

Constraint Handling Rules - The Story So Far

Thom Frühwirth

University of Ulm, Germany

www.informatik.uni-ulm.de/pm/mitarbeiter/fruehwirth/

Abstract

Rule-based programming experiences renaissance due to its applications in areas such as Business Rules, Semantic Web, Computational Biology, Verification and Security. Executable rules are used in declarative programming languages, in program transformation and analysis, and for reasoning in artificial intelligence applications.

Constraint Handling Rules (CHR) [6, 8, 11] is a concurrent committed-choice constraint logic programming language consisting of guarded rules that transform multi-sets of atomic formulas (constraints) into simpler ones until exhaustion. CHR was initially developed for solving constraints, but has matured into a general-purpose concurrent constraint language over the last decade, because it can embed many rule-based formalisms and describe algorithms in a declarative way. The clean semantics of CHR facilitates non-trivial program analysis and transformation.

Categories and Subject Descriptors D.1.6 [Programming Techniques]: Logic Programming; D.3.2 [Programming Languages]: Language Classifications—Constraint and logic languages, Constraint Handling Rules; D.3.3 [Programming Languages]: Language Constructs and Features—Constraints

General Terms Algorithms, Languages, Theory

Keywords Constraint Programming, Constraint Solving, Computational Logic, Concurrency, Executable Specification, Rule-Based Programming, Program Analysis, Applications

CHR in a Nutshell

CHR programs consist of two main kinds of rules: **Simplification rules** replace constraints by simpler constraints while preserving logical equivalence, e.g., $X \leq Y, Y \leq X \Leftrightarrow X = Y$. **Propagation rules** add new constraints that are logically redundant but may cause further simplification, e.g., $X \leq Y, Y \leq Z \Rightarrow X \leq Z$. Together with $X \leq X \Leftrightarrow \text{true}$, these three rules encode the axioms of a partial order relation. The rules compute its transitive closure and replace \leq by equality $=$ whenever possible. For example, $A \leq B, B \leq C, C \leq A$ will be simplified into $A = B, A = C$.

Direct **ancestors of CHR** are logic programming, constraint logic programming [9] and concurrent committed-choice logic programming [10] languages. Like these languages, CHR has an operational semantics and a declarative semantics that are closely

related. Other influences were the chemical abstract machine [4], term rewriting systems, and, of course, production rule systems.

CHR is appealing for **computational logic**, because logical theories are usually specified by implications and logical equivalences, corresponding to propagation and simplification rules. On the meta-level, given the transformation rules for deduction in a calculus, inference rules map to propagation rules and replacement rules to simplification rules.

The use of CHR as a **general purpose programming language** is justified by the following observation: Given a state transition system, its transition rules can readily be expressed with simplification rules. In this way, dynamics and changes (e.g., updates) can be modelled, possibly triggered by events and handled by actions (that are all represented by atomic constraints). In such applications, conjunctions of constraints are best regarded as interacting collections of concurrent agents or processes. The standard declarative semantics based on first order predicate logic is likely to break down, but linear logic does the job [5].

Rule-based programming languages have the stigma of inefficiency. The paper [12] introduces CHR machines, analogous to RAM and Turing machines. It shows that these machines can simulate each other in polynomial time, thus establishing that CHR is Turing-complete and, more importantly, that every algorithm can be implemented in CHR with **best known time and space complexity**, something that is not known to be possible in other pure declarative programming languages like Prolog. These results hold in practice, as optimal and elegant implementations of algorithms like union-find, shortest paths and Fibonacci heaps have shown.

CHR libraries exist for most Prolog systems, several for Java and Haskell. Standard constraint systems as well as novel ones such as temporal, spatial, or description logic constraints have been implemented, many programs are available online. Besides constraint solvers, **applications of CHR** can be found in computational logic¹, in agent programming, multi-set rewriting and production rule systems. The several hundred publications [11] mentioning CHR cover such diverse applications as type system design for Haskell, time tabling for universities, optimal sender placement, computational linguistics, spatio-temporal reasoning, chip card verification, semantic web information integration, and decision support for cancer diagnosis.

One advantage of a declarative programming language is the ease of **program analysis**. Techniques for termination and time complexity, as well as confluence and operational equivalence of CHR have been investigated.

Since CHR is Turing-complete, **termination** is undecidable. For simplification rules, techniques from term rewriting can be adapted, for propagation rules from deductive databases. Both kinds of rules in one program can make termination proofs hard. From a termination order, an upper bound for the time **complex-**

Copyright is held by the author/owner(s).

PPDP'06 July 10–12, 2006, Venice, Italy.
ACM 1-59593-388-3/06/0007.

¹ Integrating deduction and abduction, bottom-up and top-down execution, forward and backward chaining, tabulation and integrity constraints.

