


The Fun of Programming

Oege de Moor,
Jeremy Gibbons
and Geraint Jones



Contents

Fun with phantom types
R. Hinze

5

Fun with phantom types

R. Hinze

Haskell is renowned for its many extensions to the Hindley-Milner type system (type classes, polymorphic recursion, rank- n types, existential types, functional dependencies—just to name a few). In this chapter we look at yet another extension. I can hear you groaning but this is quite a mild extension and one that fits nicely within the Hindley-Milner framework. Of course, whenever you add a new feature to a language, you should throw out an existing one (especially if the language at hand is named after a logician). Now, for this chapter we abandon type classes—judge for yourself how well we get along without Haskell’s most beloved feature.

1 Introducing phantom types

Suppose you want to embed a programming language, say, a simple expression language in Haskell. Since you are a firm believer of *static typing*, you would like your embedded language to be statically typed, as well. This requirement rules out a simple *Term* data type as this choice would allow us to freely mix terms of different types. The next idea is to parameterize the *Term* type so that *Term t* comprises only terms of type *t*. The different compartments of *Term* are then inhabited by declaring *constructors* of the appropriate types (we confine ourselves to a few basic operations):

$$\begin{aligned} \text{Zero} & \quad :: \text{Term Int} \\ \text{Succ, Pred} & \quad :: \text{Term Int} \rightarrow \text{Term Int} \\ \text{IsZero} & \quad :: \text{Term Int} \rightarrow \text{Term Bool} \\ \text{If} & \quad :: \forall a. \text{Term Bool} \rightarrow \text{Term a} \rightarrow \text{Term a} \rightarrow \text{Term a}. \end{aligned}$$

The types are essentially those of the corresponding Haskell functions except that every argument and every result type has *Term* wrapped around it. For instance, the Haskell function $\text{succ} :: \text{Int} \rightarrow \text{Int}$ corresponds to the constructor $\text{Succ} :: \text{Term Int} \rightarrow \text{Term Int}$.

This term representation meets the typing requirement: we can apply *Succ* only to an arithmetic expression; applying *Succ* to a Boolean expression results

in a type error. Unfortunately, the above signature cannot be translated into a **data** declaration (Haskell’s linguistic construct for introducing constructors). The reason is simply that all constructors of a data type must share the same result type, namely, the declared type on the left-hand side. Thus, we can assign *Zero* the type *Term t* but not *Term Int*. Of course, using the first type would defeat the purpose of the whole exercise. The only constructor that fits into the scheme is *If*, which has the desired general result type.

If only we had the means to constrain the type argument of *Term* to a certain type. Now, this is exactly what the aforementioned ‘mild’ extension allows us to do. Given this extension we declare the *Term* data type as follows:

```

data Term t = Zero                               with t = Int
             | Succ (Term Int)                   with t = Int
             | Pred (Term Int)                   with t = Int
             | IsZero (Term Int)                 with t = Bool
             | If (Term Bool) (Term a) (Term a) with t = a.

```

The **with** clause that it attached to each constructor records its type constraints. For instance, *Zero* has *Type t* with the additional constraint $t = \text{Int}$. Note that the **with** clause of the *If* constructor is not strictly necessary. We could have simply replaced *a* by *t*. Its main purpose is to illustrate that the type equation may contain type variables that do not appear on the left-hand side of the declaration. These variables can be seen as being *existentially quantified*.

Let us move on to defining an interpreter for the expression language. The interpreter takes an expression of type *Term t* to a value of type *t*. The definition proceeds by straightforward structural recursion.

```

eval          :: ∀t. Term t → t
eval (Zero)   = 0
eval (Succ e) = eval e + 1
eval (Pred e) = eval e - 1
eval (IsZero e) = eval e == 0
eval (If e1 e2 e3) = if eval e1 then eval e2 else eval e3

```

Even though *eval* is assigned the type $\forall t. \text{Term } t \rightarrow t$, each equation—with the notable exception of the last one—has a more specific type as dictated by the type constraints. As an example, the first equation has type $\text{Term Int} \rightarrow \text{Int}$ as *Zero* constrains *t* to *Int*.

The interpreter is quite noticeable in that it is *tag free*. If it receives a Boolean expression, then it returns a Boolean. By contrast, a more conventional interpreter of type $\text{Term} \rightarrow \text{Val}$ has to inject the Boolean into the *Val* data type. Conversely, when evaluating a conditional it has to untag the evaluated condition and furthermore it has to check whether the value is actually a Boolean. To make a long story short, we are experiencing the benefits of static typing. Here is a short interactive session that shows the interpreter in action (**:type** displays the type of an

expression).

```

Main> let one = Succ Zero
Main> :type one
Term Int
Main> eval one
1
Main> eval (IsZero one)
False
Main> IsZero (IsZero one)
Type error: couldn't match 'Int' against 'Bool'
Main> eval (If (IsZero one) Zero one)
1
Main> let true = IsZero Zero
Main> let false = IsZero one
Main> eval (If true true false)
True

