

# Design Patterns as Higher-Order Datatype-Generic Programs

Jeremy Gibbons

<http://www.comlab.ox.ac.uk/jeremy.gibbons/>

7th December 2010

## Abstract

Design patterns are reusable abstractions in object-oriented software. However, using current mainstream programming languages, these elements can only be expressed extra-linguistically: as prose, pictures, and prototypes. We believe that this is not inherent in the patterns themselves, but evidence of a lack of expressivity in the languages of today. We expect that, in the languages of the future, the code parts of design patterns will be expressible as reusable library components. Indeed, we claim that the languages of tomorrow will suffice; the future is not far away. All that is needed, in addition to commonly-available features, are *higher-order* and *datatype-generic* constructs; these features are already or nearly available now. We argue the case by presenting higher-order datatype-generic programs capturing ORIGAMI, a small suite of patterns for recursive data structures.

## 1 Introduction

Design patterns, as the subtitle of the seminal book [20] has it, are ‘elements of reusable object-oriented software’. However, within the confines of existing mainstream programming languages, these supposedly reusable elements can only be expressed extra-linguistically: as prose, pictures, and prototypes. We believe that this is not inherent in the patterns themselves, but evidence of a lack of expressivity in the languages of today. We expect that, in the languages of the future, the code parts of design patterns will be expressible as directly reusable library components. The benefits will be considerable: patterns may then be reasoned about, type-checked, applied and reused, just as any other abstractions can.

Indeed, we claim that the languages of tomorrow will suffice; the future is not far away. All that is needed, in addition to what is provided by essentially every programming language, are *higher-order* (parametrization by code) and *datatype-generic* (parametrization by type constructor) features. Higher-order constructs have been available for decades in functional programming languages such as ML [53] and Haskell [66]. Datatype genericity can be simulated in existing programming languages [11, 36, 59], but we already have significant experience with robust prototypes of languages that support it natively [37, 49].

We argue our case by capturing as higher-order datatype-generic programs a small subset ORIGAMI of the Gang of Four (GOF) patterns. (For the sake of rhetorical style, in this paper we equate ‘GOF patterns’ with ‘design patterns’.) These programs are parametrized along three dimensions: by the *shape* of the computation, which is determined by the shape of the underlying data, and represented by a type constructor (an operation on types); by the *element type* (a type); and by the *body* of the computation, which is a higher-order argument (a value, typically a function).

Although our presentation is in a functional programming style, we do not intend to argue here that functional programming is the paradigm of the future. Rather, we believe that functional programming languages are a suitable test-bed for experimental language features—as evidenced by parametric polymorphism, list comprehensions, and of course closures, which are all now finding their way into mainstream programming languages such as Java and C#. We expect that the evolution of programming languages will continue to follow the same trend: experimental language features will be developed and explored in small, nimble laboratory languages, and the successful experiments will eventually make their way into the outside world. In particular, we expect that the mainstream languages of tomorrow will be broadly similar to the languages of today—strongly and statically typed, object-oriented, with an underlying imperative

mindset—but incorporating additional features from the functional world—specifically, higher-order operators and datatype genericity. (Indeed, we would not be in the least surprised to see the language mainstream heading in the direction of the functional/object-oriented language Scala [57]. We have elsewhere [60] argued that Scala makes a fine language for generic programming; it supports both the familiar concepts of the OO paradigm and the higher-order and datatype-generic features we make use of in this paper.)

Thus, our main contribution in this paper is to show that the code aspects of four of the familiar Gang of Four *design patterns* [20]—namely COMPOSITE, ITERATOR, VISITOR, and BUILDER, together with some variations—can be captured as reusable library code in the form of *higher-order, datatype-generic programs*—specifically, as recursive datatypes with maps, folds, and unfolds, again with some variations; we therefore claim that higher-order and datatype-generic features are very helpful in defining flexible software components. This argument is presented in Section 6, where we show how to capture our ORIGAMI patterns as higher-order datatype-generic programs; Sections 2 to 5 cover the necessary background, on functional programming, generic programming, folds and unfolds, and design patterns, respectively, and Sections 7 to 9 discuss limitations, survey related work, and conclude.

## 2 Functional programming

We start with a brief review of the kinds of parametrization available in functional programming languages. We do this in order to emphasize what is different about datatype genericity, which we believe will be the next step in parametrization.

### 2.1 First-order, monomorphic

Functional programming is a matter of programming with expressions rather than statements, manipulating values rather than actions. For example, consider the following two datatype definitions, of lists of integers and lists of characters respectively. In each case, there are two variants: a constant for the empty list, and a binary operator constructing a non-empty list from its head and tail.

```
data ListI = NilI | ConsI Integer ListI
data ListC = NilC | ConsC Char ListC
```

Unlike in conventional imperative (and object-oriented) programming, computations do not proceed by executing actions that destructively update a state; instead, they construct values by evaluating expressions. Thus, instead of a loop repeatedly updating a running total, summing a list of integers entails the recursive evaluation of one lengthy expression. For example, the following program computes the sum of a list of integers: the sum of the empty list is zero, and the sum of a non-empty list is the head plus the sum of the tail.

```
sumI :: ListI → Integer
sumI NilI          = 0
sumI (ConsI x xs) = x + sumI xs
```

Similarly, programs to append two lists (of integers, or of characters) proceed non-destructively by recursion.

```
appendI :: ListI → ListI → ListI
appendI NilI          ys = ys
appendI (ConsI x xs) ys = ConsI x (appendI xs ys)

appendC :: ListC → ListC → ListC
appendC NilC          ys = ys
appendC (ConsC x xs) ys = ConsC x (appendC xs ys)
```

