Semi-Automated Direction-Driven Functional Conversion

EKATERINA VERBITSKAIA, JetBrains Research, Serbia

IGOR ENGEL, JetBrains Research, Germany

DANIIL BEREZUN, JetBrains Research, Netherlands

One of the most attractive applications of relational programming is inverse computation. It offers an approach to solving complex problems by transforming verifiers into solvers with relatively low effort. Unfortunately, inverse computation often suffers from interpretation overhead, leading to subpar performance compared to direct program inversion. A prior study introduced a functional conversion scheme capable of creating inversions of MINIKANREN specifications with respect to a known fixed direction. This paper expands upon it by providing a semi-automated functional conversion algorithm. Our evaluation demonstrates a significant performance improvement achieved through functional conversion.

CCS Concepts: • **Software and its engineering → Constraint and logic languages**.

Additional Key Words and Phrases: program inversion, inverse computations, relational programming, functional programming, conversion

1 INTRODUCTION

One of the most attractive applications of relational programming is *inverse computation*. It is helpful, when the program being inverted is a relational interpreter of some sort: this way an interpreter for a programming language may be used for program synthesis, a type checker — to solve type inhabitation problem and so on [3, 4]. Constructing relational interpreters out of functional implementations can be done automatically by relational conversion [5]. MINIKANREN along with relational conversion are capable of inverse computations. However, it is important to note that inverse computations exhibit lower performance compared to directly executing an inversion of the original program due to the interpretation overhead [1, 2].

Relational programs do not exist on their own: they are a part of a host program, which utilizes query results in some way. The host languages are not expected to be able to process logic variables, nondeterminism and other aspects of relational computations. The host program usually only deals with a finite subset of answers, which have been reified into a ground representation, meaning they do not include any logic variables.

When a relation is expected to produce ground answers, and the direction in which it is intended to be run is known, then it becomes possible to convert it into a function which may execute significantly faster than its relational counterpart. Performance improvement comes from reducing interpretation overhead as well as replacing expensive unifications with considerably faster equality checks, assignments and pattern matches of a host language. An informal functional conversion scheme was introduced in the paper [9]. We are building upon this research effort, presenting a semi-automatic functional conversion algorithm and implementation for a minimal core relational programming language MICROKANREN. This paper focuses on converting to the target languages of HASKELL and OCAML, although other languages can also be considered as potential target languages. Our evaluation showed performance improvement of 2.5 times for propositional formulas synthesis and up to 3 orders of magnitude improvement for relations over Peano numbers.

Authors' addresses: Ekaterina Verbitskaia, kajigor@gmail.com, JetBrains Research, Belgrade, Serbia; Igor Engel, igorengel@mail.ru, JetBrains Research, Bremen, Germany; Daniil Berezun, daniil.berezun@jetbrains.com, JetBrains Research, Amsterdam, Netherlands.

2 BACKGROUND

In this section, we give the abstract syntax of MICROKANREN version used in this paper and describe a concept of modes which was developed earlier for other logic languages.

2.1 Normal Form Abstract Syntax of MICROKANREN

To simplify the functional conversion scheme, we consider MICROKANREN relations to be in the superhomogeneous normal form used in the MERCURY programming language [7]. Converting an arbitrary MICROKANREN relation into the normal form is a simple syntactic transformation, which we omit.

In the normal form, a term is either a variable or a constructor application which is flat and linear. Linearity means that arguments of a constructor are distinct variables. To be flat, a term should not contain any nested constructors. Each constructor has a fixed arity n. Below is the abstract syntax of the term language over the set of variables V.

$$\mathcal{T}_V = V \cup \{C_n\,(x_1, \ldots, x_n) \mid x_i \in V; i \neq j \Rightarrow x_i \neq x_j\}$$

Whenever a term which does not adhere to this form is encountered in a unification or as an argument of a call, it is transformed into a conjunction of several unifications, as illustrated by the following examples:

$$C\,(x_1, x_2) \equiv C\,(C\,(y_1, y_2)\,, y_3) \Rightarrow x_1 \equiv C\,(y_1, y_2) \wedge x_2 \equiv y_3$$
$$C\,(C\,(x_1, x_2)\,, x_3) \equiv C\,(C\,(y_1, y_2)\,, y_3) \Rightarrow x_1 \equiv y_1 \wedge x_2 \equiv y_2 \wedge x_3 \equiv y_3$$
$$x \equiv C\,(y, y) \Rightarrow x \equiv C\,(y_1, y_2) \wedge y_1 \equiv y_2$$
$$add^o\,(x, x, z) \Rightarrow add^o\,(x_1, x_2, z) \wedge x_1 \equiv x_2$$

Unification in the normal form is restricted to always unify a variable with a term. We also prohibit using disjunctions inside conjunctions. The normalization procedure declares a new relation whenever this is encountered.

