Views:
A way for pattern matching to cohabit with data abstraction

Philip Wadler

Programming Research Group, Oxford University, UK
and Programming Methodology Group, Chalmers University, Sweden

Abstract

Pattern matching and data abstraction are important concepts in designing programs, but they do not fit well together. Pattern matching depends on making public a free data type representation, while data abstraction depends on hiding the representation. This paper proposes the *views* mechanism as a means of reconciling this conflict. A view allows any type to be viewed as a free data type, thus combining the clarity of pattern matching with the efficiency of data abstraction.

1 Introduction

Induction and abstraction are fundamental tools of the mathematician's trade, and equally essential to the computer scientist. Pattern matching is a language feature that supports induction, and data abstraction is a feature that supports abstraction; but unfortunately these two features do not get on well together. This paper proposes the views mechanism as a means of resolving this problem.

As an example of the conflict between pattern matching and data abstraction, consider the definition of exponentiation. Mathematicians traditionally define it as follows:

$$\begin{aligned} x^0 &= 1 \\ x^{n+1} &= x(x^n) \end{aligned} \qquad (1)$$

This definition makes it easy to prove properties of exponentiation by means of induction. Functional programming languages encourage a similar style of definition. For example, in a language like Hope [BMS80] or Miranda [Tur85] we might declare a new type

$$peano ::= Zero \mid Succ\ peano$$

and then we can write a definition that is essentially equivalent to (1):

$$\begin{aligned} power\ x\ Zero &= 1 \\ power\ x\ (Succ\ n) &= x \times power\ x\ n \end{aligned} \qquad (2)$$

This style of definition has several advantages: each case is displayed clearly as a pattern on the left-hand side of an equation; the compiler can check that no cases have been accidentally omitted; and the definitions are well-suited for proofs by structural induction [Bur69] and for program transformation [BD77].

However, there is a problem with the above definition: it specifies a particular way of representing natural numbers, as the free data type *peano*. The representation of the number seven is the data structure

$$Succ\ (Succ\ (Succ\ (Succ\ (Succ\ (Succ\ (Succ\ Zero))))))$$

Compared with the representation of integers built-in to the computer hardware, this representation is astonishingly inefficient. If we had followed the fundamental principle of *data abstraction* (or *representation hiding*) then this problem would not arise, because we would be free to implement natural numbers in any convenient way, including the built-in integer data type.

In short, pattern matching supports clear definitions and induction, but it requires that the representing type be a free data type and be visible. Data abstraction supports efficiency, but it requires that the representing type be hidden. Thus, the programmer is often faced with an unenviable choice between clarity and efficiency.

The programming language Hope finesses this problem in the special case of the natural numbers. It provides a special mechanism that allows the built-in integer type to be viewed as if it were the type *peano*. Such a useful mechanism is clearly a candidate for generalisation, as the need to view one data type as if it were another is hardly limited to this special case. This paper proposes a language mechanism called *views* as a means of satisfying this need. A view specifies how any arbitrary data type (including an abstract data type) can be viewed as a free data type. It is even possible to specify several different views of the same type.

This paper discusses views in the context of functional languages; similar ideas may be useful in imperative languages. The notation used in this paper is styled after Miranda. The essential requirement is that the host language permits the declaration of free data types. Views are useful regardless of whether eager, applicative order evaluation (as in Hope) or lazy, normal order evaluation (as in Miranda) is used; the examples in this paper work with either evaluation order.

Views as described here should not be confused with views in OBJ2 [FGJM85]. Views in OBJ2 specify homomorphisms between modules; views as described here specify isomorphisms between data types. (Joseph Goguen has suggested that the views in this paper be called "bi-views".)

The problem of extending pattern-matching to apply to non-free types is also addressed by Miranda, through the mechanism of "lawful types" [Tho86]. This mechanism is rather more limited than views. A lawful type is simply a subset of a free type, whereas views allow one to specify a correspondence between a free type and *any* desired type.

The remainder of this paper is organised as follows. Section 2 introduces views by showing how to define a view of integers. Section 3 briefly describes an alternative view of integers. Section 4 gives a simple example of views and abstract data types, in the context of two representations of complex numbers. Sections 5 through 8 demonstrate further applications of views to lists and trees. Section 9 shows two unusual uses of views. Section 10 and 11 describe how views support equational reasoning and induction. Section 12 outlines an efficient implementation method. Section 13 concludes.

2 Viewing an integer as zero or a successor

This section introduces the views mechanism by defining a view of the built-in integer type, *int*, that is analogous to the free data type *peano* discussed in the introduction.

Here is the definition of the view:

$$
\begin{aligned}
\textbf{view } int ::= &\ Zero \mid Succ\ int \\
\textbf{in } n \quad &= \quad Zero, \qquad\qquad \text{if } n = 0 \\
&= \quad Succ\,(n - 1), \quad \text{if } n > 0 \\
\textbf{out } Zero \quad &= \quad 0 \\
\textbf{out }(Succ\ n) \quad &= \quad n + 1
\end{aligned}
$$

The first line introduces two new names, *Zero* and *Succ*, which may appear in terms (on the right-hand side of equations) and in patterns (on the left-hand side of equations).

The **in** and **out** clauses are similar to function definitions. The **in** clause defines a function to apply to an *int* to get a *Zero* or *Succ*; it is used when *Zero* or *Succ* appear in a pattern on the left-hand side of an equation. The **out** clause defines a function to apply to a *Zero* or *Succ* to get an *Int*; it is used when *Zero* or *Succ* appear on the right-hand side of an equation.