## 2.2 First-order, polymorphic

The attentive reader will note that the definitions above of *appendI* and *appendC* are essentially identical: their monomorphic types are over-specific. Abstracting from their differences allows us to capture their commonalities. The kind of abstraction that is required is parametrization by type. This can be done both in the programs and in the datatypes they manipulate, yielding *parametrically polymorphic* datatypes and functions. The datatype declaration defines a type constructor *List*, which when applied to an element type (such as *Integer*) denotes another type (in this case, lists of integers):

```
data List a = Nil | Cons a (List a)
sum  :: List Integer → Integer
sum Nil          = 0
sum (Cons x xs) = x + sum xs
```

Now we no longer need two monomorphic datatypes *ListI* and *ListC*; both can be obtained by instantiating the polymorphic datatype *List a*. Consequently, we do not need two different monomorphic append functions either; a single polymorphic function suffices.

```
append :: List a → List a → List a
append Nil      ys = ys
append (Cons x xs) ys = Cons x (append xs ys)
```

Here, for later reference, is another example of a parametrically polymorphic function: the function *concat* concatenates a list of lists (of elements of some type) to a single long list (of elements of the same type).

```
concat :: List (List a) → List a
concat Nil          = Nil
concat (Cons xs xss) = append xs (concat xss)
```

## 2.3 Higher-order, list-specific

Each of the three programs *sum*, *append*, and *concat* traverses its list argument in exactly the same way. Abstracting from their differences allows us to capture a second kind of commonality, namely the pattern of recursion. The kind of parametrization that is required is parametrization by a program; doing so yields *higher-order* programs. The common pattern is called a ‘fold’; it has two cases, one for each variant of the datatype, and makes identical recursive calls on each recursive substructure.

```
foldL :: b → (a → b → b) → List a → b
foldL n c Nil          = n
foldL n c (Cons x xs) = c x (foldL n c xs)
```

Instances of *foldL* replace the list constructors *Nil* and *Cons* with the supplied arguments *n* and *c*:

```
sum      = foldL 0 (+)
append xs ys = foldL ys Cons xs
concat   = foldL Nil append
```

For more about higher-order programming, see any textbook on functional programming [65, 6].

## 2.4 Higher-order, tree-specific

Now, suppose one also had a polymorphic datatype of binary trees (here, externally labelled—that is, with elements at the tips and not at the binary nodes):

```
data Btree a = Tip a | Bin (Btree a) (Btree a)
```

A similar process would lead one to abstract the natural pattern of computation on these trees as another higher-order function:

$$\begin{aligned} \text{foldB} &:: (a \rightarrow b) \rightarrow (b \rightarrow b \rightarrow b) \rightarrow \text{Btree } a \rightarrow b \\ \text{foldB } t \ b \ (\text{Tip } x) &= t \ x \\ \text{foldB } t \ b \ (\text{Bin } xs \ ys) &= b \ (\text{foldB } t \ b \ xs) \ (\text{foldB } t \ b \ ys) \end{aligned}$$

For example, instances of *foldB* reflect a tree, and flatten it to a list, in both cases replacing the tree constructors *Tip* and *Bin* with supplied arguments:

$$\begin{aligned} \text{reverse} &:: \text{Btree } a \rightarrow \text{Btree } a \\ \text{reverse} &= \text{foldB } \text{Tip } \text{nib} \ \mathbf{where} \\ &\quad \text{nib } xs \ ys = \text{Bin } ys \ xs \\ \text{flatten} &:: \text{Btree } a \rightarrow \text{List } a \\ \text{flatten} &= \text{foldB } \text{wrap } \text{append} \ \mathbf{where} \\ &\quad \text{wrap } x = \text{Cons } x \ \text{Nil} \end{aligned}$$

## 2.5 Datatype-generic

We have seen that each kind of parametrization allows some recurring patterns to be captured. Parametric polymorphism unifies commonality of computation, abstracting from variability in irrelevant typing information. Higher-order functions unify commonality of program shape, abstracting from variability in some of the details.

But what about the two higher-order, polymorphic programs *foldL* and *foldB*? We can see that they have something in common: both replace constructors by supplied arguments; both have patterns of recursion that follow the datatype definition, with one clause per datatype variant and one recursive call per substructure. There is evidently a recurrent pattern; but neither parametric polymorphism nor higher-order functions suffice to capture this commonality.

In fact, what differs between the two fold operators is the *shape* of the data on which they operate, and hence the shape of the programs themselves. The kind of parametrization required is by this shape; that is, by the datatype or type constructor (such as *List* or *Tree*) concerned. We call this *datatype genericity*; it allows the capture of recurring patterns in *programs of different shapes*. In Section 4 below, we explain the definition of a datatype-generic operation *fold* with the following type:

$$\text{fold} :: \text{Bifunctor } s \Rightarrow (s \ a \ b \rightarrow b) \rightarrow \text{Fix } s \ a \rightarrow b$$

Here, in addition to the type *a* of collection elements and the fold body (a function of type *s a b → b*), the shape parameter *s* varies; the type class *Bifunctor* expresses the constraints we place on its choice. The shape parameter determines the shape of the input data; for one instantiation *ListF* of *s*, the type *Fix s a* is isomorphic to *List a*, and for another instantiation *BtreeF*, the type *Fix s a* is isomorphic to *Btree a*. The same shape parameter also determines the type of the fold body, supplied as an argument with which to replace the constructors; when *s* is instantiated to *ListF*, the function *fold* specializes to *foldL*, and when *s* is *BtreeF*, it specializes to *foldB*.

