1 \wedge g_2)$.

Let $\theta'' = \theta' \circ \theta$ and let $\sigma_{LA} = \text{chr_to_la}(\sigma)$. Because θ' is a ground substitution, $\mathcal{D} \models \exists_\theta \theta'(g_1 \wedge g_2)$ implies that for all $A_i \in A^c$, $\mathcal{D} \models \theta''(A_i)$. For all positive user-defined antecedents $A_i = a(\bar{X}) \in A^u$, we have that $\theta''(a(\bar{X}, \mathbf{p})) \in \text{chr}(S)$ and so $\theta''(A_i) \in \sigma_{LA}$ and $\text{del}(\theta''(A_i)) \notin \sigma_{LA}$. For all negative user-defined antecedents $A_i = \text{del}(a(\bar{X})) \in A^u$, we have that $\theta''(a_r(\bar{X}, N)) \in \text{chr}(S)$ with $N = \mathbf{b}$ or $N = \mathbf{n}$ and so $\theta''(A_i) \in \sigma_{LA}$. We have assumed that $\theta'(r_\rho)$ is not an implied rule instance and so $\theta'(C') = \theta''(C) \not\subseteq \sigma_{LA}$.

If we again ignore rule priorities, all conditions are satisfied so that rule instance $\theta''(r)$ can fire in state σ_{LA} and it holds that $\sigma'_{LA} = \sigma_{LA} \cup \theta''(C) = \text{chr_to_la}(\sigma')$ since $\sigma' = \langle \theta'(C'), S, true, T \cup \{\langle r_\rho, id(H) \rangle\} \rangle_n$. Now we have that both the original program P and its translation $T(P)$ can fire corresponding rule instances if we ignore priorities, and so their highest priority rule instances also correspond. $\square$

Theorem 2

For every reachable CHR^{rp} state σ , if $\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$ then either $\text{chr_to_la}(\sigma) = \text{chr_to_la}(\sigma')$ or $\text{chr_to_la}(\sigma) \xrightarrow{LA}_P \text{chr_to_la}(\sigma')$.

Proof

Implied by Lemmas 1, 2 and 3. $\square$

Theorem 3

For every Logical Algorithms state σ_i and reachable CHR^{rp} state σ'_i such that $\text{chr_to_la}(\sigma'_i) = \sigma_i$, there exists a finite CHR^{rp} derivation $\sigma'_i \xrightarrow{\omega_p^*}_{T(P)} \sigma'_{i*}$ for which holds that $\text{chr_to_la}(\sigma'_{i*}) = \sigma_i$ such that if $\sigma_i \xrightarrow{LA}_P \sigma_j$ then $\sigma'_{i*} \xrightarrow{\omega_p}_{T(P)} \sigma'_j$ with $\text{chr_to_la}(\sigma'_j) = \sigma_j$ and if σ_i is a final state then σ'_{i*} is also a final state.

Proof

Implied by Lemmas 1, 2 and 3. $\square$

Given a Logical Algorithms state σ , we can use $\langle \sigma, \emptyset, \text{true}, \emptyset \rangle_1$ as initial state for the CHR^{rp} derivation. Theorem 3 is illustrated by the figure below.

3.3 Weak Bisimulation

To capture the meaning of the above correspondence results, we relate them to the notion of (weak) bisimulation. A bisimulation is a relation between the states of a labeled transition system (LTS). A relation $R \subseteq S_1 \times S_2$ between the states in S_1 and those in S_2 is a bisimulation if $p R q$ and $p \xrightarrow{\alpha} p'$ implies that $q \xrightarrow{\alpha} q'$ with $p' R q'$, and similarly, $p R q$ and $q \xrightarrow{\alpha} q'$ implies that $p \xrightarrow{\alpha} p'$ with $p' R q'$. Here, α is the label of the transition $p \xrightarrow{\alpha} p'$ from state p to state p' . If a transition from p to p' has no observable effect, it is called a *silent* transition and denoted by $p \xrightarrow{\tau} p'$. A relation $R \subseteq S_1 \times S_2$ is a *weak* bisimulation if $p R q$ and $p \xrightarrow{\alpha} p'$ implies that $q \xrightarrow{\tau^*} q_* \xrightarrow{\alpha} q'_* \xrightarrow{\tau^*} q'$ with $p' R q'_*$, and vice versa with the roles of p and q swapped. Here $p \xrightarrow{\tau^*} p'$ means p and p' are linked by zero or more silent transitions.

Let S_1 be the set of valid Logical Algorithms states for program P and let $S_2 = \{\text{chr_to_la}(\sigma) \mid \langle G, \emptyset, \text{true}, \emptyset \rangle_1 \xrightarrow{\omega_p^*}_{T(P)} \sigma \wedge G \in S_1\}$, i.e., S_2 is found by applying the chr_to_la mapping function all reachable CHR^{rp} states for program $T(P)$. We transform the state transition systems for Logical Algorithms and CHR^{rp} to labeled transition systems as follows: a Logical Algorithms transition $\sigma \xrightarrow{LA}_P \sigma'$ corresponds to an LTS transition $\sigma \xrightarrow{\alpha} \sigma'$ with $\alpha = \sigma' \setminus \sigma$, i.e., α represents the state change from σ to σ' . A CHR^{rp} transition $\sigma \xrightarrow{\omega_p}_{T(P)} \sigma'$ corresponds to an LTS transition $\text{chr_to_la}(\sigma) \xrightarrow{\alpha} \text{chr_to_la}(\sigma')$ with $\alpha = \text{chr_to_la}(\sigma') \setminus \text{chr_to_la}(\sigma)$ if this set is not empty and $\alpha = \tau$ otherwise.

Corollary 1

The equality relation between the states of S_1 and S_2 is a weak bisimulation.

4 Translating a subset of CHR^{rp} into Logical Algorithms

In the previous section, we have shown that Logical Algorithms programs can be translated into equivalent CHR^{rp} programs. In this section, we show how to do the opposite, i.e., how CHR^{rp} programs can be translated into equivalent Logical Algorithms programs. This allows us to apply the meta-complexity theorem for Logical Algorithms to the translation of these CHR^{rp} programs.

We need to impose some restrictions on the CHR^{rp} programs that can be translated. These restrictions result from the fact that the Logical Algorithms language does not have the concept of an underlying constraint solver that offers both ask and tell built-in constraints. The following two properties are required:

1. In all *reachable* states $\sigma = \langle G, S, B, T \rangle_n$: $\text{vars}(S) = \emptyset$. In words, all (stored) CHR constraints are ground.
2. All built-in constraints are comparisons; there are no built-in tell constraints.

The first property holds if the initial goal is ground and all rules are *variable restricted*, which means that all variables in the body of a rule, also appear in one of the rule heads. The second property implies that all reachable states are of the form $\langle G, S, \text{true}, T \rangle_n$, i.e., the built-in constraint store always equals *true*.

To simplify the presentation, we also assume that the priority of dynamic priority rules is determined by the arguments of its left-most head. In general, we can use the transformation schema given in Section 2.2.3 to ensure that the resulting Logical Algorithms rules have the correct syntactical form.

4.1 The Translation Schema

We now show how the rules of a CHR^{rp} program P are transformed into Logical Algorithms rules that form a program $T(P)$. To increase readability, we distinguish between simplification and simpagation rules on the one hand, and propagation rules on the other. A *simpagation* rule of the form

$$p :: r @ H_1, \dots, H_{m-1} \setminus H_m, \dots, H_n \iff g \mid B_1, \dots, B_l$$

is transformed into

$$\begin{aligned} r' @ p : H_1^{id}, \dots, H_n^{id}, \text{Alldiff}, g, \text{next_id}(Id_{next}) \Rightarrow \\ \text{del}(H_m^{id}), \dots, \text{del}(H_n^{id}), \text{del}(\text{next_id}(Id_{next})), \\ B_1^{id}, \dots, B_l^{id}, \text{next_id}(Id_{next} + l) \end{aligned}$$

where $H_i^{id} = \text{c}(\bar{X}, Id_i)$ if $H_i = \text{c}(\bar{X})$, $B_i^{id} = \text{c}(\bar{X}, Id_{next} + i - 1)$ if $B_i = \text{c}(\bar{X})$ and $\text{Alldiff} = \{(Id_i \neq Id_j) \mid \mathcal{D} \models \exists_0 H_i = H_j \wedge g\}$. The disequalities in Alldiff are between those heads that are unifiable and for which the guard is still satisfiable after this unification. The case of a simplification rule is very similar. A propagation rule of the form

$$p :: r @ H_1, \dots, H_n \implies g \mid B_1, \dots, B_l$$

is transformed into the following two rules

$$\begin{aligned} r'_1 @ p : H_1^{id}, \dots, H_n^{id}, \text{Alldiff}, g \Rightarrow \text{token}(r, [Id_1, \dots, Id_n]) \\ r'_2 @ p : H_1^{id}, \dots, H_n^{id}, \text{Alldiff}, g, \text{token}(r, [Id_1, \dots, Id_n]), \text{next_id}(Id_{next}) \Rightarrow \\ \text{del}(\text{token}(r, [Id_1, \dots, Id_n])), \text{del}(\text{next_id}(Id_{next})), \\ B_1^{id}, \dots, B_l^{id}, \text{next_id}(Id_{next} + l) \end{aligned}$$

where H_i^{id} , B_i^{id} and Alldiff are as before. The first of these rules generates a token. This token is removed by the second rule. The tokens are needed to prevent a

given rule instance from firing more than once.⁷ Note that the transformation into two rules and the use of tokens does not increase the complexity compared to the original rule, as there is only one token for each combination of rule and constraint identifiers (as well as only one `next_id/1` fact in any state).

The initial database consists of the goal (where each constraint is extended with a unique identifier) and a `next_id( $Id_{next}$ )` assertion (with Id_{next} the next free identifier).

Example 6 (Merge Sort)

The following CHR^{IP} program implements a merge sort algorithm. Its input consists of a series of n (a power of 2) `number/1` constraints. Its output is a sorted list of the numbers in the input, represented as `arrow/2` constraints, where `arrow( $X, Y$ )` indicates that X is right before Y .

```
1 :: ms1 @ arrow(X,A) \ arrow(X,B) <=> A < B | arrow(A,B).
2 :: ms2 @ merge(N,A), merge(N,B) <=> A < B | merge(2*N+1,A), arrow(A,B).
3 :: ms3 @ number(X) <=> merge(0,X).
```

Its Logical Algorithms translation is

```
ms1' @ 1 : arrow(X,A,Id1), arrow(X,B,Id2), A < B, next_id(NId) =>
    del(arrow(X,B,Id2)), del(next_id(NId)),
    arrow(A,B,NId), next_id(NId+1).
ms2' @ 2 : merge(N,A,Id1), merge(N,B,Id2), A < B, next_id(NId) =>
    del(merge(N,A,Id1)), del(merge(N,B,Id2)), del(next_id(NId)),
    merge(2*N+1,A,NId), arrow(A,B,NId+1), next_id(NId+2).
ms3' @ 3 : number(X,Id), next_id(NId) => del(number(X,Id)),
    del(next_id(NId)), merge(0,X,NId), next_id(NId+1).
```

Note that in rules `ms1` and `ms2`, the guard prevents the constraints matching the heads from being equal, and so there are no disequality constraints between the CHR constraint identifiers. In (De Koninck et al. 2007a) it is derived that the total runtime of this Logical Algorithms program is $\mathcal{O}(n \log n)$. We defer the complexity analysis of the merge sort algorithm to Section 6.1 where we analyse the CHR^{IP} implementation directly using a new meta-complexity theorem for CHR^{IP}.

Example 7 (Less-or-Equal)

To illustrate how propagation rules are dealt with, we show the translation of a rule of the `leq` program which is given further on in Example 9. The rule

```
3 :: transitivity @ leq(X,Y), leq(Y,Z) ==> leq(X,Z).
```

is translated into

⁷ In (De Koninck et al. 2007a), an erroneous translation was presented which did not use tokens, and in which a propagation rule could fire infinitely many times because the constraints in the body are assigned new identifiers each time the rule is fired.

```

transitivity'_1 @ 3 : leq(X,Y,Id_1), leq(Y,Z,Id_2), Id_1 \= Id_2 =>
    token(transitivity,[Id_1,Id_2]).
transitivity'_2 @ 3 : leq(X,Y,Id_1), leq(Y,Z,Id_2), Id_1 \= Id_2,
    token(transitivity,[Id_1,Id_2]), next_id(NId) =>
        del(token(transitivity,[Id_1,Id_2])), del(next_id(NId)),
        leq(X,Z,NId), next_id(NId+1).

