

Generic Views on Data Types

Stefan Holdermans¹, Johan Jeuring¹, Andres Löh², and Alexey Rodriguez¹

¹ Department of Information and Computing Sciences, Utrecht University
P.O.Box 80.089, 3508 TB Utrecht, The Netherlands
{stefan,johanj,alexey}@cs.uu.nl

² Institut für Informatik III, Universität Bonn
Römerstraße 164, 53117 Bonn, Germany
loeh@informatik.uni-bonn.de

Abstract. A generic function is defined by induction on the structure of types. The structure of a data type can be defined in several ways. For example, in PolyP a pattern functor gives the structure of a data type viewed as a fixed point, and in Generic Haskell a structural representation type gives an isomorphic type view of a data type in terms of sums of products. Depending on this generic view on the structure of data types, some generic functions are easier, more difficult, or even impossible to define. Furthermore, the efficiency of some generic functions can be improved by choosing a different view. This paper introduces generic views on data types and shows why they are useful. Furthermore, it shows how generic views have been added to Generic Haskell, an extension of the functional programming language Haskell that supports the construction of generic functions. The separation between inductive definitions on type structure and generic views allows us to combine many approaches to generic programming in a single framework.

1 Introduction

A generic function is defined by induction on the structure of types. Several approaches to generic programming [16, 10, 20, 19, 15] have been developed in the last decade. These approaches have their commonalities and differences:

- All the approaches provide either a facility for defining a function by induction on the structure of types, or a set of basic, compiler generated, generic functions which are used as combinators in the construction of generic functions.
- All the approaches differ on how they *view* data types. There are various ways in which the inductive structure of data types can be defined, and each approach to generic programming chooses a different one.

This paper introduces *generic views* on data types. Using generic views it is possible to define generic functions for different views on data types. Generic views provide a framework in which the different approaches to generic programming can be used and compared.

The inductive structure of types. Different approaches to generic programming view the structure of types differently:

- In PolyP [16] a data type is viewed as the fixed point of a pattern functor that has kind $* \rightarrow * \rightarrow *$. Viewing a data type as a fixed point of a pattern functor allows us to define recursive combinators such as the catamorphism and anamorphism [25], and functions that return the direct recursive children of a constructor [18]. A downside of this view on data types is that PolyP can only handle regular data types of kind $* \rightarrow *$.
- In Generic Haskell [10, 23, 22], a data type is described in terms of a top-level sums of products structural representation type. Generic functions in Generic Haskell are defined on possibly nested data types of any kind. However, because the recursive structure of data types is invisible in Generic Haskell, it is hard to define the catamorphism and children functions in a natural way, for example.

- In the ‘Scrap your boilerplate’ [20, 21] approach the generic fold is the central steering concept. The generic fold views a value of a data type as either a constructor, or as an application of a (partially applied) constructor to a value. Using the generic fold it is easy to define traversal combinators on data types, which can easily be specialized to update small parts of a value of a large data structure. A disadvantage of the boilerplate approach is that some generic functions, such as the equality and zipping functions, are harder to define. Furthermore, the approach does not naturally generalize to type-indexed data types [11, 22]. We can translate the boilerplate approach to the level of data types by defining a particular generic view.

Other approaches to representing data types can be found in the Constructor Calculus [19], and in the work of Backhouse, De Moor, and Hoogendijk, where a data type is considered to be a relator with a membership relation [3, 15].

Generic views on data types. An approach to generic programming essentially consists of two components: a facility to define recursive functions on structure types, and a view on the inductive structure of data types, which maps data types onto structure types. We call such a view on the structure of types a *generic view* (or just *view*) on data types. Wadler [30] also defines views on data types. The difference between Wadler’s views and generic views is that the former constitute a method for viewing a single data type in different ways, whereas the latter describes how the structure of a large class of data types is viewed.

Each of the above generic views on data types has its advantages and disadvantages. Some views allow the definition of generic functions that are impossible or hard to define in other approaches, other views allow the definition of more efficient generic functions. This paper

- identifies the concept of generic views as an important building block of an implementation for generic programming;
- shows that different choices of generic views have significant influence on the class of generic functions that can be expressed;
- clearly defines what constitutes a generic view, and discusses how generic views have been added to Generic Haskell;
- allows us to compare different approaches to generic programming by providing a common framework.

Views add expressiveness to a generic programming language. Generic functions still work for arbitrary data types that can be expressed in the view, but the choice between different views allows us to define more generic functions.

Organization. This paper is organized as follows. Section 2 briefly introduces generic programming in Generic Haskell. Section 3 shows by means of examples why generic views on data types are useful, and how they increase the expressiveness of a generic programming language. Section 4 defines formally what constitutes a generic view. For some of the examples of Section 3, we give the formal definition. In Section 5, we discuss how views have been added to the Generic Haskell compiler. Related work and conclusions are given in Section 6.

2 Introduction to generic programming in Generic Haskell

This section introduces generic programming in Generic Haskell. The introduction will be brief, for more information see [12, 11, 22]. Generic Haskell has slightly changed in the last couple of years, and we will use the version described in Löh’s thesis ‘Exploring Generic Haskell’ [22] (EGH) in this paper, which to a large extent has been implemented in the Coral release [23].

2.1 Type-indexed functions

A type-indexed function takes an explicit type argument, and can have behaviour that depends on the type argument. For example, suppose the unit type `Unit`, sum type `+`, and product type `×` are defined as follows,

```

data Unit = Unit
data a + b = Inl a | Inr b
data a × b = a × b.

```

We use infix types $+$ and \times and an infix type constructor \times here to ease the presentation.

The type-indexed function *collect* is used to collect values from a data structure. We define function *collect* on the unit type, sums and products, integers, characters, and lists as follows:

```

collect⟨Unit⟩ Unit    = []
collect⟨α + β⟩ (Inl a) = collect⟨α⟩ a
collect⟨α + β⟩ (Inr b) = collect⟨β⟩ b
collect⟨α × β⟩ (a × b) = collect⟨α⟩ a ++ collect⟨β⟩ b
collect⟨Int⟩ n         = []
collect⟨Char⟩ c        = []
collect⟨[α]⟩ []         = []
collect⟨[α]⟩ (a : as) = collect⟨α⟩ a ++ collect⟨[α]⟩ as.

```

The type signature of *collect* is as follows:

$$\text{collect}\langle a :: * \mid c :: * \rangle :: (\text{collect}\langle a \mid c \rangle) \Rightarrow a \rightarrow [c].$$

The type of *collect* is parameterized over two type variables. The first type variable, a , appearing to the left of the vertical bar, is a generic type variable, and represents the type of the type argument of *collect*. Type variable c , appearing to the right of a vertical bar, is called a non-generic (or parametric) type variable. Such non-generic variables appear in type-indexed functions that are parametrically polymorphic with respect to some type variables; Section 9.3 of EGH discusses the difference between non-generic variables and ordinary universally quantified type variables. The *collect* function is parametrically polymorphic in the element type of its list result. It always returns the empty list, but we will show below how to adapt it so that it collects values from a data structure. Since it always returns the empty list there is no need, but also no desire, to fix the type of the list elements. The type context $(\text{collect}\langle a \mid c \rangle) \Rightarrow$ appears in the type because *collect* is called recursively on sums and products, which means that, for example, if we want an instance of *collect* on the type $\alpha + \beta$, we need instances of *collect* on the types α and β . Thus *collect* depends on itself. The theory of dependencies and type signatures of generic functions is an integral part of dependency-style Generic Haskell, see EGH.

The type signature of *collect* can be instantiated for specific cases by the Generic Haskell compiler, yielding, for example, the types

```

collect⟨Unit⟩ :: ∀ c . Unit → [c]
collect⟨[α]⟩  :: ∀ c a . (collect⟨α⟩ :: a → [c]) ⇒ [a] → [c]

```

for the cases of the unit type and lists, respectively. The latter type can be read as “given a function $\text{collect}\langle \alpha \rangle$ of type $a \rightarrow [c]$, the expression $\text{collect}\langle [\alpha] \rangle$ is of type $[a] \rightarrow [c]$ ”.

Depending on the situation, the function $\text{collect}\langle \alpha \rangle$ can be automatically inferred by the compiler, or it can be user specified using *local redefinitions*. For example, if we only want to collect the positive numbers from a list using the function *collect*, we can write:

```

let collect⟨α⟩ x | x > 0    = [x]
                | otherwise = []
in collect⟨[α]⟩,

```

which has type $\text{Num } a \Rightarrow [a] \rightarrow [a]$. Generally, we use a *local redefinition* of a generic function to locally modify the behaviour of a generic function. Some generic functions such as *collect* only reveal their full power in the context of local redefinitions.

2.2 Structure types

Until now, it looks as if a type-indexed function such as *collect* works only on the types that appear as its type indices. To see why *collect* is in fact *generic* and works on arbitrary data types, we

give a mapping from data types to view types (or structural representation types) such as units, sums and products. It suffices to define a function on view types (and primitive or abstract types such as `Int` and `Char`) in order to obtain a function that can be applied to values of arbitrary data types. If there is no specific case for a type in the definition of a generic function, generic behaviour is derived automatically by the compiler by exploiting the structural representation.