For more about datatype genericity, see [27].

## 3 Genera of genericity

The term *generic programming* means different things to different people: to some, it means parametric polymorphism; to some, it means libraries of algorithms and data structures; to some, it means reflection and meta-programming; to some, it means polytypism. By and large, though, everyone agrees with the intention, characterized by Sheard [3], of ‘making programs more adaptable by making them more general’.

Generic programming usually manifests itself as a kind of parametrization. By abstracting from the differences in what would otherwise be separate but similar specific programs, one can make a single unified generic program. Instantiating the parameter in various ways retrieves the various specific programs, and

- (a) **generic**  
     **type** *Element* **is private** ;  
     **procedure** *Swap* (*X*, *Y* : **in out** *Element*) ;
- (b) **procedure** *Swap* (*X*, *Y* : **in out** *Element*) **is**  
     *Z* : **constant** *Element* := *X* ;  
     **begin**  
         *X* := *Y* ;  
         *Y* := *Z* ;  
     **end** *Swap* ;
- (c) **procedure** *SwapInteger* **is new** *Swap* (*Integer*) ;

Figure 1: An Ada generic subprogram: (a) declaration, (b) definition, and (c) instantiation

ideally some new ones too. The different interpretations of the term ‘generic programming’ arise from different notions of what constitutes a ‘parameter’.

Moreover, a parametrization is usually only called ‘generic’ programming if it is of a ‘non-traditional’ kind; by definition, traditional kinds of parametrization give rise only to traditional programming, not generic programming. Therefore, ‘genericity’ is in the eye of the beholder, with beholders from different programming traditions having different interpretations of the term. No doubt by-value and by-name parameter-passing mechanisms for arguments to procedures, as found in Pascal, look like ‘generic programming’ to an assembly-language programmer.

### 3.1 Parametric polymorphism

One interpretation of the term ‘generic programming’ is as embodied by Ada generics [54]. These were inspired by Liskov’s CLU language [48], and were in turn the inspiration for the C++ template mechanism, which we discuss below. Ada *generic units* encompass ‘subprograms’ (functions or procedures) and ‘packages’ (modules or abstract datatypes). These generic units can be parametrized in various ways: by types, by values, by subprograms, and by packages. As with C++, Ada generic units are templates for their non-generic counterparts; they cannot be used until they are instantiated.

For example, Figure 1 shows: (a) the declaration of a generic subprogram *Swap*, parametrized by a type *Element*; (b) the generic body of the subprogram, which makes use of the formal parameter *Element*; and (c) the instantiation of the generic unit to make a non-generic subprogram that may be used in the same way as any other subprogram.

Ada generics implement a form of parametric polymorphism: the same generic definition is applicable to all possible instantiations. This is true for the other available kinds of formal parameter, as well as for type parameters (although Cardelli and Wegner [10] would call parametrization by these other kinds of parameter *universal polymorphism*, restricting *parametric polymorphism* to mean ‘universal polymorphism with a type parameter’). For example, a value parameter could be used to parametrize a collection type by a size bound; a subprogram parameter could be used to parametrize a sorting subprogram with an ordering; a package parameter could be used to parametrize Horner’s rule for polynomial evaluation by a semiring [13], an extremal path finder by a regular algebra [4, 5], or a greedy algorithm by a matroid or greedoid structure [18, 47].

### 3.2 Generic programming and the STL

The most popular interpretation of the term ‘generic programming’ is as embodied by the C++ Standard Template Library (STL). This is an object-oriented class library providing *containers*, *iterators* and *algorithms* for many datatypes. As the name suggests, it is implemented using the C++ template mechanism, which offers similar facilities to Ada generics: class- and function templates are parametrized by type- and value parameters. (Indeed, a predecessor to the STL was an Ada library for list processing [55].) Within this

community more than any other, it is considered essential that genericity imposes no performance penalty [16, 74].

The containers in the STL are parametrically polymorphic datatypes, parametrized by the element type; these are further classified into *sequence containers* (such as *Vector*, *String* and *Deque*) and *associative containers* (such as *Set*, *Multiset* and *Map*).

Bulk access to the elements of a container type is provided by *iterators*. These are abstractions of C++ pointers to elements, and so support pointer arithmetic. They are further classified according to what pointer operations they support: *input iterators* (which can only be read from, that is, dereferenced as r-values), *output iterators* (which can only be written to, that is, dereferenced as l-values), *forwards iterators* (which can be advanced, that is, supporting increment), *backwards iterators* (which can also retreat, that is, supporting decrement), and *random access iterators* (which can move any number of steps in one operation, that is, supporting addition).

Iterators form the interface between container types and *algorithms* over data structures. STL algorithms include many general-purpose operations such as searching, sorting, filtering, and so on. Rather than operating directly on a container, an algorithm takes one or more iterators as parameters; but then the algorithm is generic, in the sense that it applies to any container that supports the appropriate kind of iterator.

In the C++ approach, the exact set of requirements on parameters (such as the iterator passed to a generic algorithm, or the element type passed to a generic container, or indeed any other template parameter) is called a *concept* [42]. It might specify the operations available on a type parameter, the laws these operations satisfy, and the asymptotic complexities of the operations in terms of time and space; the first of these can be checked at instantiation time, but the other two cannot. The C++ template mechanism provides no means explicitly to define a concept; it is merely an informal artifact rather than a formal construct. (They were dropped at a late stage from the forthcoming ‘C++0x’ language standard.)

