







# Certified Semantics for Relational Programming

Dmitry Rozplokhas<sup>1,3</sup>, Andrey Vyatkin<sup>2</sup>, and Dmitry Boulytchev<sup>2,3</sup>

<sup>1</sup> Higher School of Economics, Saint Petersburg, Russia

<sup>2</sup> Saint Petersburg State University, Saint Petersburg, Russia  
dboulytchev@math.spbu.ru

<sup>3</sup> JetBrains Research, Saint Petersburg, Russia

**Abstract.** We present a formal study of semantics for the relational programming language MINIKANREN. First, we formulate a denotational semantics which corresponds to the minimal Herbrand model for definite logic programs. Second, we present operational semantics which models interleaving, the distinctive feature of MINIKANREN implementation, and prove its soundness and completeness w.r.t. the denotational semantics. Our development is supported by a COQ specification, from which a reference interpreter can be extracted. We also derive from our main result a certified semantics (and a reference interpreter) for SLD resolution with cut and prove its soundness.

## 1 Introduction

In the context of this paper, we understand “relational programming” as a puristic form of logic programming with all extra-logical features banned. Specifically, we use MINIKANREN as an exemplary language; MINIKANREN can be seen as a logical language with explicit connectives, existentials and unification, and is mutually convertible to the pure logical subset of PROLOG.<sup>1</sup> Unlike PROLOG, which relies on SLD-resolution, most MINIKANREN implementations use a monadic *interleaving search*, which is known to be complete [15]. MINIKANREN is designed as a shallow DSL which may help to equip the host language with logical reasoning features. This design choice has been proven to be applicable in practice, and there are more than 100 implementations for almost 50 languages.

Although there already were attempts to define a formal semantics for MINIKANREN, none of them were capable of reflecting the distinctive property of MINIKANREN’s search—*interleaving* [18]. Since this distinctive search strategy is essential for the specification of the language and its extensions, the description of almost all development on miniKanren was not based on formal semantics. The introductory book on MINIKANREN [12] describes the language by means of

---

The reported study was funded by RFBR, project number 18-01-00380.

<sup>1</sup> A detailed PROLOG-to-MINIKANREN comparison can be found here: <http://minikanren.org/minikanren-and-prolog.html>.

© Springer Nature Switzerland AG 2020

B. C. d. S. Oliveira (Ed.): APLAS 2020, LNCS 12470, pp. 167–185, 2020.

[https://doi.org/10.1007/978-3-030-64437-6\\_9](https://doi.org/10.1007/978-3-030-64437-6_9)

$\mathcal{C} = \{C_i^{k_i}\}$	constructors with arities
$\mathcal{T}_X = X \cup \{C_i^{k_i}(t_1, \dots, t_{k_i}) \mid t_j \in \mathcal{T}_X\}$	terms over the set of variables $X$
$\mathcal{D} = \mathcal{T}_\emptyset$	ground terms
$\mathcal{X} = \{x, y, z, \dots\}$	syntactic variables
$\mathcal{A} = \{\alpha, \beta, \gamma, \dots\}$	semantic variables
$\mathcal{R} = \{R_i^{k_i}\}$	relational symbols with arities
$\mathcal{G} = \mathcal{T}_X \equiv \mathcal{T}_X$	unification
$\mathcal{G} \wedge \mathcal{G}$	conjunction
$\mathcal{G} \vee \mathcal{G}$	disjunction
<b>fresh</b> $\mathcal{X} . \mathcal{G}$	fresh variable introduction
$R_i^{k_i}(t_1, \dots, t_{k_i}), t_j \in \mathcal{T}_X$	relational symbol invocation
$S = \{R_i^{k_i} = \lambda x_1^i \dots x_{k_i}^i . g_i; \}$ $g$	specification

**Fig. 1.** The syntax of the source language

an evolving set of examples. In a series of follow-up papers [1, 7, 13–15, 30] various extensions of the language were presented with their semantics explained in terms of a SCHEME implementation. We argue that this style of semantic definition is fragile and not self-sufficient since it relies on concrete implementation languages’ semantics and therefore is not stable under the host language replacement. In addition, the justification of important properties of the language and specific relational programs becomes cumbersome.

In this paper, we present a formal semantics for core MINIKANREN and prove some of its basic properties. First, we define denotational semantics similar to the least Herbrand model for definite logic programs [23]; then we describe operational semantics with interleaving in terms of a labeled transition system. Finally, we prove soundness and completeness of the operational semantics w.r.t the denotational one. We support our development with a formal specification using the COQ proof assistant [4], thus outsourcing the burden of proof checking to the automatic tool and deriving a certified reference interpreter via the extraction mechanism. As a rather straightforward extension of our main result, we also provide a certified operational semantics (and a reference interpreter) for SLD resolution with cut, a new result to our knowledge; while this step brings us out of purely relational domain, it still can be interesting on its own.

## 2 The Language

In this section, we introduce the syntax of the language we use throughout the paper, describe the informal semantics, and give some examples.

The syntax of the language is shown in Fig. 1. First, we fix a set of constructors  $\mathcal{C}$  with known arities and consider a set of terms  $\mathcal{T}_X$  with constructors as functional symbols and variables from  $X$ . We parameterize this set with an alphabet of variables since in the semantic description we will need *two* kinds of variables. The first kind, *syntactic* variables, is denoted by  $\mathcal{X}$ . The second kind, *semantic* or *logic* variables, is denoted by  $\mathcal{A}$ . We also consider an alphabet of

*relational symbols*  $\mathcal{R}$  which are used to name relational definitions. The central syntactic category in the language is *goal*. In our case, there are five types of goals: *unification* of terms, conjunction and disjunction of goals, fresh variable introduction, and invocation of some relational definition. Thus, unification is used as a constraint, and multiple constraints can be combined using conjunction, disjunction, and recursion. The final syntactic category is a *specification*  $\mathcal{S}$ . It consists of a set of relational definitions and a top-level goal. A top-level goal represents a search procedure which returns a stream of substitutions for the free variables of the goal. The definition for a set of free variables for both terms and goals is conventional; as “**fresh**” is the sole binding construct the definition is rather trivial. The language we defined is first-order, as goals can not be passed as parameters, returned or constructed at run time.