```

Thinking of it, the type *Term t* is quite unusual. Though *Term* is parameterized, it is not a container type: an element of *Term Int*, for instance, is an expression that evaluates to an integer; it is not a data structure that contains integers. This means, in particular, that we cannot define a mapping function $(a \rightarrow b) \rightarrow (Term\ a \rightarrow Term\ b)$ as for many other data types. How could we possibly turn expressions of type *Term a* into expression of type *Term b*? The type *Term b* might not even be inhabited: there are, for instance, no terms of type *Term String*. Clearly, types of this characteristic deserve a special name. Since the type argument of *Term* is not related to any component, we call *Term* a *phantom type*. The purpose of this chapter is to demonstrate the usefulness and the beauty of phantom types.

Exercise 1 (Language design) Whenever we define a function that involves a phantom type, we will provide an explicit type signature. Can you imagine why? *Hint:* is it possible to infer the types of functions that involve phantom types? \square

Exercise 2 (Language design) In Haskell, constructors are introduced via **data** declarations. An alternative is to abandon the **data** construct and to introduce constructors simply by listing their signatures. Discuss the pros and cons of the two alternatives. \square

2 Generic functions

Suppose you are developing an application where the need arises to compress data to strings of bits. As it happens, you have data of many different types and you want to program a compression function that works for all of these types. This

sounds like a typical case for Haskell's type classes. Alas, I promised to do without type classes. Fortunately, phantom types offer an intriguing alternative.

The basic idea is to define a type whose elements represent types. For concreteness, assume that we need compressing functions for types built from *Int* and *Char* using the list and the pair type constructor.

```

data Type t = RInt           with t = Int
             | RChar          with t = Char
             | RList (Type a) with t = [a]
             | RPair (Type a) (Type b) with t = (a, b)

rString     :: Type String
rString     = RList RChar

```

An element *rt* of type *Type t* is a representation of *t*. For instance, *Int* is represented by *RInt*, the type *(String, Int)* is represented by *RPair rString RInt*.

Now, the compression function takes a type representation as a first argument and the to-be-compressed value as the second argument. The following interactive session illustrates the use of *compress* (note that integers require 32 bits and characters 7 bits).

```

Main> :type compress RInt
Int → [Bit]
Main> compress RInt 60
<00111100000000000000000000000000>
Main> :type compress rString
[Char] → [Bit]
Main> compress rString "Richard"
<10100101110010111100011100010111100001110100111100100110>

```

The definition of *compress* itself is straightforward: it pattern matches on the type representation and then takes the appropriate action.

```

data Bit = 0 | 1
compress :: ∀t. Type t → t → [Bit]
compress (RInt) i = compressInt i
compress (RChar) c = compressChar c
compress (RList ra) [] = 0 : []
compress (RList ra) (a : as) = 1 : compress ra a ++ compress (RList ra) as
compress (RPair ra rb) (a, b) = compress ra a ++ compress rb b

```

We assume that *compressInt* :: *Int* → [*Bit*] and *compressChar* :: *Char* → [*Bit*] are given. Consider the definition of *compress (RList ra)*. Since the list data type has two constructors, we emit one bit to distinguish between the two cases. In the case of a non-empty list, we recursively encode the head and the tail. As an aside, if we extend *compress* to data types with more than two constructors, we must ensure that the codes for the constructors have the *unique prefix property*, that is, no code

is the prefix of another code. However, we *can* use the same code for constructors of different types as compression (as well as decompression) is driven by type.

We can view *Type* as representing a *family of types* and *compress* as implementing a *family of functions*. Through the first argument of *compress* we specify which member of the family we wish to apply. Functions that work for a family of types are commonly called *generic functions*. Using a phantom type of type representations, generic functions are easy to define. Typical examples of generic functions include equality and comparison functions, pretty printers and parsers. Actually, pretty printing is quite a nice example, so let us consider this next.

In Haskell, the *Show* class takes care of converting values into string representations. We will define a variant of its *show* method building upon the pretty-printing combinators of Chapter ???. The implementation of the *Show* class is complicated by the desire to print lists of characters different from lists of other types: a list of characters is shown using string syntax whereas any other list is shown as a comma-separated sequence of elements enclosed in square brackets. Using type representations we can easily single out this special case by supplying an additional equation.

```

pretty                :: ∀t. Type t → t → Doc
pretty (RInt) i       = prettyInt i
pretty (RChar) c      = prettyChar c
pretty (RList RChar) s = prettyString s
pretty (RList ra) []  = text "[]"
pretty (RList ra) (a : as) = block 1 (text "[" ⟨ pretty ra a ⟨ prettyL as)
  where prettyL []    = text "]"
        prettyL (a : as) = text "," ⟨ line ⟨ pretty ra a ⟨ prettyL as
pretty (RPair ra rb) (a, b) = block 1 (text "(" ⟨ pretty ra a ⟨ text ","
                                     ⟨ line ⟨ pretty rb b ⟨ text ")")

block                :: Int → Doc → Doc
block i d            = group (nest i d)

```

Here, *prettyInt* :: *Int* → *Doc*, *prettyChar* :: *Int* → *Doc*, and *prettyString* :: *String* → *Doc* are predefined functions that pretty print integers, characters, and strings, respectively.