The complete abstract syntax of the MICROKANREN language used in this paper is presented in figure 1.

$$
\begin{array}{llll}
\mathcal{D}_V^N & : & R_n\,(x_1, \ldots, x_n) = \mathbf{Disj}_V, x_i \in V & \text{normalized relation definition} \\
\mathbf{Disj}_V & : & \bigvee (c_1, \ldots, c_n)\,, c_i \in \mathbf{Conj}_V & \text{normal form} \\
\mathbf{Conj}_V & : & \bigwedge (g_1, \ldots, g_n)\,, g_i \in \mathbf{Base}_V & \text{normal conjunction} \\
\mathbf{Base}_V & : & V \equiv \mathcal{T}_V & \text{flat unification} \\
& | & R_n\,(x_1, \ldots, x_n)\,, x_i \in V, i \neq j \Rightarrow x_i \neq x_j & \text{flat call}
\end{array}
$$

Fig. 1. Abstract syntax of MICROKANREN in the normal form

2.2 Modes

A mode generalizes the concept of a direction; this terminology is commonly used in the conventional logic programming community. In its most primitive form, a mode specifies which arguments of a relation will be known at runtime (input) and which are expected to be computed (output). Several logic programming languages have mode systems used for optimizations, with MERCURY standing out among them. MERCURY[1] is a modern functional-logic programming language with a complicated mode system capable not only of describing directions, but also specifying if a relation in the given mode is deterministic, among other things [6].

---

[1]Website of the MERCURY programming language: https://mercurylang.org/

Given an annotation for a relation, mode inference determines modes of each variable of the relation. For some modes, conjunctions in the body of a relation may need reordering to ensure that consumers of computed values come after the producers of said values so that a variable is never used before it is bound to some value. In this project, we employed the least complicated mode system, in which variables may only have an *in* or *out* mode. A mode maps variables of a relation to a pair of the initial and final instantiations. The mode *in* stands for $g \to g$, while *out* stands for $f \to g$. The instantiation f represents an unbound, or *free*, variable, when no information about its possible values is available. When the variable is known to be *ground*, its instantiation is g.

In this paper, we call a pair of instantiations a mode of a variable. figure 2 shows examples of the normalized MICROKANREN relations with modes inferred for the forward and backward directions. We use superscript annotation for variables to represent their modes visually. Notice the different order of conjuncts in the bodies of the add^o relation in different modes.

$$
\begin{aligned}
&\textbf{let } \mathrm{double}^o \ x^{g \to g} \ r^{f \to g} = \\
&\quad \mathrm{addo}^o \ x_1^{g \to g} \ x_2^{g \to g} \ r^{f \to g} \ \wedge \\
&\quad x_1^{g \to g} \equiv x_2^{g \to g}
\end{aligned}
$$

$$
\begin{aligned}
&\textbf{let rec } \mathrm{add}^o \ x^{g \to g} \ y^{g \to g} \ z^{f \to g} = \\
&\quad (x^{g \to g} \equiv 0 \ \wedge \ y^{g \to g} \equiv z^{f \to g}) \ \vee \\
&\quad (x^{g \to g} \equiv S \ x_1^{f \to g} \ \wedge \\
&\quad \ \mathrm{add}^o \ x_1^{g \to g} \ y^{g \to g} \ z_1^{f \to g} \ \wedge \\
&\quad \ z^{f \to g} \equiv S \ z_1^{g \to g})
\end{aligned}
$$

(a) Forward direction

$$
\begin{aligned}
&\textbf{let } \mathrm{double}^o \ x^{f \to g} \ r^{g \to g} = \\
&\quad \mathrm{addo}^o \ x_1^{f \to g} \ x_2^{f \to g} \ r^{g \to g} \ \wedge \\
&\quad x_1^{g \to g} \equiv x_2^{g \to g}
\end{aligned}
$$

$$
\begin{aligned}
&\textbf{let rec } \mathrm{add}^o \ x^{f \to g} \ y^{f \to g} \ z^{g \to g} = \\
&\quad (x^{f \to g} \equiv 0 \ \wedge \ y^{f \to g} \equiv z^{g \to g}) \ \vee \\
&\quad (z^{f \to g} \equiv S \ z_1^{g \to g} \ \wedge \\
&\quad \ \mathrm{add}^o \ x_1^{f \to g} \ y^{f \to g} \ z_1^{g \to g} \ \wedge \\
&\quad \ x^{f \to g} \equiv S \ x_1^{g \to g})
\end{aligned}
$$

(b) Backward direction

Fig. 2. Normalized doubling and addition relations with mode annotations

3 FUNCTIONAL CONVERSION FOR MICROKANREN

In this section, we describe the functional conversion algorithm. The reader is encouraged to first read the paper [9] on the topic, which introduces the conversion scheme on a series of examples.

