The Computational Power and Complexity of Constraint Handling Rules

JON SNEYERS, TOM SCHRIJVERS, and BART DEMOEN
K.U. Leuven

Constraint Handling Rules (CHR) comprise a high-level rule-based programming language which is increasingly used for general-purpose programming. We introduce the CHR machine, a model of computation based on the operational semantics of CHR. Its computational power and time complexity properties are compared to those of the well-understood Turing machine and Random Access Memory machine. This allows us to prove the interesting result that every algorithm can be implemented in CHR with the best known time and space complexity. We also investigate the practical relevance of this result and the constant factors involved. Finally we expand the scope of the discussion to other (declarative) programming languages.

Categories and Subject Descriptors: D.3.2 [**Programming Languages**]: Language Classifications—*Constraint and logic languages*; F.1.1 [**Computation by Abstract Devices**]: Models of Computation; F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; D.3.4 [**Programming Languages**]: Processors—*Compilers, optimization*

General Terms: Languages, Algorithms, Theory, Performance, Experimentation

Additional Key Words and Phrases: Constraint Handling Rules, complexity, constant factors

ACM Reference Format:

Sneyers, J., Schrijvers, T., and Demoen, B. 2009. The computational power and complexity of constraint handling rules. *ACM Trans. Program. Lang. Syst.* 31, 2, Article 8 (January 2009), 42 pages. DOI = 10.1145/1462166.1462169 <http://doi.acm.org/10.1145/1462166.1462169>

A preliminary version of this article was presented at the *2nd Workshop on Constraint Handling Rules (CHR'05)*.

J. Sneyers is funded by a Ph.D. grant of the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen). This work is partially based on work that was supported in part by the Research Foundation—Flanders (FWO-Vlaanderen) through projects G.0144.03 and G.0160.02. T. Schrijvers is a post-doctoral researcher of the Research Foundation—Flanders (FWO-Vlaanderen).

Authors' address: J. Sneyers, T. Schrijvers, and B. Demoen, Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200A, B-3001 Heverlee, Belgium; email: {jon,toms,bmd}@cs.kuleuven.be.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2009 ACM 0164-0925/2009/01-ART8 \$5.00 DOI 10.1145/1462166.1462169 <http://doi.acm.org/10.1145/1462166.1462169>

ACM Transactions on Programming Languages and Systems, Vol. 31, No. 2, Article 8, Pub. date: January 2009.

8:2 • J. Sneyers et al.

1. INTRODUCTION

Constraint Handling Rules (CHR) comprise a high-level programming language based on multiheaded committed-choice rules, introduced by Frühwirth [2009, 1998, 1992]. Originally designed for writing constraint solvers, CHR is now used as a general-purpose programming language. When new programming languages are introduced, sooner or later the question arises as to whether classical algorithms can be implemented in an efficient and elegant way. Schrijvers and Frühwirth [2006] and Sneyers et al. [2006a] have discussed case studies: They describe efficient and elegant CHR implementations of, respectively, Tarjan's union-find algorithm and Dijkstra's shortest path algorithm, indicating an affirmative answer. In contrast, it is not clear whether these algorithms can be implemented with optimal complexity in pure Prolog.

This article tackles the preceding question in a more general way. The major contribution of this article is the proof of a general affirmative answer: We demonstrate that it is possible to implement *any* algorithm in CHR in an efficient way, that is, with the best known time and space complexity. Our approach is as follows. We introduce a new model of computation, the CHR machine, and compare it with the well-known Turing machine and RAM machine models. We show how to program the CHR machine to efficiently simulate both other machines, in particular the RAM machine, which is faster than the Turing machine. We then investigate how to simulate a CHR machine efficiently on a RAM machine using an optimizing CHR compiler. This results in a general complexity metatheorem. Finally we apply this metatheorem to the RAM machine simulator. Thus we conclude that existing CHR compilation techniques suffice to implement the RAM machine simulator efficiently. As a result, *all* RAM machine programs (so also every known algorithm) can be translated to CHR programs with the same complexity. Figure 1 gives an overview of the relation between the three models of computation. Note that the $\tilde{O}((TS^{m+1})^4)$ bound for simulating a CHR machine on a Turing machine can be obtained by simulating the CHR machine on a RAM machine, and then simulating that RAM machine on a Turing machine.

Notational Conventions. Prolog list notation is used to denote sequences: $[h|T] = [h] ++ T$, where $++$ denotes sequence concatenation. Σ^* denotes the set of all sequences of elements of Σ . The notation $vars(t)$ denotes the set of variables used in term t . The abbreviation $\exists_{\{v_1, \dots, v_n\}}$ means $\exists v_1 \dots \exists v_n$. The notation $\exists_V F$ is used to denote $\exists_{vars(F) \setminus V}$, that is, quantifying over all variables not in V . Finally, $\tilde{O}(T)$ denotes $\mathcal{O}(TL)$, where L is polylogarithmic in the variables of T . For example, a complexity of $\mathcal{O}(n^2 \log^2 n \log \log n)$ could be denoted by $\tilde{O}(n^2)$. Note that a function which is $\tilde{O}(n^k)$ is also $\mathcal{O}(n^{k+\epsilon})$ for arbitrarily small $\epsilon > 0$.

Overview. In the next section we briefly recapture syntax and semantics of CHR. Section 3 defines the three models of computation: Turing machines, RAM machines, and CHR machines. In Section 4, we investigate the theoretical relation between CHR machines and the other computational models. In Section 5 we discuss compilation of CHR programs, linking the theoretical CHR machines and practical CHR systems. This concludes the proof of our main result. Section 6 investigates, experimentally, the practical relevance of the result and the

Fig. 1. Overview. An arrow between A and B indicates that A can be simulated on B . Labels indicate the relevant section and the time complexity of simulating a T -time, S -space A on B .

constant factors of the complexities. In Section 7 we try to “port” the results to other declarative languages. Section 8 concludes this work.

2. CONSTRAINT HANDLING RULES

In this section we introduce the syntax (Section 2.1) and semantics (Section 2.2) of Constraint Handling Rules. Gentler introductions to CHR can be found in Frühwirth [2009, 1998], Schrijvers [2005], and Sneyers et al. [2009]. Readers already familiar with CHR and its theoretical operational semantics may skip to Section 2.3, in which some new terminology is introduced.

2.1 Syntax

We assume the reader to be familiar with some basic notions of constraint logic programming [Marriott and Stuckey 1998; Frühwirth and Abdennadher 2003].

CHR is embedded in a host language $\mathcal{H}$ that provides data types and a number of predefined constraints. These constraints are called host language constraints or *built-in* constraints. The typical host language of CHR is Prolog. Its host language constraints are unification, Prolog built-ins, and other (user-defined) Prolog predicates. Its data types are Prolog variables and terms.

CHR constraint symbols are drawn from the set of predicate symbols, denoted by a functor/arity pair. CHR constraints, also called constraint atoms or constraints for short, are atoms constructed from these symbols and the data types offered by the host language.

We assume the host language to offer at least one data type that can be used as an identifier; for instance, Prolog variables (without unification). We assume the built-in constraint theory (denoted by $\mathcal{D}_{\mathcal{H}}$) to define at least the basic constraints *true* and *fail* and the *ask*-versions of syntactic equality (“==”) and inequality (“≠”).

We denote the set of multisets of (CHR and host language) constraints for a given CHR program $\mathcal{P}$ and host language $\mathcal{H}$ by the symbol $\mathcal{G}_{\mathcal{P}}^{\mathcal{H}}$.

8:4 • J. Sneyers et al.

```

loop   @ upto(N)  ⇔ N > 1 | prime(N), N1 is N - 1, upto(N1).
stop   @ upto(1)  ⇔ true.
absorb @ prime(A) \ prime(B) ⇔ B mod A = 0 | true.

```

Fig. 2. The CHR (Prolog) program PRIMES, a prime number sieve.

A CHR program $\mathcal{P}$ consists of an ordered sequence of CHR rules. There are three different kinds of rules. A *simplification* rule is of the form

$$c_1, \dots, c_n \iff g \mid d_1, \dots, d_m.$$

A *propagation* rule is of the form

$$c_1, \dots, c_n \implies g \mid d_1, \dots, d_m.$$

A *simpagation* rule is of the form

$$c_1, \dots, c_l \setminus c_{l+1}, \dots, c_n \iff g \mid d_1, \dots, d_m.$$

Here $l, m, n > 0$. The sequence, or conjunction, $c_1, \dots, c_n$ are CHR constraint atoms; together they are called the *head* or *head constraints* of the rule. A rule with n head constraints is named an *n-headed* rule and when $n > 1$, it is a *multiheaded* rule. A rule is optionally preceded by *name* @ where *name* is a term. No two rules may have the same name. Rules without a name get an unique implicit name.

All the head constraints of a simplification rule and the head constraints $c_{l+1}, \dots, c_n$ of a simpagation rule are called *removed* head constraints. The other head constraints (all heads of a propagation rule and $c_1, \dots, c_l$ of a simpagation rule) are called *kept* head constraints.

An *occurrence number* is associated with every head constraint. Head constraints are numbered per functor/arity pair, starting from 1, from the first rule to the last rule, removed heads before kept heads, from left to right.

The conjunction $d_1, \dots, d_n$ consists of CHR constraints and host language constraints; it is called the *body* of the rule. The part of the rule between the arrow and the body, g , is called the *guard*. It is a conjunction of host language ask constraints. The guard “ $g \mid$ ” is optional; if omitted, it is considered to be “*true* $\mid$ ”.

For simplicity, we sometimes consider both simplification and propagation rules as special cases of a simpagation rule. We often use the following notation to denote any kind of rule.

$$H_k \setminus H_r \iff g \mid B$$

If H_k is empty, then the rule is a *simplification* rule. If H_r is empty, then the rule is a *propagation* rule. At least one of H_r and H_k must be nonempty.

Example. Figure 2 lists a simple CHR (Prolog) program called PRIMES, a CHR variant of the Sieve of Eratosthenes. Dating back to at least 1992 [Frühwirth 1992], this is one of the very first examples where CHR is used as a general-purpose programming language. Given a query of the form “**upto**(n)”, where n is a positive integer, it computes all prime numbers up to n . The first rule (*loop*) does the following: If $n > 1$, it simplifies **upto**(n) to **upto**($n - 1$)

and adds a **prime**(n) constraint. The second rule handles the case for $n = 1$, removing any **upto**(1) constraint. Note that removing a constraint is done by simplifying it to the built-in constraint *true*. The third and most interesting rule (*absorb*) is a simpagation rule. If there are two **prime**/1 constraints **prime**(A) and **prime**(B), such that B is a multiple of A, the latter constraint is removed. The effect of the *absorb* rule is that all nonprimes are eventually removed. As a result, if the rules are applied exhaustively, the remaining constraints correspond exactly to the prime numbers up to n .

2.2 The ω_t Operational Semantics

In this section we present the operational semantics ω_t of CHR, sometimes also called *theoretical* or *high-level* operational semantics.

The ω_t semantics is formulated as a state transition system. Transition rules define the relation between an execution state and its subsequent execution state.

Definition 2.1. An execution state σ is a tuple $\langle \mathbb{G}, \mathbb{S}, \mathbb{B}, \mathbb{T} \rangle_n$. The goal $\mathbb{G} \in \mathcal{G}_P^{\mathcal{H}}$ is a multiset of constraints to be rewritten to solved form. The CHR constraint store $\mathbb{S}$ is a set of identified CHR constraints that can be matched with rules in the program $\mathcal{P}$. An *identified* CHR constraint $c\#i$ is a CHR constraint c associated with some unique integer i , the *constraint identifier*. This number serves to differentiate among copies of the same constraint. We introduce the functions $chr(c\#i) = c$ and $id(c\#i) = i$, and extend them to sequences and sets of identified CHR constraints in the obvious manner (e.g. $id(\mathbb{S}) = \{c\#i \in \mathbb{S}\}$). Note that $chr(\mathbb{S})$ is a multiset although $\mathbb{S}$ is a set. The *built-in constraint store* $\mathbb{B}$ is the conjunction of all built-in constraints that have been passed to the underlying solver. This abstracts the internal representation used by the host language. Its actual meaning depends on the host language $\mathcal{H}$. The *propagation history* $\mathbb{T}$ is a set of tuples, each recording the identities of the CHR constraints that fired a rule, and the name of the rule itself. This is necessary to prevent trivial nontermination for propagation rules: A propagation rule is allowed to fire on a set of constraints only if the constraints have not been used to fire the same rule before. Finally, the counter $n \in \mathbb{N}$ represents the next free integer that can be used to number a CHR constraint. We use $\sigma, \sigma_0, \sigma_1, \dots$ to denote execution states and Σ^{CHR} to denote the set of all execution states.

Transitions are defined by the binary relation $\rightarrow_{\mathcal{P}} : \Sigma^{\text{CHR}} \rightarrow \Sigma^{\text{CHR}}$ shown in Figure 3. Execution proceeds by exhaustively applying the transition rules, starting from an initial state. We define $\rightarrow_{\mathcal{P}}^*$ as the transitive closure of $\rightarrow_{\mathcal{P}}$.

2.3 Derivations

Definition 2.2. Given an *initial goal* (or *query*) $\mathbb{G} \in \mathcal{G}_P^{\mathcal{H}}$, the *initial state* is $initstate(\mathbb{G}) = \langle \mathbb{G}, \emptyset, true, \emptyset \rangle_1$. The set of initial states is denoted by $\Sigma^{\text{init}} \subset \Sigma^{\text{CHR}}$.

Definition 2.3. A *final state* $\sigma_f = \langle \mathbb{G}, \mathbb{S}, \mathbb{B}, \mathbb{T} \rangle_n$ is an execution state for which no transition applies: $\neg \exists \sigma \in \Sigma^{\text{CHR}} : \sigma_f \rightarrow_{\mathcal{P}} \sigma$. In a *failure* state, the underlying solver $\mathcal{H}$ can prove $\mathcal{D}_{\mathcal{H}} \models \neg \exists_{\emptyset} \mathbb{B}$ and such states are always final. A *successful*

8:6 • J. Sneyers et al.

1. **Solve.** $\langle \{c\} \uplus G, S, B, T \rangle_n \mapsto_P \langle G, S, c \wedge B, T \rangle_n$
where c is a built-in constraint and $\mathcal{D}_{\mathcal{H}} \models \exists_{\emptyset} B$.
2. **Introduce.** $\langle \{c\} \uplus G, S, B, T \rangle_n \mapsto_P \langle G, \{c \# n\} \cup S, B, T \rangle_{n+1}$
where c is a CHR constraint and $\mathcal{D}_{\mathcal{H}} \models \exists_{\emptyset} B$.
3. **Apply.** $\langle G, H_1 \cup H_2 \cup S, B, T \rangle_n \mapsto_P \langle C \uplus G, H_1 \cup S, \theta \wedge B, T \cup \{h\} \rangle_n$
where $\mathcal{P}$ contains a (renamed apart) rule of the form $r @ H'_1 \setminus H'_2 \iff g \mid C$
and θ is a matching substitution such that $\text{chr}(H_1) = \theta(H'_1)$ and $\text{chr}(H_2) = \theta(H'_2)$
and $h = (r, \text{id}(H_1), \text{id}(H_2)) \notin T$ and $\mathcal{D}_{\mathcal{H}} \models (\exists_{\emptyset} B) \wedge (B \rightarrow \exists_B(\theta \wedge g))$.

Fig. 3. The transition rules of the theoretical operational semantics ω_t , defining $\mapsto_P$.

final state is a final state that is not a failure state (i.e. $\mathcal{D}_{\mathcal{H}} \models \exists_{\emptyset} B$). The set of final states is denoted by $\Sigma^{\text{final}} \subset \Sigma^{\text{CHR}}$.