Exercise 3 Implement generic equality *eq* :: ∀*t*. *Type t* → *t* → *t* → *Bool* and a generic comparison function *compare* :: ∀*t*. *Type t* → *t* → *t* → *Ordering*. □

Exercise 4 Families of type-indexed functions can be implemented either using type classes or using type representations. Discuss differences and commonalities of the two approaches. □

Exercise 5 Implement a function *uncompress* :: ∀*t*. *Type t* → [*Bit*] → *t* that uncompresses a bit string. *Hint*: use tupling (see IFPH, Section 7.3). Implement a generic parser *parse* that converts a string to a value. The function *parse* should at least be able to read strings that were generated by *pretty*. □

3 Dynamic values

Even a programming language such as Haskell cannot guarantee the absence of run-time errors using static checks only. For instance, when we communicate with the environment, we have to check dynamically whether the imported values have the expected types. In this section we show how to embed dynamic checking in a statically typed language.

To this end we introduce a *universal data type*, the type *Dynamic*, which encompasses all static values (whose types are representable). To inject a static value into the universal type we bundle the value with a representation of its type.

```
data Dynamic = Dyn (Type t) t
```

It is important to note that the type variable *t* is existentially quantified: a dynamic value is a pair consisting of a type representation of *Type t* and a value of type *t* for *some* type *t*. The type *Dynamic* looks attractive but on a second thought we note a small deficiency: we can form a list of dynamic values but we cannot turn this list into a dynamic value itself, simply because the type *Dynamic* is not representable. This is, however, easily remedied: we simply add *Dynamic* to *Type t*.

```
data Type t = ...
           | RDyn with t = Dynamic
```

Note that *Type* and *Dynamic* are now defined by mutual recursion.

Dynamic values and generic functions go well together. In a sense, they are complementary concepts. It is not too difficult, for instance, to extend the generic functions of the previous section so that they also work for dynamic values (see Exercise 7 and 8): a dynamic value contains a type representation, which a generic function requires as a first argument. The following interactive session illustrates the use of dynamics and generics (note that the identifier *it* always refers to the previously evaluated expression).

```
Main> let ds = [Dyn RInt 60, Dyn rString "Bird"]
Main> :type ds
[Dynamic]
Main> Dyn (RList RDyn) ds
Dyn (RList RDyn) [Dyn RInt 60, Dyn (RList RChar) "Bird"]
Main> compress RDyn it
<010100100000111100000000000000000000000000001010001101000
111001011101001111001001100>
Main> uncompress RDyn it
Dyn (RList RDyn) [Dyn RInt 60, Dyn (RList RChar) "Bird"]
```

By pairing a value with its type representation we turn a static into a dynamic value. The other way round involves a dynamic check. This operation, usually termed *cast*, takes a dynamic value and a type representation and checks whether

the type representation of the dynamic value and the supplied one are identical. The equality check is defined

```

tequal                                ::  ∀t u . Type t → Type u → Maybe (t → u)
tequal (RInt) (RInt)                  =  return id
tequal (RChar) (RChar)                =  return id
tequal (RList ra1) (RList ra2)      =  liftM list (tequal ra1 ra2)
tequal (RPair ra1 rb1) (RPair ra2 rb2)
    = liftM2 pair (tequal ra1 ra2) (tequal rb1 rb2)
tequal _ _                             =  fail "cannot tequal".

```

If the test succeeds, *tequal* returns a function that allows us to transform the dynamic value into a static value of the specified type. Of course, as the types are equal, this function is necessarily the identity! Turning to the implementation of *tequal*, the functions *list* and *pair* are the mapping functions of the list and the pair type constructor. Since the equality check may fail, we must lift the mapping functions into the *Maybe* monad (using *return*, *liftM*, and *liftM2*). The cast operation simply calls *tequal* and then applies the conversion function to the dynamic value.

```

cast                                   ::  ∀t . Dynamic → Type t → Maybe t
cast (Dyn ra a) rt                     =  fmap (λf → f a) (tequal ra rt)

```

Here is a short interactive session that illustrates its use.

```

Main> let d = Dyn RInt 60
Main> cast d RInt
Just 60
Main> cast d RChar
Nothing

```

Exercise 6 Define functions that compress and uncompress type representations. *Hint*: define an auxiliary data type

```

data Rep = Rep (Type t)

