

*FreeCHR – An Algebraic Framework for Constraint Handling Rules Embeddings**

SASCHA RECHENBERGER and THOM FRÜHWIRTH

*Institute of Software Engineering and Programming Languages,
Ulm University, Ulm, Germany*

(e-mails: sascha.rechenberger@uni-ulm.de, thom.fruehwirth@uni-ulm.de)

submitted 28 March 2024; revised 4 February 2025; accepted 21 March 2025

Abstract

We introduce the framework *FreeCHR* which formalizes the embedding of *Constraint Handling Rules* (CHR) into a host language, using the concept of *initial algebra semantics* from category theory. We hereby establish a high-level implementation scheme for CHR as well as a common formalization for both theory and practice. We propose a lifting of the syntax of CHR via an endofunctor in the category **Set** and a lifting of the very abstract operational semantics of CHR into *FreeCHR*, using the free algebra, generated by the endofunctor. We give proofs for soundness and completeness with its original definition. We also propose a first abstract execution algorithm and prove correctness with the operational semantics. Finally, we show the practicability of our approach by giving two possible implementations of this algorithm in Haskell and Python. Under consideration in *Theory and Practice of Logic Programming*.

KEYWORDS: embedded domain-specific languages, rule-based programming languages, constraint handling rules, operational semantics, initial algebra semantics

1 Introduction

Constraint Handling Rules (CHR) is a rule-based programming language that is designed to be embedded into a general-purpose language. Having a CHR implementation available enables software developers to solve problems in a declarative and elegant manner. Aside from the obvious task of implementing constraint solvers (Frühwirth (2006); De Koninck *et al.* (2006)), it has been used, for example, to solve scheduling problems (Abdennadher and Marte (2000)) implement concurrent and multi-agent systems (Thielscher (2002, 2005); Lam and Sulzmann (2006, 2007)). In general, CHR is ideally suited for any problem that involves the transformation of collections of data, as programs consist of a set of rewriting rules, hiding away the process of finding suitable candidates for rule application.

* This is an extended version of a paper presented at the 2023 RuleML+RR Conference. The authors are grateful to Anna Fensel, Ana Ozaki and Ahmet Soylu, the conference program co-chairs, for encouraging us to submit the paper to the journal *Theory and Practice of Logic Programming* and for their help in overseeing the review process.

Hereby, we get a purely declarative representation of the algorithm without the otherwise necessary boilerplate code.

The literature on CHR as a formalism consists of a rich body of theoretical work concerning CHR, including a rigorous formalization of its declarative and operational semantics (Frühwirth (2009); Sneyers *et al.* (2010); Frühwirth (2015)), relations to other rule-based formalisms (e.g., Betz (2007) and Gall and Frühwirth (2018)) and results on properties like confluence (Christiansen and Kirkeby (2015)).

Implementations of CHR exist for a number of languages, such as Prolog (Schrijvers and Demoen (2004)), C (Wuille *et al.* (2007)), C++ (Barichard (2024)), Haskell (Lam and Sulzmann (2007); Chin *et al.* (2003)), JavaScript (Nogatz *et al.* (2018)) and Java (Abdennadher *et al.* (2002); Van Weert *et al.* (2005); Ivanović (2013); Wibiral (2022)).

While the implementations adhere to the formally defined operational semantics, they are not direct implementations of a common formal model. Therefore, the two aspects of CHR (formalism and programming language) are not strictly connected with each other and there is hence no strict guarantee that the results on the formalism CHR are applicable on the programming language CHR. Although, such a strict connection is probably not entirely achievable (unless we define and use everything inside a proof assistant like *Coq* or *Agda*), it is desirable to have both formal definition and implementation as closely linked as possible. Additionally, to being able to directly benefit from theoretical results, implementors of CHR embeddings and users of the CHR language can also use the formally defined properties to define more profound tests for their software.

Another apparent issue within the CHR ecosystem is that many of the implementations of CHR are currently unmaintained. Although some of them are mere toy implementations, others might have been of practical use. One example is JCHR (Van Weert *et al.* (2005)) which would be a useful tool if it was kept on par with the development of Java, especially with modern build tools like *Gradle*. Having a unified formal model from which every implementation is derived could ease contributing to implementations of CHR as it provides a strict documentation and description of the system a priori. Also, different projects might even be merged which would prevent confusion due to multiple competing, yet very similar implementations, as it can be observed in the *miniKanren* ecosystem (e.g., there exist about 20 implementations of *miniKanren* dialects only for Haskell²).

A third major issue is that many implementations, like the aforementioned JCHR or CCHR (Wuille *et al.* (2007)), are implemented as external EDSL, this is, they rely on a separate compiler which translates CHR code into code of the host language. This makes it significantly less convenient to use and hence also less likely to be adopted into the standard distribution of the host language. This is somewhat demonstrated by the *K.U.Leuven CHR system* which is implemented as a library in Prolog and distributed as a standard package with *SWI-Prolog*,³ or by the library *core.logic*⁴ which implements *miniKanren* for the LISP dialect *Clojure*.

To solve the three issues stated above, we introduce the framework *FreeCHR* which formalizes the embedding of CHR, using *initial algebra semantics*. This common concept in functional programming is used to inductively define languages and their semantics

² <https://minikanren.org/#implementations>

³ <https://www.swi-prolog.org/pldoc/man?section=chr>

⁴ see <https://github.com/clojure/core.logic>

(Hudak (1998); Johann and Ghani (2007)). FreeCHR provides both a guideline and high-level architecture to implement and maintain CHR implementations across host languages and a strong connection between the practical and formal aspects of CHR. To our knowledge, the presented execution algorithm and its implementations, although simple, will be the first implementations of CHR for which there are formal proofs of correctness. Also, by FreeCHR-instances being internal embeddings, we get basic tooling like syntax highlighting and type-checking for free (Fowler and Parsons (2011)) and are able to implement CHR as a library, without the need for any external tooling like source-to-source compilers.

Ultimately, the framework shall serve a fourfold purpose by providing

- a general guideline on how to implement a CHR system in modern high-level languages,
- a guideline for future maintenance of FreeCHR instances,
- a common framework for both formal considerations and practical implementations
- and a framework for the definition and verification of general criteria of correctness.

In this work, we will give first formal definitions of FreeCHR, upon which we will build our future work. Section 2 will provide the necessary background and intuitions. Section 3 introduces the syntax and semantics of Constraint Handling Rules and generalizes them to non-Herbrand domains. Section 4 introduces the framework FreeCHR. Section 4.1 lifts the syntax of CHR programs to a **Set**-endofunctor and introduces the free algebra, generated by that functor, Section 4.2 lifts the *very abstract* operational semantics ω_a of CHR to the *very abstract* operational semantics ω_a^* of FreeCHR and Section 4.3 proves soundness and completeness of ω_a^* with respect to ω_a .

Section 5 introduces a simple execution algorithm for a formal instance with *very abstract* operational semantics and proves their correctness. This is the second main contribution of this paper, as it gives an example use of a proven correct guideline to implement FreeCHR.

This is finally demonstrated with the two case studies in Section 6. The section provides two implementations of the defined instance in Haskell and Python. The languages were chosen, as they can be considered to be on the opposite sides of a static-dynamic spectrum of programming languages. They will hence give a good intuition of how the host language flavors FreeCHR.

Finally, Section 7 discusses the limitations of FreeCHR as presented in this paper, Section 8 discusses related work, Section 9 will give an overview over planned and ongoing future work and Section 10 concludes the paper.

2 Preliminaries

2.1 Endofunctors and F -algebras

In this section, we want to introduce endofunctors and F -algebras. Both concepts are taken from category theory and will be introduced as instances in the category of sets **Set**.

We do not assume any previous knowledge of category theory, but to readers more interested in the topic in general we recommend Milewski (2019) as introductory literature.

2.1.1 Basic definitions

The *disjoint union* of two sets A and B

$$A \sqcup B = \{l_A(a) \mid a \in A\} \cup \{l_B(b) \mid b \in B\}$$

is the union of both sets with additional labels l_A and l_B added to the elements to keep track of the origin set of each element. We will also use the labels l_A and l_B as *injection* functions $l_A : A \rightarrow A \sqcup B$ and $l_B : B \rightarrow A \sqcup B$ which construct elements of $A \sqcup B$ from elements of A or B , respectively.

For two functions $f : A \rightarrow C$ and $g : B \rightarrow C$, the function

$$[f, g] : A \sqcup B \rightarrow C$$

$$[f, g](l(x)) = \begin{cases} f(x), & \text{if } l = l_A \\ g(x), & \text{if } l = l_B \end{cases}$$

is called a *case analysis* function of the disjoint union $A \sqcup B$. It can be understood as a formal analog to a **case-of** expression. Furthermore, we define two functions

$$f \sqcup g : A \sqcup B \rightarrow A' \sqcup B' \qquad f \times g : A \times B \rightarrow A' \times B'$$

$$(f \sqcup g)(l(x)) = \begin{cases} l_{A'}(f(x)), & \text{if } l = l_A \\ l_{B'}(g(x)), & \text{if } l = l_B \end{cases} \qquad (f \times g)(x, y) = (f(x), g(y))$$

which lift two functions $f : A \rightarrow A'$ and $g : B \rightarrow B'$ to the disjoint union and the Cartesian product, respectively.

2.1.2 Endofunctors

A **Set**-endofunctor⁵ F maps all sets A to sets FA and all functions $f : A \rightarrow B$ to functions $Ff : FA \rightarrow FB$, such that $F\mathbf{id}_A = \mathbf{id}_{FA}$ and $F(g \circ f) = Fg \circ Ff$, where $\mathbf{id}_X(x) = x$ is the identity function on a set X .⁶ A signature $\Sigma = \{\sigma_1/a_1, \dots, \sigma_n/a_n\}$, where σ_i are operators and a_i their arity, generates a functor

$$F_\Sigma X = \bigsqcup_{\sigma/a \in \Sigma} X^a \qquad F_\Sigma f = \bigsqcup_{\sigma/a \in \Sigma} f^a$$

with $X^0 = \mathbb{1}$ and $f^0 = \mathbf{id}_{\mathbb{1}}$, where $\mathbb{1}$ is a singleton set. Such a functor F_Σ models *flat* (this is, not nested) terms over the signature Σ .

Example 1.

The signature

$$\Gamma = \{\mathbf{0}/0, \oplus/2, \mathbf{1}/0, \otimes/2\}$$

models two constants $\mathbf{0}$ and $\mathbf{1}$, as well as two binary operators \oplus and \otimes . It generates the functor

$$F_\Gamma X = \mathbb{1} \sqcup X \times X \sqcup \mathbb{1} \sqcup X \times X$$

$$F_\Gamma f = \mathbf{id}_{\mathbb{1}} \sqcup f \times f \sqcup \mathbf{id}_{\mathbb{1}} \sqcup f \times f$$

⁵ Since we only deal with endofunctors in **Set** we will simply call them *functors*.

⁶ We will omit the index of \mathbf{id} if it is clear from the context.

2.1.3 F -algebras

Since an endofunctor F defines the syntax of terms, an evaluation function $\alpha : FA \rightarrow A$ defines the *semantics* of terms. We call such a function α , together with its *carrier* A , an F -algebra (A, α) .

If there are two F -algebras (A, α) and (B, β) and a function $h : A \rightarrow B$, we call h an F -algebra homomorphism if and only if $h \circ \alpha = \beta \circ Fh$, that is h preserves the structure of (A, α) in (B, β) when mapping A to B . In this case, we also write $h : (A, \alpha) \rightarrow (B, \beta)$.