Definition 2.4. Given a CHR program $\mathcal{P}$, a *finite derivation* d is a finite sequence $[\sigma_0, \sigma_1, \dots, \sigma_n]$ of states where $\sigma_0 \in \Sigma^{\text{init}}$, $\sigma_n \in \Sigma^{\text{final}}$, and $\sigma_i \mapsto_P \sigma_{i+1}$ for $0 \leq i < n$. If σ_n is a failure state, we say d has *failed*, otherwise d is a *successful* derivation.

Definition 2.5. An *infinite derivation* d_{∞} is an infinite sequence $\sigma_0, \sigma_1, \dots$ of states where $\sigma_0 \in \Sigma^{\text{init}}$ and $\sigma_i \mapsto_P \sigma_{i+1}$ for $i \in \mathbb{N}$.

We use $\#d$ to denote the length of a derivation: The length of a finite derivation is the number of transitions in the sequence; the length of an infinite derivation is ∞ . A set of (finite or infinite) derivations is denoted by Δ . The set of all derivations in Δ that start with $\text{initstate}(\mathbb{G})$ is denoted by $\Delta|_{\mathbb{G}}$. We use $\Delta_{\omega_t}^{\mathcal{H}}(\mathcal{P})$ to denote the set of all derivations (in the ω_t semantics) for a given CHR program $\mathcal{P}$ and host language $\mathcal{H}$. We now define the relation $\rightsquigarrow_{\Delta} : \Sigma^{\text{init}} \rightarrow \Sigma^{\text{CHR}} \cup \{\infty\}$.

Definition 2.6. State σ_n is a Δ -output of σ_0 if $[\sigma_0, \dots, \sigma_n] \in \Delta$. We say σ_0 Δ -outputs σ_n and write $\sigma_0 \rightsquigarrow_{\Delta} \sigma_n$. If Δ contains an infinite derivation starting with σ_0 , we say σ_0 has a non terminating derivation. We denote this as $\sigma_0 \rightsquigarrow_{\Delta} \infty$.

Definition 2.7. The CHR program $\mathcal{P}$ is Δ -deterministic for input $I \subseteq \mathcal{G}_{\mathcal{P}}^{\mathcal{H}}$ if the restriction of $\rightsquigarrow_{\Delta}$ to $\text{initstate}[I]$ is a function (so for any input in I , all derivations result in the same output) and $\forall i \in I, d \in \Delta|_i$: If d is a successful derivation, then $\forall d' \in \Delta|_i : \#d = \#d'$.

In other words, a program is Δ -deterministic if all derivations starting from a given input have the same result and all successful ones have the same length. Note that if a program $\mathcal{P}$ is $\Delta_{\omega_t}^{\mathcal{H}}(\mathcal{P})$ -deterministic for all input $\mathcal{G}_{\mathcal{P}}^{\mathcal{H}}$, it is also observable confluent [Duck et al. 2006]. The notion of observable confluence does not require derivations to have the same length.

Example. Consider the CHR program of the previous example. This program (PRIMES, listed in Figure 2) is $\Delta_{\omega_t}^{\text{Prolog}}$ -deterministic for input $\{\mathbf{upto}(n) \mid n \in \mathbb{N}\}$. Although the order in which the transitions of ω_t are applied is not fixed for a given input, allowing different derivations, the derivation length and result are always the same.

Of course not every CHR program is $\Delta_{\omega_t}^{\mathcal{H}}$ -deterministic for its intended input.

$$\begin{array}{ll}
\text{reflexivity} & @ \text{leq}(X, X) \iff \text{true}. \\
\text{antisymmetry} & @ \text{leq}(X, Y), \text{leq}(Y, X) \iff X = Y. \\
\text{idempotence} & @ \text{leq}(X, Y) \setminus \text{leq}(X, Y) \iff \text{true}. \\
\text{transitivity} & @ \text{leq}(X, Y), \text{leq}(Y, Z) \implies \text{leq}(X, Z).
\end{array}$$

Fig. 4. The CHR (Prolog) program **LEQ**, a solver for the less-than-or-equal relation.

Example. The CHR program **LEQ** listed in Figure 4 is not $\Delta_{\omega_t}^{\text{Prolog}}$ -deterministic for all conjunctions of **leq**/2 constraints. Consider the input query “**leq**(A,B), **leq**(B,A)”. One derivation consists of two **Introduce** steps followed by an **Apply** step using the antisymmetry rule, followed by a **Solve** step for “A = B” and two **Apply** steps using the reflexivity rule. This derivation has length six. In another derivation, the transitivity rule is applied after the **Introduce** steps. This results in a longer derivation, or even an infinite derivation.

For the connoisseurs: **LEQ** is not even $\Delta_{\omega_r}^{\text{Prolog}}$ -deterministic (i.e., restricting the derivations to those of the refined operational semantics ω_r , refer to Section 5.2.1). Given the query “**leq**(B,C), **leq**(B,A), **leq**(A,B)”, the derivation length depends on which partner for **leq**(A,B) is tried first when the transitivity rule is applied.

3. MODELS OF COMPUTATION

Both computability theory and computational complexity theory define idealized *models of computation*. In this section we define three different models: the well-known Turing machine (Section 3.1); the RAM machine (Section 3.2), which more closely models realistic computers; and finally the CHR machine (Section 3.3), which we will use to study the CHR programming language.

3.1 Turing Machines

The Turing machine, originally introduced by Alan Turing [1936], is the prototypical computational model used in computability and complexity theory. We will use a single-tape definition which corresponds to the one given by Hopcroft et al. [2001].

Definition 3.1. A Turing machine is a 6-tuple $M = \langle Q, \Sigma, q_0, b, F, \delta \rangle$ where:

- Q is a finite set of *states*;
- Σ is a finite set of symbols, the tape *alphabet*;
- $q_0 \in Q$ is the initial state;
- $b \in \Sigma$ is the *blank* symbol;
- $F \subseteq Q$ is the set of *accepting* final states; and
- $\delta : Q \times \Sigma \mapsto Q \times \Sigma \times \{\text{left}, \text{right}\}$. is a partial function.

The function δ is called the *transition function*, “left” represents left shift and “right” represents right shift.

A Turing machine M operates on an infinite tape of cells. Each cell contains a symbol of the tape alphabet Σ . The tape is assumed to be arbitrarily extendible to the left and the right. A head is positioned on a particular cell of the tape, can read and write a symbol in that cell, and can move left and right.

8:8 • J. Sneyers et al.

Operation starts in the initial state q_0 on a tape which contains a finite string of symbols (called the *input*), and the head is positioned on the leftmost input symbol. Execution proceeds by considering the current state q and the symbol s that is under the head. Then, one of the following holds.

- Either (q, s) is a member of the domain of δ and $\delta(q, s) = (q', s', X)$. The effect then is that the current state of M changes to q' , the head overwrites the value s in the cell under it with s' , and next the head either moves to the left or the right depending on whether $X = \text{left}$ or $X = \text{right}$.
- Or (q, s) is not part of the domain of δ . Execution stops. If $q \in F$, the input is *accepted*, otherwise it is *rejected*.

Formally, we represent an execution state of a Turing machine $M = \langle Q, \Sigma, q_0, b, F, \delta \rangle$ as a 4-tuple $\sigma^{\mathbb{M}} = \langle q, P, s, N \rangle$, representing the *current state* $q \in Q$, the *current symbol* $s \in \Sigma$, the *previous symbols* P , and the *next symbols* N , which are sequences of symbols. We use $\sigma^{\mathbb{M}_0}, \sigma^{\mathbb{M}_1}, \dots$ to denote Turing machine execution states, and $\Sigma^{\mathbb{M}}(M)$ to denote the set of all execution states of M . Given an input $I = [f|N]$ (a sequence of symbols), the initial execution state is $\sigma^{\mathbb{M}_i} = \langle q_0, [], f, N \rangle$. The transition rules are as follows.

- (1) $\langle q, [p|P], s, N \rangle \rightarrow \langle q', P, p, [s'|N] \rangle$ if $\delta(q, s) = (q', s', \text{left})$
- (2) $\langle q, P, s, [n|N] \rangle \rightarrow \langle q', [s'|P], n, N \rangle$ if $\delta(q, s) = (q', s', \text{right})$
- (3) $\langle q, [], s, N \rangle \rightarrow \langle q', [], b, [s'|N] \rangle$ if $\delta(q, s) = (q', s', \text{left})$
- (4) $\langle q, P, s, [] \rangle \rightarrow \langle q', [s'|P], b, [] \rangle$ if $\delta(q, s) = (q', s', \text{right})$

The machine operates by exhaustively applying the transition rules on the initial execution state. Note that the size of execution states is unbounded but finite. In execution state $\sigma^{\mathbb{M}_i} = \langle q, P, s, N \rangle$, the contents of the tape consists of the sequence $\text{tape}(\sigma^{\mathbb{M}_i}) = \text{reverse}(P) ++ [s] ++ N$, where *reverse* reverses a sequence (i.e., $\text{reverse}([a_1, a_2, \dots, a_{n-1}, a_n]) = [a_n, a_{n-1}, \dots, a_2, a_1]$). If the final execution state is of the form $\langle q_f, P, s, N \rangle$, the Turing machine has *accepted* the input I if $q_f \in F$. If $q_f \in Q \setminus F$, the input is *rejected*. The *output* of the Turing machine is the tape contents in the final execution state if the input is accepted, and undefined if the input is rejected. Given a Turing machine M and input I , we denote the corresponding derivation by $\text{deriv}_M(I)$. If the machine terminates on input I , we denote the output by $M(I) = \text{tape}(\sigma^{\mathbb{M}_f})$, where $\sigma^{\mathbb{M}_f}$ is the last state in $\text{deriv}_M(I)$.

3.2 Random Access Memory Machines

The Random Access Memory (RAM) machine closely models the basic features of traditional sequential computers. In the literature many variations of the RAM have been considered. We investigate two different RAM machines. The first is a RAM machine with simple Peano arithmetic operations and the second has the standard arithmetic operations as they are implemented on today's computers.

Common architecture. A RAM machine consists of three components: the Central Processing Unit (CPU), the program, and the Random Access Memory (RAM). The memory consists of an infinite number of cells, or registers, each labeled with a natural number which is called its *address*. If a register is *initialized*, it contains a *value*, which is an integer number. We use $A, A_1, A_2, \dots$ to represent the memory addresses and $\llbracket A \rrbracket$ to denote the value of the register at address A . By convention, particular memory cells are initialized in advance and contain the input, while other (or the same) memory cells are meant to contain the output. Without loss of generality we assume that the program is written such that all additional registers are initialized (using the **init** instruction) before they are used.¹

The program consists of a sequence of instructions. The program instructions are labeled with successive natural numbers. We use $L, L_1, \dots$ to denote program instruction labels. The CPU follows a fetch-and-execute cycle. It has a program counter PC that is initialized to the first program instruction label. This program counter contains the label of the next program instruction to be executed. The CPU fetches the instruction and performs the corresponding operations.² This involves setting the program counter to the next instruction, by default the successor of the current address. Table I lists the instructions supported by the standard RAM machine. The Peano-arithmetic RAM machine uses a subset of these instructions.

Definition 3.2. A *Peano-arithmetic RAM machine* consists of a program and a working memory as described previously. The program instructions are **inc**, **dec**, **clr**, **jmp**, **cjmp**, and **halt** (see Table I).

This corresponds to the definition used by Savage [1998]. Indirect addressing is not supported, so all registers that are used in a program are supposed to be initialized in advance. All copying, addition, and subtraction has to be done by repeated use of the **inc** and **dec** instructions. This makes the Peano-arithmetic RAM less practical, as actual computers do provide instructions for addition and subtraction. However, pure CHR without built-in constraints does not provide any arithmetic functionality, either. Hence, arithmetic must be encoded, for example, using Peano arithmetic.

Definition 3.3. A *standard RAM machine* (or standard-arithmetic RAM) consists of a program and a working memory as described earlier. The program instruction set is given in Table I.

¹A program that does not initialize registers before it uses them can be rewritten in the following way. We may assume all registers that are used directly in operands to be initialized in advance (these can be easily found by inspecting the program). We let the program instructions that use indirect addressing (**imv** and **mvi**) be preceded by an **init** instruction. For example, the instruction "**mvi** a b " should be rewritten to "**init** b ; **mvi** a b ". This transformation at most doubles the program size and the number of instructions executed.

²If an illegal instruction is encountered, the machine halts. Examples of illegal instructions are division by zero, jump to a nonexistent label, instructions referring to registers with a negative address, etc.

8:10 • J. Sneyers et al.

Table I. Instruction Set of the RAM

Instruction			Effect
inc	A		$\llbracket A \rrbracket \leftarrow \llbracket A \rrbracket + 1$
dec	A		$\llbracket A \rrbracket \leftarrow \llbracket A \rrbracket - 1$
clr	A		$\llbracket A \rrbracket \leftarrow 0$
jmp	L		$PC \leftarrow L$
cjmp	A	L	$PC \leftarrow L$ if $\llbracket A \rrbracket = 0$
halt			Halt execution of the RAM.
init	A		Initialize the register at address $\llbracket A \rrbracket$ to zero.
cnst	B	A	$\llbracket A \rrbracket \leftarrow B$
add	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket A_1 \rrbracket + \llbracket A_2 \rrbracket$
sub	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket A_1 \rrbracket - \llbracket A_2 \rrbracket$
mul	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket A_1 \rrbracket * \llbracket A_2 \rrbracket$
div	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket A_1 \rrbracket / \llbracket A_2 \rrbracket$ if $\llbracket A_2 \rrbracket \neq 0$
mov	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket A_2 \rrbracket$
imv	A_2	A_1	$\llbracket A_1 \rrbracket \leftarrow \llbracket \llbracket A_2 \rrbracket \rrbracket$
mvi	A_2	A_1	$\llbracket \llbracket A_1 \rrbracket \rrbracket \leftarrow \llbracket A_2 \rrbracket$

This instruction set is similar to the definition of Aho et al. [1975], and resembles more closely actual computers. The **inc**, **dec**, and **clr** instructions are redundant: They can be implemented using **add**, **sub**, and **cnst**. Without loss of generality, we assume the instructions that refer to two (**add**, **sub**, **mul**, **div**, **mov**) or three (**imv**, **mvi**) memory cells refer to two or three *different* cells. For example, the instruction “**add a a**” should be rewritten to “**mov a t ; add t a**” where t is some temporary register.

3.3 CHR Machines

We introduce a model of computation to investigate the computability and complexity properties of Constraint Handling Rules, which we call the CHR machine.

Definition 3.4. A *CHR machine* is a tuple $\mathcal{M} = (\mathcal{H}, \mathcal{P}, \mathcal{V})$. The host language $\mathcal{H}$ defines a built-in constraint theory $\mathcal{D}_{\mathcal{H}}$, $\mathcal{P}$ is a CHR program, and $\mathcal{V} \subseteq \mathcal{G}_{\mathcal{P}}^{\mathcal{H}}$ is a set of *valid goals*, such that $\mathcal{P}$ is a $\Delta_{\omega_t}^{\mathcal{H}}$ -deterministic CHR program for input $\mathcal{V}$. The machine takes an input query $\mathbb{G} \in \mathcal{V}$ and executes a derivation $d \in \Delta_{\omega_t}^{\mathcal{H}} |_{\mathbb{G}}$.

Terminology. If the derivation d for $\mathbb{G}$ is finite, we say the machine *terminates* with *output state* $\mathcal{M}(\mathbb{G}) = \langle \mathbb{G}', \mathbb{S}, \mathbb{B}, \mathbb{T} \rangle_n$, which is the last state of d . We will use the following notation: $\mathcal{M}_c(\mathbb{G}) = \mathbb{S}$ is the set of *output constraints*, $\mathcal{M}_b(\mathbb{G}) = \mathbb{B}$ is the *output built-in store*, and $\text{deriv}_{\mathcal{M}}(\mathbb{G}) = d$ is the *derivation* for $\mathbb{G}$. The machine *accepts* the input $\mathbb{G}$ if d is a successful derivation and *rejects* $\mathbb{G}$ if d is a failed derivation ($\mathcal{M}_b(\mathbb{G}) = \text{fail}$). If d is an infinite derivation, we say the machine *does not terminate*. A *CHR(X) machine* is a CHR machine for which the host language $\mathcal{H} = X$. We use Φ to denote *no host language*: The built-in constraint theory $\mathcal{D}_{\Phi}$ defines only the basic constraints *true* and *fail*, and syntactic equality and inequality (only to be used as an *ask*-constraint). This implies that the **Solve** transition can only be used once (to add *fail*). The only data types are Prolog-like variables (that cannot be bound) and (zero-arity) atoms. A *CHR-only machine* is a $\text{CHR}(\Phi)$ machine.