A view is well-defined only when the functions defined by the **in** and **out** clause are inverses of each other. Together, they specify an isomorphism between (a subset of) the *viewed* type and (a subset of) the *viewing* type. In this case, the isomorphism is between the natural numbers (a subset of the viewed type, *int*) and values constructed with *Succ* and *Zero* (the viewing type).

Given the above view declaration, one may write definitions such as (2) in Section 1, or the following definition of Fibonacci numbers:

$$
\begin{aligned}
fib\ Zero \quad &= \quad Zero \\
fib\,(Succ\ Zero) \quad &= \quad Succ\ Zero \\
fib\,(Succ\,(Succ\ n)) \quad &= \quad (fib\ n) + (fib\,(Succ\ n))
\end{aligned}
$$

Here, views are used on the right-hand side only for symmetry (and purposes of demonstration). It would work just as well to say *fib Zero* = 0 for the first equation. Later, we will see examples where the use of views on the right-hand side is more natural.

This view applies only to natural numbers. Any attempt to view a negative integer as a *Zero* or *Succ* (for example, by evaluating *fib* (−1)) will cause a run-time error.

It is easy to translate a program that uses views into a program that does not use views. The view definition above is equivalent to a type definition and two function definitions:

$$
\begin{aligned}
viewtype ::= &\ Zero \mid Succ\ int \\
viewin\ n \quad &= \quad Zero, \qquad\qquad \text{if } n = 0 \\
&= \quad Succ\,(n - 1), \quad \text{if } n > 0 \\
viewout\ Zero \quad &= \quad 0 \\
viewout\,(Succ\ n) \quad &= \quad n + 1
\end{aligned}
$$

The function *viewin* has type *int* → *viewtype* and the function *viewout* has type *viewtype* → *int*.

A function definition such as *power* or *fib* is translated in two steps. First, all pattern matching is translated into case expressions; see [Aug85,Wad87]. Second, calls of *viewin* and *viewout* are inserted at appropriate places. For example, the *fib* definition above is equivalent to:

```
fib m =
  case viewin m of
    Zero    ⇒  viewout Zero
    Succ m' ⇒  case viewin m' of
                 Zero   ⇒  viewout (Succ (viewout Zero))
                 Succ n ⇒  fib n + fib (viewout (Succ n))
```

Note that values of type *viewtype* appear in the program only in a very restricted way (namely, as a result of *viewin* or as an argument to *viewout*).

Any view may always be expanded out in the way outlined above. Thus, views do not require any significant change in the semantics of a functional language.

3 Another view of integers

It is possible to have more than one view of a data type. An alternative view of integers is as follows:

$$
\begin{aligned}
\textbf{view } int ::= &\ Zero \mid Even\ int \mid Odd\ int \\
\textbf{in } n \quad &= \quad Zero, & \text{if } n = 0 \\
&= \quad Even\,(n\ \text{div}\ 2), & \text{if } n > 0 \wedge n\ \text{mod}\ 2 = 0 \\
&= \quad Odd\,((n - 1)\ \text{div}\ 2), & \text{if } n > 0 \wedge n\ \text{mod}\ 2 = 1 \\
\textbf{out } Zero \quad &= \quad 0 \\
\textbf{out }(Even\ n) \quad &= \quad 2 \times n, & \text{if } 2 \times n > 0 \\
\textbf{out }(Odd\ n) \quad &= \quad 2 \times n + 1, & \text{if } 2 \times n + 1 > 0
\end{aligned}
$$

The **in** and **out** clauses again define inverse functions. Note that the constructor *Zero* appears in both views; this is permissible and unambiguous because it is given the same definition in each.

Using this view one can give a more efficient definition of exponentiation:

$$
\begin{aligned}
power\ x\ Zero \quad &= \quad 1 \\
power\ x\,(Even\ n) \quad &= \quad power\,(x \times x)\ n \\
power\ x\,(Odd\ n) \quad &= \quad x \times power\,(x \times x)\ n
\end{aligned}
$$

This expresses the traditional divide-and-conquer algorithm.

4 Viewing a complex number in cartesian and polar coordinates

This section gives a simple illustration of how abstract data types can be combined with pattern matching. The pattern matching here is extremely simple—no case analysis or recursive types are involved—but still useful.

Two well known representations of complex numbers are the polar and the cartesian. We might choose to represent complex numbers in the polar representation, and provide the cartesian as a view:

$$
\begin{aligned}
complex ::= &\ Pole\ real\ real \\
\\
\textbf{view } complex ::= &\ Cart\ real\ real \\
\textbf{in }(Pole\ r\ t) \quad &= \quad Cart\,(r \times cos\ t)\,(r \times sin\ t) \\
\textbf{out }(Cart\ x\ y) \quad &= \quad Pole\,(sqrt\,(x\char`^2 + y\char`^2))\,(atan2\ x\ y)
\end{aligned}
$$

We can then define the operations of multiplication and addition as follows:

$$
\begin{aligned}
add\,(Cart\ x\ y)\,(Cart\ x'\ y') \quad &= \quad Cart\,(x + x')\,(y + y') \\
mult\,(Pole\ r\ t)\,(Pole\ r'\ t') \quad &= \quad Pole\,(r \times r')\,(t + t')
\end{aligned}
$$

Here addition was defined in terms of the cartesian representation and multiplication in terms of the polar representation.