A special F -algebra is the *free F -algebra* $F^* = (\mu F, \mathbf{in}_F)$, for which there is a homomorphism $\llbracket \alpha \rrbracket : F^* \rightarrow (A, \alpha)$ for any other algebra (A, α) . We call those homomorphisms $\llbracket \alpha \rrbracket$ F -catamorphisms. The functions $\llbracket \alpha \rrbracket$ encapsulate structured recursion on values in μF with the semantics defined by the function α which is itself only defined on flat terms. The carrier of F^* , with $\mu F = F\mu F$, is the set of inductively defined values in the shape defined by F . The function $\mathbf{in}_F : F\mu F \rightarrow \mu F$ inductively constructs the values in μF .

Example 2.

The initial F_Γ -algebra

$$(\mu F_\Gamma, [\mathbf{0}, \oplus, \mathbf{1}, \otimes])$$

can be used to construct the nested expressions like $\mathbf{1} \oplus (\mathbf{1} \otimes \mathbf{0})$ of μF_Γ . The F_Γ algebras

$$(2, [\text{false}, \vee, \text{true}, \wedge])$$

$$(\mathbb{N}_0, [0, +, 1, \cdot])$$

give to F_Γ the semantics of Boolean and arithmetic expressions, respectively. Using the catamorphism $\llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket : \mu F_\Gamma \rightarrow 2$, we can evaluate the expression from above as

$$\begin{aligned} & \llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket (\mathbf{1} \oplus (\mathbf{1} \otimes \mathbf{0})) \\ &= \llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket (\mathbf{1}) \vee \llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket (\mathbf{1} \otimes \mathbf{0}) \\ &= \text{true} \vee \llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket (\mathbf{1}) \wedge \llbracket [\text{false}, \vee, \text{true}, \wedge] \rrbracket (\mathbf{0}) \\ &= \text{true} \vee (\text{true} \wedge \text{false}) \\ &= \text{true} \vee \text{false} \\ &= \text{true} \end{aligned}$$

Analogously, if we use $\llbracket [0, +, 1, \cdot] \rrbracket : \mu F_\Gamma \rightarrow \mathbb{N}_0$, we get

$$\begin{aligned} & \llbracket [0, +, 1, \cdot] \rrbracket (\mathbf{1} \oplus (\mathbf{1} \otimes \mathbf{0})) \\ &= \llbracket [0, +, 1, \cdot] \rrbracket (\mathbf{1}) + \llbracket [0, +, 1, \cdot] \rrbracket (\mathbf{1} \otimes \mathbf{0}) \\ &= 1 + \llbracket [0, +, 1, \cdot] \rrbracket (\mathbf{1}) \cdot \llbracket [0, +, 1, \cdot] \rrbracket (\mathbf{0}) \\ &= 1 + (1 \cdot 0) \\ &= 1 + 0 \\ &= 1 \end{aligned}$$

Technically, the constants $\mathbf{0}$, $\mathbf{1}$, false, true, 0 and 1 are supposed to be functions from the singleton set $\mathbb{1}$ to the respective carrier. Since $\mathbb{1}$ only has one element, there is essentially no difference between such a function and the constant itself. We will hence use the constant directly.

2.2 Labeled transition systems

In this section, we will lay out our definition and notation of *labeled transition systems* and notions of soundness and completeness.

Definition 1 (Labeled transition system).

A *labeled transition system (LTS)* is a structure $\langle S, A, R \rangle$, where S is the set of states called the domain of the system, A is a set of labels and $R \in S \times A \times S$ a ternary transition relation.

For an LTS $\omega = \langle S, A, \mapsto \rangle$, we write $s \xrightarrow{a} s'$, if $(s, a, s') \in (\mapsto)$ and $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \xrightarrow{a_3} \dots \xrightarrow{a_n} s_n$ for $s_0 \xrightarrow{a_1} s_1 \wedge s_1 \xrightarrow{a_2} s_2 \wedge \dots \wedge s_{n-1} \xrightarrow{a_n} s_n$. Also, we will write $s_1 \xrightarrow{a}^* s_n$ if there is an $a \in A$, such that $s_1 \xrightarrow{a} s_2 \wedge \dots \wedge s_{n-1} \xrightarrow{a} s_n$, for $n \geq 0$.

Definition 2 (θ -Soundness).

Given two labeled transition systems $\omega = \langle S, A, \mapsto \rangle$ and $\omega' = \langle S, A', \mapsto' \rangle$. For function $\theta: A \rightarrow A'$, we call ω θ -sound with respect to ω' if and only if

$$s \xrightarrow{a} s' \implies s \xrightarrow{\theta(a)} s'$$

for all $s, s' \in S$ and $a \in A$ ⁷.

θ -soundness means that any transition defined by ω is also defined by ω' , when mapping labels with θ . This is typically used to verify a more constrained system (ω) against a more general one (ω').

Dually, there exists a concept of *completeness*.

Definition 3 (θ -Completeness).

Given two labeled transition systems $\omega = \langle S, A, \mapsto \rangle$ and $\omega' = \langle S, A', \mapsto' \rangle$. For a function $\theta: A \rightarrow A'$, we call ω θ -complete with respect to ω' if and only if

$$s \xrightarrow{a} s' \longleftarrow s \xrightarrow{\theta(a)} s'$$

for all $s, s' \in S$ and $a \in A$.

θ -completeness means that ω defines all transitions that are defined by ω' , but using only the mapped labels of ω .

We will also need a slightly relaxed notion of soundness.

Definition 4 (θ -Soundness up to repeated transition).

Given two labeled transition systems $\omega = \langle S, A, \mapsto \rangle$ and $\omega' = \langle S, A', \mapsto' \rangle$. For a function $\theta: A \rightarrow A'$, we call ω θ -sound up to repeated transition with respect to ω' if and only if

⁷ We will generally just use the symbol \mapsto for the relation if the respective transition system is clear from the context.

$$s \xrightarrow{a} s' \implies s \xrightarrow{\theta(a)^*} s'$$

for all $s, s' \in S$, $a \in A$ and $l \geq 0$.

This definition says that an LTS is also θ -sound with respect to another if we can accomplish a transition in one LTS with zero or more transitions in the other.

3 CHR over non-Herbrand domains

The first implementations of CHR were embedded into the logical programming language Prolog, where terms like `3+4` or `f(a,b,c)` are not evaluated, as is the case in most other programming languages, but interpreted as themselves. This is called the *Herbrand* interpretation of terms. Since we want to embed CHR in any programming language, we need to generalize the language to non-Herbrand interpretations of terms. We will formalize this, using initial algebra semantics.

3.1 Host language

We first define a *data type* in the host language. A data type determines the syntax and semantics of terms via a functor Λ_T and an algebra τ_T . The fixed point $\mu\Lambda_T$ contains terms that are inductively defined via Λ_T and the catamorphism (τ_T) evaluates those terms to values of T .

Definition 5 (Data types).

A data type is a triple $\langle T, \Lambda_T, \tau_T \rangle$, where T is a set, Λ_T a functor and (T, τ_T) a Λ_T -algebra.

We write $t \equiv_T t'$ for $t \in \mu\Lambda_T$ and $t' \in T$ if and only if $(\tau_T)(t) = t'$.

Example 3 (Boolean data type).

The signature

$$\Sigma_2 = \{(n \leq m)/0 \mid n, m \in \mathbb{N}_0\} \cup \{(n < m)/0 \mid n, m \in \mathbb{N}_0\} \cup \{\wedge/2, \text{true}/0, \text{false}/0\}$$

defines Boolean terms⁸. Σ_2 generates the functor

$$\Lambda_2 X = \mathbb{N}_0 \times \mathbb{N}_0 \sqcup \mathbb{N}_0 \times \mathbb{N}_0 \sqcup X \times X \sqcup \mathbf{1} \sqcup \mathbf{1}$$

the fixed point, $\mu\Lambda_2$, of which is the set of valid nested Boolean terms like $(0 < 4 \wedge 4 \leq 6)$. Let $\langle \mathbf{2}, \Lambda_2, \tau_2 \rangle$, with $\mathbf{2} = \{\text{true}, \text{false}\}$, be a data type. If we assume τ_2 to implement the usual semantics for Boolean terms and comparisons, $(0 < 4 \wedge 4 \leq 6)$ will evaluate as

$$(\tau_2)(0 < 4 \wedge 4 \leq 6) = (\tau_2)(0 < 4) \wedge (\tau_2)(4 \leq 6) = \text{true} \wedge \text{true} = \text{true}$$

For a set T , both Λ_T and τ_T are determined by the host language which is captured by the next definition.

⁸ We will generally overload symbols like *false*, *true*, \wedge , \neg , \dots , if their meaning is clear from the context.

Definition 6 (Host environment).

A mapping

$$\mathcal{L}T = \langle T, \Lambda_T, \tau_T \rangle$$

where $\langle T, \Lambda_T, \tau_T \rangle$ is a data type, is called a host environment.

A host environment is implied by the host language (and by the program, the CHR program is part of) and assigns to a set T a data type, effectively determining syntax and semantics of terms that evaluate to values of T .

3.2 Embedding CHR

With the formalization of our host environment, we can define the syntax and semantics of CHR.

Definition 7 (CHR programs).

CHR programs are sets of multiset-rewriting rules of the form

$$N @ K \setminus R \iff G \mid B$$

For a set C , called the domain of the program, for which there is a data type $\mathcal{L}C = \langle C, \Lambda_C, \tau_C \rangle$, $K, R \in \mathbf{list}C$ are called the kept and removed head, respectively. The backslash (\setminus) is used to separate the kept and removed part of the head. Either K or R must be non-empty. If either head part is empty, the separator is omitted as well.

$\mathbf{list}X = \bigcup_{i \in \mathbb{N}_0} X^i$ maps a set X to the set of lists over X , with $X^0 = \varepsilon$ being the empty sequence, and functions $f: A \rightarrow B$ to functions

$$\mathbf{list}f(a_1, \dots, a_n) = (f(a_1), \dots, f(a_n))$$

The optional $G \in \mu\Lambda_2$ is called the guard. If G is omitted, we assume $G \equiv_2 \text{true}$ and omit the pipe (\mid) which separates the guard from the body as well. $B \in \mathbf{mset}\mu\Lambda_C$ is called the body. The functor \mathbf{mset} maps sets X to the set of multisets over X and functions $f: X \rightarrow Y$ to functions

$$\mathbf{mset}f(\{a_1, a_2, \dots\}) = \{f(a_1), f(a_2), \dots\}$$

N is an optional name for the rule, which is generally used for debugging and tracing. The $@$ symbol is used to separate the name of the rule from the rule itself. It is also omitted if the rule is not given a name.

The members of the kept and removed head are matched against values of the domain C . The guard G is a term that can be evaluated to a Boolean value. The body B is a multiset over terms which can be evaluated to values of C . This includes any call of functions or operators which evaluate to Boolean, or values of C , respectively.

In literature, rules with an empty kept head are called *simplification* rules and rules with an empty removed head *propagation* rules. Rules which have both a removed and kept head are called *simpagation* rules.

Definition 7 corresponds to the *positive range-restricted ground* segment of CHR which is commonly used as the target for embeddings of other (rule-based) formalisms (e.g., colored Petri nets as by Betz (2007)) into CHR (Frühwirth, 2009, Chapter 6.2). *Positive*

means that the body of the rule contains only user constraints (this is, values from C) which guarantees that computations do not fail. *Range-restricted* means that instantiating all variables of the head (K and R) will ground the whole rule. This also maintains the *groundness* of the segment of CHR which requires that the input and output of a program are ground. \mathbf{PRG}_C denotes the set of all such programs over a domain C .

Example 4 (Euclidean algorithm).

The program $\text{GCD} = \{\text{zero}, \text{subtract}\}$ ⁹

$$\text{zero} @ 0 \Leftrightarrow \emptyset$$

$$\text{subtract} @ N \setminus M \Leftrightarrow 0 < N \wedge 0 < M \wedge N \leq M \mid M - N$$

computes the greatest common divisor of a collection of natural numbers. The first rule removes all zeros from the collection. For any pair of numbers N and M greater 0 and $N \leq M$, the second rule replaces M by $M - N$. Note that we omitted the kept head and guard of the zero rule.