For example, the definition of the function *collect* on lists is superfluous: the instance generically derived for lists coincides with the specific definition of *collect* given above. To obtain this instance, the compiler needs to know the structural representation of lists, and how to convert between lists and their structural representation. We will describe these components in the remainder of this section.

The structural representation of types is expressed in terms of units, sums, products, and base types such as integers, characters, etc. For example, for the list and tree data types defined by

```
data List a  = Nil | Cons a (List a)
data Tree a b = Tip a | Node (Tree a b) b (Tree a b)
```

we obtain the following structural representations:

```
type List◦ a  = Unit + a × List a
type Tree◦ a b = a + Tree a b × b × Tree a b,
```

where we assume that \times binds stronger than $+$, and both type constructors associate to the right. Note that the representation of a recursive type is not recursive, and refers to the recursive type itself. The structural representation of a type in Generic Haskell only represents the structure of the top level of the type.

If two types are isomorphic, the corresponding isomorphisms, called embedding-projection pairs, can be stored as a pair of functions converting back and forth:

```
data EP a b = EP { from :: (a → b), to :: (b → a) }.
```

A type T and its structural representation type T° are isomorphic, witnessed by a value $conv_T :: EP\ T\ T^\circ$. For example, for the list data type we have that $conv_{List} = EP\ from_{List}\ to_{List}$, where $from_{List}$ and to_{List} are defined by

```
fromList          :: ∀ a . List a → List◦ a
fromList Nil      = Inl Unit
fromList (Cons a as) = Inr (a × as)
toList           :: ∀ a . List◦ a → List a
toList (Inl Unit) = Nil
toList (Inr (a × as)) = Cons a as.
```

The definitions of the embedding-projection pairs are automatically generated by the Generic Haskell compiler for all data types that appear in a program.

With the structural representation types and the embedding-projection pairs in place, a call to a generic function on a data type T can always be reduced to a call on type T° . Hence, if the generic function is defined for the data types occurring in the structural representations, such as `Unit`, $+$, and \times , we do not need cases for specific datatypes such as `List` or `Tree` anymore.

For primitive types such as `Int`, `Float`, `IO` or \rightarrow , no structure is available. For a generic function to work on these types, a specific case is necessary. Other types are deliberately declared abstract to hide the structure of the type. Here, the situation is the same, and a specific case is required to extend a generic function to an abstract type. The types `Unit`, $+$, and \times are abstract types as far as Generic Haskell is concerned, instead of being their own representation types. This choice implies that a missing case for one of these types is reported as a compile-time error.

2.3 Specializing generic functions

In this section we sketch how Generic Haskell specializes a generic function. Assume that *collect*, the collect function from Section 2.1, is called on the type argument `Bool`. No case is given for

Bool, so Generic Haskell considers the structural representation for Bool. The data type Bool and its structural representation are given by

```
data Bool = False | True,
type Boolo = Unit + Unit.
```

We reduce a call of $collect\langle Bool \rangle$ to a call $collect\langle Bool^o \rangle$. The translation of the latter function to Haskell code, using the cases of $collect$ for view types, is quite simple and described elsewhere (EGH,[10]). The call $collect\langle Bool^o \rangle$ is of type $Bool^o \rightarrow [c]$, whereas $collect\langle Bool \rangle$ is of type $Bool \rightarrow [c]$. So to express the call of $collect\langle Bool \rangle$ in terms of the the call of $collect\langle Bool^o \rangle$, we have to lift the isomorphism between Bool and its representation to the type of the generic function $collect$.

Given an embedding-projection pair between a type D and its structure type D^o , we can use the generic function $bimap$ to lift the isomorphism to arbitrarily complex types. Recall that $collect$ is defined in such a way that it returns the empty list for every data type, and only becomes useful when locally redefined. Similarly, $bimap$ defines the identity embedding-projection pair for each data type generically. A remarkable fact is that $bimap$ can be defined on function types. We give the cases for Unit, +, and \rightarrow as an example (see, for example, EGH for a complete definition):

```
 $bimap\langle a_1 :: *, a_2 :: * \rangle :: (bimap\langle a_1, a_2 \rangle) \Rightarrow EP\ a_1\ a_2$ 
bimap $\langle Unit \rangle = EP\ id\ id$ 
bimap $\langle \alpha + \beta \rangle =$ 
  let  $from_+$  ( $Inl\ a$ ) =  $Inl\ (from\ bimap\langle \alpha \rangle\ a)$ 
       $from_+$  ( $Inr\ b$ ) =  $Inr\ (from\ bimap\langle \beta \rangle\ b)$ 
       $to_+$  ( $Inl\ a$ ) =  $Inl\ (to\ bimap\langle \alpha \rangle\ a)$ 
       $to_+$  ( $Inr\ b$ ) =  $Inr\ (to\ bimap\langle \beta \rangle\ b)$ 
  in  $EP\ from_+\ to_+$ 
bimap $\langle \alpha \rightarrow \beta \rangle =$ 
  let  $from_{\rightarrow}\ c = from\ bimap\langle \beta \rangle \cdot c \cdot to\ bimap\langle \alpha \rangle$ 
       $to_{\rightarrow}\ c = to\ bimap\langle \beta \rangle \cdot c \cdot from\ bimap\langle \alpha \rangle$ 
  in  $EP\ from_{\rightarrow}\ to_{\rightarrow}$ .
```

Using local redefinition, we can plug in an embedding-projection pair in $bimap$ to lift the isomorphism between Bool and its representation to the type of the generic function $collect$.

```
 $collect\langle Bool \rangle = \mathbf{let}\ bimap\langle \alpha \rangle = ep_{Bool}\ \mathbf{in}\ to\ (bimap\langle \alpha \rightarrow [c] \rangle)\ collect\langle Bool^o \rangle.$ 
```

The details of why this works are omitted here. It is, however, important to realize that for generic functions that both consume and produce values of the type argument's type, both components of the embedding projection pair will be applied: a value of the original type D is transformed into D^o to be in suitable form to be passed to the function that works on the structural representation. Because the function also returns something containing values of type D^o , these values are then converted back to type D . This is the reason why the embedding-projection pair should contain an isomorphism. If it does not, a value could change simply by the conversion functions that are applied, leading to unexpected results for the user.

3 Views

We have explained how Generic Haskell defines a structural representation type plus an embedding-projection pair for any Haskell data type. A type-indexed function is generic because the embedding-projection pair is applied to the type arguments by the compiler as needed. Other approaches to generic programming use different, but still fixed representations of data types. In this section, we argue that different views improve the expressiveness of a generic programming system, because not every view is equally suitable for every generic function. In Section 4 we will give a formal definition of generic views.

3.1 Fixed points

Consider the data type `Term`, representing lambda terms, and the function `subterms` that, given a term, produces the immediate subterms.

```

data Term      = Var Variable | Abs Variable Term | App Term Term
type Variable = String
subterms      :: Term → [Term]
subterms (Var x) = []
subterms (Abs x t) = [t]
subterms (App t u) = [t, u]

```

This function is an instance of a more general pattern. The function `subtrees`, for example, produces the immediate subtrees of an external binary search tree.

```

subtrees :: ∀ a b . Tree a b → [Tree a b]
subtrees (Tip a) = []
subtrees (Node l b r) = [l, r]

```

Both `subterms` and `subtrees` retrieve the immediate children of a recursive data type’s value. Since the general pattern is clear, we would like to be able to express it as a generic function. However, Generic Haskell does not allow us to define such a function directly, due to the fact that the structure over which generic functions are inductively defined does not expose the recursive calls in a data type’s definition.

Generic Haskell’s precursor, PolyP, does give access to these recursive calls, enabling the definition of a generic function that collects the immediate recursive children of a value [18]. Generic functions in PolyP, however, are limited in the sense that they can only be applied to regular³ data types of kind $* \rightarrow *$. In particular, this precludes nested and mutually recursive data types.

Interestingly, it is possible to write a program in Generic Haskell that produces the immediate children of a value, but it requires some extra effort from the user of the program. If regular recursive data types are expressed using an explicit type-level fixed point operator:

```

data Fix f      = In (f (Fix f))
data TermBase r = VarBase Variable | AbsBase Variable r | AppBase r r
type Term'      = Fix TermBase
data TreeBase a b r = TipBase a | NodeBase r b r
type Tree' a b   = Fix (TreeBase a b),

```

then the generic function `children` can be defined with a single case for `Fix`.

```

children⟨a :: *⟩ :: (collect⟨a | a⟩) ⇒ a → [a]
children⟨Fix ϕ⟩ (In r) = let collect⟨α⟩ x = [x] in collect⟨ϕ α⟩ r

```

The `children` function depends on the `collect` function `collect` defined in Section 2. The local redefinition fixes the type of the produced list and adapts the `collect` function to construct singleton lists from the recursive components in a fixed point’s value. The function `collect` ensures that these singletons are concatenated to produce the result list.

Although this approach works fine, there is an obvious downside. The programmer needs to redefine her recursive data types in terms of `Fix`. Whenever she wants to use `children` to compute the recursive components of a value of any of the original recursive types, say `Term` or `Tree`, a user-defined bidirectional mapping from the original types to the fixed points, `Term'` and `Tree'`, has to be applied.