Functional conversion is done for a relation with a concrete fixed direction. The goal is to create a function which computes the same answers as MICROKANREN would, not necessarily in the same order. Since the search in MICROKANREN is complete, both conjuncts and disjuncts can be reordered freely: interleaving makes sure that no answers would be lost this way. Moreover, the original order of the subgoals is often suboptimal for any direction but the one which the programmer had in mind when they encoded the relation. When the relational conversion is used to create a relation, the order of the subgoals only really suits the forward direction, in which the relation is often not intended to be run (in this case, it is better to run the original function).

The mode inference results in the relational program with all variables annotated by their modes, and all base subgoals ordered in a way that further conversion makes sense. Conversion then produces functions in the intermediate language. It may then be pretty printed into concrete functional programming languages, in our case HASKELL and OCAML.

3.1 Mode Inference

We employ a simple version of mode analysis to order subgoals properly in the given direction. The mode analysis makes sure that a variable is never used before it is associated with some value.

It also ensures that once a variable becomes ground, it never becomes free, thus the value of a variable is never lost. The mode inference pseudocode is presented in listing 1.

```
1  modeInfer  (R_i (x_1,...,x_{k_i}) ≡ body)  = (R_i (x_1,...,x_{k_i}) ≡ (modeInferDisj body))
2
3  modeInferDisj (⋁(c_1,...,c_n)) = ⋁(modeInferConj c_1,..., modeInferConj c_n)
4
5  modeInferConj (⋀(g_1,...,g_n)) =
6    let (picked, theRest) = pickConjunct([g_1,...,g_n]) in
7    let moddedPicked = modeInferBase picked in
8    let moddedConjs = modeInferConj (⋀ theRest) in
9    ⋀(moddedPicked : moddedConjs)
10
11 pickConjunct goals =
12   pickGuard goals <|>
13   pickAssignment goals <|>
14   pickMatch goals <|>
15   pickCallWithGroundArguments goals <|>
16   pickUnificationGenerator goals <|>
17   pickCallGenerator goals
```

Listing 1. Mode inference pseudocode

Mode inference starts by initializing modes for all variables in the body of the given relation according to the given direction. All variables that are among arguments are annotated with their *in* or *out* modes, while all other variables get only their initial instantiations specified as f.

Then the body of the relation is analyzed (see line 1). Since the body is normalized, it can only be a disjunction. Each disjunct is analyzed independently (see line 3) because no data flow happens between them.

Analyzing conjunctions involves analyzing subgoals and ordering them. Let us first consider mode analysis of unifications and calls, and then circle back to the way we order them. Whenever a base goal is analyzed, all variables in it have some initial instantiation, and some of them also have some final instantiation. Mode analysis of a base goal boils down to making all final instantiations ground.

When analyzing a unification, several situations may occur. Firstly, every variable in the unification can be ground, as in $x^{g \to g} \equiv O$ or in $y^{g \to ?} \equiv z^{g \to ?}$ (here ? is used to denote that a final instantiation is not yet known). We call this case *guard*, since it is equivalent to checking that two values are the same.

The second case is when one side of a unification only contains ground variables. Depending on which side is ground, we call this either *assignment* or *match*. The former corresponds to assigning the value to a variable, as in $x^{f \to ?} \equiv S\, x_1^{g \to g}$ or $x^{g \to g} \equiv y^{f \to ?}$. The latter — to pattern matching with the variable as the scrutinee, as in $x^{g \to g} \equiv S\, x_1^{f \to ?}$. Notice that we allow for some variables on the right-hand side to be ground in matches, given that at least one of them is free.

The last case occurs when both the left-hand and right-hand sides contain free variables. This does not translate well into functional code. Any free logic variable corresponds to the possibly infinite number of ground values. To handle this kind of unification, we propose to use *generators* which produce all possible ground values a free variable may have.

We base our ordering strategy for conjuncts on the fact that these four different unification types have different costs. The guards are just equality checks which are inexpensive and can reduce the search space considerably. Assignments and matches are more involved, but they still take much less effort than generators. Moreover, executing non-generator conjuncts first can make some of the variables of the prospective generator ground thus avoiding generation in the end. This is the base reasoning which is behind our ordering strategy.

The function `pickConjunct` selects the base goal which is least likely to blow up the search space. The right-associative function `<|>` used in lines 12 through 16 is responsible for selecting the base goals in the order described. The function first attempts to pick a base goal with its first argument, and only if it fails, the second argument is called. As a result, `pickConjunct` first picks the first guard unification it can find (`pickGuard`). If no guard is present, then it searches for the first assignment (`pickAssignment`), and then for the match (`pickMatch`). If all unifications in the conjunction are generators, then we search for relation calls with some ground arguments (`pickCallWithGroundArguments`). If there are none, then we have no choice but selecting a generating unification (`pickUnificationGenerator`) and then a call with all arguments free (`pickCallGenerator`).