Definition 8 (C-instances of rules).

For a positive range-restricted rule

$$R @ k_1, \dots, k_n \setminus r_1, \dots, r_m \Leftrightarrow G \mid B$$

with universally quantified variables v_1, \dots, v_l , and a data type $\mathcal{LC} = \langle C, \Lambda_C, \tau_C \rangle$, we call the set

$$\begin{aligned} \Gamma_C(R) = \{ & (R @ k_1\sigma, \dots, k_n\sigma \setminus r_1\sigma, \dots, r_m\sigma \Leftrightarrow G\sigma \mid \mathbf{mset}(\llbracket \tau_C \rrbracket)(B\sigma)) \\ & \mid \sigma \text{ instantiates all variables } v_1, \dots, v_l, \\ & k_1\sigma, \dots, k_n\sigma, r_1\sigma, \dots, r_m\sigma \in C, \\ & G\sigma \in \mu\Lambda_2, \\ & B\sigma \in \mathbf{mset}\mu\Lambda_C \} \end{aligned}$$

the C -grounding of R . Analogously, for a set \mathcal{R} of rules, $\Gamma_C(\mathcal{R}) = \bigcup_{R \in \mathcal{R}} \Gamma_C(R)$ is the C -grounding of \mathcal{R} . An element $r' \in \Gamma_C(R)$ (or $\Gamma_C(\mathcal{R})$ respectively) is called a C -instance of a rule $R \in \mathcal{R}$.

A C -instance (or grounding) is obtained, by instantiating all variables and evaluating the then ground terms in the body of the rule, using the Λ_C -catamorphism $\llbracket \tau_C \rrbracket$. The functor \mathbf{mset} is used to lift $\llbracket \tau_C \rrbracket$ into the multiset.

Example 5.

Given a body $\{M - N\}$ and a substitution $\sigma = \{N \mapsto 4, M \mapsto 6\}$, the body is instantiated like

$$\mathbf{mset}(\llbracket \tau_C \rrbracket)(\{M - N\}\sigma) = \mathbf{mset}(\llbracket \tau_C \rrbracket)(\{6 - 4\}) = \{\llbracket \tau_C \rrbracket(6 - 4)\} = \{2\}$$

With Example 5, we can also easily see that if we use a data type $\mathcal{L}\mu\Lambda_C = \langle \mu\Lambda_C, \Lambda_C, \mathbf{in}_{\Lambda_C} \rangle$ we get the Herbrand interpretation of terms over C . Hence, for instance, an expression $(3 + 4) \in \mu\Lambda_{\mathbb{N}_0}$ is evaluated to itself, as it is the case in Prolog.

⁹ We will typically use the rule name as a symbol for the rule as a whole.

$$\begin{array}{c}
\frac{}{S \xrightarrow{R} S} \text{ EMPTY} \qquad \frac{r \in R \quad S \xrightarrow{\{r\}} S' \quad S' \xrightarrow{R} S''}{S \xrightarrow{R} S''} \text{ STEP} \\
\\
\frac{R@c_1, \dots, c_n \setminus c_{n+1}, \dots, c_{n+m} \Leftrightarrow G \mid B \in \Gamma_C(r) \quad \tau_2 \models G}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\{r\}} \{c_1, \dots, c_n\} \uplus B \uplus \Delta S} \text{ APPLY}
\end{array}$$

Fig. 1. Very abstract operational semantics for ground and pure CHR.

Example 6 (C -instances).

If we instantiate the rule

$$\text{subtract } @ N \setminus M \Leftrightarrow 0 < N \wedge 0 < M \wedge N \leq M \mid M - N$$

with $\sigma_1 = \{N \mapsto 4, M \mapsto 6\}$ and $\sigma_2 = \{N \mapsto 0, M \mapsto 6\}$, respectively, we get the \mathbb{N}_0 -instances

$$(\text{subtract})\sigma_1 = \text{subtract } @ 4 \setminus 6 \Leftrightarrow 0 < 4 \wedge 0 < 6 \wedge 4 \leq 6 \mid 2$$

$$(\text{subtract})\sigma_2 = \text{subtract } @ 0 \setminus 6 \Leftrightarrow 0 < 0 \wedge 0 < 6 \wedge 0 \leq 6 \mid 6$$

Both instances are elements of the \mathbb{N}_0 -grounding $\Gamma_{\mathbb{N}_0}(\text{GCD})$ of the program in Example 4.

Classically, the guard G contains constraints which are defined with respect to a constraint theory \mathcal{CT} . We typically write $\mathcal{CT} \models G$ ¹⁰ to denote that the guard is satisfiable with respect to \mathcal{CT} and $\mathcal{CT} \models \neg G$ otherwise. Since in our case $G \in \mu\Lambda_2$, \mathcal{CT} is essentially τ_2 , as it determines the semantics of Boolean terms. We thus write

$$\tau_2 \models G \iff G \equiv_2 \text{true} \qquad \tau_2 \models \neg G \iff G \equiv_2 \text{false}$$

Note that we always need a data type $\mathcal{L}2$. In Prolog, for instance, 2 corresponds to the set $\{\text{true}, \text{false}\}$, representing successful or failed computations, respectively.

Finally, the operational semantics of CHR is defined as a state transition system where the states are multisets¹¹ over the elements of C .

Definition 9 (Very abstract operational semantics of CHR).

The very abstract operational semantics of CHR programs over a domain C is given by the labeled transition system

$$\omega_a = \langle \text{mset } C, \mathbf{PRG}_C, \mapsto \rangle$$

where the transition relation (\mapsto) is defined by the inference rules in Figure 1.

Rules are applied until no more are applicable to the state, this is, we have reached a final state.

Example 7 (ω_a -transitions).

Intuitively, both

$$\tau_2 \models 0 < 4 \wedge 0 < 6 \wedge 4 \leq 6 \qquad \text{and} \qquad \tau_2 \models \neg(0 < 0 \wedge 0 < 6 \wedge 0 \leq 6)$$

¹⁰ As we only work with ground values, we do not need to quantify any variables.

¹¹ There may be some additional decoration in more refined operational semantics.

hold. Hence, we can prove the transition $\{4, 6\} \xrightarrow{\{\text{subtract}\}} \{4, 2\}$, but not $\{0, 6\} \xrightarrow{\{\text{subtract}\}} \{0, 6\}$.

The following example shows the execution of the Euclidean algorithm as a final example of the operational semantics of CHR.

Example 8 (Euclidean algorithm (cont.)).

The rules of GCD are applied until exhaustion, leaving only the greatest common divisor of all numbers of the input. For an input $\{4, 6\}$, the program will perform a sequence

$$\{4, 6\} \xrightarrow{\{\text{subtract}\}} \{4, 2\} \xrightarrow{\{\text{subtract}\}} \{2, 2\} \xrightarrow{\{\text{subtract}\}} \{2, 0\} \xrightarrow{\{\text{zero}\}} \{2\}$$

of transformations.

4 FreeCHR

The main idea of FreeCHR is to model the syntax of CHR programs as a functor CHR_C . We then use the *free* CHR_C -algebra to define the operational semantics of FreeCHR.

4.1 Syntax

We first present the fundamental definition of our work which allows us to model CHR-programs over a domain C .

Definition 10 (Syntax of FreeCHR programs).

The functor¹²

$$\begin{aligned} \text{CHR}_C D &= \text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C} \sqcup D \times D \\ \text{CHR}_C f &= \text{id} \sqcup f \times f \end{aligned}$$

describes the syntax of FreeCHR programs.

The set $\text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C}$ is the set of single rules. The kept and removed head of a rule are sequences of functions in $\text{list}2^C$ which map elements of C to Booleans, effectively checking individual values for applicability of the rule. The guard of the rule is a function in $2^{\text{list}C}$ and maps sequences of elements in C to Booleans, checking all matched values in the context of each other. Finally, the body of the rule is a function in $(\text{mset}C)^{\text{list}C}$ and maps the matched values to a multiset of newly generated values.

The set $D \times D$ represents the composition of FreeCHR programs by an execution strategy, allowing the construction of more complex programs from, ultimately, single rules.

By the structure of CHR_C , a CHR_C -algebra with carrier D is defined by two functions

$$\rho: \text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C} \longrightarrow D \quad \nu: D \times D \rightarrow D$$

¹² That CHR_C is indeed a functor can easily be verified via equational reasoning.

as $(D, [\rho, \nu])$. The free CHR_C -algebra CHR_C^* provides us with an inductively defined representation of programs which we will later use to lift the very abstract operational semantics ω_a .

Lemma 1 (Free CHR_C -algebra).

With

$$\mu\text{CHR}_C = \text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C} \sqcup \mu\text{CHR}_C \times \mu\text{CHR}_C$$

and labels/injections

$$\text{rule} : \text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C} \longrightarrow \mu\text{CHR}_C$$

$$\odot : \mu\text{CHR}_C \times \mu\text{CHR}_C \longrightarrow \mu\text{CHR}_C$$

$\text{CHR}_C^* = (\mu\text{CHR}_C, [\text{rule}, \odot])$ is the free CHR_C -algebra.

Proof.

We show that $\text{CHR}_C^* = (\mu\text{CHR}_C, [\text{rule}, \odot])$ is the free CHR_C -algebra, by constructing the CHR_C -catamorphism

$$(\llbracket [\rho, \nu] \rrbracket) : \mu\text{CHR}_C \longrightarrow (A, [\rho, \nu])$$

for any CHR_C -algebra $(A, [\rho, \nu])$. By definition of the free CHR_C -algebra,

$$(\llbracket [\rho, \nu] \rrbracket)([\text{rule}, \odot](p)) = [\rho, \nu]((\text{CHR}_C(\llbracket [\rho, \nu] \rrbracket))(p)) \quad (1)$$

needs to be true for any $p \in \mu\text{CHR}_C$.