With a *fixed-point view*, the compiler becomes capable of deriving the fixed point for any regular recursive data type and will generate and apply the required mappings automatically. The

³ A data type is *regular* if it does not contain function spaces, and if the arguments of the data type constructor on the left- and right-hand sides in its definition are the same. So the data type `Flip` defined by `data Flip a b = MkFlip a (Flip b a)` is not regular.

structure of a data type is then no longer perceived as a sum of products, but as the fixed point of a sum of products. The only thing we have to change in the definition of *children* to make use of the new view is to add the name of the view to the type signature:

$$\mathit{children}\langle a :: * \mathbf{viewed} \mathit{Fix} \rangle :: (\mathit{collect}\langle a \mid a \rangle) \Rightarrow a \rightarrow [a].$$

The definition of *children* is unchanged. For example, $\mathit{children}\langle [\mathit{Int}] \rangle [1, 2, 3]$ yields $[[2, 3]]$. The user of the function does not have to worry about defining types in terms of *Fix* any longer: the translation happens behind the scenes.

Another well-known function that can be defined using the fixed-point view is the *catamorphism* [24]. In the definition of *cata* we use a type-indexed type *AlgebraOf*, which returns the algebra of a data type: a function from the pattern functor of the data type to the result type. The details of this definition can be found in EGH, and in the forthcoming release of *Generic Haskell with views*.

The formal definition of the *Fix* view will be given in detail in Section 4.5. In the following, we will discuss examples of other views.

3.2 Balanced sums of products

Traditionally, *Generic Haskell* views the structure of data types using nested right-deep binary sums and products. The choice for such a view is rather arbitrary. A nested left-deep view or a balanced view may be just as suitable. However, the chosen view has some impact on the behaviour of certain generic programs.

The generic function *enc*, for instance, encodes values of data types as lists of bits [17, 9].

$$\begin{aligned} \mathbf{data} \text{ Bit} &= \mathit{Off} \mid \mathit{On} \\ \mathit{enc}\langle a :: * \rangle &:: (\mathit{enc}\langle a \rangle) \Rightarrow a \rightarrow [\text{Bit}] \\ \mathit{enc}\langle \mathit{Unit} \rangle \ \mathit{Unit} &= [] \\ \mathit{enc}\langle \alpha + \beta \rangle (\mathit{Inl} \ a) &= \mathit{Off} : \mathit{enc}\langle \alpha \rangle \ a \\ \mathit{enc}\langle \alpha + \beta \rangle (\mathit{Inr} \ b) &= \mathit{On} : \mathit{enc}\langle \beta \rangle \ b \\ \mathit{enc}\langle \alpha \times \beta \rangle (a \times b) &= \mathit{enc}\langle \alpha \rangle \ a \mathbin{\#} \mathit{enc}\langle \beta \rangle \ b \\ \mathit{enc}\langle \mathit{Int} \rangle \ n &= \mathit{encInt} \ n \\ \mathit{enc}\langle \mathit{Char} \rangle \ c &= \mathit{encChar} \ c \end{aligned}$$

Here, *encInt* and *encChar* denote primitive encoders for integers and characters, respectively. The interesting cases are the ones for sums where a bit is emitted for every choice that is made between a pair of constructors. In the case for products the encodings of the constituent parts are concatenated.

Applying a nested right-deep view to the type *Compass* of directions

$$\mathbf{data} \ \mathit{Compass} \ = \ \mathit{North} \mid \mathit{East} \mid \mathit{South} \mid \mathit{West},$$

gives the structure

$$\mathbf{type} \ \mathit{Compass}^\circ = \mathit{Unit} + (\mathit{Unit} + (\mathit{Unit} + \mathit{Unit})).$$

Using this structure, encoding a value with *enc* takes one bit at best (*North*) and three bits at worst (*West*). In contrast, a balanced view *Bal* on the structure, i.e.,

$$\mathbf{type} \ \mathit{Compass}_B^\circ = (\mathit{Unit} + \mathit{Unit}) + (\mathit{Unit} + \mathit{Unit}),$$

requires only two bits for any value of *Compass*.

In general, encoding requires $O(n)$ bits on average when a nested structure representation is applied, and $O(\log n)$ bits when a balanced representation is used. All we have to do (next to implementing a balanced view *Bal*) is to change the type signature of *enc* accordingly:

$$\mathit{enc}\langle a :: * \mathbf{viewed} \mathit{Bal} \rangle :: (\mathit{enc}\langle a \rangle) \Rightarrow a \rightarrow [\text{Bit}].$$

3.3 List-like sums and products

Suppose we have a generic function *show* which is of type

$$\text{show}\langle a :: * \rangle :: (\text{show}\langle a \rangle) \Rightarrow a \rightarrow \text{String}$$

and produces a human-readable string representation of a value. We want to write a function that shows only a part of a value. The part that is shown is determined by a path, of type

type Path = [Int].

Non-empty lists of type Path select a part of the data structure to print. For instance, [1] selects the second field of the top-level constructor, and [1, 0] selects the first field of the top-level constructor thereof. The function is called *showPart* and has type

$$\text{showPart}\langle a :: * \rangle :: (\text{show}\langle a \rangle, \text{showPart}\langle a \rangle) \Rightarrow \text{Path} \rightarrow a \rightarrow \text{String}.$$

Let us look at the definition of *showPart* on products:

$$\text{showPart}\langle \alpha \times \beta \rangle (0 : p) (a \times b) = \text{if } \text{null } p \text{ then } \text{show}\langle \alpha \rangle a \text{ else } \text{showPart}\langle \alpha \rangle p a.$$

If the first path element is 0, we know that the leftmost field is selected. The encoding in binary products is such that the left component is always a field of the constructor, and not an encoding of multiple fields. We can therefore test if the remainder of the path is empty: if this is the case, we show the complete field using *show*; otherwise, we show the part of the field that is selected by the tail of the current path.

$$\text{showPart}\langle \alpha \times \beta \rangle (n : p) (a \times b) = \text{showPart}\langle \beta \rangle (n - 1 : p) b$$

If the first path element is not 0, we can decrease it by one and show a part of the right component, containing the remaining fields.

There are several problems with this approach. Consider the following data type and its structural representation:

```
data Con012 a b = Con0 | Con1 a | Con2 a b
type Con012o a b = Unit + a + a × b.
```

Using the standard view of Generic Haskell a product structure is only created if there are at least two fields. If there is only one field, such as for *Con1*, the single field (here *a*) is the representation. Obviously, we then cannot use the product case of the generic function to make sure that 0 is the topmost element of the path.

We could add a check to the sum case of the function, detecting the size of the underlying product structure by calling another generic function, or by modifying the type of *showPart* to pass additional information around. However, consider a datatype *Rename* and its structural representation:

```
data Rename = R Original
type Renameo = Original.
```

The structural representation does not even contain a sum structure. Although it is possible to write *showPart* in the standard view, it is extremely tiresome to do so. The same functionality has to be distributed over a multitude of different cases, simply because the structural encoding is so irregular, and we cannot rely on sum and product structure to be present in any case.

A list-like view *List* on data types can help. For this purpose we introduce a data type without constructors and without values (except bottom).

```
data Zero
```

The type *Zero* plays the role of a neutral element for sums in the same way as *Unit* does for products. The above definition is not Haskell 98, but is supported by GHC and can be simulated in Haskell 98.

In our list-like view, the left component of a sum always encodes a single constructor, and the right component of a sum is either another sum or *Zero*. For products, the left component is

always a single field, and the right component is either another product or Unit. In particular, there is always a sum and a product structure. The data type `Con012` is encoded as follows:

```
type Con012o a b = Unit + a × Unit + a × b × Unit + Zero.
```

Now, we can define `showPart` easily:

```
showPart⟨a :: * viewed List⟩ :: (show⟨a⟩, showPart⟨a⟩) ⇒ Path → a → String
showPart⟨Unit⟩ _      Unit    = error "illegal path"
showPart⟨α × β⟩ (0 : p) (a × b) = showPart⟨α⟩      p a
showPart⟨α × β⟩ (n : p) (a × b) = showPart⟨β⟩ (n - 1 : p) b
showPart⟨Zero⟩ _      _      = error "cannot happen"
showPart⟨α + β⟩ []    x      = show⟨α + β⟩ x
showPart⟨α + β⟩ p    (Inl a) = showPart⟨α⟩ p a
showPart⟨α + β⟩ p    (Inr b) = showPart⟨β⟩ p b.
```

We have moved the check for the empty path to the sum case. We can do this because we know that every data type has a sum structure in the list-like view.

3.4 Boilerplate approach

In the ‘Scrap Your Boilerplate’ approach, Lämmel and Peyton Jones present a design pattern for writing programs that traverse data structures [20, 21]. These traversals are defined using a relatively small library that comprises two types of generic combinators: recursive traversals and type extensions. Generic functions are defined in terms of these library functions, and not by induction on the structure of types. The library functions, however, do use a particular view on data types. This section discusses this view, dubbed `Boilerplate`, and shows how to implement a traversal function based on this view. The emulation of the boilerplate approach in Generic Haskell is useful for comparing different approaches to generic programming, but it turns out to be less convenient to use than the original boilerplate library.