Definition 3.5. A *sufficiently strong* host language $\mathcal{H}$ is a host language whose built-in constraint theory $\mathcal{D}_{\mathcal{H}}$ defines at least *true*, *fail*, *==* and *!=*, the integer numbers and the arithmetic operations for addition, subtraction, multiplication, and integer division.

Clearly, most host languages are sufficiently strong. Prolog, for instance, allows arithmetic using the one-way built-in *is/2*. In CHR program listings where the host language is assumed sufficiently strong, we will use a slightly abbreviated notation. For example, if *c/1* is a CHR constraint, we write expressions like “*c(N+1)*”: a host-language-independent notation that is equivalent to “*M is N+1, c(M)*” for CHR (Prolog), to “*c(intUtil.add(N,1))*” for CHR(Java), etc.

4. COMPUTATIONAL POWER AND COMPLEXITY OF CHR MACHINES

A model of computation is called *Turing-complete* if it has the same computational power as Turing machines: Every Turing Machine can be simulated in the model and every program of the model can be simulated on a Turing machine.

First (Section 4.1) we show that the CHR-only machine has at least the computational power of the Turing machine.³ In order to prove this, it suffices to construct a CHR machine which corresponds to a universal Turing machine, namely, a CHR machine which can simulate any Turing machine. The second requirement for Turing completeness (a Turing machine which can simulate any CHR machine) is discussed in Section 5. In Section 4.2 we define time and space complexity functions for all three models of computation. Finally (Section 4.3) we investigate the complexity of CHR machines that simulate RAM machines.

4.1 Computational Power of CHR Machines

Consider the CHR program TMSIM shown in Figure 5 and the corresponding CHR-only machine $\mathcal{M}_{\mathbb{T}} = (\Phi, \text{TMSIM}, \mathcal{V}_{\mathbb{T}})$ (we postpone the definition for $\mathcal{V}_{\mathbb{T}}$ for a while). The program TMSIM simulates Turing machines.

Intuitively, the meanings of the constraints in the TMSIM program are as follows.

- delta/5** encodes the transition function δ (the Turing machine program) in the obvious way: *delta*(*q,s,q',s',d*) means that $\delta(q,s) = (q',s',d)$;
- nodelta/2** encodes the domain on which δ is undefined;
- rejecting/1** encodes the set of nonaccepting final states $Q \setminus F$;
- state/1** contains the current state;
- head/1** contains the identifier of the cell under the head;
- cell/2** represents a tape cell (the first argument is the unique identifier of the cell; the second argument is the symbol in the cell); and

³See also Sneyers [2008], in which (syntactical) subclasses of CHR are identified that are still Turing-complete.

8:12 • J. Sneyers et al.

$$\begin{aligned}
r_1 @ \text{delta}(Q_i, S_i, Q_o, S_o, \text{left}), \text{adj}(\text{LC}, C) & \setminus \text{state}(Q_i), \text{cell}(C, S_i), \text{head}(C) \\
\iff \text{LC} \setminus == \text{null} & \mid \text{state}(Q_o), \text{cell}(C, S_o), \text{head}(\text{LC}). \\
r_2 @ \text{delta}(Q_i, S_i, Q_o, S_o, \text{right}), \text{adj}(C, \text{RC}) & \setminus \text{state}(Q_i), \text{cell}(C, S_i), \text{head}(C) \\
\iff \text{RC} \setminus == \text{null} & \mid \text{state}(Q_o), \text{cell}(C, S_o), \text{head}(\text{RC}). \\
r_3 @ \text{delta}(Q_i, S_i, Q_o, S_o, \text{left}) & \setminus \text{adj}(\text{null}, C), \text{state}(Q_i), \text{cell}(C, S_i), \text{head}(C) \\
\iff \text{cell}(\text{LC}, b), \text{adj}(\text{null}, \text{LC}), \text{adj}(\text{LC}, C), & \text{state}(Q_o), \text{cell}(C, S_o), \text{head}(\text{LC}). \\
r_4 @ \text{delta}(Q_i, S_i, Q_o, S_o, \text{right}) & \setminus \text{adj}(C, \text{null}), \text{state}(Q_i), \text{cell}(C, S_i), \text{head}(C) \\
\iff \text{cell}(\text{RC}, b), \text{adj}(C, \text{RC}), \text{adj}(\text{RC}, \text{null}), & \text{state}(Q_o), \text{cell}(C, S_o), \text{head}(\text{RC}). \\
\text{fail} @ \text{nodelta}(Q_i, S_i), \text{rejecting}(Q_i), \text{state}(Q_i), & \text{cell}(C, S_i), \text{head}(C) \implies \text{fail}.
\end{aligned}$$

Fig. 5. The CHR program TMSIM, a Turing machine simulator.

—**adj**/2 encodes the order of the tape cells. The constraint **adj**(A, B) should be read: “The right neighbor of the tape cell with identifier A is the tape cell with identifier B .”

The special cell identifier **null** is used to refer to a not-yet-instantiated cell. The rules r_3 and r_4 take care of extending the tape as needed. The first four rules of TMSIM correspond to the four Turing machine transition rules.

A simulation of the execution of a Turing machine M proceeds as follows. The tape input is encoded as **cell**/2 constraints and **adj**/2 constraints. The identifier of the cell to the left of the leftmost input symbol is set to **null** and similarly for the cell to the right of the rightmost input symbol. The transition function δ of M is encoded in multiple **delta**/5 constraints. All these constraints are combined in the initial query together with the constraint **state**(q_0), where q_0 is the initial state of M and the constraint **head**(c_1), where c_1 is the identifier of the cell representing the leftmost input symbol. Every rule application of the first four rules of TMSIM corresponds directly to a Turing machine transition.

If no more (Turing machine) transitions can be made, the last rule is applicable if the current state is nonaccepting. In that case, the built-in constraint *fail* is added, which leads to a failure state. If the Turing machine ends in an accepting final state, the CHR program ends in a successful final state.

Formally, we define a function `tm_to_chr` which produces a query for TMSIM, given a Turing machine and an input tape. It is defined as follows: Given a Turing machine $M = \langle Q, \Sigma, q_0, b, F, \delta \rangle$ and an input tape $I = [i_1, \dots, i_n]$,

$$\text{tm_to_chr}(M, I) = \text{prog_to_chr}(M) \cup \{\text{state}(q_0), \text{head}(c_1)\} \cup \text{tape_to_chr}(I)$$

where `prog_to_chr`(M) is defined as follows:

$$\begin{aligned}
\text{prog_to_chr}(M) = & \{\text{delta}(q, s, q^*, s^*, d) \mid (q, s) \in Q \times \Sigma \text{ and } \delta(q, s) = (q^*, s^*, d)\} \\
& \cup \{\text{nodelta}(q, s) \mid (q, s) \in Q \times \Sigma \text{ and } \delta(q, s) \text{ is undefined}\} \\
& \cup \{\text{rejecting}(q) \mid q \in Q \setminus F\}
\end{aligned}$$

and `tape_to_chr`(I) = $\bigcup_{j=1}^n \{\text{cell}(c_j, i_j), \text{adj}(c_j, c_{j+1})\} \cup \{\text{adj}(\text{null}, c_1)\}$ where $c_{n+1} = \text{null}$ and the other c_j are unique cell identifiers. Clearly, `tm_to_chr` can be computed in $O(n + |Q \times \Sigma|)$ time: linear in the size of the input tape plus the size of the domain of the Turing machine program. We now define $\mathcal{V}_{\mathbf{T}}^{\mathbf{M}}$

as:

$\mathcal{V}_{\mathbf{T}_M} = \{\text{tm_to_chr}(M, I) \mid M \text{ is a Turing machine for which } I \text{ is an input tape}\}.$

LEMMA 4.1. $\mathcal{M}_{\mathbf{T}_M}$ is indeed a CHR machine, that is, TMSIM is a $\Delta_{\omega_t}^\Phi$ -deterministic CHR program for input $\mathcal{V}_{\mathbf{T}_M}$.

PROOF. Clearly, the rules of TMSIM maintain a valid tape representation and the invariant that there is at most one **head**/1 constraint and one **state**/1 constraint. For deterministic Turing machines the first two arguments of **delta**/5 functionally determine the other three arguments. Hence, since valid input corresponds to a deterministic Turing machine, the rules of TMSIM are mutually exclusive. As a result, the **Apply** transitions of a derivation are determined, and only the order of the **Introduce** transitions may vary. However, it can easily be verified that the order of **Introduce** transitions cannot affect the derivation result and it can only affect the derivation length for failing derivations. $\square$

We define a function chr_to_tm which returns a Turing machine execution state, given an execution state for a CHR machine: $\text{chr_to_tm} : \Sigma^{\text{CHR}} \rightarrow \Sigma^{\mathbf{T}_M}$: $\text{chr_to_tm}(\langle \mathbb{G}, \mathbb{S}, \mathbb{B}, \mathbb{T} \rangle_n) = \langle q, P, s, N \rangle$ where q , P , S , and N are such that $\mathbb{G} \uplus \mathbb{S} = \text{prog_to_chr}(M) \cup \{\mathbf{state}(q), \mathbf{head}(c), \mathbf{cell}(c, s)\} \cup \text{tape_to_chr}(\text{tape}(\langle q, P, s, N \rangle))$. It can easily be verified that this is indeed a function, and that it can be computed in time linear in the tape size.

We now show that $\mathcal{M}_{\mathbf{T}_M}$ can simulate every Turing machine.

THEOREM 4.2. For any Turing machine M and input tape I :

- (1) M terminates on input $I \Leftrightarrow \mathcal{M}_{\mathbf{T}_M}$ terminates on input $\text{tm_to_chr}(M, I)$;
- (2) M accepts $I \Leftrightarrow \mathcal{M}_{\mathbf{T}_M}$ accepts $\text{tm_to_chr}(M, I)$; and
- (3) M outputs $\text{chr_to_tm}(X)$ on input $I \Leftrightarrow \mathcal{M}_{\mathbf{T}_M}$ outputs X on input $\text{tm_to_chr}(M, I)$.

PROOF. Let $\sigma_0 = \text{initstate}(\text{tm_to_chr}(M, I))$ be the initial state of $\mathcal{M}_{\mathbf{T}_M}$ with input $\text{tm_to_chr}(M, I)$. Note that $\text{chr_to_tm}(\sigma_0) = \langle q_0, [], i_1, [i_2, \dots, i_n] \rangle$ is the initial execution state of M .

Observation 1. For every $\mathcal{M}_{\mathbf{T}_M}$ transition $\sigma_i \rightarrow \sigma_{i+1}$, either:

- (a) $\text{chr_to_tm}(\sigma_i) \rightarrow \text{chr_to_tm}(\sigma_{i+1})$ is a TM transition for M , or
- (b) $\text{chr_to_tm}(\sigma_i) = \text{chr_to_tm}(\sigma_{i+1})$.

This can be shown by induction on the number of CHR machine steps and case analysis on the transition rules of ω_t . The first case (a) holds if $\sigma_i \rightarrow \sigma_{i+1}$

8:14 • J. Sneyers et al.

by the **Apply** transition rule, where the applied rule is r_x . The second case (b) holds if $\sigma_i \rightarrow \sigma_{i+1}$ by the **Introduce** transition rule. Case (b) also holds when $\sigma_i \rightarrow \sigma_{i+1}$ by the **Apply** transition rule for the rule *fail*, which can only be followed by (a series of **Introduce** transitions followed by) a **Solve** transition which results in a failure state. Schematically, the relation between the Turing machine derivation steps and the CHR machine steps is as follows.

Observation 2. If σ_i is a final state (for $\mathcal{M}_{\text{TM}}$), then $\text{chr_to_tm}(\sigma_i)$ is a final state (for M). Furthermore, if σ_i is a successful final state, then $\text{chr_to_tm}(\sigma_i)$ is an accepting final state, and if σ_i is a failure state, then $\text{chr_to_tm}(\sigma_i)$ is nonaccepting.

Observation 3. The subderivations corresponding to one Turing machine step are finite. In other words, all chains of the form $\text{chr_to_tm}(\sigma_n) = \text{chr_to_tm}(\sigma_{n+1}) = \text{chr_to_tm}(\sigma_{n+2}) = \dots$ have a finite length. This follows from Observation 1 and the definition of the **Introduce** transition, which decrements the size of the (finite) multiset representing the goal in CHR execution states.

Observation 4. If M terminates on input I , then the corresponding $\mathcal{M}_{\text{TM}}$ derivation on input $\text{tm_to_chr}(M, I)$ is also finite. This follows from Observations 1 and 3.

The three properties follow straightforwardly from the preceding observations. $\square$

4.2 Definitions of Time and Space Complexity

We have shown that any Turing-computable function can be computed in CHR. To investigate complexity properties of CHR machines, we first define the notions of time and space complexity for the three models of computation.

4.2.1 Turing Machines.

Definition 4.3. The (worst-case) *time complexity* of a Turing machine $M = \langle Q, \Sigma, q_0, b, F, \delta \rangle$ is the maximal derivation length for inputs of a given size.

$$\mathbf{T}_{\text{TIME}_M}(n) = \max\{\#deriv_M(I) \mid I \in \Sigma^* \text{ and } |I| = n\}$$

Definition 4.4. The (worst-case) *space complexity* of a Turing machine M is the maximal tape size used in derivations for inputs of a given size.

$$\mathbf{T}_{\text{SPACE}_M}(n) = \max\{\#tape(\sigma) \mid \sigma \in deriv_M(I) \text{ and } I \in \Sigma^* \text{ and } |I| = n\}$$

4.2.2 RAM Machines.

Definition 4.5. The (worst-case) time complexity $\text{RAMTIME}_P(n)$ of a RAM machine with program P is the maximal time needed for an execution starting with some input of size n . The time needed for a program execution is the sum of the times needed for every instruction that is executed. All instructions take constant time except for the arithmetic instructions **add**, **sub**, **mul**, and **div**, which take time logarithmic in the absolute value of the numbers involved in the arithmetic operation.

Definition 4.6. The (worst-case) space complexity $\text{RAMSPACE}_P(n)$ of a RAM machine with program P is the maximum, over every execution starting with an input of size n , of the sum over all registers in the range $0, \dots, \text{maxaddr}$ of the number of bits needed to represent the maximal value that it held. Unused registers within this range are charged for one bit for the implicit value 0 stored there.

In practice, we often assume the registers to use at most a fixed number of bits. In this case, every instruction takes constant time and the space complexity is just maxaddr , the maximum register address reached in a computation.

This definition corresponds to that of Savitch [1978]. In the literature, other definitions of space complexity are often used that do not count those registers that were unused.⁴ According to such definitions, a program that uses for input size n (only) the registers with address $2^1, 2^2, 2^3, \dots, 2^n$ has a space complexity of only $\mathcal{O}(n)$; according to our definition it needs exponential space. An advantage of such definitions is that programs never use more space than time. However, in our opinion the definition of Savitch [1978] is more realistic.

In the following, we will assume RAM programs to use at least as much time as space. This is a nontrivial assumption given our definition of space complexity, but unless a program uses the registers in an unrealistically sparse way, it is not a very restrictive assumption. After all, every program can be rewritten to use memory in a dense way with only logarithmic time overhead, for example, by explicitly storing address-value pairs using a balanced search tree.