Alternatively, we might choose to represent complex numbers in the cartesian representation, and provide the polar as a view. This requires just a small variation on the previous declaration:

$$
\begin{aligned}
complex ::= &\ Cart\ real\ real \\
\\
\textbf{view } complex ::= &\ Pole\ real\ real \\
\textbf{in }(Cart\ x\ y) \quad &= \quad Pole\,(sqrt\,(x\char`^2 + y\char`^2))\,(atan2\ x\ y) \\
\textbf{out }(Pole\ r\ t) \quad &= \quad Cart\,(r \times cos\ t)\,(r \times sin\ t)
\end{aligned}
$$

The definitions of *add* and *mult* given previously are still valid under this new declaration. This shows how views can be used to hide choice of representation, while still allowing the convenience of pattern matching.

The traditional method for hiding a representation is an abstract data type. Just as abstract data types export values and functions, they should also be able to export views. It is easy to modify the abstract data type mechanism of Miranda to include views, for example:

$$
\begin{aligned}
\textbf{abstype } complex \textbf{ with} \\
complex ::= Cart\ real\ real \\
complex ::= Pole\ real\ real
\end{aligned}
$$

This abstract type only defines views, but in general an abstract type might define both views and functions; the syntax for declaring functions is shown below.

This declaration can be implemented with a polar representation and a cartesian view, or with a cartesian representation and a polar view, or with some third representation and cartesian and polar both as views. The definitions of *add* and *mult* will still be valid regardless of the representation chosen.

For comparison, consider the equivalent abstract type if views are not used:

> **abstype** *complex* **with**
> $\quad$ *xpart, ypart, rpart, tpart* :: *complex* → *real*
> $\quad$ *mkcart, mkpole* :: *real* → *real* → *complex*

The single identifier *Cart* has been replaced by three identifiers, *xpart*, *ypart*, and *mkcart*; and similarly for *Pole*. The corresponding definitions of addition and multiplication are:

$$add \; c \; c' \; = \; mkcart \; (xpart \; c + xpart \; c') \; (ypart \; c + ypart \; c')$$
$$mult \; c \; c' \; = \; mkpole \; (rpart \; c \times rpart \; c') \; (tpart \; c + tpart \; c')$$

In this example the difference is not great, but for more complicated examples the advantages of views—compactness and the ability to use pattern matching—would be more pronounced.

The example given here is slightly contrived, because in practice it would be more sensible to define the operations *add* and *mult* inside the abstraction, rather than outside it. However, it serves to show how views combine well with data abstraction, and particularly how views can reconcile the conflict between pattern matching and data abstraction mentioned in the introduction. One way to summarise this result is as follows. Traditionally, abstraction is achieved by refusing to export the representation. With views, abstraction can be achieved by exporting as many representations as desired.

5 Viewing a list backwards

Assume that lists are represented in the traditional way, so that, for example, $[1, 2]$ is taken as an abreviation for 1 Cons (2 Cons *Nil*), where the constructors *Nil* and *Cons* are defined as follows:

$$list \; \alpha \quad ::= \quad Nil \; | \; \alpha \; \textbf{Cons} \; (list \; \alpha)$$

(Here x **Cons** xs is just different syntax for *Cons* x xs.)

Of course, the Cons representation of lists is "biased" towards the first element of the list. For example, it is much easier to write a function to return the first element of the list than a function to return the last:

$$head \; (x \; \textbf{Cons} \; xs) \qquad\qquad = \; x$$
$$last \; (x \; \textbf{Cons} \; Nil) \qquad\qquad = \; x$$
$$last \; (x \; \textbf{Cons} \; (x' \; \textbf{Cons} \; xs)) \quad = \; last \; (x' \; \textbf{Cons} \; xs)$$

We can define a new view, based on constructors *Nil* and **Snoc**, that is biased in the opposite way:

> **view** *list* α ::= *Nil* | (*list* α) **Snoc** α
> $\quad$ **in** $(x \; \textbf{Cons} \; Nil)$ $\qquad\qquad = \; Nil \; \textbf{Snoc} \; x$
> $\quad$ **in** $(x \; \textbf{Cons} \; (xs \; \textbf{Snoc} \; x'))$ $\quad = \; (x \; \textbf{Cons} \; xs) \; \textbf{Snoc} \; x'$
> $\quad$ **out** $(Nil \; \textbf{Snoc} \; x)$ $\qquad\qquad = \; x \; \textbf{Cons} \; Nil$
> $\quad$ **out** $((x \; \textbf{Cons} \; xs) \; \textbf{Snoc} \; x') \; = \; x \; \textbf{Cons} \; (xs \; \textbf{Snoc} \; x')$

This allows us to view the list $[1, 2]$ as if were $(Nil \; \textbf{Snoc} \; 1) \; \textbf{Snoc} \; 2$. We can now write definitions such as

$$last \; (xs \; \textbf{Snoc} \; x) \qquad = \; x$$
$$rotleft \; (x \; \textbf{Cons} \; xs) \quad = \; xs \; \textbf{Snoc} \; x$$
$$rotright \; (xs \; \textbf{Snoc} \; x) \quad = \; x \; \textbf{Cons} \; xs$$

Here *last* is equivalent to the definition above, and, for example, *rotleft* $[1, 2, 3, 4] = [2, 3, 4, 1]$ and *rotright* $[1, 2, 3, 4] = [4, 1, 2, 3]$.

Attention is drawn to three features of the above view.

First, the constructor *Nil*, which is part of the representation, also appears in the view, just as *Zero* was shared between the two different views of integers. In general, it is permissible (and unambiguous) to let views of the same representation share any number of constructors with the representation and with each other.