In the boilerplate approach all traversals are instances of a general scheme imposed by a left-associative generic fold over constructor applications. So a type is viewed as a sum of products, where a product is either a nullary constructor, or the application of a constructor(-application) to a type. To emulate the behaviour of the generic fold, the product constructor `×` in the `Boilerplate` view is left associative as well. The right component of a product is always a single field, and the left component is either another product or Unit, similar to the `List` view from Section 3.3.

For example, the `Boilerplate` view representations of the types of lists and trees are given by:

```
type ListoBP a = Unit + (Unit × a) × List a
type TreeoBP a b = Unit × a + ((Unit × Tree a b) × b) × Tree a b.
```

Besides generic traversals such as the generic fold, the `Boilerplate` view makes use of type extensions. A type extension extends the type of a function such that it works on many types instead of a single type. To emulate type extensions, we have to be able to distinguish types by name. Therefore we use a type-indexed function equipped with cases for specific types. The remaining cases – that is, the extension – are provided in a definition on a view that does not operate on the structure of types. For this purpose, we use the identity view (`Id`), which merely wraps a data type in the `Id` data type:

```
data Id a = Id a.
```

For example, the function `addPrefixTerm` prefixes variables of a term.

```
addPrefixTerm      :: Term → Term
addPrefixTerm (Var x) = Var ("gh_" ++ x)
addPrefixTerm x      = x
```

This function is extended as follows to work on any type:

```

addPrefix⟨a :: * viewed Id⟩ :: a → a
addPrefix⟨Term⟩ x      = addPrefixTerm x
addPrefix⟨Id α⟩ (Id x) = Id x.

```

Generic Haskell’s specialisation algorithm chooses the *Term* arm when *addPrefix* is applied to that type. For all other types, the last arm is selected.

The definitions of the traversal combinators rely on the list-like character of the *Boilerplate* view. For example, the *gmapT* combinator applies a transformation argument to the immediate children of a node, traversing it in a right to left fashion. The transformation argument is a type-extended function, which in our approach is modeled by a type-indexed function.

We implement type-extended arguments to combinators as type-indexed functions bound to the combinator’s name followed by the *Par* suffix. For example, the *gmapT* combinator has the following definition (omitting the dependencies in the type, these can easily be inferred from the function definition):

```

gmapT⟨a :: * viewed Boilerplate⟩ :: a → a
gmapT⟨Unit⟩ Unit      = Unit
gmapT⟨α + β⟩ (Inl a) = Inl (gmapT⟨α⟩ a)
gmapT⟨α + β⟩ (Inr b) = Inr (gmapT⟨β⟩ b)
gmapT⟨α × β⟩ (a × b) = gmapT⟨α⟩ a × gmapTPar⟨β⟩ b.

```

The default definition of the transformation argument is the identity:

```

gmapTPar⟨a :: * viewed Id⟩ :: a → a
gmapTPar⟨Id α⟩ (Id x) = Id x.

```

The *everywhere* combinator applies a transformation to all nodes in a tree, traversing it in a bottom-up fashion. It is defined in terms of the simple non-recursive one-layer traversal combinator *gmapT*. Since we use an instance of *gmapT* for *everywhere*, we call *gmapTInst* which is the instance of *gmapT* where the parameter of *gmapT* is instantiated with *everywhere*.

```

everywhere⟨a :: * viewed Id⟩ :: a → a
everywhere⟨Id α⟩ (Id x) = Id (everywherePar⟨α⟩ (gmapTInst⟨α⟩ x))
everywherePar⟨a :: * viewed Id⟩ :: a → a
everywherePar⟨Id α⟩ (Id x) = Id x

```

The first function transforms the children by means of a call to *gmapTInst* and then applies the transformation argument *everywherePar* to the result. Generic function *gmapTInst* is the defunctionalized equivalent of the application of *gmapT* to *everywhere*:

```

gmapTInst⟨a :: * viewed Boilerplate⟩ :: a → a
gmapTInst extends gmapT
where everywhere as gmapTPar.

```

Function *gmapTInst* is defined by means of a *default case* (EGH,[7]): it behaves as *gmapT* on all types except that the dependency on *gmapTPar* is changed to a dependency on *everywhere*.

It is now trivial to write a function that adds prefixes generically by ‘applying’ *everywhere* to *addPrefix*.

```

genAddPrefix⟨a :: *⟩ :: a → a
genAddPrefix extends everywhere
where addPrefix as everywherePar.

```

Note that the case for products only recurses on the left component of a product. Since the *Boilerplate* view guarantees that all fields of a constructor are encoded as right components of products, it is easy to verify that *gmapT* does indeed define a non-recursive traversal. This simple non-recursive scheme allows us to derive several rich recursive traversal strategies from a single base combinator. These strategies are written using default cases.

The type-extension operators used in the Boilerplate approach can be defined using type-indexed functions on the `ld` view. First class type-indexed functions are not supported in Generic Haskell. We emulate application of generic combinators to type-extension operators using defunctionalization techniques [8, 28]. Defunctionalization is a standard technique to transform higher-order programs into first-order equivalents.

Because Generic Haskell lacks higher-order generic functions, these and other Scrap Your Boilerplate examples are better expressed using the standard view instead of the Boilerplate view. We believe, however, that an encoding of the Boilerplate approach within the view formalism can help to better compare it with other approaches, and improve the overall understanding of different generic programming techniques.

Hinze, Löhr and Oliveira [13] define a generic boilerplate view using generalized algebraic data types. The view uses run-time representations of types and higher order functions, and is hence closer to the original Boilerplate approach. It follows that this view represents boilerplate functions more faithfully. Generic Haskell does not use run-time representations of types, so we cannot use the same approach.

4 Generic views, formally

The previous section shows why generic views are useful. This section formally defines generic views, and presents the formal definition of the standard view and the fixed-point view. The other views mentioned in the previous section can also be defined using the formalism introduced in this section.

4.1 Notation

Throughout this section, we often use the following notation to denote repetition:

$$\begin{aligned} \{X_i\}^{i \in 1..n} &\equiv X_1 \dots X_n \\ \{X_i\};^{i \in 1..n} &\equiv X_1; \dots; X_n \end{aligned}$$

If not explicitly mentioned otherwise, such repetitions can be empty, i.e., n can be 0. We sometimes omit the range of the variable if it is irrelevant or clear from the context.

4.2 Syntax

Programs. Figure 1 shows the syntax of programs in the core language. This language is a rather standard functional language. A program consists of zero or more type declarations and a single expression: the main function.

Programs $P ::= \{D_i\}^i e$	type declarations and main expression	Expressions $e ::= x$	variable
Value declarations $d ::= x = e$	function declaration	$ C$ $ \lambda x . e$ $ (e_1 e_2)$	data constructor abstraction application
Patterns $p ::= x$	variable pattern	$ \mathbf{case} e_0 \mathbf{of} \{p_i \rightarrow e_i\}^i$ $ (\mathbf{fix} e)$	case fixed point
$ (C \{p_i\}^i)$	constructor pattern	$ \mathbf{let} \{d_i\}^i \mathbf{in} e$	let

Fig. 1. Syntax of the expression language

Types and kinds. The syntax of the type and kind language is shown in Figure 2. New types are introduced by means of **data** declarations. Such a declaration associates a type constructor with zero or more data constructors, each of which has zero or more fields. The *parameterized types* are explained below.

<p>Type declarations</p> $D ::= \mathbf{data} \ T = \{ \Lambda a_i :: \kappa_i . \}^i \{ C_j \{ t_{j,k} \}^k \}^j$ <p style="text-align: center;">algebraic data type</p> <p>Parameterized types</p> $u ::= \{ \Lambda a_i :: \kappa_i . \}^i t$ <p style="text-align: center;">type-level abstraction</p>	<p>Types</p> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">$t ::= a$</td> <td>type variable</td> </tr> <tr> <td style="padding-right: 10px;"> T</td> <td>type constructor</td> </tr> <tr> <td style="padding-right: 10px;"> $(t_1 \ t_2)$</td> <td>type application</td> </tr> <tr> <td style="padding-right: 10px;"> $\forall a :: \kappa . t$</td> <td>universal quantification</td> </tr> </table> <p>Kinds</p> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">$\kappa ::= *$</td> <td>kind of proper types</td> </tr> <tr> <td style="padding-right: 10px;"> $\kappa_1 \rightarrow \kappa_2$</td> <td>function kind</td> </tr> </table>	$t ::= a$	type variable	T	type constructor	$(t_1 \ t_2)$	type application	$\forall a :: \kappa . t$	universal quantification	$\kappa ::= *$	kind of proper types	$\kappa_1 \rightarrow \kappa_2$	function kind
$t ::= a$	type variable												
T	type constructor												
$(t_1 \ t_2)$	type application												
$\forall a :: \kappa . t$	universal quantification												
$\kappa ::= *$	kind of proper types												
$\kappa_1 \rightarrow \kappa_2$	function kind												

Fig. 2. Syntax of types and kinds

<p>Kind environments</p> $K ::= \varepsilon$ <p style="text-align: center;">empty kind environment</p> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;"> $K, a :: \kappa$</td> <td>type-variable binding</td> </tr> <tr> <td style="padding-right: 10px;"> $K, T :: \kappa$</td> <td>type-constructor binding</td> </tr> </table>	$K, a :: \kappa$	type-variable binding	$K, T :: \kappa$	type-constructor binding	<p>Type environments</p> $\Gamma ::= \varepsilon$ <p style="text-align: center;">empty type environment</p> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;"> $\Gamma, x :: t$</td> <td>variable binding</td> </tr> <tr> <td style="padding-right: 10px;"> $\Gamma, C :: t$</td> <td>data-constructor binding</td> </tr> </table>	$\Gamma, x :: t$	variable binding	$\Gamma, C :: t$	data-constructor binding
$K, a :: \kappa$	type-variable binding								
$K, T :: \kappa$	type-constructor binding								
$\Gamma, x :: t$	variable binding								
$\Gamma, C :: t$	data-constructor binding								

Fig. 3. Syntax of environments

Generic programming extensions. To facilitate generic programming, the core language is extended with parameterized type patterns and type-indexed functions with dependencies, see EGH. This extension is omitted, we only have to adapt it with the facility to specify a view in the signature of a generic function.