```

and then implement functions *compressRep* :: *Rep* → [*Bit*] and *uncompressRep* :: [*Bit*] → *Rep* that compress and uncompress elements of type *Rep*. Why do we need the auxiliary data type? □

Exercise 7 Use the results of the previous exercise to implement functions that compress and uncompress dynamic values. To compress a dynamic value, first compress the type representation and then compress the static value. Conversely, to uncompress a dynamic value first uncompress the type representation and then use the type representation to read in a static value of this type. Finally, extend the generic functions *compress* and *uncompress* to take care of dynamic values. □

Exercise 8 Implement functions that pretty print and parse dynamic values and extend the definitions of *pretty* and *parse* accordingly. □

Exercise 9 Extend the type of type representations *Type* and the dynamic type equality check *tequal* to include functional types of the form $a \rightarrow b$. \square

4 Generic traversals and queries

Let us develop the theme of Section 2 a bit further. Suppose you have to write a function that traverses a complex data structure representing a university's organisational structure, and that increases the age of a given person. The interesting part of this function, namely the increase of age, is probably dominated by the *boilerplate code* that recurses over the data structure. The boilerplate code is not only tiresome to program, it is also highly vulnerable to changes in the underlying data structure. Fortunately, generic programming saves the day as it allows us to write the traversal code once and use it many times. Before we look at an example let us first introduce a data type of persons.

```

type Name    = String
type Age     = Int
data Person = Person Name Age

```

To be able to apply generic programming techniques, we add *Person* to the type of representable types.

```

data Type t = ...
           | RPerson with t = Person

```

Now, the aforementioned function that increases the age can be programmed as follows (this is only the interesting part without the boilerplate code):

```

tick      :: Name → Traversal
tick s (RPerson) (Person n a)
  | s == n = Person n (a + 1)
tick s rt t = t

```

The function *tick s* is a so-called *traversal*, which can be used to modify data of *any* type (the type *Traversal* will be defined shortly). In our case, *tick s* changes values of type *Person* whose name equals *s*; integers, characters, lists etc are left unchanged.

The following interactive session shows the traversal *tick* in action. The combinator *everywhere*, defined below, implements the generic part of the traversal: it

applies its argument ‘everywhere’ in a given value.

```

Main> let ps = [Person "Norma" 50, Person "Richard" 59]
Main> everywhere (tick "Richard") (RList RPerson) ps
[Person "Norma" 50, Person "Richard" 60]
Main> total age (RList RPerson) it
110
Main> total sizeof rString "Richard Bird"
60

```

The second and the third example illustrate *generic queries*: *age* computes the age of a person, *sizeof* yields the size of an object (the number of occupied memory cells), *total* applies an integer query to every component of a value and sums up the results.

Turning to the implementation the type of generic traversals is given by:

type *Traversal* = $\forall t. \text{Type } t \rightarrow t \rightarrow t.$

A generic traversal takes a type representation and transforms a value of the specified type. The universal quantifier makes explicit that the function works for *all* representable types. The simplest traversal is *copy*, which does nothing.

```

copy    :: Traversal
copy rt = id

```

Traversals can be composed using the operator ‘ \circ ’, which has *copy* as its identity.

```

( $\circ$ )    :: Traversal  $\rightarrow$  Traversal  $\rightarrow$  Traversal
(f  $\circ$  g) rt = f rt  $\cdot$  g rt

```

The *everywhere* combinator is implemented in two steps. We first define a function that applies a traversal *f* to the *immediate* components of a value: $C\ t_1 \dots t_n$ is mapped to $C\ (f\ rt_1\ t_1) \dots (f\ rt_n\ t_n)$ where rt_i is the representation of t_i ’s type.

```

imap    :: Traversal  $\rightarrow$  Traversal
imap f (RInt) i      = i
imap f (RChar) c     = c
imap f (RList ra) [] = []
imap f (RList ra) (a : as) = f ra a : f (RList ra) as
imap f (RPair ra rb) (a, b) = (f ra a, f rb b)
imap f (RPerson) (Person n a) = Person (f rString n) (f RInt a)

```

The function *imap* can be seen as a ‘traversal transformer’. Note that *imap* has a so-called *rank-2 type*: it takes polymorphic functions to polymorphic functions. The combinator *everywhere* enjoys the same type.

```

everywhere, everywhere' :: Traversal  $\rightarrow$  Traversal
everywhere f             = f  $\circ$  imap (everywhere f)
everywhere' f            = imap (everywhere' f)  $\circ$  f

```

Actually, there are two flavours of the combinator: *everywhere* f applies f *after* the recursive calls (it proceeds bottom-up), whereas *everywhere'* applies f *before* (it proceeds top-down). And yes, *everywhere* and *everywhere'* have the structure of generic folds and unfolds—only the types are different (Chapter ?? treats folds and unfolds in detail).