```

Note that since in the original rule, the two heads $\text{leq}(X,Y)$ and $\text{leq}(Y,Z)$ are unifiable (and there is furthermore no guard to prevent this from happening), we have to add an explicit disequality between the constraint identifiers for these heads: $\text{Id}_1 \setminus = \text{Id}_2$.

4.2 Correspondence

In this subsection, we prove that a CHR^{IP} program and its translation to Logical Algorithms are operationally equivalent. Again we introduce a mapping function:

$$\text{la_to_chr}(\sigma) = \langle \emptyset, S, \text{true}, T \rangle_n$$

where the CHR constraint store $S = \{c(\bar{X}) \# Id \mid c(\bar{X}, Id) \in \sigma \wedge \text{del}(c(\bar{X}, Id)) \notin \sigma\}$, the propagation history $T = \{\langle R, Ids \rangle \mid \text{del}(\text{token}(R, Ids)) \in \sigma\}$, and the next free identifier n is such that $\text{next_id}(n) \in \sigma$ and $\text{del}(\text{next_id}(n)) \notin \sigma$. In the following, we consider a Logical Algorithms state σ reachable with respect to program $T(P)$ if it can be derived from an initial state consisting of CHR constraints extended with unique identifiers, and a single $\text{next_id}/1$ assertion with as argument the next free identifier. In this case, reachability amongst others implies that there can be only one (strictly) positive $\text{next_id}/1$ assertion in the database in any state, and no two CHR constraint representations share their identifier.

Theorem 4

For every reachable Logical Algorithms state σ_i it holds that if $\sigma_i \xrightarrow{LA}_{T(P)} \sigma_j$, then either it holds that $\text{la_to_chr}(\sigma_i) = \text{la_to_chr}(\sigma_j)$ or there exists a finite CHR derivation $\text{la_to_chr}(\sigma_i) = \langle \emptyset, S, \text{true}, T \rangle_n \xrightarrow{\omega_p}_P \langle C, S', \text{true}, T' \rangle_n \xrightarrow{\omega_p^*}_P \langle \emptyset, S'', \text{true}, T' \rangle_{n'} = \text{la_to_chr}(\sigma_j)$ consisting of an **Apply** transition, followed by zero or more **Introduce** transitions.

Proof

Consider a transition $\sigma_i \xrightarrow{LA}_{T(P)} \sigma_j$. The only type of transition in Logical Algorithms is the **Apply** transition which fires a rule. If $\text{la_to_chr}(\sigma_i) = \text{la_to_chr}(\sigma_j)$, then this rule must be of the form

$$r'_1 @ p : H_1^{id}, \dots, H_m^{id}, \text{Alldiff}, g \Rightarrow \text{token}(r, [Id_1, \dots, Id_m])$$

because all other types of rules either delete the representation of a CHR constraint which changes the CHR constraint store, or remove a token which results in an extended propagation history. We call the fired rule a *token generation rule*.

If $\text{la_to_chr}(\sigma_i) \neq \text{la_to_chr}(\sigma_j)$ and the rule fired is of the form

$$\begin{aligned} r' @ p : H_1^{id}, \dots, H_m^{id}, \text{Alldiff}, g, \text{next_id}(Id_{next}) \Rightarrow \\ \text{del}(H_l^{id}), \dots, \text{del}(H_m^{id}), \text{del}(\text{next_id}(Id_{next})), \\ B_1^{id}, \dots, B_o^{id}, \text{next_id}(Id_{next} + o) \end{aligned}$$

which corresponds to a simplification ($l = 1$) or simpagation ($l > 1$) rule. We further assume the case of a simpagation rule; the case of a simplification rule is similar. If $r' \in T(P)$ (with $l > 1$), then P contains a rule

$$p :: r @ H_1, \dots, H_{l-1} \setminus H_l, \dots, H_m \iff g \mid B_1, \dots, B_o$$

Since the conditions for the Logical Algorithms **Apply** transition are satisfied, there exists a ground matching substitution θ such that for each antecedent $H_i^{id} = c(\bar{X}, Id_i)$ ($1 \leq i \leq m$) it holds that $\theta(H_i^{id}) \in \sigma$ and $\text{del}(\theta(H_i^{id})) \notin \sigma$ and so by definition of the la_to_chr function, $\theta(H_i \# Id_i) \in S$ where $\text{la_to_chr}(\sigma_i) = \sigma'_i = \langle \emptyset, S, \text{true}, T \rangle_n$. For each comparison $g_i \in g$, it holds that $\mathcal{D} \models \theta(g_i)$ and so $\mathcal{D} \models \text{true} \rightarrow \exists_{\emptyset} \theta(g)$. Since r is a simpagation rule, the propagation history T does not contain any element of the form $\langle r, _ \rangle$. Ignoring priorities for the moment, all conditions are satisfied such that the rule instance $\theta(r)$ can fire in state σ'_i . We return to the issue of priorities further on.

After firing $\theta(r)$ in state σ'_i , the resulting state equals $\langle \theta(B_1 \wedge \dots \wedge B_o), S', \text{true}, T \rangle_n$ where $S' = S \setminus \{\theta(H_l \# Id_l), \dots, \theta(H_m \# Id_m)\}$. In this state, the only applicable transition is the **Introduce** transition, which is applied o times before reaching a state with an empty goal. There are $o!$ possible orders in which the introductions can be applied, but the one we need is the order in which the B_i constraints appear in the rule body. Following this order, the state resulting from the introductions equals $\sigma'_j = \langle \emptyset, S'', \text{true}, T \rangle_{(n+o)}$ where $S'' = S' \cup \{\theta(B_1) \# n, \dots, \theta(B_o) \# (n+o-1)\}$. It is easy to see that this state σ'_j equals $\text{la_to_chr}(\sigma_j)$, the state resulting from firing Logical Algorithms rule instance $\theta(r')$ in state σ_i .

If $\text{la_to_chr}(\sigma_i) \neq \text{la_to_chr}(\sigma_j)$ and the rule fired is not of the form shown above, then it must have the following form

$$\begin{aligned} r'_2 @ p : H_1^{id}, \dots, H_m^{id}, \text{Alldiff}, g, \text{token}(r, [Id_1, \dots, Id_m]), \text{next_id}(Id_{next}) \Rightarrow \\ \text{del}(\text{token}(r, [Id_1, \dots, Id_m])), \text{del}(\text{next_id}(Id_{next})), \\ B_1^{id}, \dots, B_o^{id}, \text{next_id}(Id_{next} + o) \end{aligned}$$

the corresponding CHR^{rp} rule in P looks like

$$p :: r @ H_1, \dots, H_m \implies g \mid B_1, \dots, B_o$$

Again, since the conditions for the Logical Algorithms **Apply** transition are satisfied, there exists a ground matching substitution θ such that for each antecedent $H_i^{id} = c(\bar{X}, Id_i)$ ($1 \leq i \leq m$) in rule r'_2 it holds that $\theta(H_i^{id}) \in \sigma$ and $\text{del}(\theta(H_i^{id})) \notin \sigma$ and so by definition of the la_to_chr function, $\theta(H_i \# Id_i) \in S$ where $\text{la_to_chr}(\sigma_i) = \sigma'_i = \langle \emptyset, S, \text{true}, T \rangle_n$. For each comparison $g_i \in g$, it holds that $\mathcal{D} \models \theta(g_i)$ and so $\mathcal{D} \models \text{true} \rightarrow \exists_{\emptyset} \theta(g)$. The propagation history T cannot contain $\langle r, \theta([Id_1, \dots, Id_m]) \rangle$ because by definition of the la_to_chr function this would imply that the atom

$token(r, \theta([Id_1, \dots, Id_m]))$ was deleted in some earlier state, which contradicts with the applicability of the **Apply** transition on rule instance $\theta(r'_2)$. If we again ignore the issue of priorities, all conditions are satisfied such that $\theta(r)$ can fire in state σ'_i .

After firing $\theta(r)$ in state σ'_i , the resulting state equals $\langle \theta(B_1 \wedge \dots \wedge B_o), S, true, T' \rangle_n$ where $T' = T \cup \{\langle r, [Id_1, \dots, Id_m] \rangle\}$. In this state, the only applicable transition is the **Introduce** transition, which is applied o times before reaching a state with an empty goal. Given again that these introductions are applied in the order in which the B_i constraints appear in the rule body, then the resulting state equals $\sigma'_j = \langle \emptyset, S', true, T' \rangle_{(n+o)}$ where $S' = S \cup \{\theta(B_1) \# n, \dots, \theta(B_o) \# (n + o - 1)\}$. It is again easy to see that this state σ'_j equals $la_to_chr(\sigma_j)$, the state resulting from firing Logical Algorithms rule instance $\theta(r'_2)$ in state σ_i .

This proves the theorem if we ignore priorities. Theorem 5 (see next) shows that each CHR^P rule firing has a corresponding Logical Algorithms rule firing. Under the assumption that this theorem also holds ignoring rule priorities, we have that the highest priority rule instances are the same in both programs given corresponding states and ignoring token generation rules. $\square$

Theorem 5

For every reachable CHR^P state σ_i and reachable Logical Algorithms state σ'_i with $la_to_chr(\sigma'_i) = \sigma_i$, there exists a finite Logical Algorithms derivation $\sigma'_i \xrightarrow{T(P)}^{LA} \sigma'_{i*}$ with $la_to_chr(\sigma'_{i*}) = \sigma_i$ such that if $\sigma_i = \langle \emptyset, S, true, T \rangle_n \xrightarrow{P}^{\omega_p} \langle C, S', true, T' \rangle_n \xrightarrow{P}^{\omega_p} \langle \emptyset, S'', true, T' \rangle_{n'} = \sigma_j$ where the derivation consists of a single **Apply** transition, followed by zero or more **Introduce** transitions, then $\sigma'_{i*} \xrightarrow{T(P)}^{LA} \sigma'_j$ with $la_to_chr(\sigma'_j) = \sigma_j$ and if σ_i is a final state then σ'_{i*} is also a final state.

Proof

Let there be given a reachable Logical Algorithms state σ'_i with $la_to_chr(\sigma'_i) = \sigma_i$. Because of Theorem 4, state σ_i is also reachable in CHR^P with respect to program P . Assume $\sigma_i = \langle \emptyset, S, true, T \rangle_n \xrightarrow{P}^{\omega_p} \langle C, S', true, T' \rangle_n \xrightarrow{P}^{\omega_p} \langle \emptyset, S'', true, T' \rangle_{n'} = \sigma_j$ where the derivation consists of a single **Apply** transition, followed by zero or more **Introduce** transitions, and let $\theta(r)$ be the CHR^P rule instance that fired in state σ_i . If r is simplification ($l = 1$) or simpagation ($l > 1$) rule

$$p :: r @ H_1, \dots, H_{l-1} \setminus H_l, \dots, H_m \iff g \mid B_1, \dots, B_o$$

then $\theta(H_i) \# id_i \in S$ for $1 \leq i \leq m$ with $id_i \neq id_j$ if $i \neq j$, and $\mathcal{D} \models exists_{\emptyset} \theta(g)$. Furthermore, $T(P)$ contains a rule

$$\begin{aligned} r' @ p : H_1^{id}, \dots, H_m^{id}, Alldiff, g, next_id(Id_{next}) \Rightarrow \\ del(H_l^{id}), \dots, del(H_m^{id}), del(next_id(Id_{next})), \\ B_1^{id}, \dots, B_o^{id}, next_id(Id_{next} + o) \end{aligned}$$

Now let θ' be a ground matching substitution such that $\theta'|_{vars(\theta)} = \theta$ where $\theta'_{vars(\theta)}$ is the projection of θ' on the variables in θ , and such that both $\theta'(Id_i) = id_i$ for $1 \leq i \leq m$ and $\theta'(Id_{next}) = n$. Since for $1 \leq i \leq m$, $H_i^{id} = c(\bar{X}, Id_i)$ if $H_i = c(\bar{X})$, it

holds that $\theta'(H_i^{id}) \in \sigma'_i$ and $\text{del}(\theta'(H_i^{id})) \notin \sigma'_i$. Also, $\mathcal{D} \models \exists_0 \theta(g)$ implies $\mathcal{D} \models \theta(g_i)$ for each comparison $g_i \in g$.⁸ The *Alldiff* conditions hold because $\theta'(Id_i) = \theta'(Id_j)$ implies that $i = j$. Finally, because of the reachability of state σ'_i , there is exactly one strictly positive `next_id/1` assertion in σ'_i whose argument equals n . Finally, the rule conclusion cannot be already included in the state σ'_i because it includes amongst others the deletion of at least one of the antecedents. Ignoring priorities, all conditions are satisfied such that rule instance $\theta'(r')$ can fire in state σ'_i , resulting in a state $\sigma'_j = \text{la_to_chr}(\sigma'_j)$. As stated earlier in the proof of Theorem 4, the combination of Theorems 4 and 5 without taking into account the priorities, implies that the highest priority fireable rule instances are the same in corresponding states, ignoring token generation rules.