Structure types in Haskell are declared as **type** synonyms. Type synonyms are not supported in the core language. Therefore, to describe structure types, the language contains *parameterized types*, which are essentially a nesting of type-level lambda abstractions around a type of kind $*$. Parameterized types are only used in view definitions, they cannot appear in a core-language program.

Rules. The well-formedness rules for programs, types and kinds, the kinding rules for types and the typing rules for expressions are standard. The operational semantics of the core language is omitted. More information about the core language can be found elsewhere (EGH,[14]).

4.3 Definitions

Using the notion of parameterized types, we can formalize the observation that a view comprises a collection of view types and algorithms for the generation of structure types and conversion functions. In the following definitions we will use kind environments and type environments; their syntax is defined in Figure 3.

Definition 1 (Generic View). A generic view \mathcal{V} consists of a collection of bindings for view types,

$$\text{viewtypes}_{\mathcal{V}} \equiv K; \Gamma,$$

a partial mapping from types to structure types,

$$\mathcal{V}[[D_0]]^{\text{str}} \equiv u; \{D_i\}^{i \in 1..n},$$

and, for each type in the domain of this mapping, conversions between values and structure values,

$$\mathcal{V}[[D_0]]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}}.$$

Notice that we allow the mapping from types to structure types to generate zero or more additional declarations for supporting data types. The types introduced by these declarations can be used for the generation of structure types. This is used in the fixed-point view, for example.

For a view to be useful for generic programming, we require it to have three essential properties. First, the mapping from types to structure types should preserve kinds.

Definition 2 (Kind Preservation). A generic view \mathcal{V} with

$$\text{viewtypes}_{\mathcal{V}} \equiv K_{\mathcal{V}}; \Gamma_{\mathcal{V}}$$

is kind preserving if for each well-formed declaration D_0 of a type constructor T such that $K \vdash T :: \kappa$, for which a structure type u can be derived,

$$\mathcal{V}[[D_0]]^{\text{str}} \equiv u; \{D_i\}^{i \in 1..n},$$

it follows that under kind environment K'

$$K' \equiv K, K_{\mathcal{V}} \{T_i :: \kappa_i\}^{i \in 1..n},$$

containing K , $K_{\mathcal{V}}$, and all the kinds of the D_i declarations, the supporting type declarations D_i are well-formed and the structure type u has the same kind κ as the original type T ,

$$K' \vdash u :: \kappa.$$

Furthermore, the conversion functions derived from a type declaration should be well-typed and indeed convert between values of the original type and values of the structure type, which is captured by the following definition.

Definition 3 (Well-typed Conversion). A view \mathcal{V} with

$$\text{viewtypes}_{\mathcal{V}} \equiv K_{\mathcal{V}}; \Gamma_{\mathcal{V}}$$

generates well-typed conversions if, for each well-formed declaration D_0 of a type constructor T of kind $\{\kappa_i \rightarrow\}^{i \in 1..\ell}$, for which a structure type t can be derived,

$$\mathcal{V}[[D_0]]^{\text{str}} \equiv \{\Lambda a_i :: \kappa_i .\}^{i \in 1..\ell} t; \{D_i\}^{i \in 1..n},$$

it follows that the corresponding conversion functions e_{from} and e_{to} ,

$$\mathcal{V}[[D_0]]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}},$$

take values of the original data type T to values of the structure type t and vice versa,

$$\begin{aligned} K'; \Gamma' \vdash e_{\text{from}} &:: \{\forall a_i :: \kappa_i .\}^{i \in 1..\ell} T \{a_i\}^{i \in 1..\ell} \rightarrow t \\ K'; \Gamma' \vdash e_{\text{to}} &:: \{\forall a_i :: \kappa_i .\}^{i \in 1..\ell} t \rightarrow T \{a_i\}^{i \in 1..\ell} \end{aligned}$$

under environments K' as in Definition 2 and Γ' containing the view bindings $\Gamma_{\mathcal{V}}$ and the types of the constructors from D_0 and all D_i .

Finally, the conversion functions from structure values to values should form the inverses of the corresponding functions in the opposite direction:

Definition 4 (Well-behaved Conversion). A generic view \mathcal{V} produces well-behaved conversions if, for each well-formed declaration D of a type constructor T , conversion functions e_{from} and e_{to} are generated,

$$\mathcal{V}[[D]]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}},$$

such that e_{to} is the left inverse of e_{from} with respect to function composition:

$$e_{\text{to}} (e_{\text{from}} v) \text{ evaluates to } v$$

for each value v of type T .

(Note that, for a well-behaved conversion pair, the function that takes values to structure values is injective; thus, a structure type should have at least as many elements as the corresponding original type.) Why do we want a generic view to have well-behaved conversion? Assume function gid is a generic identity function that is defined as a recursive function that traverses and copies the structure. To prove that this function is an identity, we have to ensure that the conversions that are applied during the traversal are well-behaved and do not modify the value.

Only views that possess all three of the discussed properties are considered valid:

Definition 5 (Validity). *A generic view is valid if it is kind preserving and generates well-typed, well-behaved conversions.*

Validity of a view has two important consequences. First, well-behaved conversion allows us to prove properties like the property for the generic identity function given after Definition 4. Second, for a valid view \mathcal{V} we can prove Theorem 11.2 of EGH in the context of views. This theorem states that specializing a call to a generic function on a new data type by constructing a wrapper around the call that uses the structural representation of that data type and the embedding-projection for that data type specified in the view, produces a type-correct function:

Theorem 1. *Let \mathcal{V} be a view with*

$$\begin{aligned} \mathcal{V}[[D_0]]^{\text{str}} &\equiv u; \{D_i\}^{i \in 1..n} \\ \mathcal{V}[[D_0]]^{\text{conv}} &\equiv e_{\text{from}}; e_{\text{to}}. \end{aligned}$$

For a type-indexed function x of arity $\langle r \mid s \rangle$, where all non-generic slots of x are of kind $*$, the declaration

$$\begin{aligned} &\text{let } \{ \text{bimap} \langle \beta_i \rangle = EP \ e_{\text{from}} \ e_{\text{to}} \}^{i \in 1..r}; \\ &\quad \{ \text{bimap} \langle \gamma_j \rangle = EP \ id \ id \}^{j \in 1..s} \\ &\text{in } to \ \text{bimap}(\text{base } (x \langle \{ \beta_i \}^{i \in 1..r} \mid \{ \gamma_j \}^{j \in 1..s} \rangle)) \ x(u) \end{aligned}$$

has the same type as $x \langle T \ \{ \alpha_j \}^{j \in 1..n} \rangle$. Here $\text{base } (f)$ returns the base type (see EGH) of f , i.e., the type specified for the generic function.

We will use the above declaration as the translation (or specialization) of

$$x \langle T \ \{ \alpha_j \}^{j \in 1..n} \ \text{viewed } \mathcal{V} \rangle.$$

The proof of this theorem (which is very similar to the proof of Theorem 11.1 in EGH) uses the facts that a valid view preserves kinds, and has well-typed conversions. Well-behavedness of the conversions is not necessary for proving the theorem.

4.4 The standard view

We describe the three components of a generic view for the standard Generic Haskell view \mathcal{S} of data types

View types. The view types of the standard view are given by the following declarations:

```

data Zero =
data Unit = Unit
data Sum =  $\Lambda a :: * . \Lambda b :: * . Inl \ a \mid Inr \ b$ 
data Prod =  $\Lambda a :: * . \Lambda b :: * . a \times b$ .

```

These types represent nullary sums, nullary products, binary sums, and binary products, respectively. It is easy to convert these definitions into bindings in the environments Γ and K .

Generating structure types. The algorithm that generates structural representations for data types is expressed by judgements of the forms

$$\begin{aligned} \mathcal{S}[[D_0]]^{\text{str}} &\equiv u; \{D_i\}^{i \in 1..n} \\ \mathcal{S}[[\{C_j \{t_{j,k}\}^{k \in 1..n_j}\}_1^{j \in 1..m}]]^{\text{str}} &\equiv t. \end{aligned}$$

The former express how type declarations are mapped to parameterized types and lists of supporting declarations; the latter express how a type is derived from a list of constructors. The rules are shown in Figures 4 and 5.

$$\boxed{\mathcal{S}[[D_0]]^{\text{str}} \equiv u; \{D_i\}^i}$$

$$\frac{\mathcal{S}[[\{C_j \{t_{j,k}\}^k\}_1^j]]^{\text{str}} \equiv t}{\mathcal{S}[\text{data } T = \{\Lambda a_i :: \kappa_i .\}^i \{C_j \{t_{j,k}\}^k\}_1^j]^{\text{str}} \equiv \{\Lambda a_i :: \kappa_i .\}^i t; \varepsilon}$$

Fig. 4. Representation of data types in the standard view

Type declarations are handled by the rule in Figure 4. The type parameters of a declared type constructor are directly copied to the resulting structure type. Notice that the standard view does not need auxiliary declarations.