Second, some left-hand sides in the **in** clause above contain Snoc. Matching against these will cause a recursive invocation of **in**. The **out** clause is recursive in a similar way. Recursive **in** and **out** clauses are perfectly acceptable, in the same way that recursive function definitions are.

Third, in the definition above the **in** and **out** clauses are exact inverses of one another. This can be abbreviated as follows:

> **view** *list* α ::= *Nil* | (*list* α) **Snoc** α
> $\quad$ **inout** $(x \; \textbf{Cons} \; Nil)$ $\qquad\qquad = \; Nil \; \textbf{Snoc} \; x$
> $\quad$ **inout** $(x \; \textbf{Cons} \; (xs \; \textbf{Snoc} \; x')) \; = \; (x \; \textbf{Cons} \; xs) \; \textbf{Snoc} \; x'$

As mentioned previously, a view is well-defined only when the **in** and **out** clauses define functions that are inverses of one another. In the case that a view can be defined using an **inout** clause, this property follows automatically.

Of course, the representation is still biased. For example, x Cons xs evaluates much more efficiently than xs Snoc x. Also, consider the two function definitions,

$$f_1 \; (x \; \textbf{Cons} \; (x' \; \textbf{Cons} \; xs)) = e_1 \; x \; x' \; xs$$
$$f_2 \; ((xs \; \textbf{Snoc} \; x') \; \textbf{Snoc} \; x) = e_2 \; xs \; x' \; x$$

The matching in f_1 takes constant time, while the matching in f_2 takes time proportional to the length of the input list. Indeed, the input list is traversed once to decompose the list for the outer Snoc, and then the entire input list (except for its last element) is traversed again to decompose it for the inner Snoc. The next section descibes a representation where all bias is removed, and the cons and snoc views are equally efficient.

6 The join representation of lists

As an alternative to the cons representation of lists, several researchers have suggested the following representation [Mee84,SH82]:

$$list \; \alpha \; ::= \; Nil \; | \; Unit \; \alpha \; | \; (list \; \alpha) \; \textbf{Join} \; (list \; \alpha)$$

(Interestingly, Meertens has made this suggestion for reasons of mathematical elegance, whereas Sleep and Holmström have suggested it for reasons of efficiency!) Each list now has many possible representations. For example, the list $[1, 2]$ might be represented by any of the following:

$$(Unit \; 1) \; \textbf{Join} \; (Unit \; 2)$$
$$(Nil \; \textbf{Join} \; (Unit \; 1)) \; \textbf{Join} \; ((Unit \; 2) \; \textbf{Join} \; Nil)$$
$$(Unit \; 1) \; \textbf{Join} \; ((Unit \; 2) \; \textbf{Join} \; Nil)$$

Indeed, each list has a potentially infinite number of representations, since xs and *Nil* Join xs both represent the same list.

Assuming this new representation, we can define a view that allows one to view a join list as if it were a cons list:

> **view** *list* α ::= *Nil* | α **Cons** (*list* α)
> $\quad$ **in** $(Unit \; x)$ $\qquad\qquad\qquad = \; x \; \textbf{Cons} \; Nil$
> $\quad$ **in** $(Nil \; \textbf{Join} \; xs)$ $\qquad\qquad = \; \textbf{in} \; xs$
> $\quad$ **in** $((Unit \; x) \; \textbf{Join} \; xs)$ $\qquad = \; x \; \textbf{Cons} \; xs$
> $\quad$ **in** $((xs \; \textbf{Join} \; ys) \; \textbf{Join} \; zs) \; = \; \textbf{in} \; (xs \; \textbf{Join} \; (ys \; \textbf{Join} \; zs))$
> $\quad$ **out** $(x \; \textbf{Cons} \; xs)$ $\qquad\qquad = \; (Unit \; x) \; \textbf{Join} \; xs$

The **in** clause maps all of the different ways of representing $[1, 2]$ as a join list into the same view as a cons list, namely

$$1 \; \textbf{Cons} \; (2 \; \textbf{Cons} \; Nil).$$

Conversely, the **out** clause maps this term back into a particular representation as a join list, namely

$$(Unit \; 1) \; \textbf{Join} \; ((Unit \; 2) \; \textbf{Join} \; Nil).$$

The correctness of the view depends on the equivalence between the various ways of representing a join list; otherwise, the **in** and **out** functions would not be inverses.

Note that recursion in the **in** clause above is indicated explicitly, as compared with the implicit recursion in the snoc view of the previous section.

The equation in the out clause above also appears, inverted, in the in clause. Thus, it is easy to show that the in and out clauses define inverse functions. In this case, the out clause cannot be omitted in favour of an inout clause, because a cons list can be represented in more than one way by a join list. Choosing which equations to include in the out clause is equivalent to choosing which representation to use.

A snoc view of join lists can be defined in a way completely symmetric to the definition given above. This is left as an exercise for the reader.

A join list can always be viewed as a cons list or a snoc list in time proportional to the size of the join list. (Further, unless the join list has an abundance of Nil nodes, its size will be proportional to the size of the corresponding cons and snoc lists.) Also, when the join representation is used, appending two lists requires constant time, rather than time proportional to the size of one of the lists. It is these properties that make the join representation desirable in terms of efficiency.