The STL is perhaps the best known instantiation of generic programming. Indeed, some writers have taken the STL style as the definition of generic programming; for example, Siek et al. [68] define generic programming as ‘a methodology for program design and implementation that separates data structures and algorithms through the use of abstract requirement specifications’. We feel that this is squandering a useful term on a well-established existing practice, namely good old-fashioned abstraction.

### 3.3 Metaprogramming

Another interpretation of the term ‘generic programming’ covers various flavours of *metaprogramming*, that is, the construction of programs that write or manipulate other programs. This field encompasses *program generators* such as lex and yacc, *reflection techniques* allowing a program (typically in a dynamically-typed language) to observe and possibly modify its structure and behaviour [45], *generative programming* for the automated customization, configuration and assembly of components [14], and *multi-stage programming* for partitioning computations into phases [70]. A compiler could be considered a generative metaprogram. Rather than writing machine code directly, the programmer writes meta-machine code in a high-level language, and leaves the generation of the machine code itself to the compiler.

In fact, the C++ template mechanism is surprisingly expressive, and already provides some kind of metaprogramming facility. Template instantiation takes place at compile time, so one can think of a C++ program with templates as a two-stage computation; several high-performance numerical libraries rely on templates’ generative properties [74]. The template instantiation mechanism turns out to be Turing complete; Unruh [72] demonstrated the disquieting possibility of a program whose compilation yields the prime numbers as error messages, Czarnecki and Eisenecker [14] show the Turing-completeness of the template mechanism by implementing a rudimentary Lisp interpreter as a template meta-program, and Alexandrescu [2] presents a tour-de-force of unexpected applications of templates.

### 3.4 Datatype genericity

The *Datatype-Generic Programming* project at Oxford and Nottingham [28] has yet another interpretation of the term ‘generic programming’. As the name suggests, *datatype-generic* programs are programs parametrized by a datatype or type functor. We motivated above the type declaration of the following datatype-generic ‘fold’ on multiple datatypes:

$$\begin{aligned} \text{fold} &:: \text{Bifunctor } s \Rightarrow (s \ a \ b \rightarrow b) \rightarrow \text{Fix } s \ a \rightarrow b \\ \text{fold } f &= f \cdot \text{bimap } \text{id} \ (\text{fold } f) \cdot \text{out} \end{aligned}$$

We explain the definition of this *fold* in Section 4; we stress that the function is parametrized not only by the type  $a$  of the elements of the datatype and the body  $f$  of the fold, but also by the shape  $s$  of the datatype itself.

One approach to datatype genericity is what is variously called *polytypism* [41], *structural polymorphism* [67] or *typecase* [76, 59], in which functions are defined inductively by case analysis on the structure of datatypes. This is the kind of genericity provided by Generic Haskell [37]. For example, here is a datatype-generic definition of encoding to a list of bits.

```

type Encode{*}      t = t → [Bool]
type Encode{k → l} t = ∀a. Encode{k} a → Encode{l} (t a)

encode{t :: k} :: Encode{k} t
encode{Char} c           = encodeChar c
encode{Int} n            = encodeInt n
encode{Unit} unit       = []
encode{: + :} ena enb (Inl a) = False : ena a
encode{: + :} ena enb (Inr b) = True  : enb b
encode{: * :} ena enb (a : * : b) = ena a ++ enb b

```

The generic function *encode* generates a list of booleans for any type constructed from characters and integers using sums and products. The cases for base types, sums, and products are given explicitly: by calls to *encodeChar* and *encodeInt* for those types, yielding the empty list for the unit type, prepending a boolean to indicate which injection into a sum, and concatenating the encodings of the two components of a product. The ‘obvious’ cases for type abstraction, application and recursion are generated automatically. Note that *encode* does not have a constant type, but a parametrized one; for example, the instance for products takes as two additional arguments the encoders for the two components, whereas the instance for characters needs no such arguments—*type-indexed values have kind-indexed types* [35].

Because a structurally polymorphic definition is given by case analysis in this fashion, it is an *ad-hoc* form of datatype genericity. One has the flexibility to define different behaviour in different branches, and maybe even to customize the behaviour for specific types; but consequently, there is no guarantee or check that the behaviours in different branches conform, except by type. In contrast, the definition of *fold* cited above and explained below uses a *parametric* form of datatype genericity; one has less flexibility, but instances at different types necessarily behave uniformly. Ad-hoc datatype genericity is more general than parametric; for example, it is difficult to see how to define datatype-generic encoding parametrically, and conversely, any parametric definition can be expanded into an ad-hoc one. However, parametric datatype genericity offers better prospects for reasoning, and is to be preferred when it is applicable. We consider parametric datatype genericity to be the ‘gold standard’, and in the remainder of this paper, we concentrate on parametric datatype-generic definitions where possible.

Datatype genericity is different from the other three interpretations of generic programming outlined above. It is not just a matter of parametric polymorphism, at least not in a straightforward way; for example, parametric polymorphism abstracts from the occurrence of ‘integer’ in ‘lists of integers’, whereas datatype genericity abstracts from the occurrence of ‘list’. It is not just interface conformance, as with concept satisfaction in the STL; although the latter allows *abstraction from* the shape of data, it does not allow *exploitation of* the shape of data, as required for the data compression and marshalling examples above. Finally, it is not metaprogramming: although some flavours of metaprogramming (such as reflection) can simulate datatype-generic computations, they typically do so at the expense of static checking.