Generic queries have a similar type except that they yield a value of some fixed type.

```
type Query x =  $\forall t. \text{Type } t \rightarrow t \rightarrow x$ 
```

In the rest of this section we confine ourselves to queries of type *Query Int*. Exercise 11 deals with the general case. The definition of the combinator *total* follows the model of *everywhere*. We first define a non-recursive, auxiliary function that sums up the immediate components of a value and then tie the recursive knot.

```
isum :: Query Int  $\rightarrow$  Query Int
isum f (RInt) a = 0
isum f (RChar) a = 0
isum f (RList ra) [] = 0
isum f (RList ra) (a : as) = f ra a + f (RList ra) as
isum f (RPair ra rb) (a, b) = f ra a + f rb b
isum f (RPerson) (Person s i) = f rString s + f RInt i
total :: Query Int  $\rightarrow$  Query Int
total f rt t = f rt t + isum (total f) rt t
```

It remains to define the ad-hoc queries *age* and *sizeof*.

```
age ::  $\forall t. \text{Type } t \rightarrow t \rightarrow \text{Age}$ 
age (RPerson) (Person n a) = a
age _ _ = 0
sizeof :: Query Int
sizeof (RInt) _ = 2
sizeof (RChar) _ = 2
sizeof (RList ra) [] = 0
sizeof (RList ra) (_ : _) = 3
sizeof (RPair ra rb) _ = 3
sizeof (RPerson) _ = 3
```

Using *total sizeof* we can compute the memory consumption of a data structure. Actually, the result is a conservative estimate since any sharing of subtrees is ignored. Note that the empty list consumes no memory since it need be represented only once (it can be globally shared).

Exercise 10 Prove the following properties of *imap* (which justify its name).

```
imap copy = copy
imap (f  $\circ$  g) = imap f  $\circ$  imap g
```

Does *everywhere* satisfy similar properties? \square

Exercise 11 Generalize *isum* and *total* to functions

$$icrush, everything \quad :: \quad \forall x. (x \rightarrow x \rightarrow x) \rightarrow x \rightarrow Query\ x \rightarrow Query\ x$$

such that *icrush* (+) 0 = *isum* and *everything* (+) 0 = *total*. \square

5 Normalization by evaluation

Let us move on to one of the miracles of theoretical computer science. In Haskell, one cannot show values of functional types. For reasons of computability, there is no systematic way of showing functions and any ad-hoc approach would destroy *referential transparency* (except if *show* were a constant function). For instance, if *show* yielded the text of a function definition, we could distinguish, say, quick sort from merge sort. Substituting one for the other could then possibly change the meaning of a program.

However, what we *can* do is to print the *normal form* of a function. This does not work for Haskell in its full glory, but only for a very tiny subset, the simply typed lambda calculus. Nonetheless, the ability to do that is rather surprising. Let us consider an example first. Suppose you have defined the following Haskell functions (the famous SKI combinators)

$$\begin{aligned} s &= \lambda x\ y\ z \rightarrow (x\ z)\ (y\ z) \\ k &= \lambda x\ y \rightarrow x \\ i &= \lambda x \rightarrow x \end{aligned}$$

and you want to normalize combinator expressions. The function *reify*, defined below, allows you to do that: it takes a type representation (where *b* represents the base type and ‘: \rightarrow ’ functional types) and yields the normal form of a Haskell value of this type, where the normal form is given as an element of a suitable expression data type.

```
Main> reify (b :→ b) (s k k)
Fun (λa → a)
Main> reify (b :→ (b :→ b)) (s (k k) i)
Fun (λa → Fun (λb → a))
Main> let e = (s ((s (k s)) ((s (k k)) i))) ((s ((s (k s)) ((s (k k)) i))) (k i))
Main> :type e
∀t. (t → t) → t → t
Main> reify ((b :→ b) :→ (b :→ b)) e
Fun (λa → Fun (λb → App a (App a b)))
```

The last test case is probably the most interesting one as the expression *e* is quite involved. We first use Haskell’s type inferencer to determine its type, then we call *reify* passing it a representation of the inferred type and *e* itself. And voilà:

the computed result shows that e normalizes to a function that applies its first argument twice to its second.

Now, since we want to represent simply typed lambda terms, we change the type of type representations to

```

infixr :→
data Type t = RBase           with t = Base
             | Type a :→ Type b with t = a → b
b           :: Type Base
b           = RBase.

```

Here, *Base* is the base type of the simply typed lambda calculus. We won't reveal its definition until later. To represent lambda terms we use *higher-order abstract syntax*. For instance, the lambda term $\lambda f.\lambda x.f (f x)$ is represented by the Haskell term `Fun (\f → Fun (\x → App f (App f x)))`, that is, abstractions are represented by Haskell functions.

```

data Term t = App (Term (a → b)) (Term a) with t = b
             | Fun (Term a → Term b)      with t = a → b

```

Note that since we use higher-order abstract syntax there is no need to represent variables.

