



Universiteit Utrecht

[Faculty of Science
Information and Computing Sciences]

Generic Programming for Domain Reasoners

José Pedro Magalhães¹

joint work with

Johan Jeuring^{1,2} Bastiaan Heeren²

¹ Department of Information and Computing Sciences, Utrecht University

² School of Computer Science, Open University of the Netherlands

Trends in Functional Programming
June 2–4, 2009

Overview

Domain Reasoners

Generic Programming

Evaluation

Future work and conclusions



1. Domain Reasoners



Exercise assistants I

An exercise assistant helps a student with some procedural skill:

- ▶ Turning a logical proposition into disjunctive normal form
- ▶ Solving linear equations
- ▶ Performing Gaussian elimination

There exist hundreds of exercise assistants, for many mathematical domains, logic, physics, etc.



Exercise assistants II

The screenshot shows a window titled "Exercise Assistant" with a standard macOS-style title bar. The interface is divided into several sections:

- Proposition to DNF:** A dropdown menu showing "Proposition to DNF".
- New assignment:** A button to start a new assignment.
- Assignment:** A text area containing the logical expression $((\neg r \parallel q) \wedge \neg r) \parallel (\neg(\neg r \parallel q) \wedge \neg \neg r)$.
- Enter modified term:** A text area containing the same logical expression $((\neg r \parallel q) \wedge \neg r) \parallel (\neg(\neg r \parallel q) \wedge \neg \neg r)$.
- Buttons:** A row of buttons labeled "Ready", "Hint", "Step", "Next", and "Submit". Below them is an "Undo" button and a dropdown menu.
- Derivation:** A text area on the right showing a derivation process:

```
(r -> q) <-> ~r
=> [DefImpl]
(~r || q) <-> ~r
=> [DefEquiv]
((~r || q) ^ ~r) || (~(~r || q) ^ ~~r)
```
- Progress:** A progress bar showing 2/6 steps completed.
- Feedback:** A text area with the instruction "Use rule DeMorganOr".
- Footer:** The logo for "OpenUniversiteitNederland" is visible in the bottom right corner of the window.

Also available online (demo).



Domain reasoners I

The domain reasoner is the core component of an exercise assistant:

- ▶ Track the steps the student takes
- ▶ Give hints
- ▶ Diagnose errors
- ▶ Record progress
- ▶ Show worked-out solutions



Domain reasoners II

The functionality of the domain reasoner fundamentally depends on the domain:

- ▶ The rules that hold for the domain
- ▶ The strategies for solving exercises within the domain



Example: the logic domain

```
data Logic = Logic : $\rightarrow$ : Logic -- implication
          | Logic : $\leftrightarrow$ : Logic -- equivalence
          | Logic : $\wedge$ : Logic -- conjunction (and)
          | Logic : $\vee$ : Logic -- disjunction (or)
          | Not Logic -- negation (not)
          | Var String -- variables
          | T -- true
          | F -- false
```



Example: the programming domain

```
data LExpr = Lambda String LExpr -- lambda abstraction
           | LVar String           -- variables
           | Apply LExpr LExpr    -- function application
           | Fix LExpr            -- fixpoint
           | Int Int              -- integers
           | Let Decl LExpr       -- let ... in
```

```
data Decl = String := LExpr -- variable declaration
```



2. Generic Programming



Generic behavior

Many parts of the domain reasoners can be implemented in a datatype-generic way: its behavior does not depend on the domain in question. Examples of such parts are:

- ▶ Folding
- ▶ Top-level equality
- ▶ Rewriting (which requires metavariable extension)
- ▶ Traversals
- ▶ Exercise generation



Generic programming

Generic programming techniques claim to:

- ▶ Reduce code duplication
- ▶ Make it easier to change the structure of data
- ▶ Provide implementations of many useful functions on most datatypes

We have applied generic programming techniques to our domain reasoners and compared the initial, non-generic version with two versions using different libraries for generic programming in Haskell.



Type-specific domain reasoners

We started our project with a single domain reasoner for solving linear equations. All required functionality was implemented specifically for this domain.

Rewriting, for instance, was implemented in a 449 line module:

module *EquationsRewriteAnalysis* **where**

equationsRewrite = ...

exprRewrite = ...

equationsSubst = ...

exprSubst = ...

A package for generic rewriting consists of only 370 lines of source code.



Uniplate

As the necessity to implement other domains arose, we moved to a simple generic programming library to simplify traversals and rewriting. For this we used Uniplate (Mitchell and Runciman, 2007).



Uniplate

As the necessity to implement other domains arose, we moved to a simple generic programming library to simplify traversals and rewriting. For this we used Uniplate (Mitchell and Runciman, 2007).

class *Uniplate* a **where**

uniplate :: a → ([a], [a] → a)

instance *Uniplate* Logic **where**

uniplate (l₁ :->: l₂) = ([l₁, l₂], λ [l₃, l₄] → l₃ :->: l₄)

uniplate (l₁ :<->: l₂) = ...

...