## 4 Origami programming

There is a branch of the mathematics of program construction devoted to the relationship between the structure of programs and the structure of the data they manipulate [50, 52, 7, 22]. We saw a glimpse of this field in Sections 2.3 and 2.4, with the definitions of *foldL* and *foldB* respectively: the structure of each

program reflects that of the datatype it traverses, for example in the number of clauses and the number and position of recursive references. In this section, we explore a little further. Folds are not the only program structure that reflects data structure, although they are often given unfair emphasis [30]; we outline *unfolds* and *builds* too, which are two kinds of dual (producing structured data rather than consuming it), and *maps*, which are special cases of these operators, and some simple combinations of these. The beauty of all of these patterns of computation is the direct relationship between their shape and that of the data they manipulate [39]; we go on to explain how both can be parametrized by that shape, yielding *datatype-generic* patterns of computation. Elsewhere, we have called this approach *origami programming* [23].

## 4.1 Maps on lists

Here is the polymorphic datatype of lists again.

```
data List a = Nil | Cons a (List a)
```

The ‘map’ operator for a datatype applies a given function to every element of a data structure. The (higher-order, polymorphic, but list-specific) map operator for lists is given by:

```
mapL :: (a → b) → List a → List b
mapL f Nil           = Nil
mapL f (Cons x xs) = Cons (f x) (mapL f xs)
```

Thus, mapping over the empty list yields the empty list; mapping  $f$  over a non-empty list applies  $f$  to the head and maps over the tail.

## 4.2 Folds on lists

The ‘fold’ operator for a datatype collapses a data structure down to a value. Here is the (again higher-order, polymorphic, but list-specific) fold operator for lists that we saw in Section 2.3.

```
foldL :: b → (a → b → b) → List a → b
foldL e f Nil           = e
foldL e f (Cons x xs) = f x (foldL e f xs)
```

For example, the function *filterL* (itself higher-order, polymorphic, but list-specific) takes a predicate  $p$  and a list  $xs$ , and returns the sublist of  $xs$  consisting of those elements that satisfy  $p$ .

```
filterL :: (a → Bool) → List a → List a
filterL p = foldL Nil (add p) where
  add p x xs = if p x then Cons x xs else xs
```

As we saw in Section 2.3, the functions *sum*, *append*, and *concat* are also instances of *foldL*.

## 4.3 Unfolds on lists

The ‘unfold’ operator for a datatype grows a data structure from a value. In a precise technical sense, it is the dual of the ‘fold’ operator. That duality is not so evident in the implementation for lists below, but it will become clearer with the datatype-generic version we present in Section 4.11.

```
unfoldL :: (b → Bool) → (b → a) → (b → b) → b → List a
unfoldL p f g x = if p x then Nil else Cons (f x) (unfoldL p f g (g x))
```

The unfold operates on a seed  $x$ ; if that seed satisfies  $p$ , the empty list is returned, and if not, a non-empty list is constructed, with head  $f x$  and tail generated recursively by unfolding an updated seed  $g x$ . For example, here are two instances. The function *preds* returns the list of predecessors of a (presumed non-negative) integer; the function *takeWhile* takes a predicate  $p$  and a list  $xs$ , and returns the longest initial segment of  $xs$  all of whose elements satisfy  $p$ .



```

preds :: Integer → List Integer
preds = unfoldL (0 ≡) id pred where
  pred n = n - 1

takeWhile :: (a → Bool) → List a → List a
takeWhile p = unfoldL (firstNot p) head tail where
  firstNot p Nil = True
  firstNot p (Cons x xs) = not (p x)

```

## 4.4 Origami for binary trees

We might go through a similar exercise for a datatype of internally labelled binary trees.

```

data Tree a = Empty | Node a (Tree a) (Tree a)

```

The ‘map’ operator applies a given function to every element of a tree.

```

mapT :: (a → b) → Tree a → Tree b
mapT f Empty = Empty
mapT f (Node x xs ys) = Node (f x) (mapT f xs) (mapT f ys)

```

The ‘fold’ operator collapses a tree down to a value.

```

foldT :: b → (a → b → b → b) → Tree a → b
foldT e f Empty = e
foldT e f (Node x xs ys) = f x (foldT e f xs) (foldT e f ys)

```

For example, the function *inorder* collapses a tree down to a list.

```

inorder :: Tree a → List a
inorder = foldT Nil glue
glue x xs ys = append xs (Cons x ys)

```

The ‘unfold’ operator grows a tree from a value.

```

unfoldT :: (b → Bool) → (b → a) → (b → b) → (b → b) → b → Tree a
unfoldT p f g h x = if p x then Empty
                  else Node (f x) (unfoldT p f g h (g x)) (unfoldT p f g h (h x))

```

For example, the Calkin–Wilf tree, the first few levels of which are illustrated in Figure 2, contains each of the positive rationals exactly once [1, 31]:

```

cwTree :: Tree Rational
cwTree = unfoldT (const False) frac left right (1, 1) where
  frac (m, n) = m/n
  left (m, n) = (m, m + n)
  right (m, n) = (n + m, n)

```

(Here, *const a* is the function that always returns *a*.)