We now informally describe how relational search works. As we said, a goal represents a search procedure. This procedure takes a *state* as input and returns a stream of states; a state (among other information) contains a substitution that maps semantic variables into the terms over semantic variables. Then five types of scenarios are possible (depending on the type of the goal):

- Unification “ $t_1 \equiv t_2$ ” unifies terms  $t_1$  and  $t_2$  in the context of the substitution in the current state. If terms are unifiable, then their MGU is integrated into the substitution, and a one-element stream is returned; otherwise the result is an empty stream.
- Conjunction “ $g_1 \wedge g_2$ ” applies  $g_1$  to the current state and then applies  $g_2$  to each element of the result, concatenating the streams.
- Disjunction “ $g_1 \vee g_2$ ” applies both its goals to the current state independently and then concatenates the results.
- Fresh construct “**fresh**  $x . g$ ” allocates a new semantic variable  $\alpha$ , substitutes all free occurrences of  $x$  in  $g$  with  $\alpha$ , and runs the goal.
- Invocation “ $R_i^{k_i}(t_1, \dots, t_{k_i})$ ” finds a definition for the relational symbol  $R_i^{k_i} = \lambda x_1 \dots x_{k_i} . g_i$ , substitutes all free occurrences of a formal parameter  $x_j$  in  $g_i$  with term  $t_j$  (for all  $j$ ) and runs the goal in the current state.

We stipulate that the top-level goal is preceded by an implicit “**fresh**” construct, which binds all its free variables, and that the final substitutions for these variables constitute the result of the goal evaluation.

Conjunction and disjunction form a monadic [32] interface with conjunction playing role of “**bind**” and disjunction the role of “**plus**”. In this description, we swept a lot of important details under the carpet—for example, in actual implementations the components of disjunction are not evaluated in isolation, but both disjuncts are evaluated incrementally with the control passing from one disjunct to another (*interleaving*) [18]; the evaluation of some goals can be additionally deferred (via so-called “*inverse- $\eta$ -delay*”) [13]; instead of streams the implementation can be based on “*ferns*” [8] to defer divergent computations, etc. In the following sections, we present a complete formal description of relational semantics which resolves these uncertainties in a conventional way.

As an example consider the following specification. For the sake of brevity we abbreviate immediately nested “**fresh**” constructs into the one, writing “**fresh**  $x y \dots . g$ ” instead of “**fresh**  $x$ . **fresh**  $y$ .  $\dots . g$ ”.

```

appendo = λ x y xy .
  ((x ≡ Nil) ∧ (xy ≡ y)) ∨
  (fresh h t ty .
    (x ≡ Cons (h, t)) ∧
    (xy ≡ Cons (h, ty)) ∧
    (appendo t y ty));

reverso = λ x xr .
  ((x ≡ Nil) ∧ (xr ≡ Nil)) ∨
  (fresh h t tr .
    (x ≡ Cons (h, t)) ∧
    (appendo tr (Cons (h, Nil)) xr) ∧
    (reverso t tr));

reverso x x

```

Here we defined<sup>2</sup> two relational symbols—“`appendo`” and “`reverso`”,—and specified a top-level goal “`reverso x x`”. The symbol “`appendo`” defines a relation of concatenation of lists—it takes three arguments and performs a case analysis on the first one. If the first argument is an empty list (“`Nil`”), then the second and the third arguments are unified. Otherwise, the first argument is deconstructed into a head “`h`” and a tail “`t`”, and the tail is concatenated with the second argument using a recursive call to “`appendo`” and additional variable “`ty`”, which represents the concatenation of “`t`” and “`y`”. Finally, we unify “`Cons (h, ty)`” with “`xy`” to form a final constraint. Similarly, “`reverso`” defines relational list reversing. The top-level goal represents a search procedure for all lists “`x`”, which are stable under reversing, i.e. palindromes. Running it results in an infinite stream of substitutions:

```

α ↦ Nil
α ↦ Cons (β0, Nil)
α ↦ Cons (β0, Cons (β0, Nil))
α ↦ Cons (β0, Cons (β1, Cons (β0, Nil)))
...

```

where “ $\alpha$ ” is a *semantic* variable, corresponding to “`x`”, “ $\beta_i$ ” are free semantic variables. Therefore, each substitution represents a set of all palindromes of a certain length.

### 3 Denotational Semantics

In this section, we present a denotational semantics for the language we defined above. We use a simple set-theoretic approach analogous to the least Herbrand model for definite logic programs [23]. Strictly speaking, instead of developing it from scratch we could have just described the conversion of specifications into definite logic form and took their least Herbrand model. However, in that case, we would still need to define the least Herbrand model semantics for definite logic programs in a certified way. In addition, while for this concrete language the conversion to definite logic form is trivial, it may become less trivial for its extensions (with, for example, nominal constructs [7]) which we plan to do in future.

<sup>2</sup> We respect here a conventional tradition for MINIKANREN programming to superscript all relational names with “<sup>o</sup>”.

We also must make the following observations. First, building inductive denotational semantics in a conventional way amounts to constructing a complete lattice and a monotone function and taking its least fixed point [31]. As we deal with a first-order language with only monotonic constructs (conjunction/disjunction) these steps are trivial. Moreover, we express the semantics in COQ, where all well-formed inductive definitions already have proper semantics, which removes the necessity to justify the validity of the steps we perform. Second, the least Herbrand model is traditionally defined as the least fixed point of a transition function (defined by a logic program) which maps sets of ground atoms to sets of ground atoms. We are, however, interested in *relational* semantics which should map a program into  $n$ -ary relation over ground terms, where  $n$  is the number of free variables in the topmost goal. Thus, we deviate from the traditional route and describe the denotational semantics in a more specific way.

To motivate further development, we first consider the following example. Let us have the following goal:

$$\mathbf{x} \equiv \text{Cons } (\mathbf{y}, \mathbf{z})$$

There are three free variables, and solving the goal delivers us the following single answer:

$$\alpha \mapsto \text{Cons } (\beta, \gamma)$$

where semantic variables  $\alpha$ ,  $\beta$  and  $\gamma$  correspond to the syntactic ones “ $\mathbf{x}$ ”, “ $\mathbf{y}$ ”, “ $\mathbf{z}$ ”. The goal does not put any constraints on “ $\mathbf{y}$ ” and “ $\mathbf{z}$ ”, so there are no bindings for “ $\beta$ ” and “ $\gamma$ ” in the answer. This answer can be seen as the following ternary relation over the set of all ground terms:

$$\{(\text{Cons } (\beta, \gamma), \beta, \gamma) \mid \beta \in \mathcal{D}, \gamma \in \mathcal{D}\} \subseteq \mathcal{D}^3$$

The order of “dimensions” is important, since each dimension corresponds to a certain free variable. Our main idea is to represent this relation as a set of total functions

$$f : \mathcal{A} \mapsto \mathcal{D}$$

from semantic variables to ground terms. We call these functions *representing functions*. Thus, we may reformulate the same relation as

$$\{(f(\alpha), f(\beta), f(\gamma)) \mid f \in \llbracket \alpha \equiv \text{Cons } (\beta, \gamma) \rrbracket\}$$

where we use conventional semantic brackets “ $\llbracket \bullet \rrbracket$ ” to denote the semantics. For the top-level goal, we need to substitute its free syntactic variables with distinct semantic ones, calculate the semantics, and build the explicit representation for the relation as shown above. The relation, obviously, does not depend on the concrete choice of semantic variables but depends on the order in which the values of representing functions are tupled. This order can be conventionalized, which gives us a completely deterministic semantics.