RAM machines are Turing-complete. Both RAM machines can simulate a T -time Turing machine in $\mathcal{O}(T)$ time. The main difference between the two RAM machines is the time complexity that can be achieved when simulating them on a Turing machine. According to Savage [1998], a T -time Peano-arithmetic RAM using S registers can be simulated on a Turing machine in $\mathcal{O}(ST \log^2 S)$ time. The standard RAM is also polynomially related to the Turing machine, although it is more expensive to simulate on a TM. According to Aho et al. [1975], a standard RAM machine with time complexity T can be simulated on a multitape TM in $\mathcal{O}((T \log T \log \log T)^2)$ time. Simulating a multitape TM on a single-tape TM squares the time complexity [Hartmanis and Stearns 1965], so we have the next lemma.

⁴See also van Emde Boas [1990, page 28] for a discussion of different space complexity measures.

8:16 • J. Sneyers et al.

A	B	C	D	E	$\% \text{ horizontal:}$ $\text{word}(\text{A,B,C,D,E}),$ $\text{word}(\text{I,J,K}),$ $\text{word}(\text{M,N,O,P}),$ $\text{word}(\text{Q,R,S,T,U}),$ $\Rightarrow$ $\text{solution}(\text{A,B,C,D,E,F,G,H,I,J,K,L,M,N,O,P,Q,R,S,T,U}).$	vertical: $\text{word}(\text{A,F,I}),$ $\text{word}(\text{J,M,R}),$ $\text{word}(\text{C,G,K,N,S}),$ $\text{word}(\text{O,T}),$ $\text{word}(\text{E,H,L,P,U})$
F		G		H		
I	J	K		L		
	M	N	O	P		
Q	R	S	T	U		

Fig. 6. Solver program for some crossword puzzle.

LEMMA 4.7. *Any standard RAM machine with time complexity T can be simulated on a Turing machine with time complexity $\tilde{O}(T^4)$.*

4.2.3 CHR Machines.

Definition 4.8. Given a CHR machine $\mathcal{M} = (\mathcal{H}, \mathcal{P}, \mathcal{V})$, the time function $\text{chrtime}_{\mathcal{M}}$ returns the derivation length, given a valid goal.

$$\text{chrtime}_{\mathcal{M}} : \mathcal{V} \rightarrow \mathbb{N} : \mathbb{G} \mapsto \#\text{deriv}_{\mathcal{M}}(\mathbb{G})$$

Definition 4.9. Given a CHR machine $\mathcal{M} = (\mathcal{H}, \mathcal{P}, \mathcal{V})$ and assuming that host language constraints of $\mathcal{H}$ take constant time, the (worst-case) time complexity function $\text{CHRTIME}_{\mathcal{M}}$ is defined as

$$\text{CHRTIME}_{\mathcal{M}}(n) = \max\{\text{chrtime}_{\mathcal{M}}(\mathbb{G}) \mid \mathbb{G} \in \mathcal{V} \wedge \text{inputsize}(\mathbb{G}) = n\}$$

where inputsize is a function which returns the size of a goal.

The time given by $\text{chrtime}_{\mathcal{M}}$ is the time needed by a theoretical CHR machine. The next section deals with the relation between this theoretical time and practical CHR implementations. This is important because the definition of $\text{chrtime}_{\mathcal{M}}$ does not correspond to the reality of CHR implementations, as the following example illustrates.

Example. A CHR machine can solve a crossword puzzle in time linear in the number of words, which does not seem to be achievable on a RAM machine. Given a query containing n words from a dictionary as word/k constraints (word length k), the CHR machine with the program of Figure 6 returns all s solutions for a crossword puzzle in time $\mathcal{O}(n + s)$.

In many cases we simply use the number of conjuncts in the goal $\mathbb{G}$ as its size $\text{inputsize}(\mathbb{G})$. In other cases, valid goals consist of just one constraint and the goal size is determined by its arguments.

Example. Consider the CHR machine $\mathcal{M}_P = (\text{Prolog}, \text{PRIMES}, \{\text{upto}(n) \mid n \in \mathbb{N}\})$ (refer to Figure 2). An appropriate goal size function is given by $\text{inputsize}(\text{upto}(n)) = n$. Every derivation starting with $\text{upto}(n)$ consists of $n - 1$ **Solve** steps (one for every $\text{is}/2$ built-in constraint), $2n - 1$ **Introduce** steps (n $\text{upto}/1$ constraints and $n - 1$ **prime}/1 constraints), and $n + \text{nonprimes}(n) < 2n$ **Apply** steps, where $\text{nonprimes}(n)$ is the number of composite numbers between 2 and n (loop is applied $n - 1$ times, stop once, and $\text{absorb nonprimes}(n)$ times). As a result, the time complexity $\text{CHRTIME}_{\mathcal{M}_P}(n) < 5n - 2$ is $\mathcal{O}(n)$.**

$i(L, L_o, \text{inc}, A) \setminus$	$m(A, X), c(L)$	$\iff m(A, Z), s(Z, X), c(L_o).$
$i(L, L_o, \text{dec}, A) \setminus$	$m(A, X), s(X, Y), c(L)$	$\iff m(A, Y), c(L_o).$
$i(L, L_o, \text{clr}, A) \setminus$	$m(A, X), c(L)$	$\iff m(A, \text{zero}), c(L_o).$
$i(L, L_o, \text{jmp}, A) \setminus$	$c(L)$	$\iff c(A).$
$i(L, L_o, \text{cjmp}, A, J),$	$m(A, \text{zero}) \setminus c(L)$	$\iff c(J).$
$i(L, L_o, \text{cjmp}, A, J),$	$m(A, X), s(X, _) \setminus c(L)$	$\iff c(L_o).$
$i(L, L_o, \text{halt}) \setminus$	$c(L)$	$\iff \text{true}.$

Fig. 7. The CHR program PARAMSIM, a simulator of Peano-arithmetic RAM machines.

Definition 4.10 (State Size Function). The state size function is

$$\text{SIZE} : \Sigma^{\text{CHR}} \rightarrow \mathbb{N} : \langle G, S, B, T \rangle_n \mapsto \text{size}(G) + \text{size}(S) + \text{size}(B) + \text{size}(T)$$

where for sets X , $\text{size}(X) = \sum_{x \in X} |x|$ and the size $|x|$ is the usual term size.

Definition 4.11. Given a CHR machine $\mathcal{M} = (\mathcal{H}, \mathcal{P}, \mathcal{V})$, the *space function* $\text{chrspac}_{\mathcal{M}}$ returns the worst-case execution state size, given an initial goal.

$$\text{chrspac}_{\mathcal{M}}(G) = \max\{\text{SIZE}(\sigma) \mid \sigma \in \text{deriv}_{\mathcal{M}}(G)\}$$

Definition 4.12. Given a CHR machine $\mathcal{M} = (\mathcal{H}, \mathcal{P}, \mathcal{V})$, the (worst-case) space complexity function $\text{CHRSPACE}_{\mathcal{M}}$ is defined as follows.

$$\text{CHRSPACE}_{\mathcal{M}}(n) = \max\{\text{chrspac}_{\mathcal{M}}(G) \mid G \in \mathcal{V} \wedge \text{inputsize}(G) = n\}$$

Example. Consider, as in the previous example, the CHR machine $\mathcal{M}_P$ and the goal size function $\text{inputsize}(\text{upto}(n)) = n$. If we assume all integers can be represented in a fixed number of bits, the space complexity $\text{CHRSPACE}_{\mathcal{M}_P}$ is $\mathcal{O}(n)$. Note that if the constraint size is bounded, the size of the constraint store is asymptotically dominated by the number of **Introduce** steps, and the size of the built-in store is dominated by the number of **Solve** steps.

4.3 Complexity of CHR Machines

We now show that the CHR machine is at least as efficient as the RAM machine. Specifically, the CHR-only machine can simulate any PA-RAM machine with the same time complexity, and $\text{CHR}(\mathcal{H})$ machines can simulate standard RAM machines with the same time and space complexity if $\mathcal{H}$ is a sufficiently strong host language.

THEOREM 4.13. *A CHR-only machine $\mathcal{M}_{\text{PARAM}}$ exists which can simulate, in $\mathcal{O}(T + P + I)$ time and $\mathcal{O}(P + J)$ space, a T -time, S -space Peano-arithmetic RAM machine with a program of P lines, where I is the sum of the values of the input registers and J is the maximal sum of the values of all registers during the entire computation.*

PROOF. The proof roughly corresponds to that of Theorem 4.2. Consider the CHR-only program PARAMSIM, listed in Figure 7, and the corresponding CHR-only machine $(\Phi, \text{PARAMSIM}, \mathcal{V}_{\text{PARAM}})$.

The mapping from CHR states to RAM states is as follows. Memory cells are represented as $m(A, V)$ constraints, where A is the address and V refers to the value. If V is the atom `zero`, the value is zero. Otherwise, there is a

8:18 • J. Sneyers et al.

successor constraint $\mathbf{s}(V, W)$ which expresses that $V = W + 1$. Again, W can be either zero or there is another successor constraint $\mathbf{s}(W, X)$, etc. For example, to represent that the register r_1 contains the value 3, the following conjunction of constraints could be used: $\mathbf{m}(r_1, N_3)$, $\mathbf{s}(N_3, N_2)$, $\mathbf{s}(N_2, N_1)$, $\mathbf{s}(N_1, \text{zero})$.

The mapping from a RAM program and input to an initial CHR goal is now obvious. We define $\mathcal{V}_{\text{PARAM}}$ to be the image of this mapping. The RAM program is encoded as $\mathbf{i}/\{3, 4, 5\}$ constraints; the first argument represents the label (or line number), the second argument is the label of the next program line, the other arguments represent the actual instruction. The input memory cells are represented as before. The program counter is set to the label of the first line of the program by adding a corresponding initial $\mathbf{c}/1$ constraint.

After the initialization, one rule is applied for every step of the RAM. Applying a rule causes at most three constraints to be inserted. Therefore, the number of CHR steps is bounded by four times the number of RAM steps, plus P steps to **Introduce** the $\mathbf{i}/\{3, 4, 5\}$ constraints. Initializing the input memory cells to their values means inserting one $\mathbf{m}/2$ constraint for every input and I $\mathbf{s}/2$ constraints. The number of registers (and hence $\mathbf{m}/2$ constraints) is fixed for a given program (remember we don't have indirection), so the total time complexity of the CHR machine is $\mathcal{O}(T + P + I)$. The CHR store contains P $\mathbf{i}/\{3, 4, 5\}$ constraints; since the number of $\mathbf{m}/2$ constraints is fixed and the number of $\mathbf{s}/2$ constraints is bounded by J , the space complexity of the CHR machine is $\mathcal{O}(P + J)$. $\square$

The PARAMSIM program does not use any host language arithmetic operations. Its unary number representation causes an exponential space penalty. If host language arithmetic is available, a similar program can be written which uses only $\mathcal{O}(P + S)$ space by directly storing the value of a register in the second argument of the $\mathbf{m}/2$ constraints.

We can also simulate the more realistic standard-arithmetic RAM machine in CHR. However, if we want to do this without a harsh complexity penalty, we need host language support to handle the integers.

THEOREM 4.14. *For any sufficiently strong host language $\mathcal{H}$, a $\text{CHR}(\mathcal{H})$ machine $\mathcal{M}_{\text{RAM}}$ exists which can simulate, in $\mathcal{O}(T + P + S)$ time and $\mathcal{O}(S + P)$ space, a T -time, S -space standard RAM machine with a program of P lines.*

PROOF. As in Theorem 4.13, for the simulator program of Figure 8. The RAM memory representation is simpler since integer numbers are available. The representation of a RAM machine program and memory is as before, except that we assume input registers to be in a continuous range $0, \dots, m$ and a constraint $\mathbf{maxm}(m)$ is added to the initial CHR goal. The auxiliary constraint $\mathbf{maxm}(m)$ indicates that the current highest initialized register address is m . When a register with a higher address n is initialized, the auxiliary constraint $\mathbf{initm}/3$ is used to initialize all addresses in the range $m + 1, \dots, n$. In addition to the $\mathcal{O}(P)$ time to **Introduce** the encoded program and the $\mathcal{O}(T)$ time to simulate it (again one CHR rule application per program instruction), the simulator needs $\mathcal{O}(S)$ time to initialize memory. $\square$

$i(L, \text{init}, A),$	$m(A, B), \text{maxm}(M) \setminus c(L)$	$\iff \text{initm}(M+1, B, L).$
	$\text{initm}(A, B, L) \iff A \leq B \mid m(A, 0), \text{initm}(A+1, B, L).$	
	$\text{initm}(A, B, L), m(B, X) \iff A > B \mid m(B, 0), \text{maxm}(B), c(L+1).$	
$i(L, \text{cst}, B, A) \setminus$	$m(A, X), c(L)$	$\iff m(A, B), c(L+1).$
$i(L, \text{add}, B, A),$	$m(B, Y) \setminus m(A, X), c(L)$	$\iff m(A, X+Y), c(L+1).$
$i(L, \text{sub}, B, A),$	$m(B, Y) \setminus m(A, X), c(L)$	$\iff m(A, X-Y), c(L+1).$
$i(L, \text{mul}, B, A),$	$m(B, Y) \setminus m(A, X), c(L)$	$\iff m(A, X*Y), c(L+1).$
$i(L, \text{div}, B, A),$	$m(B, Y) \setminus m(A, X), c(L)$	$\iff m(A, X/Y), c(L+1).$
$i(L, \text{mov}, B, A),$	$m(B, Y) \setminus m(A, X), c(L)$	$\iff m(A, Y), c(L+1).$
$i(L, \text{imv}, B, A),$	$m(B, C), m(C, Y) \setminus m(A, X), c(L)$	$\iff m(A, Y), c(L+1).$
$i(L, \text{mvi}, B, A),$	$m(B, Y), m(A, C) \setminus m(C, X), c(L)$	$\iff m(C, Y), c(L+1).$
$i(L, \text{jmp}, A) \setminus$	$c(L)$	$\iff c(A).$
$i(L, \text{cjmp}, A, J),$	$m(A, 0) \setminus c(L)$	$\iff c(J).$
$i(L, \text{cjmp}, A, J),$	$m(A, X) \setminus c(L)$	$\iff X \neq 0 \mid c(L+1).$
$i(L, \text{halt}) \setminus$	$c(L)$	$\iff \text{true}.$

Fig. 8. The CHR program SRAMSIM, a simulator of standard RAM machines.

Note that for a given, fixed RAM program which uses at least as much time as space, P is a constant and S is $\mathcal{O}(T)$, so the CHR simulator SRAMSIM takes $\mathcal{O}(T)$ time and $\mathcal{O}(S)$ space.

5. IMPLEMENTING CHR MACHINES ON RAM MACHINES

The previous section dealt with the complexity properties of the theoretical CHR machine. In this section we investigate the complexity achievable in practice: on a RAM machine. CHR compilers convert CHR programs to (RAM machine) executable code (Section 5.1). We prove a general result that gives a bound on the complexity of executing a CHR machine on a RAM machine (Section 5.2). We then apply it to the RAM machine simulator (Section 5.3). This allows us to conclude that every algorithm can be implemented efficiently in CHR.

5.1 Compilation of CHR

Holzbaur and Frühwirth [1999] pioneered efficient implementations of Constraint Handling Rules by compilation to host language (Prolog) code. The resulting host language code can be executed on RAM machines by an interpreter or a further compilation step. RAM machines can obviously simulate CHR machines since CHR implementations exist.

LEMMA 5.1. *The RAM machine can simulate the CHR machine.*

Because CHR machines can simulate Turing machines (Theorem 4.2), and Turing machines can simulate CHR machines (Lemma 4.7 and the preceding), we get the following rather unsurprising result.