The function *reify* takes a Haskell value of type t to an expression of type $Term\ t$. It is defined by induction over the structure of types, that is, it is driven by the type representation of t . Let us consider functional types first. In this case, *reify* has to turn a value of type $a \rightarrow b$ into an expression of type $Term\ (a \rightarrow b)$. The constructor *Fun* constructs terms of this type, so we are left with converting an $a \rightarrow b$ value to a $Term\ a \rightarrow Term\ b$ value (unfortunately, *Term* does not give rise to a mapping function). Suppose that there is a transformation of type $Term\ a \rightarrow a$ available. Then we can reflect a $Term\ a$ to an a , apply the given function, and finally reify the resulting b to a $Term\ b$. In other words, to implement *reify* we need its converse, as well. Turning to the base case, this means that we require functions of type $Base \rightarrow Term\ Base$ and $Term\ Base \rightarrow Base$. Fortunately, we are still free in the choice of the base type. An intriguing option is to set *Base* to the *fixed point* of *Term*.

```

newtype Base = In{ out :: Term Base }

```

Then the isomorphisms *out* and *In* constitute the required functions. Given these prerequisites we can finally define *reify* and its inverse *reflect*.

```

reify           :: ∀t. Type t → (t → Term t)
reify (RBase) v = out v
reify (ra :→ rb) v = Fun (\x → reify rb (v (reflect ra x)))
reflect         :: ∀t. Type t → (Term t → t)
reflect (RBase) e = In e
reflect (ra :→ rb) e = λx → reflect rb (App e (reify ra x))

```

Exercise 12 Implement a *show* function for *Term t*. *Hint*: augment the expression type *Term t* by an additional constructor *Var* of type *String* \rightarrow *Term t*. \square

6 Functional unparsing

Can we program C's *printf* function in a statically typed language such as Haskell? Yes, we can, provided we use a tailor-made type of format directives (rather than a string). Here is an interactive session that illustrates the puzzle (we renamed *printf* to *format*).

```

Main> :type format (Lit "Richard")
String
Main> format (Lit "Richard")
"Richard"
Main> :type format Int
Int  $\rightarrow$  String
Main> format Int 60
"60"
Main> :type format (String :^: Lit " is " :^: Int)
String  $\rightarrow$  Int  $\rightarrow$  String
Main> format (String :^: Lit " is " :^: Int) "Richard" 60
"Richard is 60"

```

The format directive *Lit s* means emit *s* literally. The directives *Int* and *String* instruct *format* to take an additional argument of the types *Int* and *String* respectively, which is then shown. The operator $^:$ is used to concatenate two directives.

The type of *format* depends on its first argument, the format directive. This is something we have already seen a number of times: the type of *compress*, for instance, depends on its first argument, the type representation. Of course, the dependence here is slightly more involved. Yet, this smells like a case for phantom types.

The format directive can be seen as a binary tree of type representations: *Lit s*, *Int*, *String* form the leaves, $^:$ constructs the inner nodes. The type of *format* is essentially obtained by linearizing the binary tree mapping, for instance, *String :^: Lit " is " :^: Int* to *String* \rightarrow *Int* \rightarrow *String*.

Before tackling the puzzle proper it is useful to reconsider flattening binary trees (see IFPH, Section 7.3.1). To avoid the repeated use of the expensive $^:$ operation, one typically defines an auxiliary function that makes use of an accu-

mulating parameter.

```

data Btree a      = Leaf a | Fork (Btree a) (Btree a)
flatten           :: ∀a . Btree a → [a]
flatten t        = flatcat t []
flatcat          :: ∀a . Btree a → [a] → [a]
flatcat (Leaf a) as = a : as
flatcat (Fork tl tr) as = flatcat tl (flatcat tr as)

```

The auxiliary function *flatcat* linearizes the given tree and additionally appends the accumulator to the result.

Now, this technique can be mirrored on the type level using a two-argument phantom type.

```

data Dir x y = Lit String           with y = x
              | Int                  with y = Int → x
              | String                with y = String → x
              | Dir y1 y2 :^: Dir x y1 with y = y2

```

The first argument corresponds to the accumulating parameter and the second to the overall result. The binary tree is implicitly given by the value constructor. Forming a functional type in a **with** clause corresponds to consing an element to a list. The major difference to the definition of *flatcat* is that *Dir* employs a relational style! In fact, with a little bit of imagination you can read the **data** declaration as a relational program (see also Chapter ??).