Once one conjunct is picked, it is analyzed (see line 7). The picked conjunct may instantiate new variables, thus this information is propagated onto the rest of the conjuncts. Then the rest of the conjuncts is mode analyzed as a new conjunction (see line 8). If any new modes for any of the relations are encountered, they are also mode analyzed.

It is worth noticing that any relation can generate infinitely many answers. We cannot judge the relation to be such generator solely by its mode: for example, the addition relation in the mode $\text{add}^o\ x^{g \to g}\ y^{f \to g}\ z^{f \to g}$ generates an infinite stream, while $\text{add}^o\ x^{f \to g}\ y^{f \to g}\ z^{g \to g}$ does not.

3.2 Conversion into Intermediate Representation

To represent nondeterminism, our functional conversion uses the basis of MICROKANREN — the stream data structure. A relation is converted into a function with n arguments which returns a stream of m-tuples, where n is the number of the input arguments, and m — the number of the output arguments of the relation. Since stream is a monad, functions can be written elegantly in HASKELL using do-notation (see figure 4). We use an intermediate representation which draws inspiration from HASKELL's do-notation, but can then be pretty-printed into other functional languages. The abstract syntax of our intermediate language is shown in figure 3. The conversion follows quite naturally from the modded relation and the syntax of the intermediate representation.

$$
\begin{array}{rcll}
\mathcal{F}_V & = & \textbf{Sum}\,[\mathcal{F}_V] & \text{concatenation of streams} \\
& | & \textbf{Bind}\,[([V]\,,\mathcal{F}_V)] & \text{monadic bind for streams} \\
& | & \textbf{Return}\,[\mathcal{T}_V] & \text{return of a tuple of terms} \\
& | & \textbf{Guard}\,(V, V) & \text{equality check} \\
& | & \textbf{Match}_V\,(\mathcal{T}_V, \mathcal{F}_V) & \text{match a variable against a pattern} \\
& | & R_n([V]\,, [G]) & \text{function call} \\
& | & \textbf{Gen}_G & \text{generator}
\end{array}
$$

Fig. 3. Abstract syntax of the intermediate language $\mathcal{F}$

A body of a function is formed as an interleaving concatenation of streams (**Sum**), each of which is constructed from one of the disjuncts of the relation. A conjunction is translated into a sequence of bind statements (**Bind**): one for each of the conjuncts and a return statement (**Return**) in the

end. A bind statement binds a tuple of variables (or nothing) with values taken from the stream in the right-hand side.

A base goal is converted into a guard (**Guard**), match (**Match**), or function call, depending on the goal's type. Assignments are translated into binds with a single return statement on the right. Notice, that a match only has one branch. This branch corresponds to a unification. If the scrutinee does not match the term it is unified with, then an empty stream is returned in the catch-all branch. If a term in the right-hand side of a unification has both *out* and *in* variables, then additional guards are placed in the body of the branch to ensure the equality between values bound in the pattern and the actual ground values.

Generators (**Gen**) are used for unifications with free variables on both sides. A generator is a stream of possible values for the free variables, and it is used for each variable from the right-hand side of the unification. The variable from the left-hand side of the unification is then simply assigned the value constructed from the right-hand side. Our current implementation works with an untyped deeply embedded MICROKANREN, in which there is not enough information to produce generators automatically. We decided to delegate the responsibility to provide generators to the user: a generator for each free variable is added as an argument of the relation. When the user is to call the function, they have to provide the suitable generators.

4 EXAMPLES

In this section, we provide some examples which demonstrate mode analysis and conversion results.

4.1 Multiplication Relation

Figure 4 shows the implementation of the multiplication relation mul^o, the mode analysis result for mode $\text{mul}^o\ x^{f \to g}\ y^{g \to g}\ z^{g \to g}$, and the results of functional conversion into HASKELL and OCAML.

Note that the unification comes last in the second disjunct. This is because before the two relation calls are done, both variables in the unification are free. Our version of mode inference puts the relation calls before the unification, but the order of the calls depends on their order in the original relation. There is nothing else our mode inference uses to prefer the order presented in the figure over the opposite: $\text{mul}^o\ x_1^{f \to g}\ y^{g \to g}\ z_1^{f \to g} \wedge \text{add}^o\ y^{g \to g}\ z_1^{g \to g}\ z^{g \to g}$. However, it is possible to derive this optimal order, if determinism analysis is employed: $\text{add}^o\ y^{g \to g}\ z_1^{f \to g}\ z^{g \to g}$ is deterministic while $\text{mul}^o\ x_1^{f \to g}\ y^{g \to g}\ z_1^{f \to g}$ is not. Putting nondeterministic computations first makes the search space larger, and thus should be avoided if another order is possible.