Case 1 ($p = \text{rule}(k, r, g, b)$).

$$\begin{aligned} & (\llbracket [\rho, \nu] \rrbracket)([\text{rule}, \odot](k, r, g, b)) = [\rho, \nu]((\text{CHR}_C(\llbracket [\rho, \nu] \rrbracket))(k, r, g, b)) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(\text{rule}(k, r, g, b)) = [\rho, \nu](\text{id}(k, r, g, b)) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(\text{rule}(k, r, g, b)) = [\rho, \nu](k, r, g, b) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(\text{rule}(k, r, g, b)) = \rho(k, r, g, b) \end{aligned}$$

Case 2 ($p = p_1 \odot p_2$).

$$\begin{aligned} & (\llbracket [\rho, \nu] \rrbracket)([\text{rule}, \odot](p_1, p_2)) = [\rho, \nu]((\text{CHR}_C(\llbracket [\rho, \nu] \rrbracket))(p_1, p_2)) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(p_1 \odot p_2) = [\rho, \nu]((\llbracket [\rho, \nu] \rrbracket \times \llbracket [\rho, \nu] \rrbracket)(p_1, p_2)) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(p_1 \odot p_2) = [\rho, \nu](\llbracket [\rho, \nu] \rrbracket(p_1), \llbracket [\rho, \nu] \rrbracket(p_2)) \\ \Leftrightarrow & \llbracket [\rho, \nu] \rrbracket(p_1 \odot p_2) = \nu(\llbracket [\rho, \nu] \rrbracket(p_1), \llbracket [\rho, \nu] \rrbracket(p_2)) \end{aligned}$$

Therefore, (1) holds if and only if

$$\begin{aligned} & (\llbracket [\rho, \nu] \rrbracket) : \mu\text{CHR}_C \longrightarrow A \\ & \llbracket [\rho, \nu] \rrbracket(\text{rule}(k, r, g, b)) = \rho(k, r, g, b) \\ & \llbracket [\rho, \nu] \rrbracket(p_1 \odot p_2) = \nu(\llbracket [\rho, \nu] \rrbracket(p_1), \llbracket [\rho, \nu] \rrbracket(p_2)) \end{aligned}$$

$$\begin{array}{c}
\frac{}{S \xrightarrow{p} S} \text{ PASS/FINAL} \qquad \frac{S \xrightarrow{p_i} S' \quad S' \xrightarrow{p_1 \odot \dots \odot p_i \odot \dots \odot p_l} S''}{S \xrightarrow{p_1 \odot \dots \odot p_i \odot \dots \odot p_l} S''} \text{ STEP}_i \\
\\
\frac{k_1(c_1) \wedge \dots \wedge k_n(c_n) \wedge r_1(c_{n+1}) \wedge \dots \wedge r_m(c_{n+m}) \wedge g(c_1, \dots, c_{n+m}) \equiv_2 \text{true}}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\text{rule}(k, r, g, b)} \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{m+n}) \uplus \Delta S} \text{ APPLY}
\end{array}$$

Fig. 2. *Very abstract* operational semantics of freeCHR.

for all CHR_C -algebras $(A, [\rho, \nu])$. Hence, there is a unique homomorphism

$$(\llbracket \rho, \nu \rrbracket) : \text{CHR}_C^* \rightarrow (A, [\rho, \nu])$$

for any CHR_C -algebra $(A, [\rho, \nu])$, making CHR_C^* the free CHR_C -algebra. \square

The free CHR_C -algebra corresponds to the definition of abstract syntax trees of programs, while the catamorphism $(\llbracket \alpha \rrbracket)$ corresponds to an interpretation that preserves the semantics of α .

We can easily see that \odot is associative up to isomorphism.¹³ We thus will not explicitly write parentheses and generally use chained expressions like $p_1 \odot \dots \odot p_l$ for some $l \in \mathbb{N}$.

Example 9 (Euclidean algorithm (cont.)).

The program $\text{gcd} = \text{zero} \odot \text{subtract}$ with

$$\begin{aligned}
\text{zero} &= \text{rule}(\varepsilon, (\lambda n. n = 0), (\lambda n. \text{true}), (\lambda n. \emptyset)) \\
\text{subtract} &= \text{rule}((\lambda n. 0 < n), (\lambda m. 0 < m), (\lambda n \ m. n \leq m), (\lambda n \ m. \{m - n\}))
\end{aligned}$$

implements the euclidean algorithm, as defined in Example 4. λ -abstractions are used for ad-hoc definitions of functions.

4.2 Operational semantics

We now lift the *very abstract* operational semantics ω_a of CHR to the very abstract operational semantics ω_a^* of FreeCHR. We assume that our programs are defined over a domain C , where there is a data type $\mathcal{LC} = \langle C, \Lambda_C, \tau_C \rangle$. Like ω_a , ω_a^* is defined as a labeled transition system, where states are multisets over elements of C .

Definition 11 (Very abstract operational semantics ω_a^*).

The very abstract operational semantics of FreeCHR is defined as the labeled transition system

$$\omega_a^* = \langle \text{mset } C, \mu\text{CHR}_C, \mapsto \rangle$$

where the transition relation (\mapsto) is defined by the inference rules in Figure 2.

The rule PASS/FINAL states that a program is always allowed to do nothing to a state. STEP_{*i*} states that we can derive a transition from S to S'' , if we can transition from S to S' with the i -th program in the composition $p_1 \odot \dots \odot p_l$ (for $1 \leq i \leq l$) and from S' to S'' with the whole composition. The idea is that, without loss of generality p_i is a

¹³ By $\text{assoc}(a, (b, c)) = ((a, b), c)$ and $\text{assoc}^{-1}((a, b), c) = (a, (b, c))$.

rule that is applied to the current state, whereafter execution is continued. APPLY is the translation of the original APPLY-rule of ω_a . k_i and r_j denote the i -th and j -th elements of the sequences k and r .

We want to demonstrate the *very abstract* operational semantics of FreeCHR on a small example.

Example 10 (Euclidean algorithm (cont.)).

Given the program gcd of Example 9 and the initial state $\{12, 9\}$, we can prove the derivation

$$\{12, 9\} \xrightarrow{\text{gcd}} \{3\}$$

as follows. We first use the STEP transition to apply the subtract rule to the initial state.

$$\frac{\frac{\{12, 9\} \xrightarrow{\text{subtract}} \{3, 9\} \quad \{3, 9\} \xrightarrow{\text{zero} \odot \text{subtract}} \{3\}}{\{12, 9\} \xrightarrow{\text{zero} \odot \text{subtract}} \{3\}} \text{STEP}_2$$

and prove that this transition is valid

$$\frac{0 < 9 \wedge 0 < 12 \wedge 9 \leq 12 \equiv \text{true}}{\{12, 9\} \uplus \emptyset \xrightarrow{\text{subtract}} \{9\} \uplus \{3\} \uplus \emptyset} \text{APPLY}$$

We repeat this process for the transitions $\{3, 9\} \xrightarrow{\text{gcd}} \{3, 6\}$, $\{3, 6\} \xrightarrow{\text{gcd}} \{3, 3\}$ and $\{3, 3\} \xrightarrow{\text{gcd}} \{3, 0\}$. At this point, we have to apply the zero rule to remove the 0 and use the PASS/FINAL transition to prove the empty transition $\{3\} \xrightarrow{\text{gcd}} \{3\}$.

$$\frac{\frac{0 = 0 \wedge \text{true} \equiv \text{true}}{\{3, 0\} \uplus \emptyset \xrightarrow{\text{zero}} \{3\} \uplus \emptyset \uplus \emptyset} \text{APPLY} \quad \frac{\frac{\{3, 0\} \xrightarrow{\text{zero}} \{3\} \quad \{3\} \xrightarrow{\text{zero} \odot \text{subtract}} \{3\}}{\{3, 0\} \xrightarrow{\text{zero} \odot \text{subtract}} \{3\}} \text{STEP}_2 \quad \text{PASS/FINAL}}$$

With this done, we have fully proven the derivation.

4.3 Soundness and completeness of ω_a^*

To prove the soundness and completeness of ω_a^* with respect to ω_a , we first need to embed FreeCHR into the *positive range-restricted ground* segment of CHR. This is a common approach to relate rule-based formalisms to CHR (Frühwirth, 2009, Chapter 6).

Definition 12 (Embedding FreeCHR into CHR).

Let

$$\begin{aligned} \Theta : \mu\text{CHR}_C &\longrightarrow \mathbf{PRG}_C \\ \Theta(\text{rule}(k, r, g, b)) &= \{R @ v_1, \dots, v_n \setminus v_{n+1}, \dots, v_{n+m} \Leftrightarrow G | b(v_1, \dots, v_{n+m})\} \\ \Theta(p_1 \odot \dots \odot p_l) &= \Theta(p_1) \cup \dots \cup \Theta(p_l) \end{aligned}$$

be the function embedding FreeCHR programs into the positive range-restricted ground segment of CHR, with universally quantified variables v_1, \dots, v_{n+m} , R a uniquely generated rule name and

$$G = k_1(v_1) \wedge \dots \wedge r_m(v_{n+m}) \wedge g(v_1, \dots, v_{n+m})$$

Example 11.

If we apply the embedding Θ to the program *gcd* (including the evaluation of function applications¹⁴), we get the program

$$\begin{aligned} \text{zero} @ \varepsilon \setminus n &\Leftrightarrow n = 0 \wedge \text{true} \mid \emptyset \\ \text{subtract} @ n \setminus m &\Leftrightarrow 0 < n \wedge 0 < m \wedge n \leq m \mid m - n \end{aligned}$$

with universally quantified variables n and m . The program corresponds to the head-normalization (Duck (2005)) of the program in Example 4.

We now prove Θ -soundness of ω_a^* with respect to ω_a , this is, if we can prove a derivation under ω_a^* for a program p , we can prove it under ω_a for $\Theta(p)$.

Theorem 1 (Θ -Soundness of ω_a^*).

ω_a^* is Θ -sound with respect to ω_a , this is, for $S, S' \in \mathbf{mset} \ C$ and $p \in \mu\text{CHR}_C$,

$$S \xrightarrow{p} S' \implies S \xrightarrow{\Theta(p)} S'$$

Proof.

We prove soundness by induction over the inference rules of ω_a^* .

Induction Base Case (PASS/FINAL).

Follows directly from EMPTY.

Induction Base Case (APPLY).

Given a proof

$$\frac{G \equiv_2 \text{true}}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{p} \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S} \text{APPLY}$$

with $r = \text{rule}((k_1, \dots, k_n), (r_1, \dots, r_m), g, b)$ and

$$G = k_1(t_1) \wedge \dots \wedge k_n(t_n) \wedge r_1(t_{n+1}) \wedge \dots \wedge r_m(t_{n+m}) \wedge g(t_1, \dots, t_{n+m})$$

there is a C -instance

$$R @ c_1, \dots, c_n \setminus c_{n+1}, \dots, c_{n+m} \Leftrightarrow G \mid b(c_1, \dots, c_{n+m}) \in \Gamma_C(\Theta(p))$$

such that $\tau_2 \models G$. By Definition 12 we know that $\{p'\} = \Theta(p)$. We can thus prove $S \xrightarrow{\Theta(p)} S'$ by

$$\frac{R @ c_1, \dots, c_n \setminus c_{n+1}, \dots, c_{n+m} \Leftrightarrow G \mid b(c_1, \dots, c_{n+m}) \in \Gamma_C(\Theta(p')) \quad \tau_2 \models G}{\underbrace{\{c_1, \dots, c_{n+m}\} \uplus \Delta S}_{=S} \xrightarrow{\{p'\}} \underbrace{\{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S}_{=S'}} \text{APPLY}$$

Induction Step (STEP_{*i*})

Given a proof

$$\frac{S \xrightarrow{p_i} S' \quad S' \xrightarrow{p} S''}{S \xrightarrow{p} S''} \text{STEP}_i$$

¹⁴ this is, β -reductions

with $p = p_1 \odot \dots \odot p_i \odot \dots \odot p_l$, for $1 \leq i \leq l$, we assume

$$\forall i \in 1, \dots, l. S \xrightarrow{p_i} S' \implies S \xrightarrow{\Theta(p_i)} S' \quad (2)$$

$$S' \xrightarrow{p} S'' \implies S' \xrightarrow{\Theta(p)} S'' \quad (3)$$

as induction hypotheses.

Since we can, without loss of generality, assume that $p_i = \text{rule}(\dots)$, we know from Definition 12 that $\{p'\} = \Theta(p_i) \subseteq \Theta(p)$. By induction hypothesis (2) we know from $S \xrightarrow{p_i} S'$ that $S \xrightarrow{\Theta(p_i)} S'$. By induction hypothesis (3) we know from $S' \xrightarrow{p} S''$ that $S' \xrightarrow{\Theta(p)} S''$. We can thus prove $S \xrightarrow{\Theta(p)} S''$ by

$$\frac{p' \in \Theta(p) \quad S \xrightarrow{\{p'\}} S' \quad S' \xrightarrow{\Theta(p)} S''}{S \xrightarrow{\Theta(p)} S''} \text{STEP}$$

□

We established that we can prove any derivation we can prove with a program p under ω_a^* , with a program $\Theta(p)$ under ω_a . We also want to prove Θ -completeness, this is, we can prove any derivation with a program $\Theta(p)$ under ω_a with a program p under ω_a^* .

Theorem 2 (Θ -completeness of ω_a^*).