Now we implement this idea. First, for a representing function

$$f : \mathcal{A} \rightarrow \mathcal{D}$$

we introduce its homomorphic extension

$$\bar{f} : \mathcal{T}_{\mathcal{A}} \rightarrow \mathcal{D}$$

which maps terms to terms:

$$\begin{aligned} \bar{f}(\alpha) &= f(\alpha) \\ \bar{f}(C_i^{k_i}(t_1, \dots, t_{k_i})) &= C_i^{k_i}(\bar{f}(t_1), \dots, \bar{f}(t_{k_i})) \end{aligned}$$

Let us have two terms  $t_1, t_2 \in \mathcal{T}_{\mathcal{A}}$ . If there is a unifier for  $t_1$  and  $t_2$  then, clearly, there is a substitution  $\theta$  which turns both  $t_1$  and  $t_2$  into the same *ground* term (we do not require  $\theta$  to be the most general). Thus,  $\theta$  maps (some) variables into ground terms, and its application to  $t_{1(2)}$  is exactly  $\bar{\theta}(t_{1(2)})$ . This reasoning can be performed in the opposite direction: a unification  $t_1 \equiv t_2$  defines the set of all representing functions  $f$  for which  $\bar{f}(t_1) = \bar{f}(t_2)$ .

We will use the conventional notions of pointwise modification of a function  $f[x \leftarrow v]$  and substitution  $g[t/x]$  of a free variable  $x$  with a term  $t$  in a goal (or a term)  $g$ .

For a representing function  $f : \mathcal{A} \rightarrow \mathcal{D}$  and a semantic variable  $\alpha$  we define the following *generalization* operation:

$$f \uparrow \alpha = \{f[\alpha \leftarrow d] \mid d \in \mathcal{D}\}$$

Informally, this operation generalizes a representing function into a set of representing functions in such a way that the values of these functions for a given variable cover the whole  $\mathcal{D}$ . We extend the generalization operation for sets of representing functions  $\mathfrak{F} \subseteq \mathcal{A} \rightarrow \mathcal{D}$ :

$$\mathfrak{F} \uparrow \alpha = \bigcup_{f \in \mathfrak{F}} (f \uparrow \alpha)$$

Now we are ready to specify the semantics for goals (see Fig. 2). We've already given the motivation for the semantics of unification: the condition  $\bar{f}(t_1) = \bar{f}(t_2)$  gives us the set of all (otherwise unrestricted) representing functions which “equate” terms  $t_1$  and  $t_2$ . Set union and intersection provide a conventional interpretation for disjunction and conjunction of goals. In the case of a relational invocation we unfold the definition of the corresponding relational symbol and substitute its formal parameters with actual ones.

The only non-trivial case is that of “**fresh**  $x . g$ ”. First, we take an arbitrary semantic variable  $\alpha$ , not free in  $g$ , and substitute  $x$  with  $\alpha$ . Then we calculate the semantics of  $g[\alpha/x]$ . The interesting part is the next step: as  $x$  can not be free in “**fresh**  $x . g$ ”, we need to generalize the result over  $\alpha$  since in our model the semantics of a goal specifies a relation over its free variables. We introduce some nondeterminism by choosing arbitrary  $\alpha$ , but we can prove that with different choices of free variable the semantics of a goal does not change.

$$\begin{array}{ll}
 \llbracket t_1 \equiv t_2 \rrbracket & = \{\mathfrak{f} : \mathcal{A} \rightarrow \mathcal{D} \mid \bar{\mathfrak{f}}(t_1) = \bar{\mathfrak{f}}(t_2)\} & [\text{UNIFY}_D] \\
 \llbracket g_1 \wedge g_2 \rrbracket & = \llbracket g_1 \rrbracket \cap \llbracket g_2 \rrbracket & [\text{CONJ}_D] \\
 \llbracket g_1 \vee g_2 \rrbracket & = \llbracket g_1 \rrbracket \cup \llbracket g_2 \rrbracket & [\text{DISJ}_D] \\
 \llbracket \mathbf{fresh} \ x . g \rrbracket & = (\llbracket g[\alpha/x] \rrbracket) \uparrow \alpha, \alpha \notin FV(g) & [\text{FRESH}_D] \\
 \llbracket R(t_1, \dots, t_k) \rrbracket & = \llbracket g[t_1/x_1, \dots, t_k/x_k] \rrbracket, \text{ where } R = \lambda x_1 \dots x_k . g & [\text{INVOKE}_D]
 \end{array}$$

**Fig. 2.** Denotational semantics of goals

**Lemma 1.** *For any goal  $\mathbf{fresh} \ x . g$ , for any two variables  $\alpha$  and  $\beta$  which are not free in this goal, if  $\mathfrak{f} \in \llbracket g[\alpha/x] \rrbracket$ , then for any representing function  $\mathfrak{f}'$ , such that*

1.  $\mathfrak{f}'(\beta) = \mathfrak{f}(\alpha)$
  2.  $\forall \gamma : \gamma \neq \alpha \wedge \gamma \neq \beta, \mathfrak{f}'(\gamma) = \mathfrak{f}(\gamma)$
- it is true that  $\mathfrak{f}' \in \llbracket g[\beta/x] \rrbracket$ .*

The proof turned out to be the most cumbersome among all others in the case where  $g$  is a nested **fresh** construct. In that case, we have to constructively build two representing functions (including an intermediate one for an intermediate goal) by pointwise modification. The details of this proof can be found in the extended version of the paper.<sup>3</sup>

We can prove the following important *closedness condition* for the semantics of a goal  $g$ .

**Lemma 2 (Closedness condition).** *For any goal  $g$  and two representing functions  $\mathfrak{f}$  and  $\mathfrak{f}'$ , such that  $\mathfrak{f}|_{FV(g)} = \mathfrak{f}'|_{FV(g)}$ , it is true, that  $\mathfrak{f} \in \llbracket g \rrbracket \Leftrightarrow \mathfrak{f}' \in \llbracket g \rrbracket$ .*

In other words, representing functions for a goal  $g$  restrict only the values of free variables of  $g$  and do not introduce any “hidden” correlations. This condition guarantees that our semantics is closed in the sense that it does not introduce artificial restrictions for the relation it defines.

## 4 Operational Semantics

In this section we describe the operational semantics of MINIKANREN, which corresponds to the known implementations with interleaving search. The semantics is given in the form of a labeled transition system (LTS) [17]. From now on we assume the set of semantic variables to be linearly ordered ( $\mathcal{A} = \{\alpha_1, \alpha_2, \dots\}$ ).