Now assume that in the CHR^{IP} state σ_i , a rule instance $\theta(r)$ fires where r is a propagation rule:

$$p :: r @ H_1, \dots, H_m \Longrightarrow g \mid B_1, \dots, B_o$$

In this case the Logical Algorithms translation $T(P)$ contains the following rules:

$$\begin{aligned} r'_1 @ p : H_1^{id}, \dots, H_n^{id}, \text{Alldiff}, g &\Rightarrow \text{token}(r, [Id_1, \dots, Id_n]) \\ r'_2 @ p : H_1^{id}, \dots, H_n^{id}, \text{Alldiff}, g, \text{token}(r, [Id_1, \dots, Id_n]), \text{next_id}(Id_{next}) &\Rightarrow \\ &\text{del}(\text{token}(r, [Id_1, \dots, Id_n])), \text{del}(\text{next_id}(Id_{next})), \\ &B_1^{id}, \dots, B_l^{id}, \text{next_id}(Id_{next} + l) \end{aligned}$$

A similar analysis as above shows that there exists a matching substitution θ' with $\theta'|_{\text{vars}(\theta)} = \theta$ and both $\theta'(Id_i) = id_i$ for $1 \leq i \leq m$ and $\theta'(Id_{next}) = n$, such that rule instance $\theta'(r'_1)$ can fire (ignoring priorities) if $\text{token}(r, [id_1, \dots, id_n]) \notin \sigma'_i$ and $\theta'(r'_2)$ otherwise. If $\theta'(r'_1)$ fires then the resulting state $\sigma'_{i*} = \sigma'_i \cup \{\text{token}(r, [id_1, \dots, id_n])\}$ and clearly $\text{la_to_chr}(\sigma'_{i*}) = \text{la_to_chr}(\sigma'_i)$. Moreover, in state σ'_{i*} , rule instance $\theta'(r'_2)$ can fire and for the resulting state σ'_j it holds that $\text{la_to_chr}(\sigma'_j) = \sigma_j$. If already $\text{token}(r, [id_1, \dots, id_n]) \in \sigma'_i$ then the same reasoning holds with $\sigma'_i = \sigma'_{i*}$.

Finally, assume that CHR^{IP} state σ_i is a final state. If σ'_i is not a final Logical Algorithms state, then because of Theorem 4, the only applicable rules are those that do not change the result of the `la_to_chr` function. Only the token generation rules satisfy this property. Since they only generate tokens and these tokens do not appear in their antecedents, these rules can fire only finitely many times before a final Logical Algorithms state σ_{i*} is reached. $\square$

5 Implementing CHR^{IP} , the Logical Algorithms way

This section presents a new implementation for CHR^{IP} , based on the implementation proposal for Logical Algorithms presented in (Ganzinger and McAllester 2002), as well as on the scheduling algorithm presented in (De Koninck 2007). The purpose of this implementation is not to replace our existing CHR^{IP} implementation as presented in (De Koninck et al. 2007), but to support a new meta-complexity theorem

⁸ Because $\theta(g)$ is ground, there is no existential quantification.

for CHR^{rp} , based on the result for Logical Algorithms, and extended towards the full CHR^{rp} language. This includes in particular support for non-ground constraints and a built-in constraint theory. We note that a better worst case complexity for certain operations is not always worthwhile in practice due to larger constant factors in the average case. Also, the proposed implementation may not always achieve a better complexity than the existing implementation. The main purpose remains to have a relatively straightforward way to derive for a given CHR^{rp} program, a bound that is guaranteed to be an upper bound for at least the implementation proposed. Since the meta-complexity result is insensitive to constant factors, we can present the new implementation as a source-to-source transformation to regular CHR.

The proposed implementation consists of the compilation of the CHR^{rp} rules of the input program into regular CHR rules in which matching is made explicit, combined with a scheduler module that is responsible for the execution control. The implementation is correct if it is executed according to the refined operational semantics of CHR (Duck et al. 2004), which describes the execution strategy followed by most current CHR implementations. We have based our implementation on the high-level implementation proposal for Logical Algorithm of (Ganzinger and McAllester 2002), extended where necessary to support general built-in constraints. By using a CHR implementation with advanced indexing support, like for example the K.U.Leuven CHR system (Schrijvers and Demoen 2004), our implementation also offers strong complexity guarantees that facilitate a new meta-complexity theorem for CHR^{rp} , similar to the one for Logical Algorithms (see Section 6). In the following, we make use of Prolog as CHR's host language, but the implementation can easily be adapted to work with a different host language.

5.1 Overview

The implementation is based on a form of lazy (on-demand) matching with retainment of previously computed partial matches. It combines the concept of alpha and beta memories from the RETE algorithm (Forgy 1982), with lazy matching as for example implemented by the LEAPS algorithm (Miranker et al. 1990).⁹ The basic idea is as follows. A new constraint can function both as a single headed partial or full match, and as an extension of an existing partial match into either a new (larger) partial match or a full match. In order to extend partial matches, all previously computed matches are stored. A scheduler decides which partial match is extended with which constraint, or which full match has its corresponding rule instance fired. More details on the scheduler are given in Section 5.3.

First, to simplify the presentation, we propose an alternative syntax for CHR^{rp} rules. An *intermediate form* CHR^{rp} rule looks as follows:

$$p :: r @ s_1 A_1, \dots, s_n A_n \iff B$$

where $s_i \in \{+, -, ?\}$ and A_i is an atom for $1 \leq i \leq n$. If $s_i = +$ or $s_i = -$ then

⁹ Most current CHR systems, including the K.U.Leuven CHR system and the CHR^{rp} system of (De Koninck et al. 2007), use a variant of the LEAPS algorithm for rule matching.

A_i must be a CHR constraint and if $s_i = ?$ then A_i must be a built-in constraint. An intermediate form CHR^{rp} rule corresponds to a regular CHR^{rp} rule as follows: a term $+A$ corresponds to a kept head A , a term $-A$ corresponds to a removed head A , and a term $?A$ corresponds to a conjunct of the rule guard. The main advantage of the intermediate form is that it supports specifying a join order for the heads, as well as an evaluation order for the guards. In particular it supports specifying the evaluation of part of the guard after having computed only a partial rule match. The intermediate form gives us the same syntactical flexibility as exists in the Logical Algorithms language where comparisons are interleaved with the (kept and removed) user-defined antecedents.

Consider, in general, a simpagation rule of the form

$$p :: r @ H_1, \dots, H_i \setminus H_{i+1}, \dots, H_n \iff g \mid B$$

where the guard g is a conjunction of atomic guards $g_1, \dots, g_m$. We can rewrite this rule in intermediate form syntax (amongst others) as follows:

$$p :: r @ +H_1, \dots, +H_i, -H_{i+1}, \dots, -H_n, ?g_1, \dots, ?g_m \iff B$$

In the following, we assume that all rules have the following form

$$p :: r @ \pm H_1, ?g_1, \pm H_2, ?g_2, \dots, \pm H_n, ?g_n \iff B$$

where $\pm$ means $+$ or $-$. Each g_i ($1 \leq i \leq n$) can be a conjunction of primitive built-in constraints, and can in particular also be equal to *true*. The transformation from regular CHR^{rp} syntax to intermediate form syntax can be done automatically using the above transformation schema, or by hand.

Using terminology similar to that of (Ganzinger and McAllester 2002), we refer to a partial match, matching the heads $H_1, \dots, H_i$ and satisfying the partial guard $g_1 \wedge \dots \wedge g_{i-1}$, as a *suspended strong prefix firing*. If also the partial guard g_i is satisfied, we speak of a *regular* (or *non-suspended*) strong prefix firing. A constraint matching the next head H_{i+1} is called a *prefix extension* of such a (regular) strong prefix firing. A prefix firing that consists of all heads is (also) called a (suspended or regular) *rule firing*.¹⁰ Every prefix firing contains the left-most head and hence determines the rule priority. In our implementation, we assume that all guards are monotone, i.e., once they are entailed by the built-in constraint store, they remain entailed in any later state. This is in fact required by the CHR operational (and declarative) semantics, although most current CHR systems also support non-monotone (impure) guards like for example `var/1` in CHR on top of Prolog.

5.2 Program-Dependent Part

The program-dependent part of our implementation (i.e., the part that depends on the actual program to be implemented) consists of rules for

¹⁰ A rule firing actually means a rule instance that is fireable. To avoid confusion, we refer to the actual firing of such a rule firing as *firing a rule instance*.

- generating a representation for CHR constraint occurrences and deleting them when the represented constraint is removed;
- generating and scheduling constraints representing prefix firings, prefix extensions and rule firings and deleting them when a constituent constraint is removed;
- matching prefix firings with prefix extensions, firing rule instances, and managing suspended prefix and rule firings.

The different types of rules of the program-dependent part are illustrated by using a running example program, namely Dijkstra's shortest path algorithm, already given in the Logical Algorithms language in Example 2 and given here in CHR^{IP} intermediate form syntax. To illustrate non-trivial head matching, we have added a rule `d1` that removes simple loops from the input graph.

1 :: d1 @ -e(V,_,V), ?true	<=> true.
1 :: d2 @ +source(V), ?true	<=> dist(V,0).
1 :: d3 @ -dist(V,D ₁), ?true, +dist(V,D ₂), ?(D ₂ < D ₁)	<=> true.
D + 2 :: d4 @ +dist(V,D), ?true, +e(V,C,U), ?true	<=> dist(U,D+C).

5.2.1 Constraint Occurrence Representation

Although CHR^{IP} constraints and CHR constraints obviously have the same syntax and semantics (i.e., multi-set semantics with non-monotone deletion), we introduce a new representation for them to allow unambiguous reference, reduce work in case of constraint reactivation, and support the efficient deletion of those prefix firings, prefix extensions, and rule firings in which they participate (see further). For each CHR^{IP} constraint of predicate c/n , we create a set of unique *occurrence representations* $c_occ_i/(n+1)$, one for each occurrence of the predicate in a rule head. The arguments of a $c_occ_i/(n+1)$ constraint consist of the arguments of the original c/n constraint, together with a unique constraint identifier that is shared by all occurrence representations. This identifier is an uninstantiated variable as long as the constraint is in the store and is instantiated the moment that the constraint is to be deleted. For each user-defined constraint predicate c/n with m occurrences, the occurrence representations are generated using rules of the following form.

$$c(X_1, \dots, X_n) \text{ <=> } c_occ_1(X_1, \dots, X_n, Id), \dots, c_occ_m(X_1, \dots, X_n, Id).$$

For the example program, these rules look as follows.

source(V) <=> source_occ_1(V, Id).
dist(V,D) <=> dist_occ_1(V,D, Id), dist_occ_2(V,D, Id), dist_occ_3(V,D, Id).
e(V,C,U) <=> e_occ_1(V,C,U, Id), e_occ_2(V,C,U, Id).