7   Viewing a list of pairs as a pair of lists

The function that converts a pair of lists into a list of pairs is defined as follows:

$$zip\ (Nil, Nil)\qquad\qquad =\quad Nil$$
$$zip\ (a\ \text{Cons}\ as, b\ \text{Cons}\ bs)\quad =\quad (a, b)\ \text{Cons}\ zip\ (as, bs)$$

(The pair notation is just another syntax for constructors; think of (a, b) as equivalent to $Pair\ a\ b$.) For example,

$$zip\ ([1, 2, 3], [\text{'a'}, \text{'b'}, \text{'c'}]) = [(1, \text{'a'}), (2, \text{'b'}), (3, \text{'c'})]$$

Very often, it is also necessary to decompose a list of pairs into a pair of lists. This is usually done by an idiom such as the following:

$$f\ cs\quad =\quad e\ as\ bs$$
$$\text{where}\ as\quad =\quad [a\ |\ (a, b) \leftarrow cs]$$
$$bs\quad =\quad [b\ |\ (a, b) \leftarrow cs]$$

(This uses *list comprehension* notation; see [Tur81,Wad87].)

Clearly, zip defines an isomorphism: given a list of pairs cs there is always a unique pair of lists as and bs such that $zip\ (as, bs) = cs$. Thus, we can discard the definition of zip given above, and instead define a view Zip of pairs of lists:

$$\text{view}\ list\ (\alpha, \beta) ::= Zip\ (list\ \alpha, list\ \beta)$$
$$\text{inout}\ Nil\qquad\qquad =\quad Zip\ (Nil, Nil)$$
$$\text{inout}\ ((a, b)\ \text{Cons}\ Zip\ (as, bs))\quad =\quad Zip\ (a\ \text{Cons}\ as, b\ \text{Cons}\ bs)$$

(Here the type of the pair (a, b) is written (α, β).) Now we can write $Zip\ (as, bs)$ in place of $zip\ (as, bs)$ on the right-hand side of equations. We can also write $Zip\ (as, bs)$ on the left-hand sides of equations. For example, the idiom given above becomes

$$f\ (Zip\ (as, bs))\quad =\quad e\ as\ bs$$

which is somewhat shorter.

8   Two representations of trees

One common representation of trees is the following:

$$tree\ \alpha\quad ::=\quad Leaf\ \alpha\ |\ Branch\ (tree\ \alpha)\ (tree\ \alpha)$$

Another common representation uses what is called the "spine" (mixing an anatomical metaphor with an arboreal one). Spine trees, together with the isomorphism that relates them to branch trees, can be conveniently described by the following view:

$$\text{view}\ tree\ \alpha ::= Spine\ \alpha\ (list\ (tree\ \alpha))$$
$$\text{inout}\ (Leaf\ x)\qquad\qquad =\quad Spine\ x\ Nil$$
$$\text{inout}\ (Branch\ (Spine\ x\ xts)\ xt)\quad =\quad Spine\ x\ (xts\ \text{Snoc}\ xt)$$

For example, the branch tree

$$Branch\ (Branch\ (Leaf\ \text{"f"})\ (Leaf\ \text{"a"}))\ (Leaf\ \text{"b"})$$

is equivalent to the spine tree

$$Spine\ \text{"f"}\ [Spine\ \text{"a"}\ [\ ], Spine\ \text{"b"}\ [\ ]].$$

Of course, one could use spine trees as the underlying representation with branch trees as the view, or hide the representation in an abstract type and provide both branch and spine as views.

9   Two other uses of views

This section presents two rather unusual uses of views. It is not clear whether these uses should be considered good style, but they do demonstrate the power of the view mechanism.

Occasionally, it is convenient to both match an argument against a pattern, and to refer to it by a single name. (One might say we wish to "eat our argument and have it too".) Hope provides the as construct for this purpose. For example, one might write:

$$factorial\ (n\ \text{as}\ Zero)\qquad =\quad 1$$
$$factorial\ (n\ \text{as}\ Succ\ n')\quad =\quad n \times factorial\ n'$$

Surprisingly, as can be defined as a view:

$$\text{view}\ \alpha ::= \alpha\ \text{as}\ \alpha$$
$$\text{in}\ x\qquad\qquad =\quad x\ \text{as}\ x$$
$$\text{out}\ (x\ \text{as}\ x')\quad =\quad x,\qquad \text{if}\ x = x'$$

This defines an isomorphism between any type α, and the subset of the viewing type α as α where the left and right arguments of the constructor are equal. It is unlikely that one would want to use as on the right-hand side in an equation, but the out clause is necessary for the view to be well-defined.

One may even use views in place of predicates. For example, one might define:

$$\text{view}\ int ::= EvenP\ int\ |\ OddP\ int$$
$$\text{in}\ n\qquad\qquad =\quad EvenP\ n,\quad \text{if}\ n\ \text{mod}\ 2 = 0$$
$$=\quad OddP\ n,\quad \text{if}\ n\ \text{mod}\ 2 = 1$$
$$\text{out}\ EvenP\ n\quad =\quad n,\qquad \text{if}\ n\ \text{mod}\ 2 = 0$$
$$\text{out}\ OddP\ n\quad =\quad n,\qquad \text{if}\ n\ \text{mod}\ 2 = 1$$

Then we can write

$$f\ (EvenP\ n)\quad =\quad e_1\ n$$
$$f\ (OddP\ n)\quad =\quad e_2\ n$$

instead of

$$f\ n\quad =\quad e_1\ n,\quad \text{if}\ n\ \text{mod}\ 2 = 0$$
$$=\quad e_2\ n,\quad \text{if}\ n\ \text{mod}\ 2 = 1$$

Replacing conditions by patterns may occasionally be clearer, particularly if many functions test the same condition. It may also improve clarity when a function has many arguments, and the test of the condition interacts with pattern matching for the other arguments. From the out clause, it can be seen that the term $EvenP\ n$ is equivalent to n with the additional assertion that n is even. Conceivably, it might be useful to use $EvenP$ and $OddP$ on the right-hand sides of equations, as a way of documenting that certain conditions hold.