We introduce the notion of substitution

$$\sigma : \mathcal{A} \rightarrow \mathcal{T}_{\mathcal{A}}$$

as a (partial) mapping from semantic variables to terms over the set of semantic variables. We denote  $\Sigma$  the set of all substitutions,  $Dom(\sigma)$ —the domain for a substitution  $\sigma$ ,  $\mathcal{VRan}(\sigma) = \bigcup_{\alpha \in Dom(\sigma)} \mathcal{FV}(\sigma(\alpha))$ —its range (the set of all free variables in the image).

<sup>3</sup> The extended version of this paper is available at <https://arxiv.org/abs/2005.01018>.

The *non-terminal states* in the transition system have the following shape:

$$S = \mathcal{G} \times \Sigma \times \mathbb{N} \mid S \oplus S \mid S \otimes \mathcal{G}$$

As we will see later, an evaluation of a goal is separated into elementary steps, and these steps are performed interchangeably for different subgoals. Thus, a state has a tree-like structure with intermediate nodes corresponding to partially-evaluated conjunctions (“ $\otimes$ ”) or disjunctions (“ $\oplus$ ”). A leaf in the form  $\langle g, \sigma, n \rangle$  determines a goal in a context, where  $g$  is a goal,  $\sigma$  is a substitution accumulated so far, and  $n$  is a natural number, which corresponds to a number of semantic variables used to this point. For a conjunction node, its right child is always a goal since it cannot be evaluated unless some result is provided by the left conjunct.

The full set of states also include one separate terminal state (denoted by  $\diamond$ ), which symbolizes the end of the evaluation.

$$\hat{S} = \diamond \mid S$$

We will operate with the well-formed states only, which are defined as follows.

**Definition 1.** *Well-formedness condition for extended states:*

- $\diamond$  is well-formed;
- $\langle g, \sigma, n \rangle$  is well-formed iff  $\mathcal{FV}(g) \cup \text{Dom}(\sigma) \cup \mathcal{VRan}(\sigma) \subseteq \{\alpha_1, \dots, \alpha_n\}$ ;
- $s_1 \oplus s_2$  is well-formed iff  $s_1$  and  $s_2$  are well-formed;
- $s \otimes g$  is well-formed iff  $s$  is well-formed and for all leaf triplets  $\langle -, -, n \rangle$  in  $s$  it is true that  $\mathcal{FV}(g) \subseteq \{\alpha_1, \dots, \alpha_n\}$ .

Informally the well-formedness restricts the set of states to those in which all goals use only allocated variables.

Finally, we define the set of labels:

$$L = \circ \mid \Sigma \times \mathbb{N}$$

The label “ $\circ$ ” is used to mark those steps which do not provide an answer; otherwise, a transition is labeled by a pair of a substitution and a number of allocated variables. The substitution is one of the answers, and the number is threaded through the derivation to keep track of allocated variables.

The transition rules are shown in Fig. 3. The first two rules specify the semantics of unification. If two terms are not unifiable under the current substitution  $\sigma$  then the evaluation stops with no answer; otherwise, it stops with the most general unifier applied to a current substitution as an answer.

The next two rules describe the steps performed when disjunction or conjunction is encountered on the top level of the current goal. For disjunction, it schedules both goals (using “ $\oplus$ ”) for evaluating in the same context as the parent state, for conjunction—schedules the left goal and postpones the right one (using “ $\otimes$ ”).

The rule for “**fresh**” substitutes bound syntactic variable with a newly allocated semantic one and proceeds with the goal.



The rule for relation invocation finds a corresponding definition, substitutes its formal parameters with the actual ones, and proceeds with the body.

The rest of the rules specify the steps performed during the evaluation of two remaining types of the states—conjunction and disjunction. In all cases, the left state is evaluated first. If its evaluation stops, the disjunction evaluation proceeds with the right state, propagating the label (SUMSTOP and SUMSTEP), and the conjunction schedules the right goal for evaluation in the context of the returned answer (PRODSTOPANS) or stops if there is no answer (PRODSTOP).

$$\begin{array}{l}
 \langle t_1 \equiv t_2, \sigma, n \rangle \xrightarrow{\circ} \diamond, \nexists \text{ mgu}(t_1\sigma, t_2\sigma) \quad [\text{UNIFYFAIL}] \\
 \langle t_1 \equiv t_2, \sigma, n \rangle \xrightarrow{(\text{mgu}(t_1\sigma, t_2\sigma)\circ\sigma, n)} \diamond \quad [\text{UNIFYSUCCESS}] \\
 \langle g_1 \vee g_2, \sigma, n \rangle \xrightarrow{\circ} \langle g_1, \sigma, n \rangle \oplus \langle g_2, \sigma, n \rangle \quad [\text{DISJ}] \\
 \langle g_1 \wedge g_2, \sigma, n \rangle \xrightarrow{\circ} \langle g_1, \sigma, n \rangle \otimes g_2 \quad [\text{CONJ}] \\
 \langle \text{fresh } x.g, \sigma, n \rangle \xrightarrow{\circ} \langle g[\alpha_{n+1}/x], \sigma, n+1 \rangle \quad [\text{FRESH}] \\
 \frac{R_i^{k_i} = \lambda x_1 \dots x_{k_i}. g}{\langle R_i^{k_i}(t_1, \dots, t_{k_i}), \sigma, n \rangle \xrightarrow{\circ} \langle g[t_1/x_1 \dots t_{k_i}/x_{k_i}], \sigma, n \rangle} \quad [\text{INVOKE}] \\
 \frac{s_1 \xrightarrow{\circ} \diamond}{(s_1 \oplus s_2) \xrightarrow{\circ} s_2} \quad [\text{SUMSTOP}] \\
 \frac{s_1 \xrightarrow{r} \diamond}{(s_1 \oplus s_2) \xrightarrow{r} s_2} \quad [\text{SUMSTOPANS}] \\
 \frac{s \xrightarrow{\circ} \diamond}{(s \otimes g) \xrightarrow{\circ} \diamond} \quad [\text{PRODSTOP}] \\
 \frac{s \xrightarrow{(\sigma, n)} \diamond}{(s \otimes g) \xrightarrow{\circ} \langle g, \sigma, n \rangle} \quad [\text{PRODSTOPANS}] \\
 \frac{s_1 \xrightarrow{\circ} s'_1}{(s_1 \oplus s_2) \xrightarrow{\circ} (s_2 \oplus s'_1)} \quad [\text{SUMSTEP}] \\
 \frac{s_1 \xrightarrow{r} s'_1}{(s_1 \oplus s_2) \xrightarrow{r} (s_2 \oplus s'_1)} \quad [\text{SUMSTEPANS}] \\
 \frac{s \xrightarrow{\circ} s'}{(s \otimes g) \xrightarrow{\circ} (s' \otimes g)} \quad [\text{PRODSTEP}] \\
 \frac{s \xrightarrow{(\sigma, n)} s'}{(s \otimes g) \xrightarrow{\circ} (\langle g, \sigma, n \rangle \oplus (s' \otimes g))} \quad [\text{PRODSTEPANS}]
 \end{array}$$

**Fig. 3.** Operational semantics of interleaving search

The last four rules describe *interleaving*, which occurs when the evaluation of the left state suspends with some residual state (with or without an answer). In the case of disjunction the answer (if any) is propagated, and the constituents of the disjunction are swapped (SUMSTEP, SUMSTEPANS). In the case of conjunction, if the evaluation step in the left conjunct did not provide any answer, the evaluation is continued in the same order since there is still no information to

proceed with the evaluation of the right conjunct (PRODSTEP); if there is some answer, then the disjunction of the right conjunct in the context of the answer and the remaining conjunction is scheduled for evaluation (PRODSTEPANS).