5.2.2 RETE Memory Constraints

Regular and suspended prefix firings as well as prefix extensions are represented as CHR constraints. We call them RETE memory constraints because they coincide

with the alpha and beta memories of the RETE algorithm. The RETE memory constraints contain all arguments of their constituent CHR constraints, as well as their identifiers. Each RETE memory constraint moreover has its own unique identifier. We use the following functors for RETE memory constraints:

- r_pf_i for a regular (non-suspended) prefix firing of rule r , consisting of i heads, and $r_pf_i_suspended$ for its suspended version
- r_pe_i for a prefix extension, consisting of the $i + 1^{\text{th}}$ head of rule r
- r_rf for a (regular) rule firing of rule r and $r_rf_suspended$ for its suspended version.

If in a rule r , the partial guard after the i^{th} head equals *true*, then there is no suspended version of the i -headed prefix firings of r , or of its rule firings if r is an i -headed rule. In the example program, the following prefix firings, prefix extensions and rule firings are defined:

- $d1_rf/4$
- $d2_rf/3$
- $d3_pf_1/4$, $d3_pe_1/3$, $d3_rf/6$ and $d3_rf_suspended/6$
- $d4_pf_1/4$, $d4_pe_1/4$ and $d4_rf/7$

5.2.3 Suspended Prefix and Rule Firings

Suspended prefix and rule firings are converted into regular prefix and rule firings as soon as the relevant part of the guard is entailed. If on the other hand this partial guard is disentailed, the suspended prefix or rule firing is removed. Given a rule in intermediate form syntax

$$p :: r @ \pm H_1, ?g_1, \pm H_2, ?g_2, \dots, \pm H_n, ?g_n \iff B$$

we generate the following rules:

- For each i -headed suspended prefix firing:

$$\begin{aligned}
 &r_pf_i_suspended(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \iff \\
 &\quad g_i \mid r_pf_i(X_1, \dots, X_m, Id_1, \dots, Id_i, SId), \\
 &\quad \text{schedule_pf}(r_i(Y_1, \dots, Y_l), p, SId). \\
 &r_pf_i_suspended(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \iff \backslash+ g_i \mid \text{true}.
 \end{aligned}$$
 where $Y_1, \dots, Y_l$ are those variables in $X_1, \dots, X_m$ that also appear in H_{i+1}
- For each rule firing:

$$\begin{aligned}
 &r_rf_suspended(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \iff g_n \mid \\
 &\quad r_rf(X_1, \dots, X_m, Id_1, \dots, Id_n, SId), \text{ schedule_rf}(p, SId). \\
 &r_rf_suspended(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \iff \backslash+ g_n \mid \text{true}.
 \end{aligned}$$

Note that if g_i or g_n equals *true*, then we can apply unfolding to replace occurrences of respectively $r_pf_i_suspended/(m + i + 1)$ and $r_rf_suspended/(m + n + 1)$ by the bodies of the corresponding rules above (see (Tacchella et al. 2007)). After this unfolding step, some of the above rules may be removed. In the example program, only a rule firing of rule **d3** can be suspended.

```

d3_rf_suspended(V,D1,D2,Id1,Id2,SId) <=>
    D2 < D1 | d3_rf(V,D1,D2,Id1,Id2,SId), schedule_rf(1,SId).
d3_rf_suspended(V,D1,D2,Id1,Id2,SId) <=> \+ (D2 < D1) | true.

```

In second rule, $\backslash + (C)$ is a safe approximation of the negation of constraint C , i.e., it is only entailed if constraint C cannot possibly hold. In the Prolog context, the built-in negation as failure can be used.

Suspended constraints are attached to all guarded variables so that they are reactivated whenever one of these variables is affected by a built-in constraint. We assume that both attaching and detaching can be done in constant time, although certain current CHR implementations like the K.U.Leuven CHR system do not support detaching in constant time.

5.2.4 Scheduling

Each constraint occurrence corresponds to a (potentially suspended) rule firing if it is the only head of a single headed rule, a (potentially suspended) prefix firing if it is the first head of a multi-headed rule, and a prefix extension in all other cases. A conversion between constraint occurrence and rule firing, prefix firing or prefix extension is made as soon as the constraint in question matches with the head. If such a match is shown to be impossible, the constraint occurrence is discarded. Let there be given a head constraint $c(X_1, \dots, X_n)$. The following function is used to construct a head match.

$$\text{head_match}([X|\bar{X}]) = \begin{cases} \langle [X|\bar{Y}], g \rangle & \text{if } X \text{ is a variable and } X \notin \text{vars}(\bar{X}) \\ \langle [Y|\bar{Y}], (Y = X) \wedge g \rangle & \text{otherwise} \end{cases}$$

where $\langle \bar{Y}, g \rangle = \text{head_match}(\bar{X})$

$$\text{head_match}([]) = \langle [], \text{true} \rangle$$

Now, for each rule in intermediate form syntax

$$p :: r @ \pm H_1, ?g_1, \pm H_2, ?g_2, \dots, \pm H_n, ?g_n \iff B$$

and for $1 \leq i \leq n$ we generate the rules below where $H_i = c(X'_1, \dots, X'_n)$ is the j^{th} occurrence of the user-defined constraint predicate c/n , $\langle [X_1, \dots, X_n], g \rangle = \text{head_match}([X'_1, \dots, X'_n])$, and $\{Y_1, \dots, Y_m\} = \text{vars}(H_i) \setminus \text{vars}(\{H_1, \dots, H_{i-1}\})$.

- If $i = n = 1$:

$$c_occ_j(X_1, \dots, X_n, Id) \iff g \mid r_rf_suspended(Y_1, \dots, Y_m, Id, SId).$$

$$c_occ_j(X_1, \dots, X_n, Id) \iff \backslash + g \mid \text{true}.$$
- If $i = 1$ and $n > 1$:

$$c_occ_j(X_1, \dots, X_n, Id) \iff g \mid r_pf_1_suspended(Y_1, \dots, Y_m, Id, SId).$$

$$c_occ_j(X_1, \dots, X_n, Id) \iff \backslash + g \mid \text{true}.$$
- Otherwise, if $i > 1$:

$$c_occ_j(X_1, \dots, X_n, Id) \iff g \mid$$

$$r_pe_i - 1(Y_1, \dots, Y_m, Id, SId), \text{ schedule_pe}(r_i - 1(Z_1, \dots, Z_l), SId).$$

$$c_occ_j(X_1, \dots, X_n, Id) \iff \backslash + g \mid \text{true}.$$

where $\{Z_1, \dots, Z_l\} = \text{vars}(H_i) \cap \text{vars}(\{H_1, \dots, H_{i-1}\})$.

In the above, if $g = \text{true}$ then the second rule of each pair of rules can be discarded. The suspended prefix and rule firings can sometimes be replaced by regular prefix and rule firings by unfolding (see Section 5.2.3).

In the example program, only the first occurrence of the $e/3$ constraint has a non-trivial head match (the first and last argument must be the same). All prefix and rule firings are followed by the trivial guard true and so we only generate regular prefix and rule firings. They are scheduled using the `schedule_pf/3` and `schedule_rf/2` predicates.

```

source_occ_1(V,Id) <=> d2_rf(V,Id,SIId), schedule_rf(1,SIId).

dist_occ_1(V,D,Id) <=> d3_pf_1(V,D,Id,SIId), schedule_pf(d3_1(V),1,SIId).
dist_occ_2(V,D,Id) <=> d3_pe_1(D,Id,SIId), schedule_pe(d3_1(V),SIId).
dist_occ_3(V,D,Id) <=> d4_pf_1(V,D,Id,SIId), schedule_pf(d4_1(V),D+2,SIId).

e_occ_1(V,C,U,Id) <=> V = U | d1_rf(V,C,Id,SIId), schedule_rf(1,SIId).
e_occ_1(V,C,U,Id) <=> \+ (V = U) | true.
e_occ_2(V,C,U,Id) <=> d4_pe_1(C,U,Id,SIId), schedule_pe(d3_1(V),SIId).
    
```

Prefix firings and extensions are scheduled using a key containing their shared variables. For example for the prefix firings consisting of the first head of rule `d3` and the corresponding prefix extensions consisting of the second head of the same rule, the key equals `d3_1(V)`.

Similar to the suspended prefix and rule firings, the constraint occurrences are attached to all guarded variables. We again assume that both attaching and detaching can be done in constant time.

5.2.5 Matching and Firing

The scheduler initiates the firing of a rule instance by asserting a `fire/1` constraint, and the matching of a prefix firing with a prefix extension by asserting a `match/2` constraint. These constraints have as arguments the identifiers of the corresponding RETE memory constraints. After matching a prefix firing with a prefix extension, a new suspended prefix or rule firing is generated. For a given n -headed rule r with $n > 1$ and for $1 \leq i \leq n - 2$, we generate the following rule

$$r_pf_i(X_1, \dots, X_m, Id_1, \dots, Id_i, SIId_1), r_pe_i(X_{m+1}, \dots, X_l, Id_{i+1}, SIId_2) \setminus \\ \text{match}(SIId_1, SIId_2) \text{ <=> } Id_{i+1} \setminus == Id_1, \dots, Id_{i+1} \setminus == Id_i \mid \\ r_pf_i + 1_suspended(X_1, \dots, X_l, Id_1, \dots, Id_{i+1}).$$

and similarly for $i = n - 1$:

$$r_pf_n - 1(X_1, \dots, X_m, Id_1, \dots, Id_{n-1}, SIId_1), r_pe_n - 1(X_{m+1}, \dots, X_l, Id_n, SIId_2) \setminus \\ \text{match}(SIId_1, SIId_2) \text{ <=> } Id_n \setminus == Id_1, \dots, Id_n \setminus == Id_{n-1} \mid \\ r_rf_suspended(X_1, \dots, X_l, Id_1, \dots, Id_n).$$

A rule firing of an n -headed rule r with body B is fired as follows:

$$r_rf_i(X_1, \dots, X_m, Id_1, \dots, Id_n, SIId), \text{ fire}(SIId) \text{ <=> } \\ Id_{r(1)} = \text{dead}, \dots, Id_{r(l)} = \text{dead}, B.$$

where $r(1), \dots, r(l)$ are the indices of the removed heads of the rule (if any). We furthermore add the following rules at the end of the code, to make sure the CHR compiler detects that the `match/2` and `fire/1` constraints are never to be stored.

```
match(_,_) <=> true.
fire(_) <=> true.
```

For the example program, the generated code is as follows:

```
d1_rf(V,C,Id,SId), fire(SId) <=> Id = dead.
d2_rf(V,Id,SId), fire(SId) <=> dist(V,0).
d3_pf_1(V,D1,Id1,SId1), d3_pe_1(D2,Id2,SId2) \ match(SId1,SId2) <=>
    Id2 \== Id1 | d3_rf_suspended(V,D1,D2,Id1,Id2,SId).
d3_rf(V,D1,D2,Id1,Id2,SId), fire(SId) <=> Id1 = dead.
d4_pf_1(V,D,Id1,SId1), d4_pe_1(C,U,Id2,SId2) \ match(SId1,SId2) <=>
    Id2 \== Id1 | d4_pf(V,D,C,U,Id1,Id2,SId), schedule_rf(D+2,SId).
d4_rf(V,D,C,U,Id1,Id2,SId), fire(SId) <=> dist(U,D+C).

match(_,_) <=> true.
fire(_) <=> true.
```

5.2.6 Clean-up

Whenever a constraint's identifier variable is instantiated, its occurrence representations, as well as those RETE memory constraints in which it participates, are removed. The rules look as follows.