10   Equational reasoning

In order for a language feature to be useful, it must be easy to reason about programs containing that feature. Views have been carefully designed to support two important proof techniques, equational reasoning and induction. These are discussed in this section and the next.

Equational reasoning is a principle of such supreme importance that it goes by many names: referential transparency, the rule of Leibniz, and more plainly "substituting equals for equals". As a very simple example, given the function definition

$$last\ (x\ \text{Cons}\ Nil)\qquad\qquad =\quad x \qquad\qquad (1)$$
$$last\ (x\ \text{Cons}\ (x'\ \text{Cons}\ xs))\quad =\quad last\ (x'\ \text{Cons}\ xs) \qquad (2)$$

equational reasoning is sufficient to calculate the value of $last\ [\text{'a'}, \text{'b'}]$,

as follows:

$$last\ (\text{`a' Cons (`b' Cons } Nil)) \ = \ last\ (\text{`b' Cons } Nil) \qquad \text{(by 2)}$$
$$= \ \text{`b'} \qquad \text{(by 1)}$$

A key principle in the design of views is that all equations in the in and out clauses of the view can be used just like any other equations for equational reasoning. Thus, given the view

$$\textbf{view } list\ \alpha ::= Nil \mid (list\ \alpha)\ \text{Snoc } \alpha$$
$$\text{inout } (x \text{ Cons } Nil) \qquad\quad = \ Nil \text{ Snoc } x \qquad (3)$$
$$\text{inout } (x \text{ Cons } (xs \text{ Snoc } x')) \ = \ (x \text{ Cons } xs) \text{ Snoc } x' \qquad (4)$$

and the definition

$$last\ (xs \text{ Snoc } x) \ = \ x \qquad (5)$$

we may now calculate as follows:

$$last\ (\text{`a' Cons (`b' Cons } Nil))$$
$$= \ last\ (\text{`a' Cons } (Nil \text{ Snoc `b'})) \qquad \text{(by 3)}$$
$$= \ last\ ((\text{`a' Cons } Nil) \text{ Snoc `b'}) \qquad \text{(by 4)}$$
$$= \ b \qquad \text{(by 5)}$$

Of course, the main value of equational reasoning is not in calculating values but in performing proofs. Given the definitions

$$rotleft\ (x \text{ Cons } xs) \ = \ xs \text{ Snoc } x \qquad (6)$$
$$rotright\ (xs \text{ Snoc } x) \ = \ x \text{ Cons } xs \qquad (7)$$

we may prove

$$rotleft\ (rotright\ xs) = xs$$

for a non-empty finite list xs, by simply observing that

$$rotleft\ (rotright\ (x \text{ Cons } xs) \ = \ rotleft\ (xs \text{ Snoc } x) \qquad \text{(by 6)}$$
$$= \ x \text{ Cons } xs \qquad \text{(by 7)}$$

Both the definitions and the proof are rather more involved if views are not used.

Equational reasoning is valid for all definitions that use pattern matching over free data types. A view establishes an isomorphism between (a subset of) the viewed data type and (a subset of) a free data type. Thus, equational reasoning is also valid for all definitions that use pattern matching over views.

However, some caution is required, because the view may imply additional conditions that the program must satisfy. For example, the polar view of complex numbers in Section 4 is valid only if for all angles t_1 and t_2 the equation

$$Pole\ 0\ t_1 = Pole\ 0\ t_2$$

is consistent with the rest of the program. This is required because all points of the form $Pole\ 0\ t$ map into the same cartesian representation, $Cart\ 0\ 0$.

As a more extended example, consider the view of join lists as cons lists. A look at this view will reveal that it establishes the following conditions on join lists:

$$Nil \text{ Join } xs \ = \ xs$$
$$xs \text{ Join } (ys \text{ Join } xs) \ = \ (xs \text{ Join } ys) \text{ Join } xs$$

That is, Join must have Nil as a (left) identity and be associative. The programmer must verify that every definition containing Join is consistent with these conditions. It is also desirable that definitions be consistent with the condition

$$xs \text{ Join } Nil \ = \ xs$$

but this is not required by the view.

An example of a satisfactory definition is

$$length\ Nil \qquad\quad = \ 0$$
$$length\ (Unit\ x) \qquad = \ 1$$
$$length\ (xs \text{ Join } ys) \ = \ length\ xs + length\ ys$$

This definition establishes a homomorphism mapping Join onto + and Nil onto 0. It is valid because + has 0 as an identity and is associative, and so the desired properties are preserved. (For further discus-

sion homomorphisms and operations on lists, the reader is referred to [Mee84,Bir86].)

An example of an unsatisfactory definition is

$$silly\ Nil \qquad\quad = \ 1$$
$$silly\ (Unit\ x) \qquad = \ 2$$
$$silly\ (xs \text{ Join } ys) \ = \ silly\ xs + silly\ ys$$

This definition does not preserve the desired properties, because 1 is not a right identity of +. Thus, from the condition that $xs = Nil \text{ Join } xs$ we could derive

$$2 = silly\ (Unit\ \text{`a'}) = silly\ (Nil \text{ Join } (Unit\ \text{`a'})) = 3$$

It is impossible to reason equationally about a program containing this definition, which is just as well since it is indeed a silly program.