The introduced transition system is completely deterministic: there is exactly one transition from any non-terminal state. There is, however, some freedom in choosing the order of evaluation for conjunction and disjunction states. For example, instead of evaluating the left substate first, we could choose to evaluate the right one, etc. This choice reflects the inherent non-deterministic nature of search in relational (and, more generally, logical) programming. Although we could introduce this ambiguity into the semantics (by replacing specific rules for disjunctions and conjunctions evaluation with some conditions on it), we want an operational semantics that would be easy to present and easy to employ to describe existing language extensions (already described for a specific implementation of interleaving search), so we instead base the semantics on one canonical search strategy. At the same time, as long as deterministic search procedures are sound and complete, we can consider them “equivalent”.<sup>4</sup>

It is easy to prove that transitions preserve well-formedness of states.

**Lemma 3.** (*Well-formedness preservation*) *For any transition  $s \xrightarrow{l} \hat{s}$ , if  $s$  is well-formed then  $\hat{s}$  is also well-formed.*

A derivation sequence for a certain state determines a *trace*—a finite or infinite sequence of answers. The trace corresponds to the stream of answers in the reference MINIKANREN implementations. We denote a set of answers in the trace for state  $\hat{s}$  by  $\mathcal{T}r_{\hat{s}}$ .

We can relate sets of answers for the partially evaluated conjunction and disjunction with sets of answers for their constituents by the two following lemmas.

**Lemma 4.** *For any non-terminal states  $s_1$  and  $s_2$ ,  $\mathcal{T}r_{s_1 \oplus s_2} = \mathcal{T}r_{s_1} \cup \mathcal{T}r_{s_2}$ .*

**Lemma 5.** *For any non-terminal state  $s$  and goal  $g$ ,  $\mathcal{T}r_{s \otimes g} \supseteq \bigcup_{(\sigma, n) \in \mathcal{T}r_s} \mathcal{T}r_{\langle g, \sigma, n \rangle}$ .*

These two lemmas constitute the exact conditions on definition of these operators that we will use to prove the completeness of an operational semantics.

We also can easily describe the criterion of termination for disjunctions.

**Lemma 6.** *For any goals  $g_1$  and  $g_2$ , substitution  $\sigma$ , and number  $n$ , the trace from the state  $\langle g_1 \vee g_2, \sigma, n \rangle$  is finite iff the traces from both  $\langle g_1, \sigma, n \rangle$  and  $\langle g_2, \sigma, n \rangle$  are finite.*

These simple statements already allow us to prove two important properties of interleaving search as corollaries: the “fairness” of disjunction—the fact that the trace for disjunction contains all the answers from both streams for disjuncts—and the “commutativity” of disjunctions—the fact that swapping two disjuncts (at the top level) does not change the termination of the goal evaluation.

<sup>4</sup> There still can be differences in observable behavior of concrete goals under different sound and complete search strategies. For example, a goal can be refutationally complete [6] under one strategy and non-complete under another.

## 5 Equivalence of Semantics

Now we can relate two different kinds of semantics for MINIKANREN described in the previous sections and show that the results given by these two semantics are the same for any specification. This will actually say something important about the search in the language: since operational semantics describes precisely the behavior of the search and denotational semantics ignores the search and describes what we *should* get from a mathematical point of view, by proving their equivalence we establish the *completeness* of the search, which means that the search will get all answers satisfying the described specification and only those.

$$\begin{aligned}
 \llbracket \diamond \rrbracket &= \emptyset \\
 \llbracket \langle g, \sigma, n \rangle \rrbracket &= \llbracket g \rrbracket \cap \llbracket \sigma \rrbracket \\
 \llbracket s_1 \oplus s_2 \rrbracket &= \llbracket s_1 \rrbracket \cup \llbracket s_2 \rrbracket \\
 \llbracket s \otimes g \rrbracket &= \llbracket s \rrbracket \cap \llbracket g \rrbracket
 \end{aligned}$$

**Fig. 4.** Denotational semantics of states

But first, we need to relate the answers produced by these two semantics as they have different forms: a trace of substitutions (along with the numbers of allocated variables) for the operational one and a set of representing functions for the denotational one. We can notice that the notion of a representing function is close to substitution, with only two differences:

- representing functions are total;
- terms in the domain of representing functions are ground.

Therefore we can easily extend (perhaps ambiguously) any substitution to a representing function by composing it with an arbitrary representing function preserving all variable dependencies in the substitution. So we can define a set of representing functions that correspond to a substitution as follows:

$$\llbracket \sigma \rrbracket = \{ \bar{f} \circ \sigma \mid f : \mathcal{A} \mapsto \mathcal{D} \}$$

And the *denotational analog* of operational semantics (a set of representing functions corresponding to the answers in the trace) for a given state  $\hat{s}$  is then defined as the union of sets for all substitutions in the trace:

$$\llbracket \hat{s} \rrbracket_{op} = \cup_{(\sigma, n) \in \mathcal{T}r_{\hat{s}}} \llbracket \sigma \rrbracket$$

This allows us to state theorems relating the two semantics.

**Theorem 1 (Operational semantics soundness).** *If indices of all free variables in a goal  $g$  are limited by some number  $n$ , then  $\llbracket \langle g, \epsilon, n \rangle \rrbracket_{op} \subseteq \llbracket g \rrbracket$ .*

It can be proven by nested induction, but first, we need to generalize the statement so that the inductive hypothesis is strong enough for the inductive step.

To do so, we define denotational semantics not only for goals but for arbitrary states. Note that this definition does not need to have any intuitive interpretation, it is introduced only for the proof to go smoothly. The definition of the denotational semantics for extended states is shown on Fig. 4. The generalized version of the theorem uses it.