For constructors, we distinguish five cases. The first rule, (str-std-1), represents empty constructor lists with *Zero*. The next three cases handle singleton lists of constructors. Fieldless constructors are, by rule (str-std-2), represented by nullary products. Rule (str-std-3) represents a unary constructors by the type of its field. If a constructor has two or more fields, rule (str-std-4) generates a product type and recurses. Finally, lists that contain two or more constructors are represented by a recursively built sum (str-std-5).

$$\boxed{\mathcal{S}[[\{C_j \{t_{j,k}\}^k\}_1^j]]^{\text{str}} \equiv t}$$

$$\frac{}{\mathcal{S}[[\varepsilon]]^{\text{str}} \equiv \text{Zero}} \quad (\text{str-std-1}) \quad \frac{}{\mathcal{S}[[C]]^{\text{str}} \equiv \text{Unit}} \quad (\text{str-std-2}) \quad \frac{}{\mathcal{S}[[C \ t]]^{\text{str}} \equiv t} \quad (\text{str-std-3})$$

$$\frac{n \in 2.. \quad \mathcal{S}[[C \ \{t_k\}^{k \in 2..n}]]^{\text{str}} \equiv t'_2}{\mathcal{S}[[C \ \{t_k\}^{k \in 1..n}]]^{\text{str}} \equiv \text{Prod } t_1 \ t'_2} \quad (\text{str-std-4})$$

$$\frac{m \in 2.. \quad \mathcal{S}[[\{C_j \ \{t_{j,k}\}^{k \in 1..n_j}\}_1^{j \in 2..m}]]^{\text{str}} \equiv t_2 \quad \mathcal{S}[[C_1 \ \{t_{1,k}\}^{k \in 1..n_1}]]^{\text{str}} \equiv t_1}{\mathcal{S}[[\{C_j \ \{t_{j,k}\}^{k \in 1..n_j}\}_1^{j \in 1..m}]]^{\text{str}} \equiv \text{Sum } t_1 \ t_2} \quad (\text{str-std-5})$$

Fig. 5. Representation of constructors in the standard view

Generating conversions. The rules for generating conversion functions are shown in Figures 6 and 7 and are of the forms

$$\begin{aligned} \mathcal{S}[[D_0]]^{\text{conv}} &\equiv e_{\text{from}}; e_{\text{to}} \\ \mathcal{S}[[\{C_j \{t_{j,k}\}^k\}_j^j]]^{\text{conv}} &\equiv \{p_{\text{from},j}\}_j^j; \{p_{\text{to},j}\}_j^j, \end{aligned}$$

i.e., type declarations give rise to pairs of conversion functions, while lists of data constructors give rise to pairs of patterns.

$$\boxed{\mathcal{S}[[D]]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}}}$$

$$\frac{\mathcal{S}[[\{C_j \{t_{j,k}\}^k\}_j^j]]^{\text{conv}} \equiv \{p_{\text{from},j}\}_j^j; \{p_{\text{to},j}\}_j^j \quad e_{\text{from}} \equiv \lambda x . \mathbf{case} \ x \ \mathbf{of} \ \{p_{\text{from},j} \rightarrow p_{\text{to},j}\}_j^j \quad e_{\text{to}} \equiv \lambda x . \mathbf{case} \ x \ \mathbf{of} \ \{p_{\text{to},j} \rightarrow p_{\text{from},j}\}_j^j}{\mathcal{S}[[\mathbf{data} \ T = \{\Lambda a_i :: \kappa_i . \}_i \{C_j \{t_{j,k}\}^k\}_j^j]]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}}}$$

Fig. 6. Conversions for data types in the standard view

The rule in Figure 6 constructs a ‘from’ function that matches values of the original type against a list of patterns. If a value is successfully matched against a certain pattern, a structure value is produced by using a complementary pattern; hence, here, we make use of the fact that the pattern language is just a subset of the expression language. A ‘to’ function is created by inverting the patterns.

The pairs of pattern lists are generated using the rules for constructor lists. These rules are analogous to the rules for generating structure types from constructor lists.

If there are no constructors, there are no patterns either (conv-std-1). Rule (conv-std-2) associates a single constructor with the value *Unit*. Rule (conv-std-3) associates unary constructors with variables that correspond to their field values. If a constructor has two or more fields, rule (conv-std-4) associates the corresponding variables to product patterns. Finally, if the list of constructors has two or more elements, rule (conv-std-5) applies; it prefixes the patterns with the injection constructors *Inl* and *Inr*.

$$\boxed{\mathcal{S}[[\{C_j \{t_{j,k}\}^k\}_j^j]]^{\text{conv}} \equiv \{p_{\text{from},j}\}_j^j; \{p_{\text{to},j}\}_j^j}$$

$$\frac{}{\mathcal{S}[[\varepsilon]]^{\text{conv}} \equiv \varepsilon; \varepsilon} \quad (\text{conv-std-1}) \quad \frac{}{\mathcal{S}[[C \ t]]^{\text{conv}} \equiv C \ x; x} \quad (\text{conv-std-3})$$

$$\frac{}{\mathcal{S}[[C]]^{\text{conv}} \equiv C; \text{Unit}} \quad (\text{conv-std-2}) \quad \frac{n \in 2.. \quad \{x_1 \neq x_k\}^{k \in 2..n} \quad \mathcal{S}[[C \ \{t_k\}^{k \in 2..n}]]^{\text{conv}} \equiv C \ \{x_k\}^{k \in 2..n}; p_{\text{to}}}{\mathcal{S}[[C \ \{t_k\}^{k \in 1..n}]]^{\text{conv}} \equiv C \ \{x_k\}^{k \in 1..n}; x_1 \times p_{\text{to}}} \quad (\text{conv-std-4})$$

$$\frac{m \in 2.. \quad \mathcal{S}[[\{C_j \{t_{j,k}\}^{k \in 1..n_j}\}_j^{j \in 2..m}]]^{\text{conv}} \equiv \{p_{\text{from},j}\}_j^{j \in 2..m}; \{p_{\text{to},j}\}_j^{j \in 2..m} \quad \mathcal{S}[[C_1 \ \{t_{1,k}\}^{k \in 1..n_1}]]^{\text{conv}} \equiv p_{\text{from},1}; p_{\text{to},1}}{\mathcal{S}[[\{C_j \{t_{j,k}\}^{k \in 1..n_j}\}_j^{j \in 1..m}]]^{\text{conv}} \equiv \{p_{\text{from},j}\}_j^{j \in 1..m}; \text{Inl} \ p_{\text{to},1} \ \{ \text{Inr} \ p_{\text{to},j}\}_j^{j \in 2..m}} \quad (\text{conv-std-5})$$

Fig. 7. Conversions for constructors in the standard view

Theorem 2. *The standard view is valid.*

A partial proof of this theorem is obtained from the kind-checkers of Generic Haskell and Haskell, and the type-checker of Haskell. Well-behaved conversion is proved by means of a standard induction argument over the construction of the conversions.

4.5 The fixed-point view

An essential aspect of the fixed-point view is the automatic derivation of pattern functors.

Given a declaration D_1 of type T , a declaration D_2 for the pattern functor $\text{ptr}(T)$ is generated by the rule in Figure 8, which takes the form

$$\llbracket D_1 \rrbracket^{\text{ptr}} \equiv D_2.$$

The metafunction ptr produces a unique name for the functor. The definition of $\text{ptr}(T)$ follows the structure of T , replacing all recursive calls by an extra type argument.

$$\boxed{\llbracket D_1 \rrbracket^{\text{ptr}} \equiv D_2}$$

$$D \equiv \text{data ptr}(T) = \{\Lambda a_i :: * . \}^i \Lambda a_{\ell+1} :: * . \{\text{ptr}(C_j) \{t'_{j,k}\}^k\}_j^j$$

$$\frac{\{\Lambda a_{\ell+1} \neq a_i\}^i \quad \{\{t'_{j,k} \equiv [a_{\ell+1} / T \{a_i\}^i] t_{j,k}\}^k\}_j^j}{\llbracket \text{data } T = \{\Lambda a_i :: * . \}^i \{C_j \{t_{j,k}\}^k\}_j^j \rrbracket^{\text{ptr}} \equiv D}$$

Fig. 8. Pattern functors

View types. The sole view type of the fixed-point view is Fix :

$$\text{data Fix} = \Lambda \varphi :: * \rightarrow * . \text{In } (\varphi (\text{Fix } \varphi)).$$

Generating structure types. The rule for generating structure types for \mathcal{F} , which has the form

$$\mathcal{F} \llbracket D_0 \rrbracket^{\text{str}} \equiv u; \{D_i\}_{i \in 1..n}.$$

is given in Figure 9. This rather straightforward rule states that a structure type is derived by applying the type Fix to a partially applied pattern functor. The declaration of the pattern functor is emitted as a supporting declaration. Note that the parameters of the original data type are restricted to kind $*$, excluding higher-order kinded types from the view domain. The need for this restriction will be explained later.