- For the i^{th} occurrence representation for constraint predicate c/n :
 $c_occ_i(X_1, \dots, X_n, Id) \lt;=> \text{nonvar}(Id) \mid \text{true}.$
- For an i -headed suspended prefix firing of rule r :
 $r_pf_i_suspended(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \lt;=> \text{nonvar}(Id_i) \mid \text{true}.$
 $\dots$
 $r_pf_i_suspended(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \lt;=> \text{nonvar}(Id_i) \mid \text{true}.$
- For an i -headed regular prefix firing of rule r :
 $r_pf_i(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \lt;=> \text{nonvar}(Id_i) \mid \text{remove_pf}(SId).$
 $\dots$
 $r_pf_i(X_1, \dots, X_m, Id_1, \dots, Id_i, SId) \lt;=> \text{nonvar}(Id_i) \mid \text{remove_pf}(SId).$
- For a prefix extension of an i -headed prefix firing of rule r :
 $r_pe_i(X_1, \dots, X_m, Id, SId) \lt;=> \text{nonvar}(Id) \mid \text{remove_pe}(SId).$
- For a suspended rule firing of an n -headed rule r :
 $r_rf_suspended(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \lt;=> \text{nonvar}(Id_1) \mid \text{true}.$
 $\dots$
 $r_rf_suspended(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \lt;=> \text{nonvar}(Id_n) \mid \text{true}.$
- For a regular rule firing of an n -headed rule r :
 $r_rf(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \lt;=> \text{nonvar}(Id_1) \mid \text{remove_rf}(SId).$
 $\dots$
 $r_rf(X_1, \dots, X_m, Id_1, \dots, Id_n, SId) \lt;=> \text{nonvar}(Id_n) \mid \text{remove_rf}(SId).$

The predicates `remove_pf/1`, `remove_pe/1` and `remove_rf/1` remove respectively a prefix firing, prefix extension and rule firing from the schedule. The following clean-up rules are generated for the example program.

```

source_occ_1(V,Id) <=> nonvar(Id) | true.
dist_occ_1(V,D,Id) <=> nonvar(Id) | true.
dist_occ_2(V,D,Id) <=> nonvar(Id) | true.
dist_occ_3(V,D,Id) <=> nonvar(Id) | true.
e_occ_1(V,C,U,Id) <=> nonvar(Id) | true.
e_occ_2(V,C,U,Id) <=> nonvar(Id) | true.

d1_rf(V,C,Id,SId) <=> nonvar(Id) | remove_rf(SId).
d2_rf(V,Id,SId) <=> nonvar(Id) | remove_rf(SId).
d3_pf_1(V,D1,Id1,SId) <=> nonvar(Id1) | remove_pf(SId).
d3_pe_1(D2,Id2,SId) <=> nonvar(Id2) | remove_pe(SId).
d3_rf(V,D1,D2,Id1,Id2,SId) <=> nonvar(Id1) | remove_rf(SId).
d3_rf(V,D1,D2,Id1,Id2,SId) <=> nonvar(Id2) | remove_rf(SId).
d4_pf_1(V,D,Id1,SId) <=> nonvar(Id1) | remove_pf(SId).
d4_pe_1(C,U,Id2,SId) <=> nonvar(Id2) | remove_pe(SId).
d4_rf(V,D,C,U,Id1,Id2,SId) <=> nonvar(Id1) | remove_rf(SId).
d4_rf(V,D,C,U,Id1,Id2,SId) <=> nonvar(Id2) | remove_rf(SId).

d3_rf_suspended(V,D1,D2,Id1,Id2,SId) <=> nonvar(Id1) | true.
d3_rf_suspended(V,D1,D2,Id1,Id2,SId) <=> nonvar(Id2) | true.
    
```

5.3 Program-Independent Part: the Scheduler

The scheduler implements the `schedule_rf/2`, `schedule_pf/3` and `schedule_pe/2` predicates. It furthermore implements the `execute/0` predicate which retrieves and executes the highest priority scheduled task. This task either is the firing of a rule instance by asserting a `fire/1` constraint, or the matching of a prefix firing with a prefix extension by asserting a `match/2` constraint. The `execute/0` predicate recursively calls itself until no more tasks are scheduled. It is first called after processing the initial goal.

For the implementation of the scheduler, we use a variant of the scheduling algorithm presented in (De Koninck 2007). This algorithm can be used to maintain which prefix firings are still to match with which prefix extensions. It is roughly based on the $\mathcal{W}(r, t)$ data structures used in (Ganzinger and McAllester 2002). Such a data structure consists of a series (implemented as a linear linked list) of *prefix blocks*, which are sets of prefix firings and (apart from the last one) are associated with a prefix extension.

The semantics of the $\mathcal{W}(r, t)$ data structure is that the prefix firings of a given prefix block are still to match with the prefix extension associated to it, as well as with all prefix extensions associated to subsequent prefix blocks. The last prefix block has no associated prefix extension, and represents those prefix firings that have been matched with all prefix extensions and hence are passive (or *completed* using the terminology of (Ganzinger and McAllester 2002)). Whenever a prefix extension is deleted, its prefix block is merged with the next prefix block.

There is one $\mathcal{W}(r, t)$ data structure for each prefix length of each rule and for each combination of arguments shared between a prefix firing and prefix extension. Each prefix block is represented as a (local) priority queue whose items are the block's prefix firing. The highest priority item of each prefix block, together with its associated prefix extension, is also represented in a global priority queue. This prefix block representative is updated whenever the highest priority prefix firing of the prefix block is removed, a new prefix firing has the highest priority, or the associated prefix extension is removed. The global priority queue furthermore contains a representative for each rule firing. The reason for using two layers of priority queues is to reduce the amount of work needed when the prefix firings of a prefix block all become passive due to a prefix extension removal. It is the global priority queue that determines the next task to perform, i.e., matching a prefix firing with a prefix extension, or firing a rule instance.

In the context of CHR^{P} , built-in constraint (in particular equality constraints) on the arguments shared between a prefix firing and extension, may require merging of $\mathcal{W}(r, t)$ data structures. The data structure of (De Koninck 2007) supports schedule merges in quasi constant time. The most notable difference with the $\mathcal{W}(r, t)$ data structure of (Ganzinger and McAllester 2002) is that the prefix blocks form a circular linked list. Using this representation, merging schedules consists of cross-linking the circular lists and reactivating the prefix firings that were passive before the merge. Special care is taken to prevent both that a prefix firing is being matched with the same prefix extension more than once, and that a prefix firing 'misses' a prefix extension.

One consequence of using a circular linked list instead of a linear one to represent the prefix blocks, is that it is unclear (or more precisely, too expensive to decide) which prefix firings become passive whenever a prefix extension is deleted. Therefore, this decision is postponed until the scheduler tries to match the prefix firing with the next prefix extension in line. For complexity reasons, it is important that all prefix firings that have simultaneously been reactivated, and have not been matched with a prefix extension since this reactivation, are simultaneously made passive in time independent of the number of prefix firings affected. In (De Koninck 2007), a so-called *element schedule* based on a stack is proposed to support this. In our context, we need an element schedule that is based on priority queues. It works as follows.

We use three types of priority queues. The first one is a single *global* priority queue which contains an item for each rule firing, for each *active* prefix firings that either has not been passive before or has been matched with at least one prefix extension since its last activation, and finally, for each set of prefix firings that have been simultaneously activated and have not been matches with a prefix extension since. A second type of priority queues is called a *local* queue and represents the above mentioned sets of prefix firings. Finally, the third type of queues is the *passive* queue which contains an item for each passive (completed) prefix firing. There is one passive queue for each schedule. Essentially, we again use two layers of priority queues. Whenever a set of previously passive prefix firings, represented as a passive priority queue, is reactivated because of a new prefix extension or because of a

Figure 1. Example schedule with global, local and passive priority queues

schedule merge, this passive priority queue becomes a local priority queue and has a representative inserted into the global priority queue. If such a representative is the highest priority item in the global priority queue, and an `execute/0` call is made, then the highest priority prefix firing of the represented local priority is removed and dealt with as an ordinary prefix firing. The representatives of local priority queues are updated (and potentially removed) similarly to how this is done in the $\mathcal{W}(r, t)$ data structure of (Ganzinger and McAllester 2002).

Example 8

Figure 1 illustrates the prefix blocks, the different types of priority queues, and their contents. The global queue, which is shared by all schedules, contains the rule firings RF_1 and RF_2 , the prefix firings PF_1 , PF_4 , PF_5 and PF_8 (the last of which belongs to another schedule), and the local queue representative LQ_1 . The represented local queue contains the prefix firings PF_2 and PF_3 which are by definition also in the same prefix block. The schedule's passive queue contains the prefix firings PF_6 and PF_7 . The schedule has two prefix blocks, which are associated with respectively the prefix extensions PE_1 and PE_2 .

Using our approach, the cost of deleting items from the global priority queue can be amortized to one of the following events: a new rule firing, a new prefix firing, a new prefix extension (for each representative of a local priority queue), or a match between a prefix firing and a prefix extension (which corresponds to either a new larger prefix firing, or a rule firing).

In (Ganzinger and McAllester 2002), retrieving the schedule for a given prefix firing or prefix extension is done by hashing. In our approach, we use a variant of hashing, which we call *non-ground hashing* and which consists of first replacing all variables by a unique identifier, and then using the resulting (ground) term for hashing. Unifications may require rehashing the affected keys and potentially also the merging of schedules.

5.4 Priority Queues

A priority queue or heap is a data structure that contains a set of prioritized items and supports the following operations: inserting and removing an item, finding a highest priority item and merging with another queue. The implementation proposal in (Ganzinger and McAllester 2002) suggests the use of two types of priority queues, one for the fixed priorities, where each of the supported operations takes constant time, and Fibonacci heaps for the dynamic priorities.

Fibonacci heaps (Fredman and Tarjan 1987) are a type of priority queue that offer $\mathcal{O}(1)$ amortized time insertion, heap merging and finding a highest priority item, and $\mathcal{O}(\log n)$ amortized time item removal with n the number of items in the queue. It is suggested in (Ganzinger and McAllester 2002) that by using only one node per priority, using linked lists to represent the items that share this priority, the item removal cost can be reduced to $\mathcal{O}(\log N)$ with N the number of distinct priorities. However, this increases the cost of heap merging from $\mathcal{O}(1)$ for a single merge operation to a total cost of $\mathcal{O}(n \log N)$ for merging heaps when there are n items in total and N distinct priorities (as is shown in an Appendix of (De Koninck et al. 2007a)). A CHR implementation of Fibonacci heaps is described in (Sneyers et al. 2006a). It can easily be extended to support multiple heaps that can be merged and to use only one node for each distinct priority per heap.

6 A New Meta-Complexity Result for CHR^{P}

In this section, we give a new meta-complexity result for CHR^{P} . It extends the result via translation to Logical Algorithms, by also supporting built-in constraints and non-ground CHR constraints. We make the following assumptions:

- Hash tables support $\mathcal{O}(1)$ insertion, removal, and retrieval of all elements that match a given (ground) key.
- The inverse of the Ackermann function ($\alpha(n)$) is a constant.

The first assumption is also made in (Ganzinger and McAllester 2002) and holds on average as long as the hash function is good enough. The second assumption is needed for our scheduling data structure (De Koninck 2007) which internally makes use of the union-find algorithm. The inverse of the Ackermann function is positive and less than 5 for all practical purposes.

We start by looking at the complexity of the different operations supported by our scheduler.