Functional conversions in both languages are similar, modulo the syntax. The HASKELL version employs do-notation, while we use let-syntax in the OCAML code. Both are syntactic sugar for monadic computations over streams. We use the following convention to name the functions: we add a suffix to the relation's name whose length is the same as the number of the relation's arguments. The suffix consists of the letters I and O which denote whether the argument in the corresponding position is *in* or *out*. The function msum uses the interleaving function mplus to concatenate the list of streams constructed from disjuncts. To check conditions, we use the function guard which fails the monadic computation if the condition does not hold. Note that even though patterns for the variable x0 in the function addoIOI are disjunct in two branches, we do not express them as a single pattern match. Doing so would improve readability, but it does not make a difference when it comes to the performance, according to our evaluation.

4.2 The Mode of Addition Relation which Needs a Generator

Consider the example of the addition relation in mode $\text{add}^o\ x^{g \to g}\ y^{f \to g}\ z^{f \to g}$ presented in figure 5. The unification in the first disjunct of this relation involves two free variables. We use a generator

```
let rec mulᵒ x y z = conde [
  (x ≡ 0 ∧ z ≡ 0);
  (fresh (x₁ z₁)
    (x ≡ S x₁ ∧
     addᵒ y z₁ z ∧
     mulᵒ x₁ y z₁)) ]
```

(a) Implementation in MINIKANREN

$$
\textbf{let rec } \mathrm{mul}^o \ x^{f \to g} \ y^{g \to g} \ z^{g \to g} =
$$
$$
(x^{f \to g} \equiv 0 \ \wedge \ z^{g \to g} \equiv 0) \ \vee
$$
$$
(\mathrm{add}^o \ y^{g \to g} \ z_1^{f \to g} \ z^{g \to g} \ \wedge
$$
$$
\mathrm{mul}^o \ x_1^{f \to g} \ y^{g \to g} \ z_1^{g \to g} \ \wedge
$$
$$
x^{f \to g} \equiv S \ x_1^{g \to g})
$$

(b) Mode inference result

```
muloOII x1 x2 = msum
  [ do { let {x0 = O}
       ; guard (x2 == O)
       ; return x0 }
  , do { x4 ← addoIOI x1 x2
       ; x3 ← muloOII x1 x4
       ; let {x0 = S x3}
       ; return x0 } ]
addoIOI x0 x2 = msum
  [ do { guard (x0 == O)
       ; let {x1 = x2}
       ; return x1 }
  , do { x3 ← case x0 of
             { S y3 → return y3
             ; _ → mzero }
       ; x4 ← case x2 of
             { S y4 → return y4
             ; _ → mzero }
       ; x1 ← addoIOI x3 x4
       ; return x1 } ]
```

(c) Functional conversion into HASKELL

```
let rec muloOII x1 x2 = msum
  [ ( let* x0 = return O in
      let* _ = guard (x2 = O) in
      return x0 )
  ; ( let* x4 = addoIOI x1 x2 in
      let* x3 = muloOII x1 x4 in
      let* x0 = return (S x3) in
      return x0 ) ]
and addoIOI x0 x2 = msum
  [ ( let* _ = guard (x0 = O) in
      let* x1 = return x2 in
      return x1 )
  ; ( let* x3 = match x0 with
         | S y3 → return y3
         | _ → mzero in
      let* x4 = match x2 with
         | S y4 →   return y4
         | _ → mzero in
      let* x1 = addoIOI x3 x4 in
      return x1 ) ]
```

(d) Functional conversion into OCAML

Fig. 4. Multiplication relation

gen_addoIIO_x2 to generate a stream of ground values for the variable z which is passed into the function addIIO as an argument. It is up to the user to provide a suitable generator. One of the possible generators which produces all Peano numbers in order and an example of its usage are presented in figure 5b.

The generators which produce an infinite stream should be inverse eta-delayed in OCAML and other non-lazy languages. Otherwise, the function would not terminate trying to eagerly produce all possible ground values before using any of them.

It is possible to automatically produce generators from the data type of a variable, but it is currently not implemented, as we work with an untyped version of MICROKANREN.

5 EVALUATION

To evaluate our functional conversion scheme, we implemented the proposed algorithm in HASKELL. We compared execution time of several OCANREN relations in different directions against their

```
let rec add^o x^{g→g} y^{f→g} z^{f→g} =
  (x^{g→g} ≡ 0 ∧ y^{f→g} ≡ z^{f→g}) ∨
  (x^{g→g} ≡ S x_1^{f→g} ∧
   add^o x_1^{g→g} y^{f→g} z_1^{f→g} ∧
   z^{f→g} ≡ S z_1^{g→g})
```

(a) Mode inference result

```
genNat = msum
  [ return 0
  , do { x ← genNat
       ; return (S x) } ]
runAddoIIO x = addoIIO x genNat
```

(b) Generator of Peano numbers

```
addoIOO x0 gen_addoIOO_x2 = msum
  [ do { guard (x0 == 0)
       ; (x1, x2) ← do { x2 ← gen_addoIOO_x2 ; return (x2, x2) }
       ; return (x1, x2) }
  , do { x3 ← case x0 of { S y3 → return y3 ; _ → mzero }
       ; (x1, x4) ← addoIOO x3 gen_addoIOO_x2
       ; let {x2 = S x4} ; return (x1, x2) } ]
```

(c) Functional conversion

Fig. 5. Addition relation when only the first argument is *in*

functional counterparts in the OCaml language. Here we showcase two relational programs and their conversions. The implementation of the functional conversion[2] as well as the execution code[3] can be found on Github.