COROLLARY 5.2. *The CHR machine is Turing-complete.*

5.2 Complexity of the Compiled Code

We now examine the practically achievable complexity of simulating a CHR machine on a RAM machine. We consider existing CHR compilers. One of the most optimizing CHR compilers currently available is the K.U. Leuven CHR

8:20 • J. Sneyers et al.

system.⁵ It was developed by Schrijvers and Demoen [2004], building on the system developed by Holzbaur and Frühwirth [1999]. For more details on optimizing CHR compilation, see Holzbaur et al. [2005], Schrijvers [2005], Duck [2005], and Sneyers et al. [2009].

First we give an overview of some notions of optimizing compilation that are crucial for time and space efficiency. Given the refined operational semantics, the join ordering strategy, functional dependencies, and efficient constraint store operations, we define the *dependency rank* of a constraint occurrence, which determines the time complexity of executing a CHR machine. We then discuss memory reuse techniques that have an impact on the space complexity. Finally, we formulate a general complexity metatheorem.

5.2.1 The Refined Operational Semantics ω_r . Like most other CHR implementations, the K. U. Leuven CHR system implements the refined operational semantics ω_r , first formalized by Duck et al. [2004]. Since the ω_r semantics instantiates the abstract ω_t semantics, every derivation performed by the K. U. Leuven CHR system corresponds to a valid ω_t derivation. Note that the reverse is not true: In general, not every ω_t derivation corresponds to a ω_r derivation.

Essential in the ω_r semantics is the notion of an *active* constraint. Query and body constraints are introduced from left to right. Once a constraint is introduced (or triggered by a **Solve** transition), it becomes active: Its occurrences in the program are tried, in textual order. For every occurrence, the corresponding rule is tried by looking up matching *partner* constraints.⁶

5.2.2 Join Ordering. CHR compilers implement a strategy to pick the order of partner constraint lookups, called the *join ordering* [Holzbaur et al. 2005]. Given a CHR program $\mathcal{P}$, the join ordering induces, for every head constraint occurrence c of $\mathcal{P}$, an order $<_c^{\mathcal{P}}$ on its partners. The join ordering strategy of the K.U. Leuven CHR compiler is based on the algorithm introduced in Holzbaur et al. [2005].

5.2.3 Functional Dependencies. Duck and Schrijvers [2005] formally defined the (set semantics) functional dependency property and they also proposed an abstract interpretation-based analysis to detect this property. Informally, for a given CHR store, a constraint has a *set semantics functional dependency* on certain *key* arguments if there are no two instances of the constraint with the same key arguments. We omit the formal definition, referring to Duck and Schrijvers [2005]. Functional dependencies can be enforced using simpagation rules. For example, to make sure that **c/4** has a functional dependency on the combination of its first and third arguments, a program could start with the rule “**c**(A,.,B,.) \ **c**(A,.,B,.) $\iff$ true”.

⁵The K.U. Leuven CHR Web site: <http://www.cs.kuleuven.be/~toms/CHR/>.

⁶The partner constraints of an occurrence c_i in a rule with a head of the form $c_1, \dots, c_k \setminus c_{k+1}, \dots, c_n$ are simply all the *other* head constraints in that rule: $c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n$.

5.2.4 Constraint Store Operations. Naive implementations of CHR use Prolog lists to represent the constraint store. This allows insertion in $\mathcal{O}(1)$ time, but lookup and removal take $\mathcal{O}(n)$ time in the worst case (where n is the number of constraints in the store). Optimizing CHR compilers like the K.U. Leuven CHR system use more advanced data structures to implement the constraint store. The join ordering determines which combinations of key arguments (lookup patterns) are used for partner constraint lookups. The constraint store is implemented in such a way that for every combination of constraint arguments that is used as a lookup pattern, an index is maintained, for instance, using hash-tables. As a result, all constraint store operations can be done in $\mathcal{O}(1)$ (amortized) time.

5.2.5 Determined Partners. To obtain a stronger bound on the complexity of finding suitable partner constraints, we introduce the notions of determined partners and dependency rank. These notions depend both on the power of the static analysis in the CHR compiler and on the dynamic program behavior. The former influences the quality of the join order strategy, while the latter determines the functional dependencies in derivations starting with valid goals $\mathcal{W} \subseteq \mathcal{G}_P^H$.

Definition 5.3. Given a join ordering strategy $\prec$, a CHR program $\mathcal{P}$, and a set of valid goals $\mathcal{W}$, we say an occurrence c is *determined* by an (active) occurrence a if, for all execution states σ that occur in a derivation $d \in \Delta_{\omega}^H|_{\mathbb{G}}$ for some valid goal $\mathbb{G} \in \mathcal{W}$, the following holds: If σ is of the form $\langle [a\#i : j | \mathbb{A}], \mathbb{S}, \mathbb{B}, \mathbb{T} \rangle_n$ (i.e., the occurrence subprocedure for the j th occurrence of constraint a is about to be executed), then a set semantic functional dependency for c holds in state σ , where the key arguments of c are fixed by a and all partners x for which $x \prec_a^P c$.

Definition 5.4. The *dependency rank* of an (active) occurrence a is the number of nondetermined partner constraints of a .

We first illustrate the preceding definitions. We assume the join ordering strategy $\prec$ used in the K. U. Leuven CHR system, the program SRAMSIM, and the set of valid goals corresponding to the valid RAM machine instances.

Example. Consider the third rule of SRAMSIM.

$$\mathbf{i}(\mathbf{L}, \mathbf{add}, \mathbf{B}, \mathbf{A}), \mathbf{m}(\mathbf{B}, \mathbf{Y}) \setminus \mathbf{m}(\mathbf{A}, \mathbf{X}), \mathbf{c}(\mathbf{L}) \iff \mathbf{m}(\mathbf{A}, \mathbf{X}+\mathbf{Y}), \mathbf{c}(\mathbf{L}+1).$$

With respect to the active occurrence $\mathbf{c}(\mathbf{L})$, the following join ordering is computed: $\mathbf{c}(\mathbf{L}) \prec_{\mathbf{c}(\mathbf{L})}^{\text{SRAMSIM}} \mathbf{i}(\mathbf{L}, \mathbf{add}, \mathbf{B}, \mathbf{A}) \prec_{\mathbf{c}(\mathbf{L})}^{\text{SRAMSIM}} \mathbf{m}(\mathbf{B}, \mathbf{Y}) \prec_{\mathbf{c}(\mathbf{L})}^{\text{SRAMSIM}} \mathbf{m}(\mathbf{A}, \mathbf{X})$. Given the first argument $\mathbf{L}$ there can be only one $\mathbf{i}/4$, which means the first arguments of the $\mathbf{m}/2$ constraints are known. Again, for any given first argument, there can be only one $\mathbf{m}/2$. These functional dependencies are not enforced explicitly by simpagation rules; they are implied by the set of valid goals and the rules. Since all partners are determined by $\mathbf{c}(\mathbf{L})$, the dependency rank of $\mathbf{c}(\mathbf{L})$ is zero.

The reader can verify that for all occurrences of $\mathbf{c}/1$ in the rules of SRAMSIM, the dependency rank is zero. In general, the dependency rank is of course

8:22 • J. Sneyers et al.

not always zero. Also, in one rule, different occurrences may have a different dependency rank.

Example. Consider the following rule.

$$\mathbf{m}(A,B), \mathbf{m}(B,C), \mathbf{m}(D,E), \mathbf{m}(E,F), \mathbf{m}(G,H) \iff \dots$$

With respect to the active occurrence $a = \mathbf{m}(A,B)$, we get the following join ordering: $\mathbf{m}(A,B) <_a \mathbf{m}(B,C) <_a \mathbf{m}(D,E) <_a \mathbf{m}(E,F) <_a \mathbf{m}(G,H)$. Assuming that the first argument of $\mathbf{m}/2$ uniquely determines the constraint, the dependency rank of $\mathbf{m}(A,B)$ is two: $\mathbf{m}(D,E)$ and $\mathbf{m}(G,H)$ contribute to the dependency rank. The dependency rank of $\mathbf{m}(B,C)$ is three.

5.2.6 Space Reuse. Every time a CHR constraint is removed, its representation in memory becomes garbage. If this garbage is not collected, we may get a space complexity which is not linear in the size of the constraint store. Using garbage collection we get the right space complexity. However, in most Prolog systems, garbage collection has a time complexity linear in the number of live cells. This may result in a severe time complexity penalty.

In earlier work [Sneyers et al. 2006b] we have tackled this problem by introducing memory reuse techniques called *in-place updates* and *suspension reuse*, inspired by compile-time garbage collection [Mazur 2004]. These optimizations improve the space complexity by eliminating garbage, with only a small constant factor worst-case time overhead. The basic idea of suspension reuse is to store the representation of a removed constraint in a cache. Later, when a new constraint has to be inserted, a representation from the cache is used to build the new constraint representation. In-place updates are a special case where both the removal and the insertion are in the same rule, eliminating the need for an intermediate cache.

LEMMA 5.5. *Using suspension reuse with unlimited cache size, the following holds: If during a particular execution, the maximal number of constraints in the store is S_{max} , then at any point in the execution, $S + C \leq S_{max}$, where S is the number of constraints in the store and C is the number of elements in the cache.*

PROOF. Execution consists of a sequence of insertion and removal operations. We denote the store size after the i th operation with S_i , the cache size with C_i , and their sum with $M_i = S_i + C_i$. Initially both the store and the cache are empty: $S_0 = C_0 = 0$. Insertion increments the store size and decrements the cache size if it is not already empty (otherwise it simply remains empty). Removal decrements the store size and increments the cache size. We proceed by induction on the sequence length. For zero-length sequences the property holds trivially. Assuming the property holds for any sequence of length n , we show that it also holds for sequences of length $n + 1$. Because of the induction hypothesis, we have the following.

$$\forall x \leq n : M_x \leq \max_{i \leq n} S_i \leq \max_{i \leq n+1} S_i = S_{max}$$

We now only have to show that $M_{n+1} \leq S_{max}$. If the last operation in the sequence is a removal, then $M_{n+1} = (S_n - 1) + (C_n + 1) = M_n \leq S_{max}$. Assuming the

last operation is an insertion, there are two cases: If $C_n > 0$, then $M_{n+1} = (S_n + 1) + (C_n - 1) = M_n \leq S_{max}$; otherwise $C_n = 0$ and $M_{n+1} = S_{n+1} \leq S_{max}$. $\square$

With an unlimited cache size, no constraint representation ever becomes garbage (from the point of view of the underlying host language in which the CHR system is implemented): All constraint representations are alive, either because they are in the store or because they are in the cache. The previous lemma shows that the right space complexity can be achieved for all CHR programs, without having to resort to runtime garbage collection in the host language.

5.2.7 General Complexity Result. We now state the main result of this section.

THEOREM 5.6. *Given a CHR program $\mathcal{P}$ and a ω_l derivation d of length T which has a corresponding ω_r derivation, for which the maximal store size is S , m is the maximum dependency rank of the active occurrences in $\mathcal{P}$, and p is the number of propagation rule applications in d ; assuming the host language constraints used in the guards and bodies of the rules of $\mathcal{P}$ can be evaluated in constant time; the K. U. Leuven CHR system compiles $\mathcal{P}$ to Prolog code which has, for the given derivation d , a time complexity $\mathcal{O}(TS^{m+1})$ and a space complexity $\mathcal{O}(S + p)$.*

PROOF. Assume the derivation d consists of s **Solve** steps, i **Introduce** steps, and $a = p + r$ **Apply** steps: $T = s + i + a$. The cost of finding a match for an active occurrence is $\mathcal{O}(S^m)$ since this process basically boils down to nested iteration over the constraints in the store, where the nesting depth is the dependency rank. Indeed, determined partners only contribute a constant factor to this cost. Checking and extending the propagation history can be done in constant time if the history is implemented as a hash-table. The **Apply** steps take $\mathcal{O}(a)$ time plus the time to find a match, which we attribute to the **Introduce** and **Solve** transitions. The **Introduce** steps take $\mathcal{O}(iS^m)$ time: Matches are tried for every occurrence of the introduced constraint, and for a fixed CHR program, the number of occurrences is bound by a constant. Every **Solve** step potentially triggers all the $\mathcal{O}(S)$ constraints in the store, so the **Solve** steps may take up to $\mathcal{O}(sS^{m+1})$ time. Since s , i , and a are all $\mathcal{O}(T)$, the total time complexity is $\mathcal{O}(TS^{m+1})$. The $\mathcal{O}(S + p)$ space complexity can be achieved using suspension reuse with unlimited cache size, as shown in Lemma 5.5. $\square$

THEOREM 5.7. *If in the previous theorem, the CHR program is ground (i.e., all constraint arguments are ground), then $\mathcal{O}(TS^m)$ time complexity can be achieved.*

PROOF. In ground programs constraints are never triggered. This reduces the complexity of a single **Solve** transition to a constant. $\square$

Formulating this result in terms of CHR machines and RAM machines, we get the following corollary.

8:24 • J. Sneyers et al.

COROLLARY 5.8. *A ground CHR machine without propagation rules, with time complexity T and space complexity S , can be simulated on a RAM machine with time complexity $\mathcal{O}(TS^m)$ and space complexity $\mathcal{O}(S)$, where m is the maximum dependency rank of the active occurrences in the program of the CHR machine.*

5.3 Complexity of SRAMSIM

Now we use the general result of Corollary 5.8 to analyze the time and space complexity of the RAM simulation (compiler-generated code) of the CHR machine $\mathcal{M}_{\text{RAM}}$, which itself simulates RAM machines.

LEMMA 5.9. *The CHR machine $\mathcal{M}_{\text{RAM}}$ with program SRAMSIM, with time complexity T and space complexity S can be simulated on a RAM machine with time complexity $\mathcal{O}(T)$ and space complexity $\mathcal{O}(S)$.*

PROOF. It can be easily verified that the dependency rank for all occurrences of $\mathbf{c}/1$ is zero, given the join ordering strategy used in the K. U. Leuven CHR system (refer to the first example in Section 5.2.5). For the other occurrences this is slightly less straightforward. Valid goals have exactly one $\mathbf{c}/1$ constraint. All rules remove one $\mathbf{c}/1$ constraint and optionally insert another one in the body (indirectly in the case of the rule for the **init** instruction). Hence there is never more than one $\mathbf{c}/1$ constraint; in other words, $\mathbf{c}/1$ has a set semantics functional dependency on the empty key. If the join order strategy does the lookup of $\mathbf{c}/1$ first, the remaining partners become determined.

Since the dependency rank is zero for all occurrences, applying Corollary 5.8 (with $m = 0$) results in the desired complexities. $\square$

In practice we will add a rule like “ $\mathbf{c}(_) \iff \text{fail}$ ” at the very end of the SRAMSIM program. In the refined semantics, for valid goals, this rule is never applied. However, it allows the *never stored* optimization [Holzbaur et al. 2005] to deduce that all partners of $\mathbf{c}/1$ are *passive*. This implies that only the dependency ranks of the occurrences of $\mathbf{c}/1$ have to be considered.

We conclude that “everything can be done efficiently in CHR”.

COROLLARY 5.10. *For every (RAM machine) algorithm which uses at least as much time as it uses space, a CHR program exists which can be executed in the K. U. Leuven CHR system, with time and space complexity within a constant from the original complexities.*

PROOF. Consider any algorithm which can be expressed as a RAM machine program with a program of P lines, and let its time and space complexities be T and S , respectively. Because of Theorem 4.14, a CHR (Prolog) machine $\mathcal{M}_{\text{RAM}}$ with program SRAMSIM exists, which simulates that RAM machine in $\mathcal{O}(S)$ space and $\mathcal{O}(T)$ time, since P is a constant and S is $\mathcal{O}(T)$. Now, because of Lemma 5.9, executing SRAMSIM in the K. U. Leuven CHR system also takes $\mathcal{O}(T)$ time and $\mathcal{O}(S)$ space. $\square$

$$\begin{array}{ll}
i(L, \mathbf{init}, A), m(A, B), m(B, -) \setminus c(L) & \Longleftrightarrow c(L+1). \\
i(L, \mathbf{init}, A), m(A, B) \setminus c(L) & \Longleftrightarrow m(B, 0), c(L+1).
\end{array}$$

Fig. 9. Alternative implementation of the **init** instruction, using the ω_r semantics.