## 4.5 Aside: ad-hoc polymorphism

For another example of an unfold to trees, consider the function *grow* that generates a binary search tree from a list of elements whose type supports an ordering.

```

grow :: Ord a ⇒ List a → Tree a
grow = unfoldT isNil head littles bigs

```




Figure 2: The first few levels of the Calkin–Wilf tree

```

littles (Cons x xs) = filterL (≤ x) xs
bigs (Cons x xs) = filterL (not · (≤ x)) xs

```

The ‘*Ord a ⇒*’ is a Haskell *type class context*, expressing a kind of *ad-hoc polymorphism*: the function *grow* is defined only for those element types *a* that support an ordering. Element ordering is ad hoc in the sense that the  $\leq$  operator (and hence also *grow*) is not defined for all element types *a*; moreover, for those types on which  $\leq$  is defined, it may be defined in quite different ways. In contrast, a *parametrically polymorphic* type, such as the type *Tree a → List a* of the function *inorder* in Section 4.4, is defined for all element types *a*; moreover, it is defined ‘in the same way’ for all types, so that there is a *parametricity* property stating coherence between the different instances [77]. A similar effect can be achieved using OO generic classes [58]. (In fact, Haskell’s type classes provide an alternative to ad-hoc polymorphism that does not involve sacrificing parametricity: the dictionary-passing translation [34] allows programs depending on a type constraint such as ordering still to be defined parametrically, by extracting the necessary operations from a supplied dictionary; the coherence property of the whole program becomes conditional a similar coherence property for the operations on which it depends.) The type class *Ord* might be defined in Haskell as follows:

```

class Ord a where
  (≤) :: a → a → Bool

```

(In fact, the definition in the standard library is more complex than this; but this will serve for illustration.) Various types are instance of the type class, by virtue of supporting a comparison operation:

```

instance Ord Integer where
  (m ≤ n) = isNonNegative (n - m)

```

We explain the type class mechanism here, because we will use it again later.

## 4.6 Hylomorphisms

An unfold followed by a fold is a common pattern of computation [52]; the unfold generates a data structure, and the fold immediately consumes it. For example, here is a (higher-order, polymorphic, but list-specific) hylomorphism operator for lists, and an instance for computing factorials: first generate the predecessors of the input using an unfold, then compute the product of these predecessors using a fold.

```

hyloL :: (b → Bool) → (b → a) → (b → b) → c → (a → c → c) → b → c
hyloL p f g e h = foldL e h · unfoldL p f g

fact :: Integer → Integer
fact = hyloL (0 ≡) id pred 1 (*)

```

With lazy evaluation, the intermediate data structure (the list of predecessors, for *fact*) is not computed all at once. It is produced on demand, and each demanded cell consumed immediately. In fact, the intermediary can be *deforested* altogether—the following equation serves as an equivalent definition of *hyloL*.

```

hyloL :: (b → Bool) → (b → a) → (b → b) → c → (a → c → c) → b → c
hyloL p f g e h x = if p x then e else h (f x) (hyloL p f g e h (g x))

```

A similar definition can be given for binary trees:

$$\begin{aligned} \text{hyloT} &:: (b \rightarrow \text{Bool}) \rightarrow (b \rightarrow a) \rightarrow (b \rightarrow b) \rightarrow (b \rightarrow b) \rightarrow c \rightarrow (a \rightarrow c \rightarrow c \rightarrow c) \rightarrow b \rightarrow c \\ \text{hyloT } p \ f \ g_1 \ g_2 \ e \ h \ x &= \mathbf{if } p \ x \ \mathbf{then } e \\ &\quad \mathbf{else } h \ (f \ x) \ (\text{hyloT } p \ f \ g_1 \ g_2 \ e \ h \ (g_1 \ x)) \\ &\quad \quad (\text{hyloT } p \ f \ g_1 \ g_2 \ e \ h \ (g_2 \ x)) \end{aligned}$$

together with an instance giving a kind of quicksort—albeit not a very quick one, since it is not in-place, it has a space leak, and it takes quadratic time:

$$\begin{aligned} \text{qsort} &:: \text{Ord } a \Rightarrow \text{List } a \rightarrow \text{List } a \\ \text{qsort} &= \text{hyloT } \text{isNil } \text{head } \text{littles } \text{bigs } \text{Nil } \text{glue} \end{aligned}$$

## 4.7 Short-cut fusion

Unfolds capture a highly structured pattern of computation for generating recursive data structures. There exist slight generalizations of unfolds, such as *monadic unfolds* [63, 64], *apomorphisms* [75] and *futurmorphisms* [73], but these still all conform to the same structural scheme, and not all programs that generate data structures fit this scheme [29]. Gill et al. [33] introduced an operator they called *build* for unstructured generation of data, in order to simplify the implementation and broaden the applicability of deforestation optimizations as discussed in Section 4.6.

The idea behind *build* is to allow the identification of precisely where in a program the nodes of a data structure are being generated; then it is straightforward for a compiler to fuse a following fold, inlining functions to replace those constructors and deforesting the data structure altogether. The definition of *build* is reminiscent of the continuation-passing style [15] beloved of LISP and Scheme programmers; it takes as argument a program with ‘holes’ for constructors, and plugs those holes with actual constructors.

$$\begin{aligned} \text{buildL} &:: (\forall b. b \rightarrow (a \rightarrow b \rightarrow b) \rightarrow b) \rightarrow \text{List } a \\ \text{buildL } g &= g \ \text{Nil } \text{Cons} \end{aligned}$$

The function *buildL* has a rank-two type; the argument *g* must be parametrically polymorphic in the constructor arguments, in order to ensure that all uses of the constructors are abstracted. We argued above that *unfoldL* is a dual to *foldL* in one sense; we will make that sense clear in Section 4.11. In another sense, *buildL* is *foldL*’s dual: whereas the fold deletes constructors and replaces them with something else, the build inserts those constructors.