Lemma 4 (Scheduler Costs)

Let N be the number of distinct priorities, and assume that a priority queue merge takes some abstract time T , then the schedule operations have the following amortized cost:

- $\mathcal{O}(1)$ and $\mathcal{O}(\log N)$ for each `schedule_pf/3`, `remove_pf/1`, `remove_pe/1`, `remove_rf/1` and `execute/0` operation involving respectively a static and dynamic priority rule

- $\mathcal{O}(T + 1)$ and $\mathcal{O}(T + \log N)$ for each `schedule_pe/2` operation involving respectively a static and dynamic priority rule
- $\mathcal{O}(1)$ for each schedule merge and `schedule_rf/2` operation

Proof

We only consider the costs related to the priority queue operations. The other costs are shown to be (quasi) constant in (De Koninck 2007). We now look at the different operations in detail:

- A `schedule_pf/3` call consists of inserting the new prefix firing into the global priority queue. We also account to this event, the cost of making the new prefix instance passive the first time. That operation consists of a removal from the global priority queue and an insertion into the schedule's passive queue. The total cost is $\mathcal{O}(1)$ if the element has a static priority, and $\mathcal{O}(\log N)$ if it has a dynamic priority.
- A `schedule_pe/2` call requires the insertion of a new representative for the local priority queue of reactivated prefix firings, into the global priority queue. We also take into account here, the cost of making all the reactivated prefix firings passive that have not been matched with a prefix extension since the reactivation. That operation consists of removing the representative and merging the local priority queue with the schedule's passive queue. The cost is $\mathcal{O}(T + 1)$ for a static priority rule and $\mathcal{O}(T + \log N)$ time for a dynamic priority one.
- A `schedule_rf/2` call requires an insertion into the global priority queue which takes $\mathcal{O}(1)$ time.
- A `remove_pf/1` call consists of deleting the prefix firing from the global priority queue, from a local priority queue or from a passive queue. A deletion from a local queue may moreover require an update of the global queue (removal and insertion). In total, this takes $\mathcal{O}(1)$ time for a static priority rule and $\mathcal{O}(\log N)$ time for a dynamic priority rule.
- A `remove_pe/1` call does not require any priority queue operations, and so the cost is $\mathcal{O}(1)$.
- A `remove_rf/1` call requires a removal from the global priority queue which takes $\mathcal{O}(1)$ time if it involves a static priority rule and $\mathcal{O}(\log N)$ time if it involves a dynamic priority rule.
- An `execute/0` call requires retrieval and potential removal (if the retrieved item corresponds to a rule firing, or to a prefix firing that becomes passive) of the highest priority item in the global priority queue. If the retrieved item represents a prefix firing or set of prefix firings that need to be made passive, the cost of this operation is already accounted for by a previous `schedule_pf/3` or `schedule_pe/2` operation. In such case, we call the `execute/0` call *unsuccessful*. An unsuccessful `execute/0` call is followed by another `execute/0` call until either such a call is successful, or the global priority queue is empty and thus a final state is reached. The cost of all unsuccessful `execute/0` calls can be amortized to previous events. If in case of a successful `execute/0` call,

the item retrieved from the global priority queue corresponds to the representative of a local priority queue, the operation requires a removal of the highest priority item (prefix firing) from this local queue, an insertion of the prefix firing into the global priority queue, and potentially the insertion of a new representative for the local queue into the global queue. The cost of a successful `execute/0` call therefore equals $\mathcal{O}(1)$ if it involves a static priority rule and $\mathcal{O}(\log N)$ otherwise.

- A schedule merge requires the reactivation of the passive prefix firings of the merged schedules. The cost analysis is similar to that of a `schedule_pe/2` call. Moreover, each schedule merge can be accounted for by at least one `schedule_pe/2` call as the resulting schedule contains at least one prefix extension more than each of the original schedules, and so the number of schedule merges is bounded by the number of prefix extensions. Therefore, the cost of a single schedule merge can be considered constant.

□

In the above lemma, we have made abstraction of the cost of priority queue merge operations. Such merges take place when the prefix firings in a local priority queue all become passive. In such an event, the local priority queue is merged with the schedule's passive queue. It is easy to see that the cost of merging priority queues for static priorities takes constant time per merge operation. In Section 6 a bound is given on the total cost of merging Fibonacci heaps with one node per distinct priority, given the number of items ever inserted into the heaps. The following lemma makes use of this result.

Lemma 5 (Fibonacci Heap Merging Cost)

The total cost of Fibonacci heap merges is $\mathcal{O}((P_d + A_d) \cdot \log N)$ where P_d is the number of strong prefix firings of dynamic priority rules, A_d is the number of constraints that may participate in a dynamic priority rule instance, and N is the number of distinct rule priorities.

Proof

We count the number of items ever inserted into the local and passive Fibonacci heaps, and then apply the result of Section . A local priority queue basically is the same as a passive priority queue in which items are no longer inserted. Therefore, a merge between a local queue and a passive queue can be seen as a special case of a merge between two passive queues and so we only need to consider these passive priority queues. Each item inserted in such a queue is either a prefix firing that has never been passive before, or a prefix firing that has been matched with a prefix extension at least once since its last activation. The total number of these items is $\mathcal{O}(P_d + A_d)$ because each prefix firing that has been matched with a prefix extension is by definition a strong prefix firing, and each new prefix firing either results from matching a (smaller) strong prefix firing and extension and hence corresponds to a (potentially suspended) strong prefix firing, or consists of a single head in which case it corresponds to a constraint assertion. Now given the number of items ever

inserted into the passive priority queues, the total cost of merging Fibonacci heaps hence is $\mathcal{O}((P_d + A_d) \cdot \log N)$. $\square$

We are now ready to formulate the new meta-complexity theorem.

Theorem 6

Let A_s and A_d be the number of assertions of constraints with an occurrence in respectively a static and dynamic priority rule. Let P_s and P_d be the number of strong prefix firings of respectively static and dynamic priority rules. The time complexity of a CHR^{rp} program executed using our implementation is

$$\mathcal{O}((1 + C_B^{\text{ask}}) \cdot (A_s + P_s + (A_d + P_d) \cdot \log N) + B \cdot C_B^{\text{tell}} \cdot (K + C_B^{\text{ask}} \cdot S))$$

where N is the number of distinct priorities, C_B^{ask} is the cost of evaluating a built-in ask constraint, C_B^{tell} is the cost of solving a built-in tell constraint, and B is the number of built-in tell constraints asserted in rule bodies; K is the maximum number of distinct combinations (keys) of arguments shared between prefix firings and extensions in which any given variable occurs, and S is the maximum number of suspended strong prefix firings (i.e., those that are followed by a non-trivial guard) and suspended instances of constraint occurrences (i.e., whose arguments are not mutually distinct variables) in which any given variable occurs.

Proof

Each new CHR constraint causes the creation of constraint occurrences which are converted into RETE memory constraints as soon as the implicit guard on the constraint arguments is entailed (i.e., the constraint matches the head in question). These RETE memory constraints are scheduled using `schedule_pf/3` for the single-headed prefix firings, `schedule_rf/2` for the single-headed rule firings, and `schedule_pe/2` for the prefix extensions. The total cost of these operations, including the cost of priority queue merges (for the `schedule_pe/2` calls), equals $\mathcal{O}((1 + C_b^{\text{ask}}) \cdot (A_s + (A_d + P_d) \log N))$. Each constraint deletion causes the deletion of those RETE memory constraints in which the deleted constraint participated. The total cost related to deletion therefore is $\mathcal{O}(A_s + P_s + (A_d + P_d) \log N)$. Each prefix firing is inserted into its schedule at most once and hence it can also be removed from this schedule only once (when one of its constituent constraints is removed). Those prefix firings that consist of at least two heads, correspond to a strong prefix firing as they are generated at a priority higher or equal to that of the highest priority rule firing. Thus, using Lemma 4 and including the cost of checking the relevant parts of the guard, the cost for inserting (and deleting) these prefix firings is $\mathcal{O}((1 + C_b^{\text{ask}}) \cdot (P_s + P_d \log N))$.

A built-in tell constraint is processed as follows. The keys used to identify the schedules and that are affected by the built-in constraint, are rehashed. If the built-in constraint causes two or more schedules to have the same key, these schedules are merged. The cost of rehashing is proportional to the number of affected keys and the cost of a schedule merge is constant by Lemma 4. A built-in constraint moreover requires the reactivation of the suspended prefix firings and rule firings, as well as those constraint occurrences for which it is not decided whether they match

with the corresponding head or not. The reactivated prefix and rule firings have their guard checked and are potentially scheduled as regular (non-suspended) prefix and rule firings. The reactivated constraint occurrences also have their (implicit) guard checked, and are potentially scheduled as single-headed prefix firings, single-headed rule firings, or prefix extensions. The cost of the scheduling operations was already taken into account above. The remaining cost per built-in tell constraint is $\mathcal{O}(C_b^{tell} \cdot (K + C_b^{ask} \cdot S))$. $\square$

Because the values of S and K might be difficult to determine in practice, we propose the following bounds: $S = \mathcal{O}(A_s + A_d + P_s + P_d)$ and $K = \mathcal{O}(A_s + A_d)$. We have used the cost of solving a built-in tell constraint as an upper bound on the number of variables that are affected. Note that in absence of built-in constraints, the theorem given here is essentially the same as the one for Logical Algorithms.¹¹

6.1 Examples

We illustrate the meta-complexity theorem on some examples, and compare with the results obtained by using the approach of (Frühwirth 2002b).

Example 9 (Less-or-Equal)

The less-or-equal (**leq**) program is classic CHR example. It implements a less-than-or-equal-to constraint by eventually translating it into equality constraints. A CHR^{IP} implementation of the program consists of the following rules.

```
1 :: idempotence @ leq(X,Y) \ leq(X,Y) <=> true.
2 :: reflexivity @ leq(X,X) <=> true.
2 :: antisymmetry @ leq(X,Y), leq(Y,X) <=> X = Y.
3 :: transitivity @ leq(X,Y), leq(Y,Z) ==> leq(X,Z).
```

Given an initial goal consisting of n **leq**/2 constraints where the arguments are taken from a set of n distinct variables, we derive the following values for the parameters:

- P_s : the number of strong prefix firings is $\mathcal{O}(n^2)$ for the **idempotence** rule, $\mathcal{O}(n)$ for the **reflexivity** rule, $\mathcal{O}(n^2)$ for the **antisymmetry** rule, and $\mathcal{O}(n^3)$ for the **transitivity** rule. These numbers are found by looking at the degrees of freedom for each constraint occurrence, based on the domain of the arguments, and given those arguments that are already fixed by the left-most heads. For example for the **transitivity** rule, we know that there are $\mathcal{O}(n^2)$ constraints matching the first head, and $\mathcal{O}(n)$ constraints matching the second head, given the first. Our reasoning is based on the fact that at priority 2 and lower, all **leq**/2 constraints have set semantics because of the **idempotence** rule.
- A_s : the number of **leq**/2 constraints asserted is $\mathcal{O}(n^3)$ (by the **transitivity** rule).

¹¹ The Logical Algorithms result makes use of the size of the initial database instead of the number of assertions. We have that $A_s = \mathcal{O}(|\sigma_0| + P_s + P_d)$.

- *B*: the number of built-in constraints is bounded by the number of rule firings of the **antisymmetry** rule, and hence is $\mathcal{O}(n^2)$.
- *K*: the schedule keys are the combination of **X** and **Y** in both the **antisymmetry** rule and the **idempotence** rule, and **Y** in the **transitivity** rule. There are at most $\mathcal{O}(n)$ different keys in which any given variable occurs.
- *S*: for any variable, and in a state in which a built-in constraint can be asserted, there are up to $\mathcal{O}(n)$ suspended instances of the **leq/2** occurrence in the **reflexivity** rule. There can be no suspended prefix or rule firings.
- C_b^{ask} and C_b^{tell} : the cost of evaluating a built-in ask constraint and the cost of solving a built-in tell constraint is constant (at least for the given query pattern).

Filling in these parameters in the formula given by Theorem 6 gives us a worst case time complexity of

$$\mathcal{O}((1+1) \cdot (n^2 + n^3) + (0+0) \cdot \log 1 + n^2 \cdot 1 \cdot (n+1 \cdot n)) = \mathcal{O}(n^3)$$

This corresponds to the actual worst-case complexity for an initial goal of the form

$$\{\text{leq}(X_1, X_2), \dots, \text{leq}(X_{n-1}, X_n), \text{leq}(X_n, X_1)\}$$

The approach of (Frühwirth 2002b) does not apply since the **transitivity** rule is a propagation rule and hence no suitable ranking function can be found.

Example 10 (Merge Sort)

Consider the CHR^{TP} implementation of the merge sort algorithm, first given in Example 6 (Section 4) and repeated here for easy reference.

```

1 :: ms1 @ arrow(X,A) \ arrow(X,B) <=> A < B | arrow(A,B).
2 :: ms2 @ merge(N,A), merge(N,B) <=> A < B | merge(2*N+1,A), arrow(A,B).
3 :: ms3 @ number(X) <=> merge(0,X).
```

We show that the total runtime of the algorithm is $\mathcal{O}(n \log n)$ given an initial goal consisting of n **number/1** constraints.

No new **number/1** constraints are ever asserted. Rule **ms3** converts one **number/1** constraint into one **merge/2** constraint each time it fires. The number of (strong) prefix firings for rule **ms3** hence is $\mathcal{O}(n)$. Rule **ms2** decreases the number of **merge/2** constraints by one and so it can fire $n - 1$ times. In any state, there are at most two **merge/2** constraints with the same first argument. This invariant holds in the initial state because there are no **merge/2** constraints in the initial goal and rule **ms2** can fire after each new **merge/2** constraint assertion, enforcing the invariant. Because of the invariant, the number of prefix firings for rule **ms2** is limited to $\mathcal{O}(n)$.