Now, using *Dir* we can assign *format* the type $\forall y . \text{Dir } \text{String } y \rightarrow y$: linearizing a directive *d* and plugging in *String* for the final result type, we obtain *y* as the type of *format d*. Unfortunately, we cannot define *format* directly since its type is not general enough to push the recursion through (see Exercise 13). We have to introduce an auxiliary function that takes a *continuation* and an accumulating string argument.

```

format'         :: ∀x y . Dir x y → (String → x) → (String → y)
format' (Lit s) = λcont out → cont (out ++ s)
format' (Int)   = λcont out → λi → cont (out ++ show i)
format' (String) = λcont out → λs → cont (out ++ s)
format' (d1 :^: d2) = λcont out → format' d1 (format' d2 cont) out
format         :: ∀y . Dir String y → y
format d       = format' d id ""

```

Note that $\text{format}' (d_1 :^: d_2)$ can be simplified to $\text{format}' d_1 \cdot \text{format}' d_2$, where ‘ \cdot ’ is ordinary function composition. This is not a coincidence. In fact, the type $(\text{String} \rightarrow x) \rightarrow (\text{String} \rightarrow y) = \text{MapTrans } \text{String } x y$ constitutes an arrow (see Chapter ??).

Exercise 13 Try to implement $\text{format} :: \forall y . \text{Dir } \text{String } y \rightarrow y$ directly. Where does the attempt fail? \square

Exercise 14 The function *format'* exhibits quadratic run-time behaviour. Remedy this defect. \square

Exercise 15 Instead of using a tree-like structure for format directives, we can alternatively employ a list-like structure.

```

data Dir x = End           with x = String
           | Lit String (Dir y) with x = y
           | Int (Dir y)      with x = Int  $\rightarrow$  y
           | String (Dir y)  with x = String  $\rightarrow$  y

```

Implement *format* $:: \forall x. \text{Dir } x \rightarrow x$ using this type. *Hint*: define an auxiliary function of type *format'* $:: \forall x. \text{Dir } x \rightarrow \text{String} \rightarrow x$. \square

7 A type equality type

We have seen in the previous sections that **with** clauses add considerably to the expressiveness of Haskell. Rather surprisingly, **with** clauses need not be a primitive concept, they can be simulated using polymorphic types. The resulting programs are more verbose—this is why we have used **with** clauses in the first place—but they can be readily evaluated using a Haskell 98 implementation that additionally supports existential types.

The principle idea is to represent type equations by a *type equality type*: the **data** declaration

```

data T t = ... | C t1 ... tn with t = u | ...

```

becomes

```

data T t = ... | C (u :=: t) t1 ... tn | ... ,

```

where ‘:=:’ is a binary type constructor, the type equality type. This type has the intriguing property that it is non-empty if and only if its argument types are equal.¹ Even more intriguing, its definition goes back to Leibniz. According to Leibniz, two terms are equal if one may be substituted for the other. Adapting this principle to types, we define

```

newtype a :=: b = Proof { apply ::  $\forall f. f a \rightarrow f b$  }.

```

Note that the universally quantified type variable *f* ranges over *type constructors* of kind $* \rightarrow *$. Thus, an element of *a :=: b* is a function that converts an element of type *f a* into an element of *f b* for *any* type constructor *f*. This function can be seen as constituting a *proof* of the type equality *a = b*. The identity function, for instance, serves as the proof of reflexivity.

```

refl  ::  $\forall a. a :=: a$ 
refl  = Proof id

```

¹We ignore the fact, that in Haskell every type contains the bottom element.

Since we have extended the value constructor C by an additional argument, we also have to adapt programs that use C . Every occurrence of the constructor C on the right-hand side of an equation is replaced by $C\ refl$. It is not hard to convince oneself that $C\ refl$ has indeed the right type. Occurrences on the left-hand side are treated as follows: the equation

$$f (C\ p_1 \dots p_n) = e$$

becomes

$$f (C\ p\ p_1 \dots p_n) = \mathit{apply}\ p\ e.$$

Assume that f has type $\forall t. T\ t \rightarrow F\ t$ where $F\ t$ is some type expression possibly involving t . The **with** clause associated with C dictates that e has type $F\ u$. The right-hand side of the transformed program, however, must have the type $F\ t$. The proof p of type $u :=: t$ allows us to turn e into a value of the desired type. Note that the universally quantified type variable f of the type equality type is instantiated to F .

In some cases it is necessary to guide the Haskell type inferencer so that it indeed instantiates f to F . The problem is that Haskell employs a *kinded first-order unification*. For instance, the types $Int \rightarrow [Bit]$ and $f\ Int$ are not unifiable, since the type checker reduces the type equation $((\rightarrow)\ Int)\ [Bit] = f\ Int$ to $(\rightarrow)\ Int = f$ and $[Bit] = Int$. The standard trick to circumvent this problem is to introduce a new type F' that is isomorphic to F .