The beauty of the idea is that fusion with a following fold is simple to state:

$$\text{foldL } e \ f \ (\text{buildL } g) = g \ e \ f$$

Perhaps more importantly, this fusion rule is also easy for a compiler to exploit [33].

Build operators are strictly more expressive than unfolds. For instance, it is possible to define *unfoldL* in terms of *buildL*:

$$\begin{aligned} \text{unfoldL} &:: (b \rightarrow \text{Bool}) \rightarrow (b \rightarrow a) \rightarrow (b \rightarrow b) \rightarrow b \rightarrow \text{List } a \\ \text{unfoldL } p \ f \ g \ b &= \text{buildL } (h \ b) \ \mathbf{where} \\ &\quad h \ b \ n \ c = \mathbf{if } p \ b \ \mathbf{then } n \ \mathbf{else } c \ (f \ b) \ (h \ (g \ b) \ n \ c) \end{aligned}$$

However, some functions that generate lists can be expressed as an instance of *buildL* and not of *unfoldL*; one such example is the function that computes the infinite list of multiples of an integer [29]:

$$\begin{aligned} \text{mults} &:: \text{Int} \rightarrow \text{List } \text{Int} \\ \text{mults } n &= \text{buildL } (\text{next } n \ 0) \ \mathbf{where} \\ \text{next} &:: \text{Int} \rightarrow \text{Int} \rightarrow b \rightarrow (\text{Int} \rightarrow b \rightarrow b) \rightarrow b \\ \text{next } i \ j \ n \ c &= c \ j \ (\text{next } i \ (i + j) \ n \ c) \end{aligned}$$

The disadvantage of *buildL* compared to *unfoldL* is a consequence of its unstructured approach: the former does not support the powerful *universal properties* that greatly simplify program calculation with the latter [22].

Of course, there is nothing special about lists in this regard. One can define build operators for any datatype—here is one for trees:

```
buildT :: (∀b. b → (a → b → b → b) → b) → Tree a
buildT g = g Empty Node
```

## 4.8 Datatype genericity

As we have already seen, data structure determines program structure. It therefore makes sense to abstract from the determining shape, leaving only what they have in common. We do this by defining a datatype *Fix*, parametrized both by an element type *a* of basic kind (a plain type, such as integers or strings), and by a shape type *s* of higher kind (a type constructor, such as ‘pairs of’ or ‘lists of’).

```
data Fix s a = In (s a (Fix s a))
out :: Fix s a → s a (Fix s a)
out (In x) = x
```

Equivalently, we could use a record type with a single named field, and define both the constructor *In* and the destructor *out* at once.

```
data Fix s a = In { out :: s a (Fix s a) }
```

## 4.9 Specific datatypes

The parameter *s* determines the shape; ‘*Fix*’ ties the recursive knot. Here are three instances of *Fix* using different shapes, for lists and for externally and internally labelled binary trees.

```
data ListF a b = NilF | ConsF a b
type List a = Fix ListF a

data BtreeF a b = TipF a | BinF b b
type Btree a = Fix BtreeF a

data TreeF a b = EmptyF | NodeF a b b
type Tree a = Fix TreeF a
```

Note that *Fix s a* is a recursive type. Typically, as in the three instances above, the shape *s* has several variants, including a ‘base case’ independent of its second argument. But with lazy evaluation, infinite structures are possible, and so the definition makes sense even with no such base case. For example, here is a type of infinite internally-labelled binary trees (which would suffice for the *cwTree* example in Section 4.4):

```
data ITreeF a b = INodeF a b b
type Fix ITreeF a
```

## 4.10 Bifunctors

Not all binary type constructors *s* are suitable for *Fixing*; for example, function types with the parameter appearing in *contravariant* (source) positions cause problems. It turns out that we should restrict attention to (covariant) *bifunctors*, which support a *bimap* operation ‘locating’ all the elements. We capture this constraint as a type class.

```
class Bifunctor s where
  bimap :: (a → c) → (b → d) → s a b → s c d
```

(For ‘locating the elements’ in an  $s$ -structure, one might initially think of an operation that returns two collections of ‘pointers’, one for each of the two kinds of element. Instead of pointers, we take a higher-order approach: the operation *bimap* accepts an  $s$ -structure and two functions, one for each kind of element, and returns a new  $s$ -structure in which each element has been subjected to the appropriate function. That is, instead of producing elements, *bimap* consumes element-consumers.) Technically speaking, *bimap* should satisfy some properties:

$$\begin{aligned} \mathit{bimap} \ \mathit{id} \ \mathit{id} &= \mathit{id} \\ \mathit{bimap} \ f \ g \cdot \mathit{bimap} \ h \ j &= \mathit{bimap} \ (f \cdot h) \ (g \cdot j) \end{aligned}$$

—properties which cannot be formally stated in Haskell, but which we might expect to be able to express in the languages of tomorrow [12, 71]. Datatype-genericity is thus a kind of constrained type-constructor polymorphism.