$$\boxed{\mathcal{F} \llbracket D_0 \rrbracket^{\text{str}} \equiv u; \{D_i\}_i^i}$$

$$D \equiv \text{data } T = \{\Lambda a_i :: * . \}^i \{C_j \{t_{j,k}\}^k\}_j^j$$

$$\frac{}{\text{type } \mathcal{F} \llbracket D \rrbracket^{\text{str}} \equiv \{\Lambda a_i :: * . \}^i \text{Fix } (\text{ptr}(T) \{a_i\}^i); \llbracket D \rrbracket^{\text{ptr}}}$$

Fig. 9. Representation of data types in the fixed-point view

Generating conversions. Generating conversion functions for \mathcal{F} is more involved. The algorithm is presented in Figures 10 and 11. It consists of judgements of the forms

$$\begin{aligned} \mathcal{F}[\![D_0]\!]^{\text{conv}} &\equiv e_{\text{from}}; e_{\text{to}} \\ \mathcal{F}[\![t]\!]_{T \{a_i\}^i; ep}^{\text{conv}} &\equiv ep'. \end{aligned}$$

The first form indicates that conversion functions e_{from} and e_{to} are derived based on the structure of a type declaration D_0 . The second form expresses the generation of embedding projection expressions that convert constructor fields to or from the representation type. The generation of these embedding projections is driven by the original type $T \{a_i\}^i$, the type of the field t and the embedding-projection ep for the data type, with type EP $(T \{a_i\}^i)$ (Fix (ptr(T) $\{a_i\}^i$)).

Each conversion function converts the fields using the appropriate components of the field embedding-projections. Because the embedding-projection may be recursively defined, Figure 10 uses the core language's recursive **let** construct.

In Figure 11, the conversion functions for a field are given by the application of *bimap* to the field type with occurrences of $T \{a_i\}^i$ abstracted. This application lifts the given embedding-projection to the abstracted field type in the fashion of Section 2.3. For instance, the embedding-projection of the data type *Rose*

data *Rose* $a = \text{Rose } a$ [*Rose* a]

with type EP (*Rose* a) (Fix (ptr(*Rose*) a)) is lifted for each of the two constructor fields to EP a a and EP [*Rose* a] [Fix (ptr(*Rose*) a)].

$$\boxed{\mathcal{F}[\![D]\!]^{\text{conv}} \equiv e_{\text{from}}; e_{\text{to}}}$$

$$\begin{aligned} &\{\{\mathcal{F}[\![t_{j,k}]\!]_{T \{a_i\}^i; ep}^{\text{conv}} \equiv e_{j,k}\}^k\}^j \\ e_{\text{from}} &\equiv \lambda x . \mathbf{case } x \mathbf{ of } \{C_j \{x_{j,k}\}^k \rightarrow \text{In } (\text{ptr}(C_j) \{from \ e_{j,k} \ x_{j,k}\}^k)\}^j \\ e_{\text{to}} &\equiv \lambda x . \mathbf{case } x \mathbf{ of } \{\text{In } (\text{ptr}(C_j) \{x_{j,k}\}^k) \rightarrow C_j \{to \ e_{j,k} \ x_{j,k}\}^k\}^j \\ ep' &\equiv \mathbf{let } ep = \text{EP } e_{\text{from}} \ e_{\text{to}} \mathbf{ in } ep \end{aligned}$$

$$\mathcal{F}[\![\mathbf{data } T = \{\Lambda a_i :: * . \}^i \{C_j \{t_{j,k}\}^k\}^j]\!]^{\text{conv}} \equiv from \ ep'; to \ ep'$$

Fig. 10. Conversions for data types in the fixed-point view

$$\boxed{\mathcal{F}[\![t]\!]_{T \{a_i\}^i; ep}^{\text{conv}} \equiv ep'}$$

$$\frac{t' \equiv [\alpha / T \{a_i\}^i] t}{\mathcal{F}[\![t]\!]_{T \{a_i\}^i; ep}^{\text{conv}} \equiv \mathbf{let } bimap\langle \alpha \rangle = ep \mathbf{ in } bimap\langle t' \rangle} \quad (\text{conv-fix})$$

Fig. 11. Conversions for fields in the fixed-point view

Theorem 3. *The fixed-point view is valid.*

Mutual recursion, nested data types and higher-order kinded types. In PolyP, generic functions can only be applied to regular data types. This restriction excludes mutually recursive data types,

nested data types and higher-order kinded data types from the class of data types to which a generic function can be applied. The fixed-point view has exactly the same restrictions, and is hence a faithful implementation of the view on data types used in PolyP. It is possible to adapt the fixed-point view such that higher-order kinded types can be handled, but in order to stay as close as possible to the view of data types of PolyP, we refrain from doing so.

5 Generic views in the Generic Haskell compiler

The latest version of the Generic Haskell compiler that implements views can be downloaded via svn: <https://svn.cs.uu.nl:12443/repos/Generic-Haskell/branches/GenericViews>. We have implemented (an extension of) the standard view, the fixed-point view and the list-like sums and products view. The next release of Generic Haskell will come with all the views mentioned in this paper: in addition to the views already implemented, the balanced sums and products view, and the identity and boilerplate views. In the previous version of the Generic Haskell compiler [23] we do not really use the standard view as presented here, but additionally use representation data types `Con` and `Lab` to encode information about constructors and record field labels in the data type. The presence of these data types makes it possible to write functions such as *show* and *read* that produce or consume a representation of a value and therefore rely on the names of constructors and labels. Refactoring the Generic Haskell compiler to support views has improved the structure of the compiler: all view-related components are localized in separate modules.

Although we do not see why we should stop at the six views described in this paper, we expect that these views suffice for most purposes and users. We do not expect that there will be many other useful views, and hence we do not think it is worthwhile to develop a special-purpose language for defining views in user programs. A generic view consists of a set of view types, a function that generates structure types, and a function that generates conversion functions, and it follows that such a special-purpose language for specifying views would be a complete programming language in itself.

To add a new view to Generic Haskell, the compiler has to be modified. Although this might sound scary, in practice it is rather simple. In this section we briefly describe how to modify the Generic Haskell compiler to add a new view. A new view is added to the compiler by implementing a module that contains a number of Haskell functions that directly correspond to the definition of a generic view in Section 4.

5.1 Adding a view to the Generic Haskell compiler

In the Generic Haskell compiler, a view would have the following type:

```
data View = View Name (TDecl → Maybe (LamType, [TDecl], Expr, Expr, [TDecl]))
```

A value of type `View` consists of a name, and a function that can be called on the declaration of a type synonym or a data type (a `TDecl`) to produce a parameterized structure representation type (a `LamType`), supporting type declarations and an embedding-projection pair (two *expr*'s). Views that apply to a subset of the Haskell data types can be implemented by returning *Nothing* on data type definitions that are outside of the view's domain.

Note that the result of the view-generating function in type `View` directly corresponds to the maps $\mathcal{V}[\cdot]^{\text{str}}$ and $\mathcal{V}[\cdot]^{\text{conv}}$. The collection `viewtypesv` of bindings that are required by the view must be added to the Generic Haskell Prelude, i.e., they must be available for the Generic Haskell compiler to parse.

The need for the second list of type declarations is better explained with an example. Consider the application of the *children* function (Section 3) to the `List` data type. This function definition uses the `Fix` structure representation type of `List`. That is, the type `Fix` applied to the supporting type `ptr(List)`, which is the base functor of `List`. This structure representation is not yet enough. The definition of *children* applies function *collect* to the base functor of the data type. Because *collect* is defined on the standard view, we need to generate a standard structure representation for `ptr(List)`, a pair of standard embedding projections and supporting declarations. In short, we

need a `Fix` view on `List` and a standard view on `ptr(List)`. The view components for `ptr(List)` are generated because `ptr(List)` occurs in the second list of type declarations. In general the second list of type declarations is a subset of the list of supporting declarations, which should recursively be processed by `View` functions. The actual implementation keeps track of additional information to determine which `View` function should or should not be called on these declarations, and in certain cases to avoid non-termination.

The validity of a view can only be checked to a certain extent. The compiler can verify the kind preservation and well-typed conversion properties of the view: for each structural representation and embedding-projection pair generated, kind and type checking is performed. The well-behavedness of the conversion cannot be verified by the compiler, since verifying that the composition of two arbitrary functions is the identity is an undecidable problem. This property remains a proof obligation for the implementor of the view.

A view implementor has to deal with some additional implementation details that slightly complicate the `View` data type, but that are not of direct concern for this paper. The interested reader can find more implementation information in the source distribution of Generic Haskell extended with views, in particular in the file `/src/views/README` and the view modules in the same directory.

5.2 Specialization

The specialization mechanism is independent of the actual view, see Theorem 1. For other views than the standard view, different structural representations and embedding-projection pairs are used, but the specialization procedure remains exactly the same. The only thing that has changed in the implementation of specialization within the Generic Haskell compiler is that all the references to structural representation types and embedding-projection pairs point to the view that is specified for the function in question.