**Lemma 7 (Generalized soundness).** *For any well-formed state  $\hat{s}$*

$$\llbracket \hat{s} \rrbracket_{op} \subseteq \llbracket \hat{s} \rrbracket.$$

It can be proven by the induction on the number of steps in which a given answer (more accurately, the substitution that contains it) occurs in the trace. We break the proof in two parts and separately prove by induction on evidence that for every transition in our system the semantics of both the label (if there is one) and the next state are subsets of the denotational semantics for the initial state.

**Lemma 8 (Soundness of the answer).** *For any transition  $s \xrightarrow{(\sigma, n)} \hat{s}$ ,  $\llbracket \sigma \rrbracket \subseteq \llbracket s \rrbracket$ .*

**Lemma 9 (Soundness of the next state).** *For any transition  $s \xrightarrow{l} \hat{s}$ ,  $\llbracket \hat{s} \rrbracket \subseteq \llbracket s \rrbracket$ .*

It would be tempting to formulate the completeness of operational semantics as soundness with the inverted inclusion, but it does not hold in such generality. The reason for this is that the denotational semantics encodes only the dependencies between free variables of a goal, which is reflected by the closedness condition, while the operational semantics may also contain dependencies between semantic variables allocated in **fresh** constructs. Therefore we formulate completeness with representing functions restricted on the semantic variables allocated in the beginning (which includes all free variables of a goal). This does not compromise our promise to prove the completeness of the search as MINIKANREN returns substitutions only for queried variables, which are allocated in the beginning.

**Theorem 2 (Operational semantics completeness).** *If the indices of all free variables in a goal  $g$  are limited by some number  $n$ , then*

$$\{f|_{\{\alpha_1, \dots, \alpha_n\}} \mid f \in \llbracket g \rrbracket\} \subseteq \{f|_{\{\alpha_1, \dots, \alpha_n\}} \mid f \in \llbracket \langle g, \epsilon, n \rangle \rrbracket_{op}\}.$$

Similarly to the soundness, this can be proven by nested induction, but the generalization is required. This time it is enough to generalize it from goals to states of the shape  $\langle g, \sigma, n \rangle$ . We also need to introduce one more auxiliary semantics—*step-indexed denotational semantics* (denoted by  $\llbracket \bullet \rrbracket^i$ ). It is an implementation of the well-known approach [2] of indexing typing or semantic logical relations by a number of permitted evaluation steps to allow inductive reasoning

on it. In our case,  $\llbracket g \rrbracket^i$  includes only those representing functions that one can get after no more than  $i$  unfoldings of relational calls.

The step-indexed denotational semantics is an approximation of the conventional denotational semantics; it is clear that any answer in conventional denotational semantics will also be in step-indexed denotational semantics for some number of steps.

**Lemma 10.**  $\llbracket g \rrbracket \subseteq \cup_i \llbracket g \rrbracket^i$

Now the generalized version of the completeness theorem is as follows.

**Lemma 11 (Generalized completeness).** *For any set of relational definitions, for any number of unfoldings  $i$ , for any well-formed state  $\langle g, \sigma, n \rangle$ ,*

$$\{f|_{\{\alpha_1, \dots, \alpha_n\}} \mid f \in \llbracket g \rrbracket^i \cap \llbracket \sigma \rrbracket\} \subseteq \{f|_{\{\alpha_1, \dots, \alpha_n\}} \mid f \in \llbracket \langle g, \sigma, n \rangle \rrbracket_{op}\}.$$

The proof is by the induction on number of unfoldings  $i$ . The induction step is proven by structural induction on goal  $g$ . We use Lemmas 4 and 5 for evaluation of a disjunction and a conjunction respectively, and Lemma 1 in the case of fresh variable introduction to move from an arbitrary semantic variable in denotational semantics to the next allocated fresh variable. The details of this proof may be found in the extended version of the paper.

## 6 Specification in Coq

We certified all the definitions and propositions from the previous sections using the COQ proof assistant.<sup>5</sup> The COQ specification for the most part closely follows the formal descriptions we gave by means of inductive definitions (and inductively defined propositions in particular) and structural induction in proofs. The detailed description of the specification, including code snippets, is provided in the extended version of the paper, and in this section we address only some non-trivial parts of it and some design choices.

The language formalized in COQ has a few non-essential simplifications for the sake of convenience. Specifically, we restrict the arities of all constructors to be either zero or two and require all relations to have exactly one argument. These restrictions do not make the language less expressive in any way since we can always represent a sequence of terms as a list using constructors `Nil`<sup>0</sup> and `Cons`<sup>2</sup>.

In our formalization of the language we use higher-order abstract syntax [27] for variable binding, therefore we work explicitly only with semantic variables. We preferred it to the first-order syntax because it gives us the ability to use substitution and the induction principle provided by COQ. On the other hand, we need to explicitly specify a requirement on the syntax representation, which is trivially fulfilled in the first-order case: all bindings have to be “consistent”, i.e. if

<sup>5</sup> The specification is available at <https://github.com/dboulytchev/miniKanren-coq>.

we instantiate a higher-order **fresh** construct with different semantic variables the results will be the same up to some renaming (provided that both those variables are not free in the body of the binder). Another requirement we have to specify explicitly (independent of HOAS/FOAS dichotomy) is a requirement that the definitions of relations do not contain unbound semantic variables.

To formalize the operational semantics in COQ we first need to define all preliminary notions from unification theory [3] which our semantics uses. In particular, we need to implement the notion of the most general unifier (MGU). As it is well-known [25] all standard recursive algorithms for calculating MGU are not decreasing on argument terms, so we can't define them as simple recursive functions in COQ due to the termination check failure. The standard approach to tackle this problem is to define the function through well-founded recursion. We use a distinctive version of this approach, which is more convenient for our purposes: we define MGU as a proposition (for which there is no termination requirement in COQ) with a dedicated structurally-recursive function for one step of unification, and then we use a well-founded induction to prove the existence of a corresponding result for any arguments and defining properties of MGU. For this well-founded induction, we use the number of distinct free variables in argument terms as a well-founded order on pairs of terms.

In the operational semantics, to define traces as (possibly) infinite sequences of transitions we use the standard approach in COQ—coinductively defined streams. Operating with them requires a number of well-known tricks, described by Chlipala [9], to be applied, such as the use of a separate coinductive definition of equality on streams.

The final proofs of soundness and completeness of operational semantics are relatively small, but the large amount of work is hidden in the proofs of auxiliary facts that they use (including lemmas from the previous sections and some technical machinery for handling representing functions).

## 7 Applications

In this section, we consider some applications of the framework and results, described in the previous sections.

### 7.1 Correctness of Transformations

One important immediate corollary of the equivalence theorems we have proven is the justification of correctness for certain program transformations. The completeness of interleaving search guarantees the correctness of any transformation that preserves the denotational semantics, for example:

- changing the order of constituents in conjunctions and disjunctions;
- distributing conjunctions over disjunctions and vice versa, for example, normalizing goals into CNF or DNF;
- moving fresh variable introduction upwards/downwards, for example, transforming any relation into a top-level fresh construct with a freshless body.