5.1 Evaluator of Propositional Formulas

In this example, we converted a relational evaluator of propositional formulas: see figure 6. It evaluates a propositional formula fm in the environment st to get the result u. A formula is either a boolean literal, a numbered variable, a negation of another formula, a conjunction or a disjunction of two formulas. Converting it in the direction when everything but the formula is *in* (see figure 6a), allows one to synthesize formulas which can be evaluated to the given value. The conversion of this relation does not involve any generators and is presented in figure 6b.

We ran an experiment to compare the execution time of the relational interpreter vs. its functional conversion. In the experiment, we generated from 1000 to 10000 formulas which evaluate to true and contain up to 3 variables with known values. The results are presented in figure 7. The functional conversion improved execution time of the query about 2.5 times from $724ms$ to $291ms$ for retrieving 10000 formulas.

5.2 Multiplication

In this example, we converted the multiplication relation in several directions and compared them to the relational counterparts: see figure 8. Functional conversion significantly reduced execution time in most directions.

In the forward direction, we run the query mul^o n 10 q with n in the range from 100 to 1000, and the functional conversion was 2 orders of magnitude faster: $927ms$ vs $9.4ms$ for the largest n, see figure 8a. In the direction which serves as division we run the query mul^o (n /10) q n with n ranging from 100 to 1000. Here, performance improved 3 orders of magnitude: from $24s$ to $0.17s$ for

---

[2]The repository of the functional conversion project https://github.com/kajigor/uKanren_transformations
[3]Evaluation code https://github.com/kajigor/miniKanren-func

```
let rec evalᵒ stᵍ→ᵍ fmᶠ→ᵍ uᵍ→ᵍ =
  ( fmᶠ→ᵍ ≡ Lit uᵍ→ᵍ) ∨
  ( elemᵒ zᶠ→ᵍ stᵍ→ᵍ uᵍ→ᵍ ∧
    fmᶠ→ᵍ ≡ Var zᵍ→ᵍ ) ∨
  ( notᵒ vᶠ→ᵍ uᵍ→ᵍ ∧
    evalᵒ stᵍ→ᵍ xᶠ→ᵍ vᵍ→ᵍ ∧
    fmᶠ→ᵍ ≡ Neg xᵍ→ᵍ ) ∨
  ( orᵒ vᶠ→ᵍ wᶠ→ᵍ uᵍ→ᵍ ∧
    evalᵒ stᵍ→ᵍ xᶠ→ᵍ vᵍ→ᵍ ∧
    evalᵒ stᵍ→ᵍ yᶠ→ᵍ wᵍ→ᵍ ∧
    fmᶠ→ᵍ ≡ Disj xᵍ→ᵍ yᵍ→ᵍ) ∨
  ( andᵒ vᶠ→ᵍ wᶠ→ᵍ uᵍ→ᵍ ∧
    evalᵒ stᵍ→ᵍ xᶠ→ᵍ vᵍ→ᵍ ∧
    evalᵒ stᵍ→ᵍ yᶠ→ᵍ wᵍ→ᵍ ∧
    fmᶠ→ᵍ ≡ Conj xᵍ→ᵍ yᵍ→ᵍ) ∨
```

(a) Mode inference result

```
evaloIOI x0 x2 = msum
  [ do { let {x1 = Lit x2}
       ; return x1 }
  , do { x7 ← elemoOII x0 x2
       ; let {x1 = Var x7}
       ; return x1 }
  , do { x5 ← notoOI x2
       ; x3 ← evaloIOI x0 x5
       ; let {x1 = Neg x3}
       ; return x1 }
  , do { (x5, x6) ← oroOOI x2
       ; x3 ← evaloIOI x0 x5
       ; x4 ← evaloIOI x0 x6
       ; let {x1 = Disj x3 x4}
       ; return x1 }
  , do { (x5, x6) ← andoOOI x2
       ; x3 ← evaloIOI x0 x5
       ; x4 ← evaloIOI x0 x6
       ; let {x1 = Conj x3 x4}
       ; return x1 } ]
```