Using similar reasoning it holds that in any state, there are at most two **arrow/2** constraints in the store with the same first argument. Now we define that in a given state, two numbers X_1 and X_m are connected by a *chain* of length $m - 1$ if the following constraints are in the store: **arrow**(X_1, X_2), **arrow**(X_2, X_3), ..., **arrow**(X_{m-1}, X_m). At priority 2 it holds that for each **merge**(N, X) constraint in the store, the maximal length of a chain starting in X is N . Indeed, this holds

for the initial `merge(0, ⊥)` constraints and if it holds for `merge(N, ⊥)` constraints, it also holds for `merge(2 · N + 1, ⊥)` constraints, because when such a constraint is asserted, two chains of length N are linked with an extra `arrow/2` constraint and merged by up to $2 · N$ firings of rule `ms1`. Two `merge(N, ⊥)` constraints are combined into a `merge(2 · N + 1, ⊥)` constraint, so the n `merge(0, ⊥)` constraints asserted by rule `ms3` are replaced by $n/2$ `merge(1, ⊥)` constraints, which in turn are combined into $n/4$ `merge(3, ⊥)` constraints and so on until finally 1 `merge(n − 1, ⊥)` constraint remains. The sum of all N in these `merge(N, ⊥)` constraints is $\mathcal{O}(n \log n)$. Rule `ms1` fires $\mathcal{O}(N)$ times after every new `merge(N, ⊥)` constraint assertion and because there are at most two `arrow/3` constraints with the same first argument, there are $\mathcal{O}(n \log n)$ strong prefix firings of rule `ms1`.

In conclusion, for an input database consisting of n `number/1` constraints, there are $\mathcal{O}(n \log n)$ strong prefix firings for rule `ms1`, $\mathcal{O}(n)$ for rule `ms2` and $\mathcal{O}(n)$ for rule `ms3`. Using the meta-complexity theorem, which simplifies to the one for Logical Algorithms because there are no built-in tell constraints, the total runtime is $\mathcal{O}(n \log n)$, which is also a tight complexity bound. We now compare this result with the result found by using the meta-complexity theorem of (Frühwirth 2002b).

Using a similar analysis as above, we can derive that $D = \mathcal{O}(n \log n)$ and $c_{max} = \mathcal{O}(n)$ where n is the number of `number/1` constraints in the query.¹² The cost of head matching ($\mathcal{O}_{H_r}$), guard checking ($\mathcal{O}_{G_r}$), adding built-in constraints ($\mathcal{O}_{C_r}$), and adding and removing CHR constraints ($\mathcal{O}_{B_r}$), can all be assumed constant. The number of heads n_r of a rule $r \in P$ is at most 2. Filling in these numbers, we derive a total worst case complexity of $\mathcal{O}(n^3 \log n)$, which is clearly suboptimal.

6.2 Comparison with the Logical Algorithms meta-complexity result

In (De Koninck et al. 2007a), we have presented a direct implementation of the Logical Algorithms language into CHR that satisfies the complexity requirements needed for the Logical Algorithms meta-complexity result to hold. In this subsection, we show that this implementation has become somewhat obsolete because we can achieve the same result by combining the translation from Logical Algorithms to CHR^{IP} of Section 3, with the CHR^{IP} implementation presented in Section 5. We assume here that the comparison antecedents in Logical Algorithms programs are scheduled after the corresponding user-defined antecedents in the translation, and that the guards on the mode indicators (these have the form $N \neq p$) are scheduled right after the head to which they apply.

Theorem 7

The time complexity of Logical Algorithms programs executed by first translating them into CHR^{IP} programs using the translation schema of Section 3, and then executing the resulting CHR^{IP} program using the implementation of Section 5, is $\mathcal{O}(|\sigma_0| + P_s + (P_d + A_d) \cdot \log N)$ with S_0 , P_s , P_d , A_d and N as defined in Section 2.1.

¹² In Theorem 4.2 of (Frühwirth 2002b) a worst case upper bound of $c_{max} = \mathcal{O}(c + D)$ is used, with c the number of constraints in the query, which becomes $c_{max} = \mathcal{O}(n \log n)$ in this example. The bound we use is tight, i.e., $c_{max} = \Theta(n)$.

Proof

The translation of a Logical Algorithms program P consists of two parts as defined in Section 3. The first part, denoted by $T_{S/D}(P)$, contains for each user-defined predicate a/n the following rules:

$$\begin{aligned}
 1 &:: a_r(\bar{X}, M) \setminus a(\bar{X}) \iff M \neq \mathbf{n} \mid \text{true} \\
 1 &:: a_r(\bar{X}, \mathbf{n}), a(\bar{X}) \iff a_r(\bar{X}, \mathbf{b}) \\
 2 &:: a(\bar{X}) \iff a_r(\bar{X}, \mathbf{p}) \\
 1 &:: a_r(\bar{X}, M) \setminus \text{del}(a(\bar{X})) \iff M \neq \mathbf{p} \mid \text{true} \\
 1 &:: a_r(\bar{X}, \mathbf{p}), \text{del}(a(\bar{X})) \iff a_r(\bar{X}, \mathbf{b}) \\
 2 &:: \text{del}(a(\bar{X})) \iff a_r(\bar{X}, \mathbf{n})
 \end{aligned}$$

It is easy to see that for an initial goal containing no constraints of the form $a_r(\bar{X}, M)$ and since these are the only rules that assert such a constraint, in any state it holds that if $a_r(\bar{X}, M_1) \# i_1$ and $a_r(\bar{X}, M_2) \# i_2$ are in the CHR constraint store, then $i_1 = i_2$ and $M_1 = M_2$. This implies that the number of strong prefix firings for these rules is bounded by the number of assertions of $a(\bar{X})$ or $\text{del}(a(\bar{X}))$.

The second part of the translation, denoted by $T_R(P)$, contains for each Logical Algorithms rule

$$r @ p : A_1, \dots, A_n \Rightarrow C$$

a set of rules

$$p + 2 :: r_\rho @ H \implies g_1, g_2 \mid C'$$

as shown in the translation schema of Section 3.1.2. Amongst these rules is one, say $r_{\rho'}$, with a maximal number of heads, namely as many as there are user-defined antecedents in $A_1, \dots, A_n$. Because the (implicit and explicit) guards on the mode indicators of the head constraints are scheduled as soon as they are decidable, and because the comparisons are scheduled at corresponding places, it is easy to see that the number of strong prefix firings of rule $r_{\rho'}$ is the same as the number of strong prefix firings of Logical Algorithms rule r . The other r_ρ rules are restricted versions of $r_{\rho'}$ and therefore have at most as many strong prefix firings as $r_{\rho'}$.

The assertions with occurrences in dynamic priority rules are of the form $a_r(\bar{X}, _)$. The set and deletion semantics rules ensure that the number of these assertions is the same in the original program and in its translation. Now using our new meta-complexity result for CHR^{P} (Theorem 6), we derive that the total runtime complexity of the translated program is $\mathcal{O}(|\sigma_0| + P_s + (P_d + A_d) \cdot \log N)$.¹³ $\square$

6.3 Comparison with the “As Time Goes By” approach

In Section 2.3.3 we already briefly compared the Logical Algorithms meta-complexity theorem with the theorem given by Frühwirth in (Frühwirth 2002b). In this sub-

¹³ Since $A_s = \mathcal{O}(|\sigma_0| + P_s + P_d)$.

section, we make the comparison complete by also considering built-in constraints, using the new meta-complexity theorem presented in Section 6.

Let there be given a CHR^{rp} program P in which each rule has the same (static) priority. Theorem 3 in (De Koninck et al. 2007b) states that such a CHR^{rp} program and its corresponding CHR program (which is found by removing the rule priorities) have the same derivations. Therefore, such programs are suitable for comparing the result of (Frühwirth 2002b) with the result of Theorem 6 in Section 6. In Section 2.3.3 we have already shown that the number of strong prefix firings is $\mathcal{O}(D \cdot \sum_{r \in P} c_{max}^{n_r})$ where D is the derivation length (i.e., the number of rule firings), and c_{max} is the maximal number of CHR constraints in the store in any state. The number of constraint assertions is $\mathcal{O}(c_{max} + D)$. If we assume that the initial goal does not contain any built-in constraints (as is done in (Frühwirth 2002b)), then the number of built-in constraints is $\mathcal{O}(D)$. The number of suspended prefix firings is bounded by $\mathcal{O}(\sum_{r \in P} c_{max}^{n_r})$ in any state and the number of suspended assertions by $\mathcal{O}(c_{max})$. Now, filling in these parameters in the CHR^{rp} meta-complexity result gives us that the total runtime complexity is

$$\mathcal{O}\left((1 + O_C) \cdot D \sum_{r \in P} (c_{max}^{n_r} \cdot O_{G_r})\right) \quad (3)$$

where $O_C = \sum_{r \in P} (O_{C_r})$. This formula strongly resembles the result of (Frühwirth 2002b) which, assuming the cost of head matching O_{H_r} and adding and removing CHR constraints O_{B_r} is constant, equals

$$\mathcal{O}\left(D \sum_{r \in P} (c_{max}^{n_r} \cdot O_{G_r} + O_{C_r})\right) \quad (4)$$

The difference lies in how built-in tell constraints are dealt with. In our CHR^{rp} implementation, as well as in any CHR implementation based on the refined operational semantics of CHR, a built-in tell constraint causes the constraints or matches whose variables are affected, to be reconsidered.¹⁴ Because each individual (atomic) built-in constraint is dealt with separately, this may cost more in total than the naive approach taken in (Frühwirth 2002b) in which after each rule firing, *all* constraints or matches are reconsidered *once*. So, while in certain rather exceptional cases, a naive approach to dealing with built-in tell constraints might in fact be better than the usual approach of selective reactivation (as can be seen by comparing Formulas (3) and (4)), in general we expect the latter approach to be an improvement over the naive one. Moreover, in these exceptional cases, the meta-complexity theorem of (Frühwirth 2002b) does not apply to optimized CHR implementations like the K.U.Leuven CHR system, i.e., in these cases it does not overestimate the actual worst case time complexity.

¹⁴ Which constraints are reactivated depends on the wake-up policy used for the **Solve** transition, see also (Schrijvers 2005, Section 5.4.2).

7 Conclusions

In this paper, we have investigated the relationship between the Logical Algorithms language and Constraint Handling Rules. We have presented an elegant translation schema from Logical Algorithms to CHR^{P} : CHR extended with user-definable rule priorities. The original program and its translation are shown to be essentially weakly bisimilar. However, our current CHR^{P} system (De Koninck et al. 2007) does not give the complexity guarantees needed for the Logical Algorithms meta-complexity theorem to hold via this translation.

As a first step towards applying the Logical Algorithms meta-complexity result to CHR^{P} programs, we have shown how a subclass of CHR^{P} can be translated into Logical Algorithms. By using this translation, we can directly apply the meta-complexity theorem for Logical Algorithms to the translated CHR^{P} programs. A drawback is that the CHR^{P} programs that can be translated this way, are restricted to those that do not make use of an underlying constraint solver.

In order to remedy both the limitation that the translation from Logical Algorithms to CHR^{P} does not exhibit the required complexity when executing translated Logical Algorithms programs using our CHR^{P} system, and the restriction of those CHR^{P} programs that can be translated to Logical Algorithms and hence to which the Logical Algorithms meta-complexity result can be applied, we have proposed a new implementation for the complete CHR^{P} language that gives strong complexity guarantees. The implementation is based on the high-level implementation proposal of (Ganzinger and McAllester 2002) as well as on the scheduling data structure of (De Koninck 2007), and consists of the compilation of CHR^{P} rules into (regular) CHR rules, combined with a scheduler that controls the execution. The implementation supports a new and accurate meta-complexity theorem for CHR^{P} . When combining the translation from Logical Algorithms to CHR^{P} with the new implementation, the new meta-complexity theorem implies the Logical Algorithms meta-complexity result. Moreover, it is shown that in general¹⁵ the new theorem is at least as accurate as the meta-complexity result for CHR given by Frühwirth in (Frühwirth 2002b). This is illustrated on two non-trivial examples, one of which contains both built-in constraints and propagation rules and therefore cannot be analyzed using the Logical Algorithms approach or Frühwirth's result.