A further example of an unsatisfactory definition is

$$nothead\ ((Unit\ x) \text{ Join } xs) \ = \ x$$

This definition won't do, because we have

$$(Unit\ \text{`a'}) \text{ Join } (Unit\ \text{`b'}) = Nil \text{ Join } ((Unit\ \text{`a'}) \text{ Join } (Unit\ \text{`b'}))$$

but applying $nothead$ to the left of this equation yields 'a', while applying $nothead$ to the right yields an undefined value.

One should consider it a bonus that views reveal an additional condition that definitions involving Join must be checked for. On the other hand, it is useful to be able to limit the portions of the program that must be so checked. This can be done in the usual way, by encapsulating Join within an abstract data type. Further, it is completely safe to export the cons and snoc views of lists outside of this type, because they are guaranteed to enforce the necessary conditions. For example, the definition

$$head\ (x \text{ Cons } xs) \ = \ x$$

is completely satisfactory, because the cons view necessarily respects the equivalences among various join representations of the same list.

11 Induction

It is also essential that induction should work over views. For example, in order to demonstrate that a property $P(xs)$ holds for every list xs, it is sufficient to show

1. $P(Nil)$ holds, and

2. $P(xs \text{ Snoc } x)$ holds, assuming $P(xs)$ holds.

Similarly, one can prove properties of the natural numbers by inducting over $Zero$ and $Succ$.

In general, induction over a view is valid to demonstrate properties that hold for all elements *generated* by the viewing type. For example, every finite list can be generated using Nil and Snoc, so induction over these serves to prove properties of finite lists. Similarly, every natural number can be generated using $Zero$ and $Succ$, so induction over these serves to prove properties of natural numbers. Here the induction only demonstrates properties of a subset (natural numbers) of the viewed type (integers).

A more problematic example is the unusual view of even and odd integers presented in Section 9. The set of values that can be generated using just the two constructors $EvenP$ and $OddP$ is empty, and so induction is not appropriate over this view.

12 Implementation

As explained in Section 2, each view can be expanded out into the definition of a new type, *viewtype*, and applications of conversion functions, *viewin* and *viewout*. This is workable, but not ideal. Constructors of the viewed type (such as $Zero$ and $Succ$) appear only as the result of *viewin* or the argument to *viewout*. At run-time, such constructors are always allocated storage and then immediately examined and never referenced again. It would be preferable to use a scheme that avoided

allocating such constructors altogether.

Such a scheme is possible. Instead of introducing a new type and two conversion functions, in the modified scheme a view with k constructors introduces $k + 1$ new functions (and no new types). One function acts as the equivalent of *viewin* and the associated case expression, while the remaining k functions provide an equivalent of *viewout* for each constructor.

For example, the view of integers in Section 2 translates into the following three functions:

$$
\begin{aligned}
viewcase\ s\ s\ n\ &=\ s &&\text{if } n = 0 \\
&=\ s\ (n - 1) &&\text{if } n > 0 \\
zero\ &=\ 0 \\
succ\ n\ &=\ n + 1
\end{aligned}
$$

Note that *viewcase* is designed so that if n corresponds to *Zero* then s is returned, and if n corresponds to *Succ* n' then s n' is returned. Since s is itself a function, *viewcase* can be defined only in a higher-order language, where functions may be passed to other functions.

Under this scheme, the definition of *fib* now translates as follows:

$$
\begin{aligned}
fib\ m\ =\ &viewcase \\
&zero \\
&(\lambda m'.viewcase \\
&\quad (succ\ zero) \\
&\quad (\lambda n.\ n + fib\ (succ\ n)) \\
&\quad m') \\
&m
\end{aligned}
$$

It is left to the reader to verify that this translation is equivalent to the one described in Section 2.

The new translation scheme has the desired property that it introduces no new constructors (like *Succ* and *Zero*) and hence requires no extra allocation operations at run-time. However, it has the disadvantage that application of *viewcase* may be more difficult to evaluate efficiently than the corresponding case expressions. The exact nature of this problem will depend on the evaluation method used. Here the problem will be considered in the context of a G-machine style compiler, and a solution suggested. Some familiarity with the G-machine is assumed; see [Aug84,Joh85,Pey87] for an introduction.

As an example, consider the following definition (which happens to provide a fast method for calculating Fibonacci numbers):

$$
\begin{aligned}
fibx\ a\ b\ Zero\ &=\ a \\
fibx\ a\ b\ (Succ\ n)\ &=\ fibx\ b\ (a + b)\ n
\end{aligned}
$$

Translated using the above scheme, this becomes

$$
fibx\ a\ b\ m\ =\ viewcase\ a\ (\lambda n.fibx\ b\ (a + b)\ n)\ m
$$

After lambda-lifting, this definition in turn becomes

$$
\begin{aligned}
fibx\ a\ b\ m\ &=\ viewcase\ a\ (fibx_0\ a\ b)\ m \\
fibx_0\ a\ b\ n\ &=\ fibx\ b\ (a + b)\ n
\end{aligned}
$$

At run-time, the code implementing the body of *fibx* will need to allocate heap storage to represent the application $(fibx_0\ a\ b)$. Thus, we have saved the allocation of a constructor *Zero* or *Succ* only to replace it by a potentially larger allocation (since in general storage will need to be allocated for each free variable in each argument to *viewcase*).