(b) Functional conversion

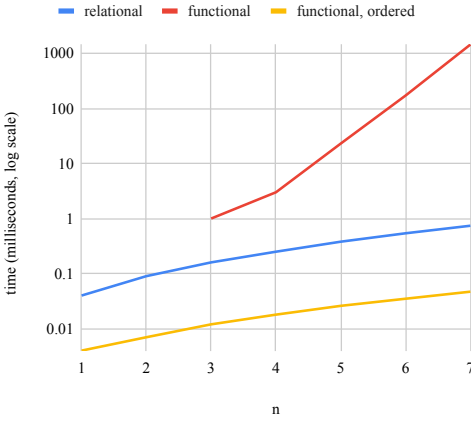Fig. 6. Evaluator of propositional formulas

Fig. 7. Execution time of the evaluators of propositional formulas, eval [**true** ; **false** ; **true**] q **true**

the largest n, see figure 8b. Even more impressive was the backward direction mulᵒ xᶠ→ᵍ yᶠ→ᵍ zᵍ→ᵍ. Querying for all 16 pairs of divisors of 1000 (mulᵒ q r 1000) took OCᴀɴʀᴇɴ about 32.9s, while the functional conversion succeeded in 1.1s.

What was surprising was the mode mulᵒ xᵍ→ᵍ yᶠ→ᵍ zᶠ→ᵍ. In this case, the functional conversion was not only worse than its relational counterpart, its performance degraded exponentially with
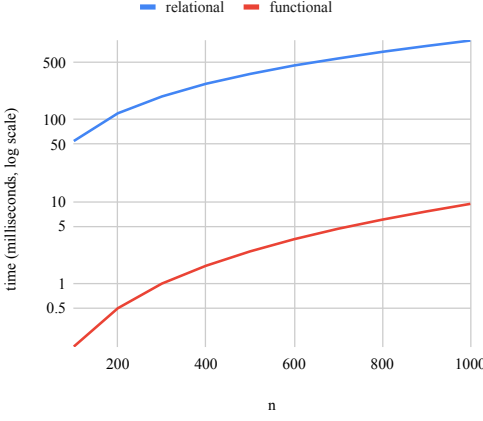
(a) Multiplication: `mulo n 10 q`
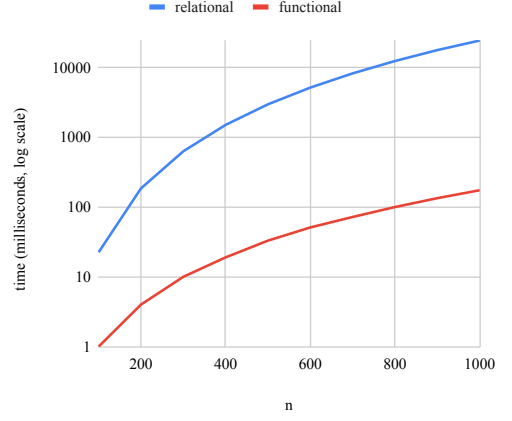
(b) Division: `mulo (n/10) q n`

$$
\begin{aligned}
&\textbf{let rec } \mathrm{mul}^o\ x^{g\to g}\ y^{f\to g}\ z^{f\to g}\ = \\
&(x^{g\to g}\ \equiv\ 0\ \wedge\ z^{f\to g}\ \equiv\ 0)\ \vee \\
&(x^{g\to g}\ \equiv\ S\ x_1^{f\to g}\ \wedge \\
&\quad \mathrm{add}^o\ y^{f\to g}\ z_1^{f\to g}\ z^{f\to g}\ \wedge \\
&\quad \mathrm{mul}^o\ x_1^{g\to g}\ y^{g\to g}\ z_1^{g\to g}\ )
\end{aligned}
$$

(d) Inefficient mode

$$
\begin{aligned}
&\textbf{let rec } \mathrm{mul}^o\ x^{g\to g}\ y^{f\to g}\ z^{f\to g}\ = \\
&(x^{g\to g}\ \equiv\ 0\ \wedge\ z^{f\to g}\ \equiv\ 0)\ \vee \\
&(x^{g\to g}\ \equiv\ S\ x_1^{f\to g})\ \wedge \\
&\quad \mathrm{mul}^o\ x_1^{g\to g}\ y^{f\to g}\ z_1^{f\to g}\ \wedge \\
&\quad \mathrm{add}^o\ y^{g\to g}\ z_1^{g\to g}\ z^{f\to g}\ )
\end{aligned}
$$

(e) Efficient mode

(c) Generation: `take n (mulo 10 q r)`

Fig. 8. Execution times of the multiplication relation

the number of answers asked. It took almost $1450ms$ to find the first 7 pairs of numbers q and r such that $10\ *\ q\ =\ r$, while OCANREN was able to execute the query in $0.74ms$ (see figure 8c). The source of this terrible behavior was the suboptimal order of the calls in the second disjunct of the mul^o relation in the corresponding mode (see figure 8d). In this case, the call $\mathrm{add}^o\ y^{f\to g}\ z_1^{f\to g}\ z^{f\to g}$ is put first, which generates all possible triples in the addition relation before filtering them by the call $\mathrm{mul}^o\ x_1^{g\to g}\ y^{g\to g}\ z_1^{g\to g}$. The other order of calls is much better (see figure 8e): it is an order of magnitude faster than its relational source. To achieve the better of these two orders automatically, we delay picking any call with all arguments free. It is not clear if these heuristics are universal.