Multirec I

However, Uniplate was too limited for all the functionality we needed:

- ▶ No generic producers
- ▶ Does not integrate nicely with type-indexed datatypes

As the number of domains grew, we moved to Multirec (Rodriguez *et al.*, 2009):

- ▶ Fixed point view on data, which we need for functionality such as type-safe rewriting
- ▶ Supports mutually recursive datatypes, which we need for some of our domains (like programming)




```
type instance PF Programming =  
    (K String :×: I LExpr) :▷: LExpr  
  :+: (K String)           :▷: LExpr  
  :+: (I LExpr :×: I LExpr) :▷: LExpr  
  :+: (I LExpr)           :▷: LExpr  
  :+: (K Int)             :▷: LExpr  
  :+: (I Decl :×: I LExpr) :▷: LExpr  
  :+: (K String :×: I LExpr) :▷: Decl
```

class *Fam* φ **where**

from :: φ ix \rightarrow ix \rightarrow PF φ I* ix

to :: φ ix \rightarrow PF φ I* ix \rightarrow ix



3. Evaluation



ISO 9126 quality model

To compare our initial, non-generic domain reasoners with the versions using Uniplate and Multirec, we have used the quality characteristics described in the ISO 9126 international standard for the evaluation of software quality:

- ▶ Functionality
- ▶ Reliability
- ▶ Usability
- ▶ Efficiency
- ▶ Maintainability
- ▶ Portability



Functionality

Using Uniplate, we can separate the functionality for generic rewriting and traversals. This leads to increased modularity and separation of concerns.

Using Multirec, we can further separate the functionality for folds, top-level equality and exercise generation.

However, Multirec does not (yet) support parametric datatypes. This means we cannot use it with our domain reasoner for systems of equations.



Reliability I

The generic domain reasoners are more reliable since they make use of more library code.

Another example of increased reliability is the way metavariables were handled:

```
class MetaVar a where  
  metaVar  :: Int → a  
  isMetaVar :: a  → Maybe Int
```

The idea was to represent metavariables as variables in the domain with a particular name.



Reliability II

For the logic domain:

instance *MetaVar* Logic **where**

metaVar *n* = *Var* ("_" ++ *show* *n*)

isMetaVar (*Var* ('_' : *s*)) | \neg (*null* *s*) \wedge *all isDigit* *s*
= *return* (*read* *s*)

isMetaVar _ = *Nothing*



Reliability II

For the logic domain:

instance *MetaVar* Logic **where**

metaVar *n* = *Var* ("_" ++ *show* *n*)

isMetaVar (*Var* ('_' : *s*)) | \neg (*null* *s*) \wedge *all isDigit* *s*
= *return* (*read* *s*)

isMetaVar _ = *Nothing*

The version with Multirec uses type-safe metavariable extension:

type Scheme *s* = K MVar :+: PF *s*

type MVar = Int

This ensures that no metavariables can escape and it works on domains that have no notion of variables.



From the perspective of a programmer who has to add a new domain, customizing our domain reasoners is now much easier. A programmer only has to specify:

- ▶ Abstract syntax
- ▶ Semantics-dependent operations:
 - ▶ Parsing and pretty-printing
 - ▶ Rewriting rules
 - ▶ Associative and commutative operators
 - ▶ The strategies for the exercises

Generic functionality comes for free.



Efficiency

The generic domain reasoners are probably less efficient than their handwritten counterparts.

We have not yet performed detailed benchmarks, but recent research (Rodriguez 2009) shows that while Uniplate performs reasonably well, Multirec can be more than 40 times slower than a handwritten solution.

Further work is required to optimize the performance of generic programming libraries.



Maintainability

Maintaining the code of the generic domain reasoners is easier, since there is less code to update. If a domain changes, the generic functionality does not have to be adapted.

The number of lines of source code decreases dramatically when using generic programming (defining the abstract syntax and essential functionality):

Handwritten	Uniplate	Multirec
496	282	214

However, advanced generic programming techniques are now used, so the maintainer needs to be a skilled functional programmer.



Portability

The original domain reasoners were developed in Haskell 98, which is a standard portable across many compilers.

For the domain reasoners with Uniplate, we used an embedded, simplified implementation of the library which is also Haskell 98.

The use of Multirec introduces dependencies on a number of advanced extensions of Haskell only available in GHC, so portability is worse.



4. Future work and conclusions



Future work I

Currently we are working on using a new version of Multirec which supports parametric mutually recursive datatypes.

This should allow us to use Multirec in the domain reasoner for systems of equations.

We still expect to report on the results of the new version, since this is a considerable limitation.



Future work II

Other improvements using the new generic framework are still possible:

- ▶ Associativity- and commutativity-aware rewriting
 - ▶ Reduce the number of rules
 - ▶ Simplify recognition of application of rules by the user
- ▶ Use of generic zipper for traversals for increased reliability
- ▶ Generic selections



Conclusions

The use of generic programming libraries in our domain reasoners has given us increased functionality, reliability, usability, and maintainability.

Using generic programming probably leads to worse performance, but this is, in our case, not noticeable for users. Using Multirec we depend on more language extensions and can only compile with GHC.

Overall, the advantages outweigh the disadvantages and the effort invested into translating to a generic approach pays off when adding multiple new domains.