Fortunately, this problem can be solved by simply rearranging the way free variables are passed into arguments to *viewcase*. In the rearranged method, all free variables of the function containing the *viewcase* are lambda-abstracted from each argument to the *viewcase* (and in the same order); these variables are then passed into the result

returned by the *viewcase*. Thus, the definition of *fibx* now translates to

$$
\begin{aligned}
fibx\ a\ b\ m\ =\ &\\
viewcase\ &(\lambda a.\lambda b.\lambda m.a)(\lambda n.\lambda a.\lambda b.\lambda m.fibx\ b\ (a + b)\ n)\ m\ a\ b\ m
\end{aligned}
$$

After lambda-lifting this becomes

$$
\begin{aligned}
fibx\ a\ b\ m\ &=\ viewcase\ fibx_1\ fibx_2\ m\ a\ b\ m \\
fibx_1\ a\ b\ m\ &=\ a \\
fibx_2\ n\ a\ b\ m\ &=\ fibx\ b\ (a + b)\ n
\end{aligned}
$$

Now the body of *fibx* need merely push pointers to $fibx_1$ and $fibx_2$ onto the stack, copy the parameter m onto the top of the stack, and perform a tail-call (that is, a jump) to *viewcase*. In turn, *viewcase* will examine the argument m; if it is zero it will perform a tail-call (jump) to $fibx_1$, and if it is positive it will push $m - 1$ onto the stack and perform a tail-call (jump) to $fibx_2$. By passing the free variables in the body of *fibx* to each argument in the same order, the rearrangement of the stack necessary at run-time is minimised. (Given the above definition of *fibx*, the current G-machine compiler would not produce quite the sequence of steps described here, but it could be modified to do so.)

This is quite acceptably efficient. Further efficiency might be gained by expanding out non-recursive function applications at compile time. (The compilers described by [HK84,FW86] perform expansion of this kind.) For example, performing expansion on the above definition of *fibx* yields

$$
\begin{aligned}
fibx\ a\ b\ m\ &=\ a, &&\text{if } m = 0 \\
&=\ fibx\ b\ (a + b)\ (m - 1), &&\text{if } m > 0
\end{aligned}
$$

which resembles the code one would write if a view had not been used. However, it is difficult to see how to perform an equivalent expansion of a recursive *viewcase*, such as the one associated with the snoc view of lists.

13 Conclusions

Designers of software are continually faced with trade-offs. Some of these trade-offs are necessary, but others can be avoided by careful design. It is particularly worrying when we are forced to choose between valuable methods such as pattern matching and data abstraction. Views move this trade-off from the "necessary" to the "avoidable" category.

After the views mechanism was defined, several unexpected uses of it emerged. These included the list-of-pairs to pair-of-lists view discussed in Section 7 and the two unusual views discussed in Section 9. No doubt many other uses of views are waiting to be discovered.

Programming languages are awash with features, and new features must be approached with caution. Views are worth consideration because they address an important need—reconciling pattern matching with data abstraction. In doing so, they also bring a new perspective. Instead of thinking of an abstract data type as something which hides the representation, with views we can think of it as something which exports as many representations as convenient.

Trademark notice. Miranda is a trademark of Research Software, Ltd.

312

References

[Aug84] L. Augustsson. A compiler for lazy ML. In *Proceedings of the 1984 ACM Symposium on Lisp and Functional Programming*, pages 218–227, Austin, 1984.

[Aug85] L. Augustsson. Compiling pattern matching. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*, Springer-Verlag, September 1985.

[BD77] R. M. Burstall and J. Darlington. A transformation system for developing recursive programs. *Journal of the ACM*, 24(1):44–67, January 1977.

[Bir86] R. S. Bird. An introduction to the theory of lists. In *Marktoberdorf Workshop on Logics of Programming*, August 1986.

[BMS80] R. Burstall, D. MacQueen, and D. Sanella. *Hope: An experimental applicative language*. Technical Report Report CSR-62-80, Edinburgh University, Computer Science Dept., 1980.

[Bur69] R. M. Burstall. Proving properties of programs by structural induction. *The Computer Journal*, 12(1), February 1969.

[FGJM85] K. Futasagi, J. A. Goguen, J.-P. Jouannaud, and J. Meseguer. Principles of OBJ2. In *ACM Symposium on Principles of Programming Languages*, pages 52–66, January 1985.

[FW86] J. Fairbairn and S. C. Wray. Code generation techniques for functional languages. In *Proceedings of the 1986 ACM Symposium on Lisp and Functional Programming*, pages 94–104, Boston, 1986.

[HK84] P. Hudak and D. Krans. A combinator-based compiler for a functional language. In *ACM Symposium on Principles of Programming Languages*, pages 121–132, January 1984.

[Joh85] T. Johnsson. Lambda lifting: transforming programs to recursive equations. In *Proceedings 1985 Conference on Functional Programming Languages and Computer Architecture*, Nancy, France, 1985.

[Mee84] L. Meertens. Algorithmics: Towards programming as a mathematical activity. In J. W. de Bakker, et. al., editors, *Mathematics and Computer Science*, North-Holland, 1984.

[Pey87] S. L. Peyton-Jones. *Implementing Functional Languages using Graph Reduction*. Prentice-Hall, 1987.

[SH82] M. R. Sleep and S. Holmström. A short note concerning lazy reduction rules of append. *Software Practice and Experience*, 12(11):1082–4, November 1982.

[Tho86] S. Thompson. Laws in miranda. In *ACM Symposium on Lisp and Functional Programming*, pages 1–12, August 1986.

[Tur81] D. A. Turner. Recursion equations as a programming language. In J. Darlington, P. Henderson, and D. Turner, editors, *Functional Programming and Its Applications*, Cambridge University Press, 1981.

[Tur85] D. A. Turner. Miranda: A non-strict functional language with polymorphic types. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*, Springer-Verlag, September 1985.

[Wad87] P. L. Wadler. Compiling pattern matching; List comprehensions. In [Pey87].