One may expect to pay *some* performance penalty for using a very high-level language like CHR, so it is comforting that at least we can always get the asymptotic complexities right.

5.4 Discussion: Register Initialization

Our definition of the space complexity of a RAM machine (see Section 4.2.2) is based on Savitch [1978]. It counts all registers in the used address range, whether or not each individual register was effectively used. In the literature, other definitions of RAM machine space complexity only take the used registers into account [van Emde Boas 1990]. If we would use such a definition, the preceding results no longer hold in general: For a program that uses the registers in a sparse way, the RAM-machine simulator SRAMSIM of Figure 8 would use more space (and thus possibly also more time) than the original RAM machine. The reason is that the simulator initializes the entire register range.

However, by relying on the refined operational semantics, we can implement the **init** instruction in a different way; see Figure 9. The auxiliary constraints **maxm**/1 and **initm**/3 are not needed in this version of the simulator. Of course, this alternative program is no longer confluent: It depends on the rule application order enforced by the refined semantics for correctness. Indeed, under ω_r semantics the second rule is applied only if the first rule cannot be applied because of the absence a corresponding **m**/2 constraint. It is not known whether checking for absence of constraint is possible in a confluent way [Van Weert et al. 2006].

In the rest of this article we will assume, for simplicity, that registers are initialized exactly once before use. In that case, we can simply use just the second rule of Figure 9, without sacrificing confluence.

5.5 Time Complexity

THEOREM 5.11. *The K. U. Leuven CHR system can simulate the standard-arithmetic RAM machine with the same time complexity.*

5.6 Space Complexity

THEOREM 5.12. *The K. U. Leuven CHR system can simulate the standard-arithmetic RAM machine with the same space complexity.*

5.7 Everything can be Done in CHR

COROLLARY 5.13. *For every algorithm, a CHR program exists which can be executed in the K. U. Leuven CHR system with the correct complexity.*

8:26 • J. Sneyers et al.

movl \$1, %eax	m(1,1),
movl \$1000000, %ecx	m(2,1000000),
movl \$0, %edx	m(3,0),
.L1: addl %eax, %edx	i(1, add , 1, 3),
subl %eax, %ecx	i(2, sub , 1, 2),
je .L5	i(3, cjmp , 2, 5),
jmp .L1	i(4, jmp , 1),
.L5:	i(5, halt),
	c(1).

Fig. 10. Assembler code and corresponding RAM machine query.

6. CONSTANT FACTORS AND EXPERIMENTAL RESULTS

In the previous sections we have shown that every algorithm can be implemented in CHR with the right asymptotic time and space complexity. However, the constant factors hidden behind asymptotic complexities could be huge and completely paralyzing in practice. In this section we investigate these constants experimentally.

6.1 Practical Relevance of the General Complexity Result

In principle, every algorithm can be implemented in CHR using the RAM simulator program. Of course, this does not result in a natural and elegant CHR program, but at least the resulting CHR program has the right time and space complexity.

Consider the following very simple C program.

```
long a=1, b=1000000, c=0;
while(b != 0) {
  c += a;
  b -= a;
}
```

This C program corresponds to the Intel assembler code shown in Figure 10 (on the lefthand side). The assembler code can also be seen as a query for the RAM machine program RAMSIMUL, as shown in the righthand side of Figure 10.

By translating assembler code to a RAM simulator query, we get a CHR program with the same asymptotic time and space complexity: Both the CHR program and the assembler code take linear time and constant space. Figure 19 lists this and some other examples of RAM programs.

Although the CHR (Prolog) RAM simulator executes such RAM programs with the correct asymptotic complexity, the execution time is about ten thousand times larger than that of the original assembler code program: The RAM simulator uses about 10 seconds while the assembler program uses 1.6 milliseconds. In other words, the computational power of a Pentium 4 is reduced to that of a Commodore 64.

Of course no sane programmer would write CHR programs in this way, and not just because of the debilitating slowdown: Such programs also lack

Table II. Different Levels of Type Information Detail in CHR Constraint Declarations

Level	Constraint Declaration
0	<code>:- chr_constraint root/2, ;/2, make/1, union/2, link/2, find/2.</code>
1	<code>:- chr_constraint root(+,+), + ~+ , make(+), union(+,+), link(+,+), find(+,-).</code>
2	<code>:- chr_constraint root(+dense_int,+int), +dense_int ~+int, make(+int), union(+int,+int), link(+int,+int), find(+int,-int).</code>

desirable properties of CHR programs (conciseness, readability, adaptability, incrementality, concurrency, etc.) that are often obtained naturally in hand-written CHR programs. Hence it remains necessary to manually construct CHR programs.

6.2 Experimental Evaluation

In the rest of this section we investigate the performance of hand-written CHR programs that implement two classical algorithms in an elegant way. In particular, we compare their performance to that of an efficient reference implementation, in the low-level language C, of the same algorithms. The goal is to obtain an estimate of the (constant factor) performance penalty for using CHR.

6.2.1 Setup. The execution times listed in the tables in this section were obtained as follows. We used GCC version 3.3.5 with the `-O4` option to compile the C programs, hProlog⁷ version 2.4.39-32 with the K. U. Leuven CHR compiler by Schrijvers and Demoen [2004], SICStus Prolog⁸ version 3.12.2 (x86-linux-glibc2.2) with CHR version 2.2 by Holzbaur and Frühwirth [1998], and JCHR⁹ version 1.3.3 by Van Weert et al. [2005], executed in Java 1.5.0.03 with the “server” virtual machine and background compilation disabled (options `-server -Xbatch`). For the C programs, we listed the user cpu time. For the CHR programs in Prolog systems, we measured user cpu time, not including garbage collection time. For the JCHR programs, we measured clock time. We made sure the test machine had a very low load, the Java heap size was set as big as possible while still fitting in RAM, and we took the best time of five runs. We tested on a Pentium 4 machine with 512MB of RAM and a cpu with a clock speed of 1.7GHz (cache size 256KB), running Debian GNU/Linux version 3.1 (sarge) with Linux kernel 2.6.15.

The K. U. Leuven CHR system allows the programmer to specify optional type information for the constraint arguments. This optional information can be declared inside the obligatory constraint declarations. The information is used for optimizing compilation. Table II illustrates the three levels of detail we consider. Detail level 2 corresponds to precise type information, level 1 corresponds to coarser groundness information, level 0 gives no information about the constraint arguments at all. The original CHR system in SICStus Prolog, by Holzbaur and Frühwirth [1998], does not have a mechanism to provide type

⁷hProlog home page: <http://www.cs.kuleuven.be/~bmd/hProlog/>.

⁸SICStus Prolog home page: <http://www.sics.se/sicstus/>.

⁹JCHR home page: <http://www.cs.kuleuven.be/~petervw/JCHR/>.

8:28 • J. Sneyers et al.

```

make(A)                 $\iff$  root(A,0).
union(A,B)               $\iff$  find(A,X), find(B,Y), link(X,Y).
A  $\leadsto$  B,              find(A,X)       $\iff$  find(B,X), A  $\leadsto$  X.
root(B,-) \ find(B,X)     $\iff$  X=B.
find(-,-)                $\iff$  fail.
link(A,A)                $\iff$  true.
link(A,B), root(A,R), root(B,S)  $\iff$  R  $\geq$  S | B  $\leadsto$  A, root(A,max(R,S+1)).
link(B,A), root(A,R), root(B,S)  $\iff$  R  $\geq$  S | B  $\leadsto$  A, root(A,max(R,S+1)).
link(-,-)                $\iff$  fail.
    
```

Fig. 11. The union-find algorithm in CHR.

Table III. Execution Times (in seconds) for Different Implementations of the Union-Find Algorithm

n	C	CHR(hProlog)			CHR(SICStus)	CHR(Java)
	2	2	1	0	0	2
1k	<0.01	0.01	0.02	1.46	2.14	0.01
4k	<0.01	0.03	0.07	32.94	33.12	0.03
8k	<0.01	0.06	0.14	112.19	119.21	0.07
16k	<0.01	0.12	0.27	367.93	465.33	0.14
64k	0.02	0.53	1.17	> 1000	> 1000	1.36
256k	0.24	2.31	5.03			7.70
1m	1.12	9.91	21.60			64.82

On n elements, n random unions are followed by n random finds.

information. In the Java CHR system, precise type declarations are obligatory, since Java is a typed language.

6.2.2 Union-Find. The classic union-find algorithm of Tarjan and van Leeuwen [1984] efficiently implements the disjoint set union operation. It is not clear whether this algorithm can be implemented with optimal complexity in other declarative languages like pure (side-effect-free) Prolog. Schrijvers and Frühwirth [2006] have shown that CHR (Prolog) allows an elegant implementation which has the optimal almost-linear time complexity. It is listed in Figure 11.

Table III lists the execution times for the CHR program in the different CHR systems. We compare these results against a very efficient C implementation.¹⁰ In order to achieve the optimal complexity, type detail level 1 is needed. The high-level CHR(hProlog) implementation is roughly 10 times slower than the corresponding direct low-level implementation in C with type detail level 2. The results are plotted in Figure 12.

To get an idea of the constant factors involved in space usage, consider the following numbers. The C program uses only one array to represent the data structure; every element takes one word (4 bytes): Positive integers represent the index of the parent, negative integers represent the rank of a root. In contrast, the CHR(hProlog) program uses 9 words to represent an element: Two arrays are used (one for roots, one for nonroots), which contain pointers to

¹⁰Written by Ariel Faigon, based on a version by Robert Sedgewick. The source code is available at <http://www.yendor.com/programming/minauto/ufind.c>.

Fig. 12. Plot of the results of the Union-find benchmark (refer to Table III).

Table IV. Execution Times (in seconds) for Different Implementations of Dijkstra's Algorithm with Fibonacci Heaps

n	C	CHR(hProlog)			CHR(SICStus)	CHR(Java)
	2	2	1	0	0	2
1k	<0.01	0.05	0.13	3.47	6.42	0.17
4k	0.01	0.26	0.57	66.24	96.89	0.65
8k	0.02	0.50	1.13	229.88	373.48	1.53
16k	0.07	1.07	2.33	717.92	> 1000	3.77
64k	0.42	4.56	9.85	> 1000		stack overflow
256k	2.06	19.85	41.12			

On random sparse graphs with n nodes and $4n$ edges.

7-word suspension terms (refer to Schrijvers [2005]: one word for the wrapper functor, two for the constraint arguments, one for the identifier, three for the state). Hence for the union-find algorithm, the CHR version uses about ten times as much space as the C version.

6.2.3 Dijkstra's Algorithm with Fibonacci Heaps. The single-source shortest path algorithm of Dijkstra [1959] can be implemented efficiently using the Fibonacci heaps data structure of Fredman and Tarjan [1987]. An implementation of this algorithm in CHR is described in Sneyers et al. [2006a], where its performance in various CHR systems is also compared to that of a C implementation¹¹ by Cherkassky et al. [1996]. The results are listed in Table IV and plotted in Figure 13.

¹¹The source code is available at: <http://www.avglab.com/andrew/soft.html>.

8:30 • J. Sneyers et al.

Fig. 13. Plot of the results of the Dijkstra benchmark (refer to Table IV).

Note the resemblance between Figure 13 and Figure 12. Again, type detail level 1 is needed to achieve optimal complexity. Without any information, the program exhibits a quadratic time complexity, because the default data structure does not allow constant time lookups of ground constraints. When groundness information is available, the optimal $\mathcal{O}(n \log n)$ time complexity is achieved.

The time gap between the hProlog CHR program (type detail level 2) and the C program is, again, a constant factor of about 10.

The difference in space usage is less pronounced than for the union-find algorithm. The CHR program uses about three times more space than the C program: For an input size of 256k nodes, the C program uses 23 megabytes, while the CHR(hProlog) program needs 63 megabytes.

6.3 Conclusion: CHR(Prolog) / C $\approx$ 10

The aforementioned results indicate that the constant time factor separating CHR(Prolog) from C is approximately 10. In terms of space usage, the constraint representation has a fixed overhead: In programs using very lightweight data representations (e.g., union-find), this results in a relatively large constant space factor (e.g., 10); in programs with more complicated representations (e.g., Fibonacci heaps), the constant space factor is smaller (e.g., 3). Future and ongoing work in CHR compiler optimization will further reduce these factors. In particular, the space overhead can be much reduced by further specializing the constraint representation.

Extrapolating from the aforesaid examples, combined with the results of the previous sections, we quite confidently state our main thesis.

Conjecture 6.1. The current state-of-the-art in CHR and Prolog systems suffices to implement any algorithm in CHR (Prolog), in a natural and high-level way, with a time and space complexity which is within a constant factor of 10 from the best-known implementation in any other language.

We expect that a CHR(C) compiler which incorporates state-of-the-art optimizations could push this constant factor significantly closer to 1. The implementation of a CHR(C) compiler is ongoing work [Wuille et al. 2007].

7. SIMILAR RESULTS FOR OTHER LANGUAGES

Among the Turing-complete languages, many have the property we have shown for $X = \text{CHR(hProlog)}$: “[E]very algorithm can be implemented in language X with the right time and space complexity.” For instance, in all imperative languages that we are aware of, it is a straightforward exercise to construct a RAM machine simulator and show that it has the right complexity. After all, the basic ingredients needed for a RAM machine simulator are directly available in most imperative languages.

However, for higher-level declarative languages, the property is far less trivial. The time and space complexity of a program depends more crucially on an optimizing compiler. Whether or not the general complexity result holds for some language depends largely on the properties of its compiler: A pathological compiler could conceivably detect the pattern of a RAM machine simulator program as a special case and produce special, hardwired output with the desired complexity properties.

In this section we briefly investigate whether some other declarative languages allow an efficient implementation of a RAM machine simulator, given the current state-of-the-art. To keep this section concise, we will only consider one well-known language for some declarative paradigms.

7.1 Sufficient Ingredients

An efficient RAM machine simulator can be implemented if the following ingredients can be implemented:

- (1) iteration which iterates n times in $\mathcal{O}(nT_s)$ time and $\mathcal{O}(S_s)$ space, where T_s and S_s are the time and space needed to evaluate the stop condition of the iteration;
- (2) the arithmetic operations, with the same complexity as the corresponding RAM machine arithmetic operations;
- (3) constant-time, zero-space *if-then-else* and evaluation of (in)equality conditions (like $X == \text{add}$ and $X \neq 0$); and
- (4) growing arrays which allow n insertions, g lookups, and s updates in $\mathcal{O}(n + g + s)$ time and $\mathcal{O}(n)$ space.

Most declarative languages do not offer iteration (the first ingredient) as a basic language construct, but many implementations convert tail recursion to iteration. The second and third ingredients are directly available in the languages we consider.

8:32 • J. Sneyers et al.

```
:- use_module(assoc).
eval(L,Prog,I,0) :- X is L, get_assoc(L,Prog,Instr),
                    (i(Instr,I,I2,X,L1) -> eval(L1,Prog,I2,0) ; 0 = I).
i(add(B,A),I,0,L,L+1) :- get_assoc(B,I,Y), get_assoc(A,I,X,0,Z), Z is X+Y.
i(sub(B,A),I,0,L,L+1) :- get_assoc(B,I,Y), get_assoc(A,I,X,0,Z), Z is X-Y.
...
i(mvi(B,A),I,0,L,L+1) :- get_assoc(B,I,X), get_assoc(A,I,C), get_assoc(C,I,_,0,X).
i(init(A),I,0,L,L+1) :- get_assoc(A,I,B), put_assoc(B,I,0,0).
i(jmp(A),I,I,_,A).
i(cjmp(B,A),I,I,L,L1) :- (get_assoc(B,I,0) -> L1 = A ; L1 is L+1).
```

Fig. 14. A simulator of RAM machines, written in Prolog.

7.2 Arrays in Declarative Languages

The remaining ingredient, an efficient dynamically growing array, is the one that seems to be the most difficult to implement.