7.1 Related Work

The time complexity of programs is in general expressed in terms of the number of elementary operations, e.g., the number of logical inferences in Prolog, function applications in a functional programming language, or rule applications in a language such as CHR. However, while in most languages, these elementary operations all take constant time, this is not the case in a language like CHR where each rule application results from a complex matching phase.

In this work, we have made a mapping from the number of elementary operations

¹⁵ Apart from some rather exceptional cases, see Section 6.3.

(like prefix and rule firings or constraint assertions) to time complexity. To the best of our knowledge, and apart from the results in (McAllester 1999; Ganzinger and McAllester 2001; Ganzinger and McAllester 2002) and (Frühwirth 2002a; Frühwirth 2002b), there is no other work with a similar goal. There are many other formalisms though in which elementary operations take more than constant time. One such formalism is term rewriting, as implemented by the Maude system (Clavel et al. 1999) or the ACD term rewriting language (Duck et al. 2006). It is known that AC matching, which is used by most of these languages, is NP-complete. Another formalism is that of production rule systems like Drools (Proctor et al. 2007) or Jess (Friedman-Hill 2007). Production rules are in many ways similar to Constraint Handling Rules. However unlike CHR, these systems are not often used as general purpose programming language, and therefore, algorithmic complexity has never been much of a concern. More work exists on the derivation of the number of elementary operations. In the context of CHR, this mostly concerns the number of rule firings, which is often derived as part of termination analysis (Frühwirth 2000b; Pilozi et al. 2007; Voets et al. 2007).

Another related topic is that of space complexity, an issue that is not dealt with in this paper. In the context of CHR, the memory reuse techniques developed in (Sneyers et al. 2006b) are crucial to achieve optimal space complexity as is shown in (Sneyers et al. 2008). The latter also introduces a space complexity meta-theorem for CHR, stating that the space complexity is $\mathcal{O}(D + p)$ where D is the derivation length and p is the number of propagation rule firings (which takes into account the size of the propagation history).

7.2 Future Work

For a previous version of this paper (De Koninck et al. 2007a), we have made an actual implementation for the Logical Algorithms language in CHR. This implementation satisfies the complexity requirements needed for the Logical Algorithms meta-complexity theorem to hold, when executed using the K.U.Leuven CHR system on top of SWI-Prolog. However, the very large constant factors and the high memory consumption makes that the system is not very useful in practice. Currently, we have no running version of the alternative implementation for CHR^{rp} presented in Section 5. The reason is that this implementation proposal is based on a similar approach as the Logical Algorithms one, and in particular the more complicated scheduler is expected to be slow in practice. However, we do intend to investigate the advantages and disadvantages of a lazy RETE based matching algorithm for CHR^(rp) compared to the LEAPS style matching that is currently used by almost all systems. A simplified version of the scheduling data structure of (De Koninck 2007) which would offer less complexity guarantees, but might be faster in the average case, could be used for this purpose.

We have already mentioned in the related work discussion that a space complexity result for our alternative implementation is currently lacking. The RETE style matching we used is in general far from optimal as far as memory usage is concerned, in particular compared to LEAPS style matching as is used by most CHR

systems. However, in the CHR context, built-in constraints may require maintaining a propagation history which in the worst case requires as much memory as the alpha and beta memories in RETE matching. Therefore, it would be interesting to more formally compare both styles of matching in terms of memory consumption in the context of CHR.

Acknowledgements

The author would like to thank Tom Schrijvers and Bart Demoen for their helpful and insightful comments on a previous version of this paper. This research is funded by a Ph.D. grant of the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

References

- CHRISTIANSEN, H. 2005. CHR grammars. *Theory and Practice of Logic Programming* 5, 4-5, 467–501.
- CLAVEL, M., DURÁN, F., EKER, S., LINCOLN, P., MARTÍ-OLIET, N., MESEGUER, J., AND QUESADA, J. F. 1999. The Maude system. In *10th International Conference on Rewriting Techniques and Applications*, P. Narendran and M. Rusinowitch, Eds. Lecture Notes in Computer Science, vol. 1631. Springer, 240–243.
- DE KONINCK, L. 2007. Mergeable schedules for lazy matching. Tech. Rep. CW 505, Department of Computer Science, K.U.Leuven.
- DE KONINCK, L., SCHRIJVERS, T., AND DEMOEN, B. 2007a. The correspondence between the Logical Algorithms language and CHR. In *23rd International Conference on Logic Programming*, V. Dahl and I. Niemelä, Eds. Lecture Notes in Computer Science, vol. 4670. Springer, 209–223.
- DE KONINCK, L., SCHRIJVERS, T., AND DEMOEN, B. 2007b. User-definable rule priorities for CHR. In *9th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, M. Leuschel and A. Podelski, Eds. ACM Press, 25–36.
- DE KONINCK, L., STUCKEY, P. J., AND DUCK, G. J. 2007. Optimized compilation of CHR^{TP}. Tech. Rep. CW 499, Department of Computer Science, K.U.Leuven. Submitted to FLOPS’08.
- DUCK, G. J., STUCKEY, P. J., AND BRAND, S. 2006. ACD term rewriting. In *22nd International Conference on Logic Programming*, S. Etalle and M. Truszczynski, Eds. Lecture Notes in Computer Science, vol. 4079. Springer, 117–131.
- DUCK, G. J., STUCKEY, P. J., GARCÍA DE LA BANDA, M., AND HOLZBAUR, C. 2004. The refined operational semantics of Constraint Handling Rules. In *20th International Conference on Logic Programming*, B. Demoen and V. Lifschitz, Eds. Lecture Notes in Computer Science, vol. 3132. Springer, 90–104.
- DUCK, G. J., STUCKEY, P. J., AND SULZMANN, M. 2007. Observable confluence for Constraint Handling Rules. In *23rd International Conference on Logic Programming*, V. Dahl and I. Niemelä, Eds. Lecture Notes in Computer Science, vol. 4670. Springer, 224–239.
- FORGY, C. L. 1982. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence* 19, 1, 17–37.
- FREDMAN, M. L. AND TARJAN, R. E. 1987. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM* 34, 3, 596–615.

- FRIEDMAN-HILL, E. 2007. *JESS 7.0p2: The rule engine for the Java platform*. <http://herzberg.ca.sandia.gov/jess>.
- FRÜHWIRTH, T. 1998. Theory and practice of Constraint Handling Rules. *Journal of Logic Programming* 37, 1-3, 95–138.
- FRÜHWIRTH, T. 2000a. Proving termination of constraint solver programs. In *New Trends in Constraints, Joint ERCIM/Compulog Net Workshop, Paphos, Cyprus, October 1999, Selected papers*, K. R. Apt, A. C. Kakas, E. Monfroy, and F. Rossi, Eds. Lecture Notes in Computer Science, vol. 1865. Springer, 298–317.
- FRÜHWIRTH, T. 2000b. Proving termination of constraint solver programs. In *Joint ERCIM/Compulog Net Workshop on New Trends in Constraints*, K. R. Apt, A. C. Kakas, E. Monfroy, and F. Rossi, Eds. Lecture Notes in Computer Science, vol. 1865. Springer, 298–317.
- FRÜHWIRTH, T. 2002a. As time goes by: Automatic complexity analysis of simplification rules. In *8th International Conference on Principles of Knowledge Representation and Reasoning*, D. Fensel, F. Giunchiglia, D. L. McGuinness, and M.-A. Williams, Eds. Morgan Kaufmann, 547–557.
- FRÜHWIRTH, T. 2002b. As time goes by II: More automatic complexity analysis of concurrent rule programs. In *Quantitative Aspects of Programming Languages, Selected Papers*. Electronic Notes in Theoretical Computer Science, vol. 59.
- GANZINGER, H. AND MCALLESTER, D. A. 2001. A new meta-complexity theorem for bottom-up logic programs. In *1st International Joint Conference on Automated Reasoning*, R. Goré, A. Leitsch, and T. Nipkow, Eds. Lecture Notes in Computer Science, vol. 2083. Springer, 514–528.
- GANZINGER, H. AND MCALLESTER, D. A. 2002. Logical algorithms. In *18th International Conference on Logic Programming*, P. J. Stuckey, Ed. Lecture Notes in Computer Science, vol. 2401. Springer, 209–223.
- HOLZBAUR, C. AND FRÜHWIRTH, T. 1998. Constraint Handling Rules reference manual, release 2.2. Tech. Rep. TR-98-01, Österreichisches Forschungsinstitut für Artificial Intelligence, Wien.
- MCALLESTER, D. A. 1999. On the complexity analysis of static analyses. In *6th International Symposium on Static Analysis*, A. Cortesi and G. Filé, Eds. Lecture Notes in Computer Science, vol. 1694. Springer, 312–329.
- MIRANKER, D. P., BRANT, D. A., LOFASO, B., AND GADBOIS, D. 1990. On the performance of lazy matching in production systems. In *8th National Conference on Artificial Intelligence*. AAAI Press / The MIT Press, 685–692.
- PILOZZI, P., SCHRIJVERS, T., AND DE SCHREYE, D. 2007. Proving termination of CHR in Prolog: A transformational approach. In *9th International Workshop on Termination*, D. Hofbauer and A. Serebrenik, Eds. 30–33.
- PROCTOR, M., NEALE, M., FRANDSEN, M., GRIFFITH, JR., S., TIRELLI, E., MEYER, F., AND VERLAENEN, K. 2007. *Drools Documentation, Version 4.0.3*. <http://www.jboss.com/products/rules>.
- SCHRIJVERS, T. 2005. Analyses, optimizations and extensions of Constraint Handling Rules. Ph.D. thesis, K.U.Leuven, Leuven, Belgium.
- SCHRIJVERS, T. AND DEMOEN, B. 2004. The K.U.Leuven CHR system: implementation and application. In *First Workshop on Constraint Handling Rules: Selected Contributions*, T. Frühwirth and M. Meister, Eds. Ulmer Informatik-Berichte, vol. 2004-01. Universität Ulm, 1–5.
- SCHRIJVERS, T. AND FRÜHWIRTH, T. 2006. Optimal union-find in Constraint Handling Rules. *Theory and Practice of Logic Programming* 6, 1&2.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2005. The computational power and

- complexity of Constraint Handling Rules. In *2nd Workshop on Constraint Handling Rules*, T. Schrijvers and T. Frühwirth, Eds. Reports CW, vol. 421. Departement of Computer Science, K.U.Leuven, Belgium, 3–17.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006a. Dijkstra’s algorithm with Fibonacci heaps: An executable description in CHR. In *20th Workshop on Logic Programming*, M. Fink, H. Tompits, and S. Woltran, Eds. INFSYS Research Report, vol. 1843-06-02. TU Wien, 182–191.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006b. Memory reuse for CHR. In *22nd International Conference on Logic Programming*, S. Etalle and M. Truszczynski, Eds. Lecture Notes in Computer Science, vol. 4079. Springer, 72–86.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2008. The computational power and complexity of Constraint Handling Rules. <http://www.cs.kuleuven.be/~jon/> – Submitted to ACM TOPLAS.
- TACCHELLA, P., MEO, M. C., AND GABBRIELLI, M. 2007. Unfolding in CHR. In *9th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, M. Leuschel and A. Podelski, Eds. ACM Press, 179–186.
- VAN WEERT, P., SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006. Extending CHR with negation as absence. In *3rd Workshop on Constraint Handling Rules*, T. Schrijvers and T. Frühwirth, Eds. Reports CW, vol. 452. Department of Computer Science, K.U.Leuven, Belgium, 125–140.
- VOETS, D., PILOZZI, P., AND DE SCHREYE, D. 2007. A new approach to termination analysis of Constraint Handling Rules. In *4th Workshop on Constraint Handling Rules*, K. Djelloul, G. J. Duck, and M. Sulzmann, Eds. U.Porto, 77–89.