ω_a^* is Θ -complete with respect to ω_a , this is, for $S, S' \in \mathbf{mset}C$ and $p \in \mu\text{CHR}_C$,

$$S \xrightarrow{\Theta(p)} S' \implies S \xrightarrow{p} S'$$

Proof.

We prove Θ -completeness by induction over the inference rules of ω_a

Induction Base Case (EMPTY).

Follows directly from PASS/FINAL

Induction Base Case (APPLY).

Given a proof

$$\frac{R@c_1, \dots, c_n \setminus c_{n+1}, \dots, c_{n+m} \Leftrightarrow G \mid B \in \Gamma_C(\Theta(p)) \quad \tau_2 \models G}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\Theta(p)} \{c_1, \dots, c_n\} \uplus B \uplus \Delta S} \text{APPLY}$$

By definition of ω_a , we know that $\Theta(p) = \{p'\}$. Therefore, $p = \text{rule}(k, r, g, b)$, and hence $G = k_1(c_1) \wedge \dots \wedge r_m(c_{n+m}) \wedge g(c_1, \dots, c_{n+m})$ and $B = b(c_1, \dots, c_{n+m})$. From $\tau_2 \models G$ follows $G \equiv_2 \text{true}$. We can thus construct a proof

$$\frac{G \equiv_2 \text{true}}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{p} \{c_1, \dots, c_n\} \uplus B \uplus \Delta S} \text{APPLY}$$

Induction Step (STEP).

Given a proof

$$\frac{p' \in \Theta(p) \quad S \xrightarrow{\{p'\}} S' \quad S' \xrightarrow{\Theta(p)} S''}{S \xrightarrow{\Theta(p)} S''} \text{STEP}$$

By Definition 12, we know from $r' \in \Theta(p)$ that for $p = p_1 \odot \dots \odot p_i \odot \dots \odot p_l$ (with $l \geq 1$) there is a p_i such that $\{p'\} = \Theta(p_i)$. We hence assume as induction hypotheses that

$$S \xrightarrow{\{p'\}} S' \implies S \xrightarrow{p_i} S' \quad (4)$$

$$S' \xrightarrow{\Theta(p)} S'' \implies S' \xrightarrow{p} S'' \quad (5)$$

From $S \xrightarrow{\{p'\}} S'$ and induction hypothesis (4) follows $S \xrightarrow{p_i} S'$. From $S' \xrightarrow{\Theta(p)} S''$ and induction hypothesis (5) follows $S' \xrightarrow{p} S''$. We can hence construct the proof

$$\frac{S \xrightarrow{p_i} S' \quad S' \xrightarrow{p} S''}{S \xrightarrow{p} S''} \text{STEP}$$

□

With Theorem 1 and Theorem 2, we have established that, up to Θ , FreeCHR is as expressive as CHR. Analogously to classical CHR, we are now able to define more refined operational semantics for FreeCHR and prove their soundness with respect to ω_a^* , which is ongoing future work.

5 An instance with *very abstract* operational semantics

In this section, we want to present a simple execution algorithm and prove its adherence to ω_a^* by using the inductive nature of our definition of syntax and semantics of FreeCHR. This will serve both as a first, albeit simple, guideline for implementations, as well as a case study on how a FreeCHR execution algorithm can be defined and verified.

5.1 Instances

Since the operational semantics define the execution of FreeCHR programs as a state transition system, an obvious way to implement FreeCHR is by mapping programs to functions mapping one state to its successor.

Definition 13 (FreeCHR instance).

An instance of FreeCHR is a CHR_C -algebra

$$((\text{mset}C)^{\text{mset}C}, [\text{rule}, \text{compose}])$$

with

$$\begin{aligned} \text{rule} &: \text{list}2^C \times \text{list}2^C \times 2^{\text{list}C} \times (\text{mset}C)^{\text{list}C} \longrightarrow (\text{mset}C)^{\text{mset}C} \\ \text{compose} &: (\text{mset}C)^{\text{mset}C} \times (\text{mset}C)^{\text{mset}C} \longrightarrow (\text{mset}C)^{\text{mset}C} \end{aligned}$$

A FreeCHR instance hence needs to define the two functions **rule** and **compose** which implement the execution of programs and hence the operational semantics of the program. **compose** implements the application of a single rule to a state if it is possible. **compose** implements the execution strategy of, fundamentally, a sequence of rules to a state. Both are only concerned with single execution steps. We hence also need the function

$$\text{run} : (\text{mset}C)^{\text{mset}C} \longrightarrow (\text{mset}C)^{\text{mset}C}$$

which maps a function performing singular steps to a function applying this function until a final state is reached. In CHR this generally means, until no rules are applicable to the state.

Since an instance practically also defines operational semantics it makes sense to directly associate it with an LTS as well.

Definition 14 (LTS of a FreeCHR instance).

Let

$$I = ((\mathbf{mset}C)^{\mathbf{mset}C}, [\rho, \nu])$$

I implies an LTS

$$crI = \langle \mathbf{mset}C, L, \mapsto \rangle$$

where the labels L are recursively defined as

$$L = \{ \rho(k, r, g, b) \mid k, r \in \mathbf{list}(2^C), g \in 2^{\mathbf{list}C}, b \in (\mathbf{mset}C)^{\mathbf{list}C} \} \cup \{ \nu(f, g) \mid f, g \in L \}$$

and the transition relation as

$$(\mapsto) = \left\{ s \xrightarrow{f} s' \mid s \in \mathbf{mset}C, f \in L, f(s) \equiv_{\mathbf{mset}C} s' \right\}$$

With this, we can easily verify an instance against operational semantics, using the concepts from Section 2.2. We will do so from the perspective of the operational semantics, instead of from the perspective of the implementation. To verify, that an instance only performs valid transitions, we hence prove Θ -completeness

$$S \xrightarrow{p} S' \Leftarrow S \xrightarrow{(\downarrow \alpha)(p)} S'$$

Consequently, to prove that the implementation can perform all valid transitions, we prove Θ -soundness up to repeated transition

$$S \xrightarrow{p} S' \Longrightarrow S \xrightarrow{(\downarrow \alpha)(p)^*} S'$$

We need the relaxation "up to repeated transition" since the instances, and hence their respective LTS, are only defined for singular steps.

5.2 Execution algorithm

The algorithm presented in this section implements the operations required for a FreeCHR instance as non-deterministic imperative pseudo-code. The *rule* function is implemented by `RULE`, `⊙` by `COMPOSE`, in Algorithm 1.

*Definition 15 (Pure CHR-programs in **Set** (cont.)).*

Algorithm 1 shows the abstract implementation of a CHR_C -algebra

$$\mathbf{Exec}_a = ((\mathbf{mset}C)^{\mathbf{mset}C}, [\mathbf{RULE}, \mathbf{COMPOSE}])$$

The function `RULE` first checks whether there is a matching $\{c_1, \dots, c_{n+m}\} \subseteq S$, such that head constraints `kepti` applied to c_i and `removedj` applied to c_{n+j} evaluate to *true*, as well as the guard, applied to all constraints of the matching (l. 2). If such a subset

Algorithm 1 Implementation of $Exec_a$

```

1: function rule(kept, removed, guard, body,  $S$ )
2:   if  $\exists \{c_1, \dots, c_{n+m}\} \subseteq S. \text{kept}_1(c_1) \wedge \dots \wedge \text{removed}_m(c_{n+m}) \wedge \text{guard}(c_1, \dots, c_{n+m}) \equiv_2 \text{true}$ 
   then
3:     return  $(S \cup \{c_{n+1}, \dots, c_{n+m}\}) \uplus \text{body}(c_1, \dots, c_{n+m})$ 
4:   else
5:     return  $S$ 
6:   end if
7: end function
8:
9: function compose( $(p_1, \dots, p_l), S$ )
10:  if  $\exists p \in \{p_1, \dots, p_l\}. p(S) \not\equiv_{\text{mset}C} S$  then
11:    return  $p(S)$ 
12:  else
13:    return  $S$ 
14:  end if
15: end function

```

Algorithm 2 Implementation of run_f

```

1: function run( $f, S$ )
2:   repeat
3:      $S' \leftarrow S$ 
4:      $S \leftarrow f(S)$ 
5:   until  $S = S'$ 
6:   return  $S$ 
7: end function

```

exists, the constraints c_{n+j} with $1 \leq j \leq m$ are removed and the multiset resulting from applying the body to the matching will be added to S (l. 3). The function COMPOSE checks if there is a $p \in \{p_1, \dots, p_l\}$ with $p(S) \neq S$ (l. 10). If this is the case, p is applied to S (l. 11). If in either function the respective check does not succeed, S is returned unmodified (l. 5 and 13, respectively). The execution strategy implemented by COMPOSE is to non-deterministically select an applicable program and apply it to the state.

Note, that we will freely use currying on the last argument of either function, for instance, $\text{RULE}(k, r, g, b)$ is a function from $\text{mset}C$ to $\text{mset}C$, and $\text{RULE}(k, r, g, b, S)$ evaluates to an $S' \in \text{mset}C$. Also, COMPOSE will accept an arbitrary number of subprograms (but always an input state).

Finally, we need to implement the *run* function.

*Definition 16 (Pure CHR-programs in **Set** (cont.)).*

Algorithm 2 implements application until exhaustion for FreeCHR programs defined using Algorithm 1.

$\text{RUN}(f, S)$ applies the function f to the argument state S repeatedly. After each iteration, it is compared to the last state. As soon as both are equal it is assumed that no more rules could be applied. Then, the state is returned as the final state.

Note, that this algorithm would also terminate if a propagation rule with an empty body is applied. As such rules are nonsensical, we will ignore this case.

5.3 Correctness

We will now prove the correctness of \mathbf{Exec}_a . With $\mathbf{exec}_a = [\text{RULE}, \text{COMPOSE}]$, let

$$\langle \mathbf{exec}_a \rangle : \text{CHR}_C^* \rightarrow \mathbf{Exec}_a$$

be the CHR_C -catamorphism to \mathbf{Exec}_a , this is, the interpretation of programs in μCHR_C that respects the semantics of Algorithm 1.

Theorem 3.

The very abstract operational semantics ω_a^ is $\langle \mathbf{exec}_a \rangle$ -complete with respect to \mathbf{Exec}_a (Algorithm 1), this is,*

$$S \xrightarrow{\langle \mathbf{exec}_a \rangle(p)} S' \implies S \xrightarrow{p} S'$$

Proof.

To prove the theorem, we will first discern two basic cases.

Case 3 ($S = S'$).

This case is trivial by rule PASS/FINAL.

Case 4 ($S \neq S'$).

We prove this case via induction over the structure of $p \in \text{CHR}_C$.

Induction Base Case ($p = \text{rule}(k, r, g, b)$).

The transition

$$S \xrightarrow{\langle \mathbf{exec}_a \rangle(\text{rule}(k, r, g, b))} S'$$

implies that

$$\begin{aligned} & \langle \mathbf{exec}_a \rangle(\text{rule}(k, r, g, b))(S) \\ &= \text{RULE}(k, r, g, b, S) \\ &\equiv S' \end{aligned}$$

From Algorithm 1 follows that the condition in line 2 must be true, since there were changes to the state. This implies that there is a subset $\{c_1, \dots, c_{n+m}\} \subseteq S$ for which

$$k_1(c_1) \wedge \dots \wedge k_n(c_n) \wedge r_1(c_{n+1}) \wedge \dots \wedge r_m(c_{n+m}) \wedge g(c_1, \dots, c_{n+m}) \equiv_2 \text{true}$$

Hence

$$S = \{c_1, \dots, c_{n+m}\} \uplus \Delta S$$

and

$$(S \sqcup \{c_{n+1}, \dots, c_{n+m}\}) \uplus b(c_1, \dots, c_{n+m}) = \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_n) \uplus \Delta S$$

From this, we can construct a proof

$$\frac{k_1(c_1) \wedge \dots \wedge k_n(c_n) \wedge r_1(c_{n+1}) \wedge \dots \wedge r_m(c_{n+m}) \wedge g(c_1, \dots, c_{n+m}) \equiv_2 \text{true}}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\text{rule}(k,r,g,b)} \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_n) \uplus \Delta S} \text{APPLY}$$

Induction Step ($p = p_1 \odot \dots \odot p_n$).