7.2.1 Logic Programming Languages. We use the term *pure Prolog* to denote the Prolog language as described by Clocksin and Mellish [1984], without the *assert* and *retract* built-ins. Clearly, if nonpure Prolog extensions (i.e., global variables, mutable terms, *assert/retract*) are allowed, there is a RAM machine simulator implementation: Consider, for instance, the code the CHR(hProlog) compiler generates for SRAMSIM. To our knowledge, there is no Prolog system which allows an efficient pure Prolog implementation of dynamically growing arrays. Association lists, available in many Prolog systems as a standard module called *assoc*, can be used instead. The implementation of *assoc* that is used in hProlog is based on an implementation by Mats Carlsson (which was based on an implementation by Richard O’Keefe) based on AVL-trees [Adelson-Velsky and Landis 1962]. Lookup, insertion, and update take $\mathcal{O}(\log n)$ worst-case time. Figure 14 lists a Prolog program which uses association lists to implement a RAM simulator. We used the same version of hProlog as the one we used for testing CHR.

Mercury [Somogyi et al. 1996] is a strongly typed high-performance logic programming language. The Mercury system includes the array module, but the procedures in this module are written in the target languages (C, C#, and Java). In an experimental development branch of Mercury, compile-time garbage collection (CTGC) has been added by Mazur [2004]. This allows automatic *structure reuse* in a large class of Mercury programs. Perhaps CTGC allows a reasonably efficient pure Mercury implementation of growing arrays: Using AVL-trees and with in-place updates thanks to CTGC, it seems feasible to perform n insertions, g lookups, and s updates in $\mathcal{O}((n + g + s) \log n)$ time and $\mathcal{O}(n)$ space. We have not tested this experimentally, since CTGC is not yet available in the main release of the Mercury system.

7.2.2 Functional Programming Languages. Haskell [Hudak et al. 2007] is a modern typed, lazy, purely functional language. Most Haskell systems include the *Data.Array* module in their standard libraries. This module efficiently implements arrays, but it is not implemented in Haskell itself. The fastest pure


```

module Ramsimul where
import Data.IntMap
data Instr = Add Int Int | Sub Int Int | ... | Cjmp Int Int | Halt
type Program = IntMap Instr
type Memory = IntMap Int
eval :: Int -> Program -> Memory -> Memory
eval pc prog mem =
  case (prog!pc) of
    Add rx ry  -> eval (pc + 1) prog (insert ry ((mem!ry) + (mem!rx)) mem)
    Sub rx ry  -> eval (pc + 1) prog (insert ry ((mem!ry) - (mem!rx)) mem)
    ...
    Mvi rx ry  -> eval (pc + 1) prog (insert (mem!ry) (mem!rx) mem)
    Init rx    -> eval (pc + 1) prog (insert (mem!rx) 0 mem)
    Jmp l      -> eval l prog mem
    Cjmp r l   -> eval (if (mem!r) == 0 then l else pc + 1) prog mem
    Halt      -> mem

```

Fig. 15. A simulator of RAM machines, written in Haskell.

Haskell implementation of arrays we could find is available in the standard `Data.IntMap` module, which is based on an implementation of Patricia trees [Morrison 1968] by Okasaki and Gill [1998]. This data structure allows memory cell lookups, updates, and insertion (initialization) in $\mathcal{O}(\min(n, W))$ time, where W is the number of bits in an `Int`. On our test platform, $W = 32$ so the operations can be considered to be constant time. However, since the updates are not done in-place, the space complexity is $\mathcal{O}(n + s)$ instead of $\mathcal{O}(n)$. Figure 15 lists the Haskell program we have tested. We used the Glasgow Haskell Compiler¹² version 6.6 with option “-O2”. We have tested both a “lazy” version (listed in Figure 15) and a “strict” version. The latter naively forces all lazy thunks immediately to weak head normal form (WHNF); it differs from the lazy version on two accounts only. Firstly, the fields of the `Instr` datatype are declared to be strict. Secondly, each function application $f e$ is transformed into `let x = e in x 'seq' f x`, which forces the subexpression e to WHNF before evaluating the main expression $f e$.

7.2.3 Term-Rewrite Systems. Maude¹³ [Clavel et al. 2002] is a system for declarative programming in rewriting logic. It features efficient rewriting of terms with associative-commutative (AC) operators using the *stripper-collector matching* algorithm of Eker [2003]. Figure 16 lists the Maude program we have tested (in Maude version 2.2). This program is directly derived from the CHR rules. As in CHR, the collection data structures and the operations on them are implicit. In this sense CHR and Maude are higher-level languages than Prolog and Haskell.

Unfortunately, for the rules listed in Figure 16, the current implementation of Maude is not able to use its most efficient matching algorithm. By making the data structure operations more explicit (using the `Map{Int,Int}` module) we obtain a more efficient program. It is listed in Figure 17.

¹²GHC home page: <http://www.haskell.org/ghc/>.

¹³Maude home page: <http://maude.cs.uiuc.edu/>.

8:34 • J. Sneyers et al.

```

mod RAMSIMUL is protecting INT .
  op c : Int -> State .
  op m : Int Int -> State .
  op _ : State State -> State [assoc comm] .
  ops add sub ... cjmp halt : -> Instr .
  rl c(L) i(L,add,B,A) m(B,Y) m(A,X) => c(L + 1) i(L,add,B,A) m(B,Y) m(A,X + Y) .
  rl c(L) i(L,sub,B,A) m(B,Y) m(A,X) => c(L + 1) i(L,sub,B,A) m(B,Y) m(A,X - Y) .
  ...
  rl c(L) i(L,init,A) m(A,B) => c(L + 1) i(L,init,A) m(A,B) m(B,0) .
  rl c(L) i(L,jmp,A) => c(A) i(L,jmp,A) .
  rl c(L) i(L,cjmp,B,A) m(B,0) => c(A) i(L,cjmp,B,A) m(B,0) .
  crl c(L) i(L,cjmp,B,A) m(B,Y) => c(L + 1) i(L,cjmp,B,A) m(B,Y) if Y /= 0 .
  rl c(L) i(L,halt) => i(L,halt) .
endm

sorts State Instr .
op i : Int Instr Int Int -> State .
op i : Int Instr Int -> State .
op i : Int Instr -> State .
vars B A C L Y X : Int .

```

Fig. 16. A simulator of RAM machines, written in Maude.

```

mod RAMSIMUL2 is protecting INT .
  op i : Instr Int Int -> Int .
  op i : Instr Int -> Int .
  op i : Instr -> Int .
  ops add sub ... cjmp halt : -> Instr .
  op machine : Int Map{Int,Int} Map{Int,Int} -> Map{Int,Int} .
  op execute : Int Int Map{Int,Int} Map{Int,Int} -> Map{Int,Int} .
  eq machine(L,P,M) = execute(P[L],L,P,M) .
  eq execute(i(add,B,A),L,P,M) = machine(L + 1,P,insert(A,(M[A])+(M[B]),M)) .
  ...
  eq execute(i(init,A),L,P,M) = machine(L + 1,P,insert(M[A],0,M)) .
  eq execute(i(jmp,A),L,P,M) = machine(A,P,M) .
  ceq execute(i(cjmp,B,A),L,P,M) = machine(A,P,M) if (M[B]) = 0 .
  eq execute(i(cjmp,B,A),L,P,M) = machine(L + 1,P,M) [owise] .
  eq execute(i(halt),L,P,M) = M .
endm

protecting MAP{Int,Int}.
sorts Instr .
vars B A C L Y X : Int .
vars P M : Map{Int,Int} .

```

Fig. 17. A more efficient RAM machine simulator in Maude, using explicit lookups and updates.

7.2.4 Rule Engines. Jess¹⁴ [Friedman-Hill 2003] is considered one of the fastest rule engines. Like its ancestor CLIPS [Giarratano and Riley 1994], it uses the RETE algorithm of Forgy [1982]. Figure 18 lists the Jess program we have tested (in Jess version 7.0).

In a sense, Jess is higher level than Prolog and Haskell because the data structures are implicit. It is lower level than CHR and Maude (as in the Maude program of Figure 16) because the data structure *operations* (assert, retract, and modify) are explicit. Moreover, as far as we know, Jess does not have the join reordering optimization. We have picked the best possible order of rule heads in the simulator program of Figure 18. If the heads are written in a different order, performance will suffer. This is another sense in which Jess can be considered to be a lower-level language than CHR; thanks to automatic join ordering, CHR programmers do not have to worry about the order of the heads in multiheaded rules.

¹⁴Jess home page: <http://www.jessrules.com/>.

```
(deftemplate c (slot at))
(deftemplate m (slot at) (slot val))
(deftemplate i2 (slot at) (slot instr))
(deftemplate i3 (slot at) (slot instr) (slot op))
(deftemplate i4 (slot at) (slot instr) (slot op1) (slot op2))
(defrule add
  ?C <- (c (at ?L))          (i4 (at ?L) (instr add) (op1 ?B) (op2 ?A))
  ?M <- (m (at ?A) (val ?X))  (m (at ?B) (val ?Y))
  => (modify ?M (val (+ ?X ?Y))) (modify ?C (at (+ ?L 1)))
)
...
(defrule ini
  ?C <- (c (at ?L))          (i3 (at ?L) (instr init) (op ?A))
  (m (at ?A) (val ?B))
  => (assert (m (at ?B) (val 0))) (modify ?C (at (+ ?L 1)))
)
(defrule j
  ?C <- (c (at ?L))          (i3 (at ?L) (instr jmp) (op ?A))
  => (modify ?C (at ?A))
)
(defrule do_cjump
  ?C <- (c (at ?L))          (i4 (at ?L) (instr cjmp) (op1 ?R) (op2 ?A))
  (m (at ?R) (val 0))
  => (modify ?C (at ?A))
)
(defrule no_cjump
  ?C <- (c (at ?L))          (i4 (at ?L) (instr cjmp) (op1 ?R) (op2 ?A))
  (m (at ?R) (val ~0))
  => (modify ?C (at (+ ?L 1)))
)
(defrule halt
  ?C <- (c (at ?L))          (i2 (at ?L) (instr halt))
  => (retract ?C)
)
```

Fig. 18. A simulator of RAM machines, written in Jess.

In general, the eager matching RETE algorithm is asymptotically slower than the lazy matching LEAPS algorithm [Miranker et al. 1990] on which the CHR execution mechanism is based. In the specific case of the RAM machine simulator program of Figure 18, this disadvantage of RETE does not emerge.

7.3 Experimental Results

Table V lists the execution times of running the RAM simulator benchmarks of Figure 19 in different RAM simulator implementations. Garbage collection complicates accurate measurement of memory usage, which is why no memory results are given. However, the runtimes include time spent in garbage collection. In this way, space complexity is taken into account. Some benchmarks run out of memory (indicated by “mem”) when the maximal heap size is set to a value slightly lower than the amount of memory available on the test machine (512Mb).

8:36 • J. Sneyers et al.

Table V. Execution Times (in seconds) for the RAM Simulator Benchmarks of Figure 19

Bench	n	CHR (Prolog)			Prolog	Haskell		Maude		Jess
		2	1	0		Lazy	Strict	Naive	Map	
Loop	10^4	0.09	0.32	0.32	0.05	0.05	0.02	1.01	0.37	3.48
	10^5	0.88	3.23	3.10	0.44	0.41	0.21	10.09	3.48	24.07
	10^6	8.70	32.42	29.93	4.46	5.40	2.05	100.43	35.32	231.13
	10^7	86.95	324.91	298.11	42.88	mem	20.41	987.28	371.54	2266.33
MFib	10^3	0.03	0.07	2.20	0.05	0.04	0.02	32.00	0.24	3.08
	10^4	0.26	0.69	434.09	0.60	1.19	0.16	4772.70	3.69	10.22
	10^5	2.61	7.39	—	7.40	mem	1.82	—	52.12	98.72
	10^6	26.76	76.43	—	89.37	—	19.69	—	1650.49	mem
NLoop	2^7	0.18	0.44	4.02	0.26	0.17	0.08	57.05	1.08	4.75
	2^8	0.71	1.72	26.81	1.10	0.73	0.34	419.54	4.68	19.59
	2^9	2.79	6.86	192.55	4.83	9.43	1.16	3677.66	21.51	77.81
	2^{10}	11.16	27.57	1475.77	20.49	mem	4.82	—	111.13	295.20
	2^{11}	44.28	108.87	11620.42	90.04	—	20.37	—	530.60	1258.71
	2^{12}	178.22	434.48	—	380.78	—	85.27	—	2604.68	4997.12
	2^{13}	709.89	1744.55	—	1678.73	—	356.29	—	11474.24	20147.51
	2^{14}	2864.56	7050.31	—	7272.05	—	1500.36	—	51573.71	—

For different implementations of the RAM simulator.

i(1,add,1,3) i(2,sub,1,2) i(3,cjmp,2,5) i(4,jmp,1) i(5,halt) m(1,1) m(2,n) m(3,0) c(1)	i(1,init,3) i(2,imv,1,6) i(3,imv,2,7) i(4,mul,6,7) i(5,mvi,7,3) i(6,add,5,1) i(7,add,5,2) i(8,add,5,3) i(9,sub,5,4) i(10,cjmp,4,12) i(11,jmp,1) i(12,halt)	m(1,8) m(2,9) m(3,10) m(4,n) m(5,1) m(6,0) m(7,0) m(8,1) m(9,1) c(1)	i(1,mov,2,3) i(2,add,1,3) i(3,init,3) i(4,mvi,4,3) i(5,sub,4,3) i(6,mov,3,5) i(7,sub,1,5) i(8,cjmp,5,10) i(9,jmp,3) i(10,mov,2,3) i(11,add,1,3) i(12,imv,3,5) i(13,add,4,5) i(14,mvi,5,3) i(15,sub,4,3)	i(16,mov,3,5) i(17,sub,1,5) i(18,cjmp,5,20) i(19,jmp,12) i(20,sub,4,2) i(21,cjmp,2,23) i(22,jmp,10) i(23,halt) m(1,10) m(2,n) m(3,0) m(4,1) m(5,0) c(1)
(a) Loop	(b) MFib		(c) NLoop	

Fig. 19. Example RAM machine simulator queries, used for benchmarking.

The first benchmark, Loop, performs $\mathcal{O}(n)$ updates on only three memory cells. The space usage is constant in CHR, Prolog, Maude, and strict Haskell, but not in lazy Haskell: In this example, lazy evaluation creates $\mathcal{O}(n)$ lazy thunks so it needs $\mathcal{O}(n)$ space. The time complexity is linear in all systems. The Jess program is the slowest: It takes more than twice the time of the naive Maude program. The naive Maude program is roughly three times slower than the one that uses Map, which is about as fast as CHR. The CHR program with type detail level 2 is almost four times faster than the CHR programs with type detail level 1 or 0, about twice as slow as the Prolog version and four times slower than strict Haskell.

In the second benchmark, MFib, $\mathcal{O}(n)$ memory cells are used. Naive Maude and CHR without type information do not get the time complexity right: The

Fig. 20. Plot of the results of the NLoop benchmark (refer to Table V).

$\mathcal{O}(n)$ lookups seem to take $\tilde{\mathcal{O}}(n^2)$ time. Prolog, Haskell, and Maude with Map get the time complexity almost right: $\tilde{\mathcal{O}}(n)$ instead of $\mathcal{O}(n)$. Both Prolog and Haskell use too much space. In the case of Prolog and strict Haskell, garbage collection takes an increasingly higher proportion of the runtime. In lazy Haskell, garbage collection does not even prevent running out of memory. CHR with type detail level 1 and Jess get the right time and space complexity; Jess is more than ten times slower and also uses significantly more space.