$$\mathbf{newtype}\ F'\ a = In\{ out :: F\ a \}$$

The equation then becomes

$$f (C\ p\ p_1 \dots p_n) = (out \cdot \mathit{apply}\ p \cdot In)\ e.$$

Turning back to the type equality type it is interesting to note that it has all the properties of an *congruence relation*. We have already seen that it is reflexive. It is furthermore symmetric, transitive, and congruent. Here are programs that

implement the respective proofs.

```

newtype Flip f a b   = Flip { unFlip :: f b a }
symm                :: ∀ a b . (a :=: b) → (b :=: a)
symm p              = unFlip (apply p (Flip refl))
trans                :: ∀ a b c . (a :=: b) → (b :=: c) → (a :=: c)
trans p q           = Proof (apply q · apply p)
newtype List f a    = List { unList :: f [a] }
list                :: ∀ a b . (a :=: b) → ([a] :=: [b])
list p              = Proof (unList · apply p · List)
newtype Pair1 f b a = Pair1 { unPair1 :: f (a, b) }
newtype Pair2 f a b = Pair2 { unPair2 :: f (a, b) }
pair                :: ∀ a b c d . (a :=: c) → (b :=: d) → ((a, b) :=: (c, d))
pair p1 p2         = Proof (unPair2 · apply p2 · Pair2
                             · unPair1 · apply p1 · Pair1)

```

Again, we have to introduce auxiliary data types to direct Haskell's type inferencer. As an example, the proof of symmetry works as follows. We first specialize the given proof of $(a :=: b) = (\forall f . f a \rightarrow f b)$ setting f to $(:=: a)$. We obtain a function of type $(a :=: a) \rightarrow (b :=: a)$, which is then passed *refl* to yield the desired proof of $b :=: a$.

Before we conclude, let us briefly revise the type equality check *tequal* of Section 3. Recall that *tequal* returns a conversion function of type $t \rightarrow u$ that allows us to transform dynamic values into static values. A far more flexible approach is to replace $t \rightarrow u$ by $t :=: u$, so that we can transform a t to a u in *any* context.

```
tequal :: ∀ t u . Type t → Type u → Maybe (t :=: u)
```

The changes to the definition of *tequal* are simple: we have to replace *id* by *refl*, and the mapping functions *pair* and *list* by the congruence proofs of the same name.

Exercise 16 Extend the above transformation to cover multiple type arguments and multiple type equations. \square

Exercise 17 Define conversion functions *from* $:: \forall a b . (a :=: b) \rightarrow (a \rightarrow b)$ and *to* $:: \forall a b . (a :=: b) \rightarrow (b \rightarrow a)$. Try to implement them from scratch. \square

Exercise 18 We have defined congruence proofs for the list and the pair type constructor. Generalize the construction to an arbitrary n -ary data type not necessarily being a functor. \square

8 Chapter notes

This chapter is based on a paper by Cheney and Hinze [2], which shows how to combine generics and dynamics in a type-safe manner. The term *phantom type*

was coined by Leijen and Meijer [8] to denote parameterized types that do not use their type argument.

There is an abundance of work on generic programming, see, for instance, [6, 5]. For a gentle introduction to the topic the interested reader is referred to [1].

Section 4 draws from a paper by Lämmel and Peyton Jones [7]. Sections 5 and 6 adopt two pearls by Danvy, Rhiger and Rose [4] and by Danvy [3], respectively. An alternative approach to unparsing is described by Hinze [?].

Acknowledgement

I would like to thank Andres Löh for his helpful and immediate feedback on a draft version of this chapter.

References

- [1] Roland Backhouse, Patrik Jansson, Johan Jeuring, and Lambert Meertens. Generic Programming — An Introduction —. In S. Doaitse Swierstra, Pedro R. Henriques, and Jose N. Oliveira, editors, *3rd International Summer School on Advanced Functional Programming, Braga, Portugal*, volume 1608 of *Lecture Notes in Computer Science*, pages 28–115. Springer-Verlag, Berlin, 1999.
- [2] James Cheney and Ralf Hinze. A lightweight implementation of generics and dynamics. In Manuel M.T. Chakravarty, editor, *Proceedings of the 2002 ACM SIGPLAN Haskell Workshop*, pages 90–104. ACM Press, October 2002.
- [3] Olivier Danvy. Functional unparsing. *J. Functional Programming*, 8(6):621–625, November 1998.
- [4] Olivier Danvy, Morten Rhiger, and Kristoffer H. Rose. Normalization by evaluation with typed abstract syntax. *J. Functional Programming*, 11(6):673–680, November 2001.
- [5] Ralf Hinze. A new approach to generic functional programming. In Thomas W. Reps, editor, *Proceedings of the 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'00), Boston, Massachusetts, January 19-21*, pages 119–132, January 2000.
- [6] Patrik Jansson and Johan Jeuring. PolyP—a polytypic programming language extension. In *Conference Record 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'97), Paris, France*, pages 470–482. ACM Press, January 1997.
- [7] Ralf Lämmel and Simon Peyton Jones. Scrap your boilerplate: a practical approach to generic programming. Available from <http://research.microsoft.com/~simonpj/papers/hmap/>, 2002.

- [8] Daan Leijen and Erik Meijer. Domain-specific embedded compilers. In *Proceedings of the 2nd Conference on Domain-Specific Languages*, pages 109–122, Berkeley, CA, October 1999. USENIX Association.