As the induction hypothesis, we assume

$$\forall p \in \text{CHR}_C. S \xrightarrow{\langle \text{exec}_a \rangle(p)} S' \implies S \xrightarrow{p} S'$$

The transition

$$S \xrightarrow{\langle \text{exec}_a \rangle(p_1 \odot \dots \odot p_n)} S'$$

implies that

$$\langle \text{exec}_a \rangle(p_1 \odot \dots \odot p_n)(S) = \text{COMPOSE}(p_1, \dots, p_n, S) \equiv_{\text{mset } C} S'$$

From Algorithm 1, especially line 10, we know that there must hence be a p_i with $1 \leq i \leq n$ such that $\langle \text{exec}_a \rangle(p_i)(S) \equiv_{\text{mset } C} S'$, or in terms of the transition relation

$$S \xrightarrow{\langle \text{exec}_a \rangle(p_i)} S'$$

By induction hypothesis, we then have a proof for $S \xrightarrow{p_i} S'$. We can thus construct a proof

$$\frac{S \xrightarrow{p_i} S' \quad \frac{S' \xrightarrow{p_1 \odot \dots \odot p_n} S'}{\text{STEP}_i} \quad \text{PASS/FINAL}}{S \xrightarrow{p_1 \odot \dots \odot p_n} S'}$$

□

Theorem 4.

The very abstract operational semantics ω_a^* is $\langle \text{exec}_a \rangle$ -sound up to repeated transition with respect to \mathbf{Exec}_a (Algorithm 1), this is,

$$S \xrightarrow{\langle \text{exec}_a \rangle(p)}^* S' \iff S \xrightarrow{p} S'$$

Proof.

We prove the theorem by induction over the inference rules of ω_a^* .

Induction Base Case (PASS/FINAL).

Trivial, since $S \xrightarrow{\langle \text{exec}_a \rangle(p)}^* S$ is always true for zero transition steps

Induction Base Case (APPLY).

Given a proof

$$\frac{k_1(c_1) \wedge \dots \wedge k_n(c_n) \wedge r_1(c_{n+1}) \wedge \dots \wedge r_m(c_{n+m}) \wedge g(c_1, \dots, c_{n+m}) \equiv_2 \text{true}}{\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\text{rule}(k,r,g,b)} \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S} \text{APPLY}$$

We want to show that

$$\{c_1, \dots, c_{n+m}\} \uplus \Delta S \xrightarrow{\langle \text{exec}_a \rangle(\text{rule}(k,r,g,b))} \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S$$

or respectively

$$\begin{aligned}
& \langle \mathbf{exec}_a \rangle (\text{rule}(k, r, g, b)) (\{c_1, \dots, c_{n+m}\} \uplus \Delta S) \\
&= \text{RULE}(k, r, g, b, \{c_1, \dots, c_{n+m}\} \uplus \Delta S) \\
&= \{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S
\end{aligned}$$

Since by assumption the condition in line 2 will be *true*. From line 3 we then know that

$$\begin{aligned}
& \text{RULE}(k, r, g, b, \{c_1, \dots, c_{n+m}\} \uplus \Delta S) \\
&\equiv \left(\{c_1, \dots, c_{n+m}\} \uplus \Delta S \cup \{c_{n+1}, \dots, c_{n+m}\} \right) \uplus b(c_1, \dots, c_{n+m})
\end{aligned}$$

which can be simplified to $\{c_1, \dots, c_n\} \uplus b(c_1, \dots, c_{n+m}) \uplus \Delta S$.

Induction Step (STEP_{*i*})

Given a proof

$$\frac{S \xrightarrow{p_i} S' \quad S' \xrightarrow{p_1 \odot \dots \odot p_l} S''}{S \xrightarrow{p_1 \odot \dots \odot p_l} S''} \text{STEP}_i$$

with $1 \leq i \leq l$. We assume

$$S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_i)}^* S' \Leftarrow S \xrightarrow{p_i} S' \quad (6)$$

$$S' \xrightarrow{\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)}^* S'' \Leftarrow S' \xrightarrow{p_1 \odot \dots \odot p_l} S'' \quad (7)$$

as induction hypotheses. We want to show that

$$S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)}^* S''$$

From the premises of the given proof we can infer $S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_i)}^* S'$ and $S' \xrightarrow{\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)}^* S''$ by induction hypotheses (6) and (7), respectively. By the definition of COMPOSE we know that, since $\langle \mathbf{exec}_a \rangle (p_i)(S) \equiv S'$, S' is also a possible outcome of $\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)(S)$. From $S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_i)}^* S'$, we can hence assume $S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)}^* S'$. Therefrom $S \xrightarrow{\langle \mathbf{exec}_a \rangle (p_1 \odot \dots \odot p_l)}^* S''$ follows directly. \square

With Theorem 3 and Theorem 4 we established that Algorithm 1 correctly implements the *very abstract operational semantics* ω_a^* . The theorems also imply that \mathbf{Exec}_a is an equivalent representation of ω_a^* .

6 Implementation case studies

We now present two implementations that are based on the execution algorithm described in Section 5 as case studies. The source code of the implementations is also available on GitHub.¹⁵

¹⁵ <https://gist.github.com/SRechenberger/d5e1eb875ae72ce5cafe6ea1c5b3ee38>

The purpose of this section is twofold: first, we show how the abstract execution algorithm can be used as a guide to implement FreeCHR. Second, we show how the chosen host language flavors the implementation of FreeCHR both in the back end, as well as in the front end.

The presented implementations have two limitations to be noted: first, they solve any source of non-determinism in the original algorithm. Whenever the original algorithm *chooses* an element from a (multi) set, the implementations will sequentially search for them. Especially is the state, in both implementations, a list instead of a multiset. Also, the rules of a program will be traversed sequentially to find an applicable one. It is, though, easy to see that they will only perform transitions Algorithm 1 would perform as well. They are, in other words, concretizations of \mathbf{Exec}_a .

Second, as they are based on the *very abstract* operational semantics, they do not yet account for trivial non-termination due to repeated application of propagation rules (this is, rules with an empty removed head) (Frühwirth (2009)). Hence, a program with such rules will not terminate if any of them is at some point applicable. This problem is solved with the *abstract* operational semantics by recording a *propagation history*. It saves applied constraints (by a unique identifier) as well as the rule applied to them. If a propagation rule is to be applied, it is first checked if the combination was already executed. If this is the case, the rule will not fire again.

The consequence for the presented implementations is that they only work on a subset of CHR/FreeCHR without propagation rules.

6.1 Host languages

To sufficiently demonstrate how FreeCHR can be implemented in different programming languages and how the respective host language flavors the implementation, we chose two programming languages of very different paradigms: *Haskell* and *Python*.

Haskell is a statically typed functional programming language with a powerful type system. Python is a dynamically typed imperative language with plenty of practical applications.

One might consider both languages to be on the opposite of a static-dynamic/functional-imperative spectrum as well. They will hence serve well for this case study, as most languages are featurewise somewhere in between. Languages like Java or C# successively adopt functional features but stay imperative in their core. They also have a static type system like Haskell, but not as powerful and strict. Languages like Kotlin move further towards the functional end of the spectrum but still inherit many characteristics from Java. On the other hand, the whole LISP family of programming languages is functional but dynamically typed.

6.2 Execution algorithm

We will now give a detailed overview of the implementations, connect them to the abstract algorithm and compare them to each other.

```
1  newtype SolverStep c = SolverStep { runSolverStep :: [c] -> [c] }
```

Fig. 3. Wrapped solver function type.

6.2.1 Carrier

We first want to discuss the implementation of the carrier of the implemented CHR_C -algebra.

In Haskell, the carrier is explicitly defined as the type `SolverStep c` shown in Figure 3. The type is essentially equal to functions of type `[c] -> [c]` which is the Haskell representation of the set $(\text{mset } C)^{\text{mset } C}$. As mentioned above, we will use lists instead of multisets. This eliminates one source of non-determinism.

Wrapping the functions by defining a new type allows the implementation of type class instances (which is used for `COMPOSE`) and to discern programs that are "in construction" from such that were already passed to `run`.

The Python implementation directly models functions mapping lists to lists. Since Python is (by default) dynamically typed, there are no constraints on the type of the elements of the lists. This, on the one hand, gives more freedom to the user of the `FreeCHR` instance since it is easier to define rules on different types of constraints. This is also possible when using Haskell, but needs additional work like defining and matching on a sum type. On the other hand, in Python, one should always explicitly check if the constraints have the correct type in the head of the rule.

For example, the program in Figure 4b allows values of any type in its state and will apply the rules only if integers are encountered. Not checking the type of constraints beforehand is a serious source of runtime errors. If a rule of the `gcd` program is applied to something other than two integers, the comparison operations (l. 3 and 6) in the guard or the subtraction (l. 7) in the body will raise an exception. Such errors are impossible in the Haskell version in Figure 4a since the input type needs to be some sort of integer (`Integral`; see l. 1) which guarantees support for comparison (`Ord`) and subtraction (`Num`).¹⁶

6.2.2 RULE

The implementations of `RULE` can be seen in Figure 5. First, either function lazily computes the sequence of possible matchings (l. 5 in Figure 5a and l. 3–5 in Figure 5b). This is done using the `match` functions which are shown in Figure 6. Both generate a lazy sequence of all permutations of the state that have the same length as the head of the rule (l. 3–4 and l. 3, respectively) and select those for which the conjunction of the elementwise application (done by `zipWith` and `zip`) of the elements of the head evaluates to true, this is, those for which $k_1(c_1) \wedge \dots \wedge k_n(c_n) \wedge r_1(c_{n+1}) \wedge \dots \wedge r_m(c_{n+m}) \equiv_2 \text{true}$, for a head $[k_1, \dots, k_n, r_1, \dots, r_m]$ and a permutation $[c_1, \dots, c_{n+m}]$ of the state. Then, both implementations of `RULE` take the first of those sequences for which the guard g holds, this is, $g(c_1, \dots, c_{n+m}) \equiv_2 \text{true}$. In total, this corresponds to l. 2 in Algorithm 1. If such a matching is found, the new state is computed and returned as in l. 3 (l. 7 and 8–11). In the Python implementation, an explicit copy of the state is created to not modify the original state. If no matching is found, the state is returned unchanged (l. 8 and 6–7).

¹⁶ Both constraints are inherited via Haskell's type class `Real1`.


```

1 gcd :: Integral a => SolverStep a
2 gcd = rule [] [(<= 0)]
3       (const True)
4       (const []) <>
5       rule [(> 0)] [(> 0)]
6       (\[n, m] -> n <= m)
7       (\[n, m] -> [m - n])

```

(a) Haskell

```

1 gcd = compose(
2   rule([], [lambda n: isinstance(n, int)],
3     lambda n: n <= 0,
4     lambda _: []),
5   rule([lambda n: isinstance(n, int)], [lambda m: isinstance(m, int)],
6     lambda n, m: 0 < n <= m,
7     lambda n, m: [m-n])
8 )

```

(b) Python

Fig. 4. Greatest common divisor.

Note that both variants explicitly create and return a function after getting the rule elements passed. The implementations of COMPOSE (Figure 7) work in the same way.

Also, the Python implementation makes heavy use of the *spread operator* ($*$ before function arguments) (l. 4 and 11). In consequence, we can write the guard and body of a rule as functions that accept as many arguments as the head has patterns. In Haskell, guard and body accept a list as their single argument (compare lines 6 and 7 in either implementation in Figure 4).