The last benchmark, NLoop, does $\mathcal{O}(n^2)$ updates on $\mathcal{O}(n)$ memory cells. As expected from the results of the previous benchmark, only CHR with type detail level > 0 and Jess get the $\mathcal{O}(n^2)$ time complexity completely right. Jess is between 10 and 30 times slower than CHR. Prolog, (strict) Haskell, and Maude with Map achieve $\tilde{\mathcal{O}}(n^2)$ time complexity. Figure 20 shows a plot of the runtimes divided by the expected complexity $\mathcal{O}(n^2)$. If the resulting curve is horizontal, the expected complexity is achieved. Both in constant factors and unwanted nonconstant factors, the strict Haskell version is better than the Prolog version, which is in turn better than the Maude version.

In summary, we obtain the following findings.

- CHR without type declarations and naive Maude have the optimal space complexity, but they do not achieve the optimal time complexity.
- CHR with type detail level 1 or 2 achieves the optimal time and space complexity.
- Jess achieves the optimal time and space complexity, with a large constant factor.
- Prolog and strict Haskell have a time complexity which is within a polylogarithmic factor from optimal, and their space complexity is not optimal.

8:38 • J. Sneyers et al.

- Lazy Haskell does not get close to the optimal space complexity.
- Maude with Map has the optimal space complexity (as far as we can tell) and gets within a (large) polylogarithmic factor from the optimal time complexity.

8. DISCUSSION AND CONCLUSION

We have investigated the computational power and complexity of Constraint Handling Rules by introducing the CHR machine, a model of computation based on the operational semantics of CHR. Besides the expected result that CHR is a Turing-complete language, we have demonstrated the much stronger result that every RAM machine program can be implemented as a CHR program which has the same asymptotic time and space complexities if executed in the K. U. Leuven CHR system. In other words, the current state-of-the-art in CHR compilation allows CHR programmers to implement any algorithm with the best possible complexity. We have provided empirical evidence that (at least some) algorithms have a natural encoding in CHR, which is only ten times slower than a direct C implementation. As far as we know, CHR is the first declarative language for which the “optimal complexity” result can be demonstrated within the pure part of the language, namely without imperative extensions to the language.

8.1 Related Work

Earlier work by Frühwirth [2002, 2001] studied the time complexity of simplification rules, for naive implementations of CHR. In this approach, a suitable termination order (called a *tight ranking*) is used as an upper bound on the derivation length. Combined with a worst-case estimate of the number and cost of rule application attempts in a naive implementation, this results in a complexity metatheorem which gives a rough upper bound of the time complexity. For the RAM simulator program SRAMSIM simulating a T -time RAM machine, the upper bound predicted by Frühwirth [2002] is $\mathcal{O}(T^6)$, quite far from the $\mathcal{O}(T)$ bound of Lemma 5.9.

We have explicitly decoupled the two steps in the approach of Frühwirth [2002] by introducing the notion of a theoretical CHR machine. If suitable termination orders can be found, they can be used to show an upper bound on the complexity of the CHR machine. This is the first step. However, for programs that are nonterminating in general, like the RAM simulator, or for which no ranking can be found, other techniques have to be used to prove complexity properties. For example, we have shown that the complexity of the RAM simulating CHR machine is the same as the complexity of the simulated RAM program. The second step corresponds to the question of how efficiently a CHR machine can be executed in practice (i.e., on a RAM machine). Recent work on optimizing compilation of CHR [Schrijvers 2005; Duck 2005] has allowed us to achieve much tighter bounds.

Ganzinger and McAllester [2002] have introduced an abstract logic programming model of computation for which they have studied the time complexity properties. Their formalism is somewhat related to CHR extended with rule priorities, although their notion of deletion is different. It is not clear whether

all RAM machine programs can be implemented in their formalism with the same time and space complexity. Experimental results are not available, since there is no practical implementation of their formalism. An implementation is being developed, which works by translating the formalism to CHR with priorities [De Koninck et al. 2007].

8.2 Future Work

The TMSIM program (Figure 5) can also simulate nondeterministic Turing machines. However, it is only $\Delta_{\omega_i}^\Phi$ -deterministic for input that corresponds to a deterministic Turing machine. By dropping the determinism condition in the definition of CHR machines, a nondeterministic version of the CHR machine can be defined [Sneyers and Frühwirth 2008]. The nondeterminism can be “don’t care” (if rule application remains committed-choice) or “don’t know” (for a non-committed choice variant of CHR).

It is an open problem whether a result similar to the the linear speedup theorem [Hartmanis and Stearns 1965] can be demonstrated for CHR machines. To improve the time complexity of a CHR machine, one could try to reduce the number of **Apply** steps by combining rules, and the number of **Introduce** steps by combining constraints. It is not clear whether such a reduction is possible in general. This is somewhat related to partial evaluation techniques at the level of CHR source code. Preliminary results indicate speedups of up to two orders of magnitude for versions of the RAM simulator that were manually specialized for the input RAM program.

Although we are convinced that every algorithm can be implemented with an elegant CHR program, it remains a useful research topic to construct good CHR implementations of existing (or new!) algorithms.

The RAM simulator is a ground CHR program without propagation rules. In a sense, our result implies that nonground constraints (which may be triggered) and propagation rules (that require checking and maintaining a propagation history) are not strictly needed. However, since nonground constraints and propagation rules are widely used (especially in the traditional constraint solver programs), improving the complexity of their implementation is still very useful.

Our result implies that only the constant factors of CHR can be further improved. A promising approach to reduce these constants is to write a CHR(C) compiler which incorporates the compiler optimizations currently implemented in the K. U. Leuven CHR (Prolog) system. This is work in progress [Wuille et al. 2007].

ACKNOWLEDGMENTS

The authors thank G. J. Duck, P. Van Weert, L. De Koninck, and the anonymous reviewers for their helpful comments.

REFERENCES

ADELSON-VELSKY, G. M. AND LANDIS, E. M. 1962. An algorithm for the organization of information. *Doklady Akademii Nauk SSSR* 146, 263–266.

ACM Transactions on Programming Languages and Systems, Vol. 31, No. 2, Article 8, Pub. date: January 2009.

8:40 • J. Sneyers et al.

- AHO, A. V., HOPCROFT, J. E., AND ULLMAN, J. D. 1975. *The Design and Analysis of Computer Algorithms*. Addison-Wesley Longman, Boston.
- CHERKASSKY, B. V., GOLDBERG, A. V., AND RADZIK, T. 1996. Shortest paths algorithms: Theory and experimental evaluation. *Math. Program.* 73, 129–174.
- CLAVEL, M., DURÁN, F., EKER, S., LINCOLN, P., ET AL. 2002. Maude: Specification and programming in rewriting logic. *Theor. Comput. Sci.* 285, 2, 187–243.
- CLOCKSIN, W. F. AND MELLISH, C. S. 1984. *Programming in Prolog*. Springer.
- DE KONINCK, L., SCHRIJVERS, T., AND DEMOEN, B. 2007. The correspondence between the Logical Algorithms language and CHR. In *Proceedings of the 23rd International Conference on Logic Programming (ICLP'07)*, V. Dahl and I. Niemelä, Eds. Lecture Notes in Computer Science, vol. 4670. Springer, 209–223.
- DIJKSTRA, E. W. 1959. A note on two problems in connexion with graphs. *Numerische Math.* 1, 4, 269–271.
- DUCK, G. J. 2005. Compilation of Constraint Handling Rules. Ph.D. thesis, University of Melbourne, Victoria, Australia.
- DUCK, G. J. AND SCHRIJVERS, T. 2005. Accurate functional dependency analysis for Constraint Handling Rules. In *Proceedings of the 2nd Workshop on Constraint Handling Rules (CHR'05)*, T. Schrijvers and T. Frühwirth, Eds., 109–124.
- DUCK, G. J., STUCKEY, P. J., GARCÍA DE LA BANDA, M., AND HOLZBAUR, C. 2004. The refined operational semantics of Constraint Handling Rules. In *Proceedings of the 20th International Conference on Logic Programming (ICLP'04)*, B. Demoen and V. Lifschitz, Eds. Lecture Notes in Computer Science, vol. 3132, 90–104.
- DUCK, G. J., STUCKEY, P. J., AND SULZMANN, M. 2006. Observable confluence for Constraint Handling Rules. In *Proceedings of the 3rd Workshop on Constraint Handling Rules (CHR'06)*, T. Schrijvers and T. Frühwirth, Eds., 61–76.
- EKER, S. 2003. Associative-Commutative rewriting on large terms. In *Proceedings of the Rewriting Techniques and Applications (RTA'03)*, R. Nieuwenhuis, Ed. Lecture Notes in Computer Science, vol. 2706. Springer, 14–29.
- FORGY, C. L. 1982. Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artif. Intell.* 19, 1, 17–37.
- FREDMAN, M. L. AND TARJAN, R. E. 1987. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM* 34, 3, 596–615.
- FRIEDMAN-HILL, E. 2003. *Jess in Action: Java Rule-Based Systems*. Manning Publications.
- FRÜHWIRTH, T. 1992. Constraint simplification rules. Tech. rep. ECRC-92-18, European Computer-Industry Research Centre, Munich, Germany. July.
- FRÜHWIRTH, T. 1998. Theory and practice of Constraint Handling Rules. *J. Logic Program.* 37, 1–3, 95–138.
- FRÜHWIRTH, T. 2001. As time goes by II: More automatic complexity analysis of concurrent rule programs. In *Proceedings of the Quantitative Aspects of Programming Languages (QAPL'01)*, A. D. Pierro and H. Wiklicky, Eds. Electronic Notes in Theoretical Computer Science, vol. 59, 3. Elsevier, 185–206.
- FRÜHWIRTH, T. 2002. As time goes by: Automatic complexity analysis of concurrent rule programs. In *Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, D. Fensel et al., Eds. Morgan Kaufmann, 547–557.
- FRÜHWIRTH, T. 2009. *Constraint Handling Rules*. Cambridge University Press. To appear.
- FRÜHWIRTH, T. AND ABDENNADHER, S. 2003. *Essentials of Constraint Programming*. Cognitive Technologies. Springer.
- GANZINGER, H. AND MCALLESTER, D. A. 2002. Logical algorithms. In *Proceedings of the 18th International Conference on Logic Programming (ICLP'02)*. Lecture Notes in Computer Science, vol. 2401. Springer 209–223.
- GIARRATANO, J. C. AND RILEY, G. 1994. *Expert Systems: Principles and Programming*. PWS Publishing, Boston, MA.
- HARTMANIS, J. AND STEARNS, R. E. 1965. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.* 117, 285–306.
- HOLZBAUR, C. AND FRÜHWIRTH, T. 1998. CHR reference manual. Tech. rep. TR-98-01, Österreichisches Forschungsinstitut für Artificial Intelligence, Vienna, Austria.

The Computational Power and Complexity of Constraint Handling Rules • 8:41

- HOLZBAUR, C. AND FRÜHWIRTH, T. 1999. Compiling constraint handling rules into prolog with attributed variables. In *Proceedings of the 1st International Conference on Principles and Practice of Declarative Programming (PPDP'99)*, G. Nadathur, Ed. Lecture Notes in Computer Science, vol. 1702. Springer, 117–133.
- HOLZBAUR, C., GARCÍA DE LA BANDA, M., STUCKEY, P. J., AND DUCK, G. J. 2005. Optimizing compilation of constraint handling rules in HAL. *Theory Pract. Logic Program. SI Constraint Handling Rules 5*, 4-5, 503–531.
- HOPCROFT, J. E., MOTWANI, R., AND ULLMAN, J. D. 2001. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Longman, Boston, MA.
- HUDAK, P., HUGHES, J., JONES, S. P., AND WADLER, P. 2007. A history of Haskell: Being lazy with class. In *Proceedings of the 3rd ACM SIGPLAN Conference on History of Programming Languages (HOPL-III)*. ACM Press, 1–55.
- MARRIOTT, K. AND STUCKEY, P. J. 1998. *Programming with Constraints: An Introduction*. MIT Press, Cambridge, MA.
- MAZUR, N. 2004. Compile-Time garbage collection for the declarative language Mercury. Ph.D. thesis, K. U. Leuven, Leuven, Belgium.
- MIRANKER, D. P., BRANT, D. A., LOFASO, B., AND GADBOIS, D. 1990. On the performance of lazy matching in production systems. In *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI'90)*, T. Dietterich and W. Swartout, Eds. MIT Press, Boston, MA, 685–692.
- MORRISON, D. R. 1968. PATRICIA—Practical algorithm to retrieve information coded in alphanumeric. *J. ACM* 15, 4, 514–534.
- OKASAKI, C. AND GILL, A. 1998. Fast mergeable integer maps. In *Proceedings of the Workshop on ML*, 77–86.
- SAVAGE, J. E. 1998. *Models of Computation*. Addison-Wesley Longman, Boston, MA.
- SAVITCH, W. J. 1978. The influence of the machine model on computational complexity. In *Interfaces between Computer Science and Operations Research*, J. Lenstra et al., Eds. Mathematical Centre Tracts, vol. 99. Centre for Mathematics and Computer Science, Amsterdam, 1–32.
- 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 *Proceedings of the 1st Workshop on Constraint Handling Rules (CHR'04), Selected Contributions*, T. Frühwirth and M. Meister, Eds. Ulm, Germany.
- SCHRIJVERS, T. AND FRÜHWIRTH, T. 2006. Optimal union-find in constraint handling rules. *Theory Pract. Logic Program.* 6, 1–2.
- SNEYERS, J. 2008. Turing-complete subclasses of CHR. In *Proceedings of the 24th International Conference on Logic Programming (ICLP'08)*, M. García de la Banda and E. Pontelli, Eds. Lecture Notes in Computer Science, vol. 5366. Springer, 759–763.
- SNEYERS, J. AND FRÜHWIRTH, T. 2008. Generalized CHR machines. In *Proceedings of the 5th Workshop on Constraint Handling Rules (CHR'08)*, T. Schrijvers et al., Eds.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006a. Dijkstra's algorithm with Fibonacci heaps: An executable description in CHR. In *Proceedings of the 20th Workshop on Logic Programming (WLP'06)*.
- SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006b. Memory reuse for CHR. In *Proceedings of the 22nd International Conference on Logic Programming (ICLP'06)*, S. Etalle and M. Truszczynski, Eds. Lecture Notes in Computer Science, vol. 4079. Springer, 72–86.
- SNEYERS, J., VAN WEERT, P., DE KONINCK, L., AND SCHRIJVERS, T. 2009. As time goes by: Constraint handling rules—A survey of CHR research between 1998 and 2007. *Theory Pract. Logic Program. Submitted*.
- SOMOGYI, Z., HENDERSON, F., AND CONWAY, T. 1996. The execution algorithm of Mercury, An efficient purely declarative logic programming language. *J. Logic Program.* 29, 1-3, 17–64.
- TARJAN, R. E. AND VAN LEEUWEN, J. 1984. Worst-Case analysis of set union algorithms. *J. ACM* 31, 2, 245–281.
- TURING, A. M. 1936. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.* 2, 42, 230–265.
- VAN EMDE BOAS, P. 1990. Machine models and simulations. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, J. van Leeuwen, Ed. Elsevier.

8:42 • J. Sneyers et al.

- VAN WEERT, P., SCHRIJVERS, T., AND DEMOEN, B. 2005. K. U. Leuven JCHR: A user-friendly, flexible and efficient CHR system for Java. In *Proceedings of the 2nd Workshop on Constraint Handling Rules (CHR'05)*, T. Schrijvers and T. Frühwirth, Eds., 47–62.
- VAN WEERT, P., SNEYERS, J., SCHRIJVERS, T., AND DEMOEN, B. 2006. Extending CHR with negation as absence. In *Proceedings of the 3rd Workshop on Constraint Handling Rules (CHR'06)*, T. Schrijvers and T. Frühwirth, Eds. K. U. Leuven, Department of Computer Science. Venice, Italy, 125–140.
- WUILLE, P., SCHRIJVERS, T., AND DEMOEN, B. 2007. CCHR: The fastest CHR implementation, in C. In *Proceedings of the 4th Workshop on Constraint Handling Rules (CHR'07)*, K. Djelloul et al., Eds., 123–137.

Received December 2006; revised March 2008; accepted April 2008