6 Conclusions, and related work

We have shown that generic views on data types can make generic functions both easier to write and more efficient. Generic views add expressiveness to a generic programming language. Furthermore, generic views allow us to use different generic programming styles in a single framework.

Although there are a multitude of approaches for generic programming, the idea to use multiple views on data types in a single approach is, to the best of our knowledge, original.

Using our approach to views we can express many different approaches to generic programming in a single framework. Our framework does have a limitation though: because we create a structure type from a single data type declaration, we cannot define a view that considers sets of possibly mutually recursive data types. This is not a fundamental problem: we could have adapted the framework to allow for this. However, since we know of no existing approach to generic programming that would need this extra complexity we have refrained from doing so.

The name “generic view” is derived from Wadler’s proposal to introduce views in (a predecessor of) Haskell [30]. Using one of these views, a single Haskell data type can be analyzed in a different way, by introducing additional constructors by which a value can be constructed, and on which pattern matching can be performed. A view is essentially like the introduction of an additional data type, together with the definition of conversion functions between values from the original type and values of the view type. These conversions are then transparently applied by the compiler where necessary.

Generic views are different in that they define a representation and conversions for many types at the same time. Furthermore, the representation types need not be new data types, but can be built from existing data types. Wadler’s views have the immense advantage that they can be added to the Haskell programming language relatively easily, allowing every programmer to add her own views. On the other hand, generic views have to be added to a generic programming system, such as the Generic Haskell compiler, following the guidelines described in the previous sections. However, we expect that the views we describe in this paper are sufficient for most purposes and

users, and we do not assume a user will frequently want to add a new view to the Generic Haskell compiler.

Both views and generic views have in common that the definition of a new view goes along with a proof obligation for the programmer that cannot easily be captured in a language like Haskell. The conversion between original type and view type, be it a single pair of functions such as in Wadler's proposal, or a type-indexed family of functions such as for generic views, must really witness isomorphisms, otherwise unexpected results may occur.

Since Wadler's views proposal, several variations of views have been given [5, 27, 6]. Our approach is closest to Wadler's proposal in that we also require the existence of an isomorphism between the original type and the view type.

Views have also been proposed in the context of XML and databases [1, 26]. Generic views as proposed in [29] are used to automatically convert between two given views. The generic view concept as introduced in this paper does not seem to have been investigated in this field.

The idea of using different sets of data types for inductive definitions of type-indexed functions is common in the world of dependent types [2, 4]. This corresponds to the idea of having views that work on different subsets of the Haskell data types. However, in the approaches we have seen there is no automatic conversion between syntactically definable data types as offered by the dependently typed programming language into representations as defined by the view or universe.

Acknowledgements. Our thanks go to a number of anonymous referees and Daan Leijen for several helpful comments. This research has partly been funded by the Netherlands Organization for Scientific Research (NWO) (project Generic Haskell, nr. 612.069.000).

References

1. S. Abiteboul. On views and XML. In *Proceedings of the 18th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 1–9. ACM Press, 1999.
2. Thorsten Altenkirch and Conor McBride. Generic programming within dependently typed programming. In Jeremy Gibbons and Johan Jeuring, editors, *Generic Programming: IFIP TC2/WG2.1 Working Conference on Generic Programming July 11–12, 2002, Dagstuhl, Germany*, number 115 in International Federation for Information Processing, pages 1–20. Kluwer Academic Publishers, 2003.
3. Roland Backhouse and Paul Hoogendijk. Generic properties of data types. In *Generic Programming, Advanced Lectures*, volume 2793 of *LNCS*, pages 97–132. Springer-Verlag, 2003.
4. Marcin Benke, Peter Dybjer, and Patrik Jansson. Universes for generic programs and proofs in dependent type theory. *Nordic Journal of Computing*, 10(4):265–289, 2003.
5. F. W. Burton and R. D. Cameron. Pattern matching with abstract data types. *Journal of Functional Programming*, 3(2):117–190, 1993.
6. F. Warren Burton, Erik Meijer, Patrick Sansom, Simon Thompson, and Philip Wadler. Views: an extension to Haskell pattern matching. Available from <http://www.haskell.org/development/views.html>, 1996.
7. Dave Clarke and Andres Löh. Generic Haskell, specifically. In Jeremy Gibbons and Johan Jeuring, editors, *Generic Programming: IFIP TC2/WG2.1 Working Conference on Generic Programming July 11–12, 2002, Dagstuhl, Germany*, number 115 in International Federation for Information Processing, pages 21–47. Kluwer Academic Publishers, 2003.
8. Olivier Danvy and Lasse R. Nielsen. Defunctionalization at work. Research Series RS-01-23, BRICS, Department of Computer Science, University of Aarhus, June 2001. 45 pp. Extended version of an article appearing in Third International Conference on Principles and Practice of Declarative Programming (PPDP 2001), pages 162–174.
9. Ralf Hinze. A generic programming extension for Haskell. In Erik Meijer, editor, *Proceedings of the Third Haskell Workshop*, Technical report of Utrecht University, UU-CS-1999-28, 1999.
10. Ralf Hinze. Polytypic values possess polykinded types. *Science of Computer Programming*, 43(2-3):129–159, 2002.
11. Ralf Hinze and Johan Jeuring. Generic Haskell: applications. In *Generic Programming, Advanced Lectures*, volume 2793 of *LNCS*, pages 57–97. Springer-Verlag, 2003.
12. Ralf Hinze and Johan Jeuring. Generic Haskell: practice and theory. In *Generic Programming, Advanced Lectures*, volume 2793 of *LNCS*, pages 1–56. Springer-Verlag, 2003.

13. Ralf Hinze, Andres Löh, and Bruno C. d. S. Oliveira. "Scrap Your Boilerplate" Reloaded. In *Proceedings of the Eighth International Symposium on Functional and Logic Programming (FLOPS 2006)*, LNCS. Springer-Verlag, 2006.
14. Stefan Holdermans. Generic views. Master's thesis, Institute of Information and Computing Sciences, Utrecht University, 2005.
15. Paul Hoogendijk and Oege de Moor. Container types categorically. *Journal of Functional Programming*, 10(2):191–225, 2000.
16. Patrik Jansson and Johan Jeuring. PolyP — a polytypic programming language extension. In *Conference Record of POPL '97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 470–482. ACM Press, 1997.
17. Patrik Jansson and Johan Jeuring. Polytypic compact printing and parsing. In Doaitse Swierstra, editor, *Proceedings of the 8th European Symposium on Programming, ESOP'99*, volume 1576 of LNCS, pages 273–287. Springer-Verlag, 1999.
18. Patrik Jansson and Johan Jeuring. A framework for polytypic programming on terms, with an application to rewriting. In Johan Jeuring, editor, *Workshop on Generic Programming 2000, Ponte de Lima, Portugal, July 2000*, pages 33–45, 2000. Utrecht Technical Report UU-CS-2000-19.
19. C. Barry Jay. Distinguishing data structures and functions: the constructor calculus and functorial types. In S. Abramsky, editor, *Typed Lambda Calculi and Applications: 5th International Conference TLCA 2001*, volume 2044 of LNCS, pages 217–239. Springer-Verlag, 2001.
20. Ralf Lämmel and Simon Peyton Jones. Scrap your boilerplate: a practical approach to generic programming. *ACM SIGPLAN Notices*, 38(3):26–37, 2003. Proceedings ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI 2003).
21. Ralf Lämmel and Simon Peyton Jones. Scrap more boilerplate: reflection, zips, and generalised casts. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming, ICFP 2004*, pages 244–255. ACM Press, 2004.
22. Andres Löh. *Exploring Generic Haskell*. PhD thesis, Utrecht University, September 2004.
23. Andres Löh, Johan Jeuring, Dave Clarke, Ralf Hinze, Alexey Rodriguez, and Jan de Wit. The Generic Haskell user's guide, version 1.42 (Coral). Technical Report UU-CS-2005-004, Institute of Information and Computing Sciences, Utrecht University, 2005.
24. G. Malcolm. Data structures and program transformation. *Science of Computer Programming*, 14:255–279, 1990.
25. E. Meijer, M.M. Fokkinga, and R. Paterson. Functional programming with bananas, lenses, envelopes, and barbed wire. In John Hughes, editor, *Functional Programming Languages and Computer Architecture, FPCA 1991*, volume 523 of LNCS, pages 124–144. Springer-Verlag, 1991.
26. A. Ohori and K. Tajima. A polymorphic calculus for views and object sharing. In *Proceedings of the 13th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 255–266, 1994.
27. Chris Okasaki. Views for Standard ML. In *SIGPLAN Workshop on ML*, pages 14–23, 1998.
28. John C. Reynolds. Definitional interpreters for higher-order programming languages. In *ACM '72: Proceedings of the ACM annual conference*, pages 717–740, New York, NY, USA, 1972. ACM Press.
29. C. Souza dos Santos, S. Abiteboul, and C. Delobel. Virtual schemas and bases. In *Proceedings of the International Conference on Extensive Data Base Technology (EDBT'94)*, Cambridge, UK, pages 81–94. Springer-Verlag, 1994.
30. Phil Wadler. Views: a way for pattern matching to cohabit with data abstraction. In *Conference Record of POPL '87: The 14th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 1987.