An issue specific to Python can be seen in Figure 8b. Python does only allow expressions in `lambda` functions. If any function of a rule needs to use statements, it must be defined separately.

6.2.3 COMPOSE

Figure 7 shows the implementations of COMPOSE. Figure 7a implements the algorithm as the binary operator ($<>$) of the `Semigroup` type class. First, the left subprogram will be applied to the state (l. 5). If this has no effect, the second subprogram is applied (l. 7). Otherwise, the result of the first application which must have altered the state, is returned (l. 8). One can easily see that the implementation is associative, as required by `Semigroup`. In nested expressions like $p_1 <> p_2 <> \dots <> p_n$, the algorithm will always perform a depth-first search, beginning from the leftmost subprogram, effectively traversing the program from left to right.

The same is true for the Python variant in Figure 7b. Here, we loop through the subprograms, passed to the function `compose` (l. 4–7) until a program alters the state (l. 6 & 7). Finally, `result` is returned which is either the input state (see l. 3) if the program made it through the loop without finding an applicable subprogram, or the first result of an effective application (see l. 5).

Contrary to the Haskell variant, COMPOSE is implemented as a variadic function by use of the spread operator in Python. The consequence for the user can be seen in Figure 4, where we use ($<>$) to compose the rules (l. 4) in Haskell and `compose` in Python.

```

1 rule :: Eq c => [c -> Bool] -> [c -> Bool] -> ([c] -> Bool) -> ([c] -> [c])
2   -> SolverStep c
3 rule kept removed guard body = SolverStep
4   { runSolverStep = \state ->
5     let matching = find guard (match (kept ++ removed) state)
6     in case matching of
7       Just ms -> (state \\ drop (length kept) ms) ++ body ms
8       _       -> state
9   }

```

(a) Haskell

```

1 def rule(k, r, g, b):
2     def solver(state):
3         matching = next(
4             (matching for matching in match(k+r, state) if g(*matching)),
5             None)
6         if not matching:
7             return state
8         state_copy = state.copy()
9         for c in matching[len(k):]:
10             state_copy.remove(c)
11         return b(*matching) + state_copy
12     return solver

```

(b) Python

Fig. 5. Implementations of RULE.

```

1 match :: [a -> Bool] -> [a] -> [[a]]
2 match ps as = [ as'
3   | as' <- subsequences as, length as' == length ps
4   , as' <- permutations as'
5   , and (zipWith ($) ps as')
6   ]

```

(a) Haskell

```

1 def match(pattern, cs):
2     return (perm
3         for perm in permutations(cs, len(pattern))
4         if all(p(c) for p, c in zip(pattern, perm))
5     )

```

(b) Python

Fig. 6. Lazy matching.

6.2.4 RUN

Figure 9 shows the implementations of RUN. The implementation in Figure 9a first applies the passed program to the query (also called the initial state) (l. 6). If this has no effect, the state is returned (l. 3). Otherwise, `run` is recursively called with the new state (l. 4), repeatedly applying the program until exhaustion.

```

1  instance Eq c => Semigroup (SolverStep c) where
2    (<>) :: Eq c => SolverStep c -> SolverStep c -> SolverStep c
3    solver <> solver' = SolverStep
4      { runSolverStep = \state ->
5        let state' = runSolverStep solver state
6        in if state == state'
7          then runSolverStep solver' state
8          else state'
9      }

```

(a) Haskell

```

1  def compose(*solvers):
2    def solver(constraints):
3      result = constraints
4      for s in solvers:
5        result = s(constraints)
6        if result != constraints:
7          break
8      return result
9  return solver

```

(b) Python

Fig. 7. Implementations of COMPOSE.

```

1  alldifferent :: Eq a => SolverStep a
2  alldifferent = rule [const True, const True] []
3    (\[x, y] -> x == y) (const $ error "False")

```

(a) Haskell

```

1  def false(*args):
2    raise Exception(False)
3
4  all_different = rule(
5    [lambda a: True, lambda b: True], [],
6    lambda a, b: a == b,
7    false
8  )

```

(b) Python

Fig. 8. Implementation of all_diff constraint.

The algorithm in Figure 9b accomplishes the same by the loop in lines 5 to 9, which is interrupted if the state remains unchanged after applying the program (l. 6 & 7).

Again, the use of the spread operator in Python causes a minor difference in usage. While in Haskell, we need to pass a list to `run`, in Python `run` it is a variadic function. Hence, to calculate the greatest common divisor of 6, 9 and 12, we call `run(gcd)(6, 9, 12)` in Haskell and `run gcd [6, 9, 12]` in Python and get in both cases `[3]` as the result.

```

1  run :: Eq c => SolverStep c -> [c] -> [c]
2  run solver query
3      | result == query = result
4      | otherwise      = run solver result
5  where
6      result = runSolverStep solver query

```

(a) Haskell

```

1  def run(solver):
2      def solve(*query):
3          state = list(query)
4          result = solver(state)
5          while True:
6              if result == state:
7                  break
8              state = result
9              result = solver(result)
10         return result
11     return solve

```

(b) Python

Fig. 9. Implementations of RUN.

6.3 Side effects

A major difference between the two presented implementations is the handling of side effects. Python allows for arbitrary side effects like I/O, assigning variables or raising exceptions, as can be seen in Figure 8b. Although this is also possible in Haskell (see Figure 8a), it is highly discouraged by the language. The programs raise an error if they encounter the same value in the query twice. Although the implementations allow it, this behavior is not defined by the operational semantics and domain of future work.

7 Limitations

We want to briefly discuss the current limitations of FreeCHR as it is presented in this paper.

The very abstract operational semantics ω_a for CHR is meant as a tool for formal reasoning and an abstract baseline. It does not describe the behavior of an actual implementation, since it is heavily non-deterministic. As discussed above, it does also not account for trivial non-termination due to repeated application of a propagation rule on the same constraints. Since the very abstract operational semantics for FreeCHR ω_a^* is a direct translation of ω_a , it inherits these issues.

The presented implementations are hence only defined on programs without propagation rules as they would inevitably get stuck on the repeated application of such a rule. This issue is solved by the *refined* operational semantics ω_r as presented by Duck *et al.* (2004). Translating ω_r to FreeCHR is currently ongoing work. For now, the implementations hence only serve as a proof of concept of the applicability of our framework and are

not meant for actual use.¹⁷ We therefore did also not concern ourselves with applying any optimization techniques and do hence not yet expect any reasonable performance.

The approach presented herein is strictly to be understood as the formal foundation upon which we will build our future work, as it connects our framework to the original definition of CHR on the most fundamental level.

8 Related work

8.1 Constraint handling rules

8.1.1 Embeddings

The first approach to embed CHR into a host language was via source-to-source transformation. Holzbaur and Frühwirth (2000) and Schrijvers and Demoen (2004), for example, translate CHR via Prolog's macro system. Similarly, Abdennadher *et al.* (2002) and Van Weert *et al.* (2005) use a precompiler to translate CHR programs into Java, Wuille *et al.* (2007) into C, Nogatz *et al.* (2018) into Javascript, and Barichard (2024) into C++.

Van Weert *et al.* (2008) introduces compilation schemes for imperative languages, upon which, for example, Nogatz *et al.* (2018) builds. The work of van Weert (2010) also provides compilation schemes for imperative languages, but optimizes the matching process by performing it lazily. The algorithms do not first compute a full matching and check it against the patterns and the guard but collect it sequentially.

The inherent disadvantage of an approach using source-to-source compilation is the need for a precompiler in the build chain if the host language does not have a sufficiently expressive macro system like Prolog or LISP. It comes with the cost of more sources of errors and an additional dependency that is often rather tedious to fulfill. Hanus (2015) approaches this problem by extending the Curry compiler to be able to compile CHR code. The obvious problem with this approach is that it is in most cases not viable to extend the compiler of the host language.

A relatively new approach by Ivanović (2013) was to implement CHR as an internal language in Java. This has a major advantage: CHR can now simply be imported as a library, similarly as if implemented by a macro system like by Schrijvers and Demoen (2004). Wibiral (2022) further builds upon this idea and introduces the idea of explicitly composing CHR programs of singular rules by an abstract and modular execution strategy and describing rules through anonymous functions. This was our main inspiration and FreeCHR aims to improve and generalize this idea for arbitrary languages.

8.1.2 Operational semantics

Duck *et al.* (2004) first formalized the behavior of existing implementations which were mostly derived from but not exactly true to existing formal definitions of operational semantics. The techniques used in the implementations were also generalized and improved upon by Van Weert *et al.* (2008) and van Weert (2010). Duck (2005) standardized call-based semantics for CHR, especially in logical programming languages.

¹⁷ Aside maybe from very small applications.

8.2 Algebraic language embeddings

The idea of algebraic embeddings of a domain specific language (DSL) into a host language was first introduced by Hudak (1998). The style in which the languages are embedded was later called *tagless*, as it does not use an algebraic data type, to construct an abstract syntax tree (see, e.g., Carette *et al.* (2007)). The *tags* are the constructors of the data type which defines the syntax of the language. Instead, the embedding directly defines the functions which implement the semantics of the language. This is, defining the free F -algebra versus defining the concrete F -algebras, for a functor F which defines the syntax of the language. The advantage of embedding a DSL in this way is that it does not rely on external build tools, but can instead be easily implemented and used as a library. It also enables the use of any features the host language offers, without any additional work. With an external code-to-code compiler, it would be necessary to re-implement at least the syntax of any desired host language features.

Hofer *et al.* (2008) and Hofer and Ostermann (2010) then extended this idea by using type families in order to provide more flexibility concerning the semantics of the language.

8.3 Logic and constraint based languages and formalisms

CHR was initially designed as a tool to implement constraint solvers for user-defined constraints. Hence, its domain intersects with those of *Answer Set Programming* (ASP) and *Constraint Logic Programming* (CLP). However, on the one hand, CHR can rather be understood as a tool in combination with ASP or especially CLP, as it provides an efficient language to implement solvers, than as an alternative approach. On the other hand, CHR has already exceeded its original purpose and developed more towards a general purpose language.

Another relevant logic based language is *miniKanren* (see Byrd (2009)). *miniKanren* is a family of EDSLs for relational and logic programming. There exists a myriad of implementations for plenty of different host languages as well as formal descriptions of operational semantics (see e.g., Rozplochas *et al.* (2020)). The *miniKanren* language family, though, seems to suffer from similar issues as traditional *Constraint Handling Rules* does. As there is no unified embedding scheme, there is an inherent disconnect between any implementation and the formally defined semantics. Applying an approach similar to ours might be beneficial to the *miniKanren* project as well, especially as it generally is implemented in a way similar to the methods described by Hudak (1998) or Carette *et al.* (2007).

9 Future work

With this foundational introduction of FreeCHR in place, our future work will focus on three main aspects.

First, since the execution algorithm presented here is only of limited practical use, we plan to introduce an algorithm that implements the *refined* operational semantics ω_r of

CHR. The ω_r semantics of CHR is the basis for most existing implementations. Since we already have the *very abstract* operational semantics as a baseline, we can use it as a basic anchor to verify any execution algorithm. The algorithm we intend to formalize is already implemented in Haskell¹⁸ and Python.¹⁹

Second, we currently work on translating the *refined* operational semantics and verifying them against ω_a^* , as well as the aforementioned execution algorithm. This will provide a base for theory concerning practically relevant aspects of FreeCHR.

We plan to verify both the execution algorithm implementing ω_r^* and ω_r^* itself by showing that they are valid concretization of ω_a^* .

Finally, we plan to tackle the issue of side effects that we encountered in the case study in Section 6. This can be done by generalizing the functor CHR_C to use monads to encapsulate arbitrary side effects (see Jones and Wadler (1993)). This was already done in implementations in Haskell but is not yet formalized. This is an important aspect, as it further formalizes practical aspects of everyday programming.