Note that this way we can guarantee only the preservation of results as *sets of ground terms*; the other aspects of program behavior, such as termination, may be affected by some of these transformations.<sup>6</sup>

One of the applications for these transformations is a conversion from/to PROLOG. As both languages use essentially the same fragment of first-order logic, their programs are mutually convertible. The conversion from PROLOG to MINIKANREN is simpler as the latter admits a richer syntax of goals. The inverse conversion involves the transformation into a DNF and splitting the disjunction into a number of separate clauses. This transformation, in particular, makes it possible to reuse our approach to describe the semantics of PROLOG as well. In the following sections we briefly address this problem.

## 7.2 SLD Semantics

The conventional PROLOG SLD search differs from the interleaving one in just one aspect—it does not perform interleaving. Thus, changing just two rules in the operational semantics converts interleaving search into the depth-first one:

$$\frac{s_1 \xrightarrow{o} s'_1}{(s_1 \oplus s_2) \xrightarrow{o} (s'_1 \oplus s_2)} \text{ [DISJSTEP]} \quad \frac{s_1 \xrightarrow{r} s'_1}{(s_1 \oplus s_2) \xrightarrow{r} (s'_1 \oplus s_2)} \text{ [DISJSTEPANS]}$$

With this definition we can almost completely reuse the mechanized proof of soundness (with minor changes); the completeness, however, can no longer be proven (as it does not hold anymore).

## 7.3 Cut

Dealing with the “cut” construct is known to be a cornerstone feature in the study of operational semantics for PROLOG. It turned out that in our case the semantics of “cut” can be expressed naturally (but a bit verbosely). Unlike SLD-resolution, it does not amount to an incremental change in semantics description. It also would work only for programs directly converted from PROLOG specifications.

The key observation in dealing with the “cut” in our setting is that a state in our semantics, in fact, encodes the whole current search tree (including all backtracking possibilities). This opens the opportunity to organize proper “navigation” through the tree to reflect the effect of “cut”. The details of the semantic description can be found in the extended version of the paper.

For this semantics, we can repeat the proof of soundness w.r.t. to the denotational semantics. There is, however, a little subtlety with our construction: we cannot formally prove that our semantics indeed encodes the conventional meaning of “cut” (since we do not have other semantics of “cut” to compare with).

<sup>6</sup> Possible slowdown and loss of termination after reorderings in conjunction is a famous example of this phenomenon in MINIKANREN, known as conjunction non-commutativity [6].

Nevertheless, we can demonstrate a plausible behavior using the extracted reference interpreter.

## 7.4 Reference Interpreters

Using the COQ extraction mechanism, we extracted two reference interpreters from our definitions and theorems: one for conventional MINIKANREN with interleaving search and another one for SLD search with cut. These interpreters can be used to practically investigate the behavior of specifications in unclear, complex, or corner cases. Our experience has shown that these interpreters demonstrate the expected behavior in all cases.

## 8 Related Work

The study of formal semantics for logic programming languages, particularly PROLOG, is a well-established research domain. Early works [10,16] addressed the computational aspects of both pure PROLOG and its extension with the cut construct. Recently, the application of certified/mechanized approaches came into focus as well. In particular, in one work [21] the equivalence of a few differently defined semantics for pure PROLOG is proven, and in another work [20] a denotational semantics for PROLOG with cut is presented; both works provide COQ-mechanized proofs. It is interesting that the former one also advocates the use of higher-order abstract syntax. We are not aware of any prior work on certified semantics for PROLOG which contributed a correct-by-construction interpreter. Our certified description of SLD resolution with cut can be considered as a certified semantics for PROLOG modulo occurs check in unification (which PROLOG does not have by default).

The implementation of first-order unification in dependently typed languages constitutes a well-known challenge with a number of known solutions. The major difficulty comes from the non-structural recursivity of conventional unification algorithms, which requires to provide a witness for convergence. The standard approach is to define a generally-recursive function and a well-founded order for its arguments. This route is taken in a number of works [5,19,26,28], where the descriptions of unification algorithms are given in COQ, LCF and ALF. The well-founded used there is lexicographically ordered tuples, containing the information about the number of different free variables and the sizes of the arguments. We implement a similar approach, but we separate the test for the non-matching case into a dedicated function. Thus, we make a recursive call only when the current substitution extension is guaranteed, which allows us to use the number of different free variables as the well-founded order. An alternative approach suggested by McBride [25] gives a structurally recursive definition of the unification algorithm; this is achieved by indexing the arguments with the numbers of their free variables.

The use of higher-order abstract syntax (HOAS) for dealing with language constructs in COQ was addressed in early work [11], where it was employed to



describe the lambda calculus. The inconsistency phenomenon of HOAS representation, mentioned in Sect. 6, is called there “exotic terms” there and is handled using a dedicated inductive predicate “Valid\_v”. The predicate has a non-trivial implementation based on subtle observations on the behavior of bindings. Our case, however, is much simpler: there is not much variety in “exotic terms” (for example, we do not have reductions in terms), and our consistency predicate can be considered as a limited version of “Valid\_v” for a more limited language.

The study of formal semantics for MINIKANREN is not a completely novel venture. Previously, a nondeterministic small-step semantics was described [24], as well as a big-step semantics for a finite number of answers [29]; neither uses proof mechanization and in both works the interleaving is not addressed.

The work of Kumar [22] can be considered as our direct predecessor. It also introduces both denotational and operational semantics and presents a HOL-certified proof for the soundness of the latter w.r.t. the former. The denotational semantics resembles ours but considers only queries with a single free variable (we do not see this restriction as important). On the other hand, the operational semantics is non-deterministic, which makes it impossible to express interleaving and extract the interpreter in a direct way. In addition, a specific form of “executable semantics” is introduced, but its connection to the other two is not established. Finally, no completeness result is presented. We consider our completeness proof as an essential improvement.

The most important property of interleaving search—completeness—was postulated in the introductory paper [18], and is delivered by all major implementations. Hemann et al. [15] give a proof of completeness for a specific implementation of MINIKANREN; however, the completeness is understood there as preservation of all answers during the interleaving of answer streams, i.e. in a more narrow sense than in our work since no relation to denotational semantics is established.

## 9 Conclusion and Future Work

In this paper, we presented a certified formal semantics for core MINIKANREN and proved some of its basic properties (including interleaving search completeness, disjunction fairness and commutativity), which are believed to hold in existing implementations. We also derived a semantics for conventional SLD resolution with cut and extracted two certified reference interpreters. We consider our work as the initial setup for the future development of MINIKANREN semantics.

The language we considered here lacks many important features, which are already introduced and employed in many implementations. Integrating these extensions—in the first hand, disequality constraints,—into the semantics looks a natural direction for future work. We are also going to address the problems of proving some properties of relational programs (equivalence, refutational completeness, etc.).