5.3 Deterministic Directions

Running in some directions, relations produce deterministic results. For example, any forward direction of a relation created by the relational conversion produces a single result, since it mimics the original function. The guard directions are semi-deterministic: they may fail, but if they succeed, they produce a single unit value. If the addition relation is run with one of the first two arguments *out*, it acts as subtraction and is also deterministic.

For such directions, there is no need to model nondeterminism with the Stream monad. Semi-determinism can be expressed with a Maybe monad, while deterministic directions can be converted into simple functions. Our implementation of functional conversion only restricts the computations to be monadic, it does not specify which monad to use. By picking other monads, we can achieve performance improvement. For example, using Maybe for division reduces its execution time 30 times in addition to the 2 orders of magnitude improvement from the functional conversion itself: see figure 9


Fig. 9. Execution time of division: `take n (mul q 10 1000)`

6 DISCUSSION

Our experiments indicated that the functional conversion is capable of improving performance of relational computations significantly in the known directions. The improvement stems from eliminating costly unifications in favor of the cheaper equality checks and pattern matches. Besides this, we employed some heuristics which push lower-cost computations to happen sooner while delaying higher-cost ones. It is also possible to take into account determinism of some directions and improve performance of them even more by picking an appropriate monad.

It is not currently clear if the heuristics we used are universal enough. However, it is always safe to run any deterministic computations because they never increase the search space. We believe that it is necessary to integrate determinism check in the mode analysis so that the more efficient modes such as the one presented in figure 8e could be achieved more justifiably.

We also think that further integration with specialization techniques such as partial deduction may benefit the conversion even more [8]. For example, the third argument of the propositional evaluator can be either **true** or **false** . Specializing the evaluator for these two values may help to shave off even more time.

7   CONCLUSION AND FUTURE WORK

In this paper, we described a semi-automatic functional conversion of a MICROKANREN relation with a fixed direction into a functional language. We implemented the proposed conversion and applied it to a set of relations, resulting in significant performance enhancement, as demonstrated in our evaluation. As part of the future work, we plan to augment the mode analysis with a determinism check. We also plan to integrate the functional conversion with specialization techniques such as partial deduction.

REFERENCES

[1] Sergei Abramov and Robert Glück. 2000. Combining Semantics with Non-standard Interpreter Hierarchies. In *FST TCS 2000: Foundations of Software Technology and Theoretical Computer Science*, Sanjiv Kapoor and Sanjiva Prasad (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 201–213.

[2] Sergei Abramov and Robert Glück. 2002. *Principles of Inverse Computation and the Universal Resolving Algorithm.* Springer Berlin Heidelberg, Berlin, Heidelberg, 269–295. https://doi.org/10.1007/3-540-36377-7_13

[3] William E. Byrd, Eric Holk, and Daniel P. Friedman. 2012. MiniKanren, Live and Untagged: Quine Generation via Relational Interpreters (Programming Pearl). In *Proceedings of the Annual Workshop on Scheme and Functional Programming* (Copenhagen, Denmark) *(Scheme '12)*. Association for Computing Machinery, New York, NY, USA, 8–-29. https://doi.org/10.1145/2661103.2661105

[4] Petr Lozov, Ekaterina Verbitskaia, and Dmitry Boulytchev. 2019. Relational interpreters for search problems. In *Relational Programming Workshop*. 43.

[5] Petr Lozov, Andrei Vyatkin, and Dmitry Boulytchev. 2018. Typed relational conversion. In *Trends in Functional Programming: 18th International Symposium, TFP 2017, Canterbury, UK, June 19-21, 2017, Revised Selected Papers 18*. Springer, 39–58.

[6] David Overton, Zoltan Somogyi, and Peter J Stuckey. 2002. Constraint-based mode analysis of Mercury. In *Proceedings of the 4th ACM SIGPLAN international conference on Principles and practice of declarative programming*. 109–120.

[7] Zoltan Somogyi, Fergus Henderson, and Thomas Conway. 1996. The execution algorithm of Mercury, an efficient purely declarative logic programming language. *The Journal of Logic Programming* 29, 1-3 (1996), 17–64.

[8] Ekaterina Verbitskaia, Daniil Berezun, and Dmitry Boulytchev. 2021. An Empirical Study of Partial Deduction for miniKanren. *arXiv preprint arXiv:2109.02814* (2021).

[9] Ekaterina Verbitskaia, Daniil Berezun, and Dmitry Boulytchev. 2022. On a Direction-Driven Functional Conversion. In *Relational Programming Workshop*.