10 Conclusion

In this paper, we introduced the framework FreeCHR that formalizes the embedding of CHR, using initial algebra semantics. We provided the fundamental definition of our framework which models the syntax of CHR programs over a domain C as a **Set**-endofunctor CHR_C . We defined the *very abstract* operational semantics ω_a^* , using the free CHR_C -algebra CHR_C^* and proved soundness and completeness with respect to the original *very abstract* operational semantics ω_a of CHR. We also provided a first abstract instance **Exec** _{a} of FreeCHR, proved its correctness with respect to ω_a^* and discussed deterministic implementations of the algorithms in Haskell and Python.

Hereby, we established FreeCHR as a valid representation of the original definition of CHR and a usable framework for the implementation of the embedded language in a wide range of host languages.

Moreover, are the presented execution algorithm and the implementations of it, although very simple, to our knowledge the first implementations of CHR for which there is an actual prove of correctness with respect to a formal definition of their operational semantics.

Acknowledgements

We thank the reviewers for their helpful comments and constructive feedback. We also thank our colleagues Florian Sihler and Paul Bittner for proofreading and their constructive feedback.

Competing interests

The authors declare none.

¹⁸ <https://github.com/SRechenberger/free-chr-hs>

¹⁹ <https://github.com/SRechenberger/freechr-python>

References

- ABDENNADHER, S., KRÄMER, E., SAFT, M. and SCHMAUSS, M. 2002. JACK: a java constraint kit. In *Electronic Notes in Theoretical Computer Science*, vol. 64, Elsevier, pp. 1–17. [https://doi.org/10.1016/S1571-0661\(04\)80344-X](https://doi.org/10.1016/S1571-0661(04)80344-X)
- ABDENNADHER, S. and MARTE, M. 2000. University course timetabling using constraint handling rules. In *Applied Artificial Intelligence*, vol. 14, issue 4, Taylor & Francis, pp. 311–325. <https://doi.org/10.1080/088395100117016>
- BARICHARD, V. 2024. CHR++: An efficient CHR system in C++ with don't know non-determinism. In *Expert Systems with Applications Part B*, vol. 238, Elsevier. <https://doi.org/10.1016/j.eswa.2023.121810>
- BETZ, H. 2007. Relating coloured petri nets to constraint handling rules. In *CHR 2007*, Ulm University. https://www.uni-ulm.de/fileadmin/website_uni_ulm/iui.inst.170/home/betz/betz_petri_nets_chr07.pdf
- BYRD, W. E. 2009. *Relational Programming in miniKanren: Techniques, Applications, and Implementations*. Ph.D. thesis, Indiana University, USA. <https://scholarworks.iu.edu/dspace/items/450e1b65-70da-4a38-8e73-c182818de110>
- CARETTE, J., KISELYOV, O. and SHAN, C.-c. 2007. Finally tagless, partially evaluated. In *APLAS 2007*, Z. SHAO (Ed.), Lecture Notes in Computer Science, Berlin, Heidelberg, Springer, pp. 222–238. https://doi.org/10.1007/978-3-540-76637-7_15
- CHIN, W., SULZMANN, M. and WANG, M. 2003. A type-safe embedding of constraint handling rules into haskell. Technical Report. https://www.academia.edu/download/40187865/A_type-safe_embedding_of_constraint_hand20151119-21055-1osq67f.pdf
- CHRISTIANSEN, H. and KIRKEBY, M. H. 2015. Confluence modulo equivalence in constraint handling rules. In *Logic-Based Program Synthesis and Transformation*, M. PROIETTI and H. SEKI (Eds.), Lecture Notes in Computer Science, Cham, Springer International Publishing, pp. 41–58. https://doi.org/10.1007/978-3-319-17822-6_3
- DE KONINCK, L., SCHRIJVERS, T., DEMOEN, B., FINK, M., TOMPITS, H. and WOLTRAN, S. 2006. INCLP(R) - interval-based nonlinear constraint logic programming over the reals. In *WLP 2006*, vol. 1843-06-02, Technische Universität Wien, Austria. <https://lirias.kuleuven.be/1653773>
- DUCK, G. J. 2005. *Compilation of Constraint Handling Rules*. Ph.D. thesis, University of Melbourne, Victoria, Australia. <https://www.comp.nus.edu.sg/~gregory/papers/thesis.pdf>
- DUCK, G. J., STUCKEY, P. J., DE LA BANDA, M. G. and HOLZBAUR, C. 2004. The refined operational semantics of constraint handling rules. In *ICLP 2004*, B. DEMOEN and V. LIFSCHITZ (Eds.), Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 90–104. https://doi.org/10.1007/978-3-540-27775-0_7
- FOWLER, M. and PARSONS, R. 2011. *Domain-Specific Languages*. Addison-Wesley, Upper Saddle River, NJ, ISBN 978-0-321-71294-3.
- FRÜHWIRTH, T. 2006. Complete propagation rules for lexicographic order constraints over arbitrary domains. In *CSCLP 2005*, B. HNIC, M. CARLSSON, F. FAGES and F. ROSSI (Eds.), Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 14–28. https://doi.org/10.1007/11754602_2
- FRÜHWIRTH, T. 2009. *Constraint Handling Rules*. Cambridge University Press, Cambridge, U.K. New York, ISBN 978-0-521-87776-3
- FRÜHWIRTH, T. 2015. Constraint handling rules - what else?. In *RuleML 2015*, N. BASSILIADIS, G. GOTTLÖB, F. SADRI, A. PASCHKE and R. D. ROMAN (Eds.), Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 13–34. https://doi.org/10.1007/978-3-319-21542-6_2

- GALL, D. and FRÜHWIRTH, T. 2018. An operational semantics for the cognitive architecture ACT-R and its translation to constraint handling rules. In *ACM Transactions on Computational Logic*, vol. 19, issue 3, Association for Computing Machinery (ACM), pp. 1–42. <https://doi.org/10.1145/3218818>
- HANUS, M. 2015. CHR(Curry): interpretation and compilation of constraint handling rules in curry. In *PADL 2015*, E. PONTELLI and T. C. SON (Eds.), Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 74–89. https://doi.org/10.1007/978-3-319-19686-2_6
- HOFER, C. and OSTERMANN, K. 2010. Modular domain-specific language components in scala. In *GPCE '10*. ACM Press. <https://doi.org/10.1145/1942788.1868307>
- HOFER, C., OSTERMANN, K., RENDEL, T. and MOORS, A. (2008) Polymorphic embedding of dsls. In *GPCE '08*. ACM Press. <https://doi.org/10.1145/1449913.1449935>
- HOLZBAUR, C. and FRÜHWIRTH, T. 2000. A prolog constraint handling rules compiler and run-time system. In *Applied Artificial Intelligence*, vol. 14, issue 4, Taylor & Francis, pp. 369–388. <https://doi.org/10.1080/088395100117043>
- HUDAK, P. 1998. Modular domain specific languages and tools. In *ICSR 1998*, IEEE Computer Society, USA, pp. 134–142. <https://doi.org/10.1109/ICSR.1998.685738>
- IVANOVIĆ, D. 2013. Implementing constraint handling rules as a domain-specific language embedded in Java, arXiv preprint arXiv:1308:3939. <https://doi.org/10.48550/arXiv.1308.3939>
- JOHANN, P. and GHANI, N. 2007. Initial algebra semantics is enough! In *TLCA 2007*, S. R. DELLA ROCCA (Ed.), Lecture Notes in Computer Science, Berlin, Heidelberg, Springer, pp. 207–222. https://doi.org/10.1007/978-3-540-73228-0_16
- JONES, S. P. and WADLER, P. 1993. Imperative functional programming. In *POPL '93*, ACM Press, New York, NY, USA, pp. 71–84. <http://doi.acm.org/10.1145/158511.158524>
- LAM, E. S. L. and SULZMANN, M. 2007. A concurrent constraint handling rules implementation in Haskell with software transactional memory. In *DAMP '07*. ACM Press, Nice, France, 19–24. <https://doi.org/10.1145/1248648.1248653>
- LAM, E. and SULZMANN, M. 2006. Towards agent programming in CHR. In *CHR 2006*, Department of Computer Science, KU Leuven, pp. 17–31. <https://www.cs.kuleuven.be/publicaties/rapporten/cw/CW452.pdf>
- MILEWSKI, B. 2019. *Category Theory for Programmers*. Lightning Source UK, Milton Keynes. ISBN 978-0-464-24387-8
- NOGATZ, F., FRÜHWIRTH, T. and SEIPEL, D. 2018. CHR.js: a CHR implementation in JavaScript. In *RuleML 2018*, C. BENZMÜLLER, F. RICCA, X. PARENT and D. ROMAN (Eds.), Lecture Notes in Computer Science, Springer International Publishing, Cham, pp. 131–146. https://doi.org/10.1007/978-3-319-99906-7_9
- ROZPLOKHAS, D., VYATKIN, A. and BOULYTCHEV, D. 2020. Certified semantics for relational programming. In *Programming Languages and Systems*, B. C. d. S. OLIVEIRA (Ed.), Springer International Publishing, Cham, pp. 167–185. https://doi.org/10.1007/978-3-030-64437-6_9
- SCHRIJVERS, T. and DEMOEN, B. 2004. The K.U. Leuven CHR system: implementation and application. In *CHR 2004*, Department of Computer Science, KU Leuven, pp. 1–5. <https://lirias.kuleuven.be/1653784&lang=en>
- SNEYERS, J., WEERT, P. V., SCHRIJVERS, T. and KONINCK, L. D. 2010. As time goes by: Constraint handling rules: A survey of CHR research from 1998 to 2007. In *Theory and Practice of Logic Programming*, vol. 10, issue 1, Cambridge University Press, pp. 1–47. <https://doi.org/10.1017/S1471068409990123>
- THIELSCHER, M. 2002. Reasoning about actions with CHRs and finite domain constraints. In *LP 2002*, P. J. STUCKEY (Ed.), *Lecture Notes in Computer Science*, Springer, Berlin, Heidelberg, pp. 70–84. <https://doi.org/10.1007/3-540-45619-8>

- THIELSCHER, M. 2005. FLUX: a logic programming method for reasoning agents. In *Theory and Practice of Logic Programming*, vol. 5, issue 4–5, Cambridge University Press, pp. 533–565. <https://doi.org/10.1017/S1471068405002358>
- VAN WEERT, P. 2010. Efficient lazy evaluation of rule-based programs. In *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, issue 11, pp. 1521–1534. <https://doi.org/10.1109/TKDE.2009.208>
- VAN WEERT, P., SCHRIJVERS, T., DEMOEN, B., SCHRIJVERS, T. and FRÜHWIRTH, T. 2005. K.U.Leuven JCHR: a user-friendly, flexible and efficient CHR system for java. In *CHR 2005*. Department of Computer Science, KU Leuven. <https://lirias.kuleuven.be/1654703>
- VAN WEERT, P., WUILLE, P., SCHRIJVERS, T. and DEMOEN, B. 2008. CHR for imperative host languages. In *Constraint Handling Rules: Current Research Topics*, T. SCHRIJVERS and T. FRÜHWIRTH (Eds.), Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, pp. 161–212. <https://doi.org/10.1007/978-3-540-92243-8>
- WIBIRAL, T. 2022. *JavaCHR – A Modern CHR-Embedding in Java*. Bachelor thesis, Ulm University. <http://dx.doi.org/10.18725/OPARU-43430>
- WUILLE, P., SCHRIJVERS, T. and DEMOEN, B. 2007. CCHR: The fastest CHR implementation. In *CHR 2007*, Department of Computer Science, KU Leuven, pp. 123–137. <https://lirias.kuleuven.be/1655099>