## References

1. Alvis, C.E., Willcock, J.J., Carter, K.M., Byrd, W.E., Friedman, D.P.: cKanren: miniKanren with constraints. In: Proceedings of the 2011 Annual Workshop on Scheme and Functional Programming (2011)
2. Appel, A.W., McAllester, D.A.: An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.* **23**(5), 657–683 (2001)
3. Baader, F., Snyder, W.: Handbook of automated reasoning. In: Unification Theory. Elsevier Science Publishers B. V., Amsterdam, The Netherlands (2001)
4. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development - Coq'Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. An EATCS Series. Springer (2004)
5. Bove, A.: Programming in martin-löf type theory: Unification - a non-trivial example, pp. 22–42, Department of Computer Science, Chalmers University of Technology (1999)
6. Byrd, W.E.: Relational Programming in miniKanren: Techniques, Applications, and Implementations. PhD thesis, Indiana University (2009)
7. Byrd, W.E., Friedman, D.P.:  $\alpha$ kanren: a fresh name in nominal logic programming. In: Proceedings of the 2007 Annual Workshop on Scheme and Functional Programming, pp. 79–90 (2007)
8. Byrd, W.E., Friedman, D.P., Kumar, R., Near, J.P.: A shallow Scheme embedding of bottom-avoiding streams. In: To appear in a special issue of Higher-Order and Symbolic Computation, in honor of Mitchell Wand's 60th birthday
9. Chlipala, A.: Certified Programming with Dependent Types - A Pragmatic Introduction to the Coq Proof Assistant. MIT Press, Cambridge (2013)
10. Debray, S.K., Mishra, P.: Denotational and operational semantics for PROLOG. In: Formal Description of Programming Concepts - III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference on Formal Description of Programming Concepts - III, Ebberup, Denmark, 25–28 August 1986, pp. 245–274 (1987)
11. Despeyroux, J., Felty, A., Hirschowitz, A.: Higher-order abstract syntax in Coq. In: Dezani-Ciancaglini, M., Plotkin, G. (eds.) TLCA 1995. LNCS, vol. 902, pp. 124–138. Springer, Heidelberg (1995). <https://doi.org/10.1007/BFb0014049>
12. Friedman, D.P., Byrd, W.E., Kiselyov, O.: The Reasoned Schemer. MIT Press, Cambridge (2005)
13. Hemann, J., Friedman, D.P.:  $\mu$ Kanren: a minimal functional core for relational programming. In: Proceedings of the 2013 Annual Workshop on Scheme and Functional Programming (2013)
14. Hemann, J., Friedman, D.P.: A framework for extending microKanren with constraints. In Proceedings 29th and 30th Workshops on (Constraint) Logic Programming and 24th International Workshop on Functional and (Constraint) Logic Programming, WLP 2015 / WLP 2016 / WFLP 2016, Dresden and Leipzig, Germany, 22nd September 2015 and 12–14th September 2016, pp. 135–149 (2017)
15. Hemann, J., Friedman, D.P., Byrd, W.E., Might, M.: A small embedding of logic programming with a simple complete search. In: Proceedings of the 12th Symposium on Dynamic Languages, DLS 2016, Amsterdam, The Netherlands, 1 Nov 2016, pp. 96–107 (2016)
16. Jones, N.D., Mycroft, A.: Stepwise development of operational and denotational semantics for Prolog. In: Proceedings of the 1984 International Symposium on Logic Programming, Atlantic City, New Jersey, USA, 6–9 Feb 1984, pp. 281–288 (1984)

17. Keller, R.M.: Formal verification of parallel programs. *Commun. ACM* **19**(7), 371–384 (1976)
18. Kiselyov, O., Shan, C., Friedman, D.P., Sabry, A.: Backtracking, interleaving, and terminating monad transformers: (functional pearl), pp. 192–203 (2005)
19. Kothari, S., Caldwell, J.: A machine checked model of idempotent MGU axioms for lists of equational constraints. In: *Proceedings 24th International Workshop on Unification, UNIF 2010, Edinburgh, United Kingdom, 14th July 2010*, pp. 24–38 (2010)
20. Kriener, J., King, A.: Semantics for Prolog with cut - revisited. In: *Functional and Logic Programming - 12th International Symposium, FLOPS 2014, Kanazawa, Japan, 4–6 June 2014, Proceedings*, pp. 270–284 (2014)
21. Kriener, J., King, A., Blazy, S.: Proofs you can believe. In: *proving equivalences between Prolog semantics in Coq*. In: *15th International Symposium on Principles and Practice of Declarative Programming, PPDP '13, Madrid, Spain, 16–18 Sept 2013*, pp. 37–48 (2013)
22. Kumar, R.: Mechanising aspects of miniKanren in HOL. Bachelor Thesis, The Australian National University (2010)
23. Lloyd, J.W.: *Foundations of Logic Programming*, 1st edn. Springer (1984)
24. Lozov, P., Vyatkin, A., Boulytchev, D.: Typed relational conversion. In: Wang, M., Owens, S. (eds.) *TFP 2017*. LNCS, vol. 10788, pp. 39–58. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-89719-6\\_3](https://doi.org/10.1007/978-3-319-89719-6_3)
25. McBride, C.: First-order unification by structural recursion. *J. Funct. Program.* **13**(6), 1061–1075 (2003)
26. Paulson, L.C.: Verifying the unification algorithm in LCF. *Sci. Comput. Program.* **5**(2), 143–169 (1985)
27. Pfenning, F., Elliott, C.: Higher-Order Abstract Syntax, pp. 199–208 (1988)
28. Ribeiro, R., Camarão, C.: A mechanized textbook proof of a type unification algorithm. In: Cornélio, M., Roscoe, B. (eds.) *SBMF 2015*. LNCS, vol. 9526, pp. 127–141. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-29473-5\\_8](https://doi.org/10.1007/978-3-319-29473-5_8)
29. Rozplochas, D., Boulytchev, D.: Improving refutational completeness of relational search via divergence test. In: *Proceedings of the 20th International Symposium on Principles and Practice of Declarative Programming, PPDP 2018, Frankfurt am Main, Germany, 03–05 Sept 2018*, pp. 18:1–18:13 (2018)
30. Swords, C., Friedman, D.P.: rKanren: guided search in miniKanren. In: *Proceedings of the 2013 Annual Workshop on Scheme and Functional Programming* (2013)
31. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. *Pac. J. Math.* **5**, 06 (1955)
32. Wadler, P.: Monads for functional programming. In: *Advanced Functional Programming, First International Spring School on Advanced Functional Programming Techniques, Båstad, Sweden, 24–30 May 1995, Tutorial Text*, pp. 24–52 (1995)