All datatypes made from sum and product constructors induce bifunctors. Here are instances for our three example shapes.

```
instance Bifunctor ListF where
  bimap f g NilF      = NilF
  bimap f g (ConsF x y) = ConsF (f x) (g y)

instance Bifunctor BtreeF where
  bimap f g (TipF x)  = TipF (f x)
  bimap f g (BinF y z) = BinF (g y) (g z)

instance Bifunctor TreeF where
  bimap f g EmptyF    = EmptyF
  bimap f g (NodeF x y z) = NodeF (f x) (g y) (g z)
```

The type signature of the operator *bimap* is datatype-generic, since it is parameterized by the shape  $s$  of the data:

$$\mathit{bimap} :: \text{Bifunctor } s \Rightarrow (a \rightarrow c) \rightarrow (b \rightarrow d) \rightarrow s \ a \ b \rightarrow s \ c \ d$$

However, because *bimap* is encoded as a member function of a type class, the definitions for particular shapes are examples of ad-hoc rather than parametric datatype genericity; each instance entails a proof obligation that the appropriate laws are satisfied.

It is a bit tedious to have to provide a new instance of *Bifunctor* for each new datatype shape; one would of course prefer a single datatype-generic definition. This is the kind of feature for which research languages like Generic Haskell [37] are designed, and one can almost achieve the same effect in Haskell [11, 59, 36]. One might hope that these instance definitions could in fact be inferred, in the languages of tomorrow [38]. But whatever the implementation mechanism, the result will still be ad-hoc datatype-generic: it is necessarily the case that different code is used to locate the elements within data of different shapes.

## 4.11 Datatype-generic recursion patterns

The datatype-specific recursion patterns introduced above can be made generic in the shape  $s$ , provided that this is a bifunctor.

```
map :: Bifunctor s => (a -> b) -> Fix s a -> Fix s b
map f = In . bimap f (map f) . out

fold :: Bifunctor s => (s a b -> b) -> Fix s a -> b
fold f = f . bimap id (fold f) . out

unfold :: Bifunctor s => (b -> s a b) -> b -> Fix s a
unfold f = In . bimap id (unfold f) . f

hylo :: Bifunctor s => (b -> s a b) -> (s a c -> c) -> b -> c
hylo f g = g . bimap id (hylo f g) . f
```

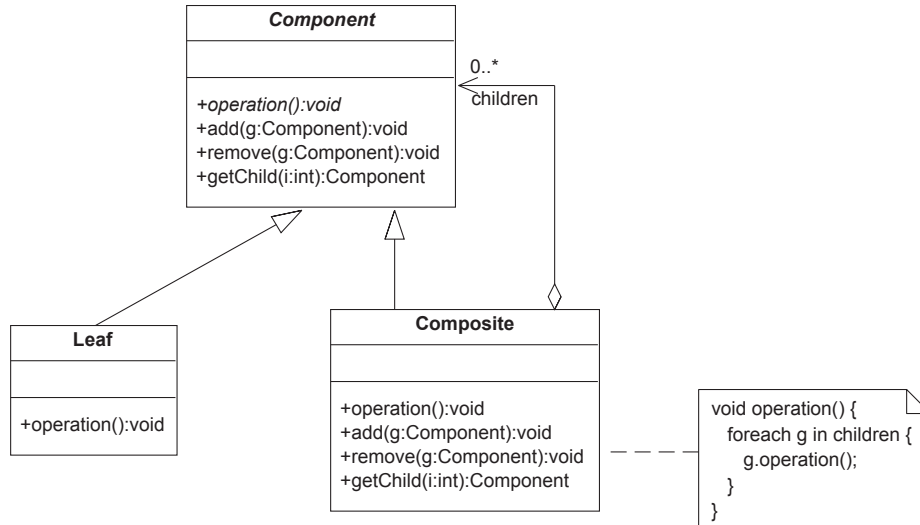


Figure 3: The class structure of the COMPOSITE pattern

$build :: Bifunctor\ s \Rightarrow (\forall b. (s\ a\ b \rightarrow b) \rightarrow b) \rightarrow Fix\ s\ a$   
 $build\ f = f\ In$

The datatype-generic definitions are surprisingly short—shorter even than the datatype-specific ones. The structure becomes much clearer with the higher level of abstraction. In particular, the duality between *fold* and *unfold* is evident from the code: compositions are reversed, and *Ins* and *outs* exchanged. .

For more about origami programming, see [22, 23].

## 5 Origami patterns

In this section we describe ORIGAMI, a little suite of patterns for recursive data structures, consisting of four of the Gang of Four design patterns [20]:

- COMPOSITE, for modelling recursive structures;
- ITERATOR, for linear access to the elements of a composite;
- VISITOR, for structured traversal of a composite; and
- BUILDER, to generate a composite structure.

These four patterns belong together. They all revolve around the notion of a hierarchical structure, represented as a COMPOSITE. One way of constructing such hierarchies is captured by the BUILDER pattern: a client application knows what kinds of part to add and in what order, but it delegates to a separate object knowledge of the implementation of the parts and responsibility for creating and holding them. Having constructed a hierarchy, there are two kinds of traversal we might perform over it: either considering it as a container of elements, in which case we use an ITERATOR for a linear traversal; or considering its shape as significant, in which case we use a VISITOR for a structured traversal.

### 5.1 Composite

The COMPOSITE pattern ‘lets clients treat individual objects and compositions of objects uniformly’, by ‘composing objects into tree structures’ [20]. The crux of the pattern is a common supertype (class *Composite* in Figure 3), of which both atomic (class *Leaf*) and aggregated (class *Composite*) objects are subtypes.

### 5.2 Iterator

The ITERATOR pattern ‘provides a way to access the elements of an aggregate object sequentially without exposing its underlying representation’ [20]. It does this by separating the responsibilities of containment

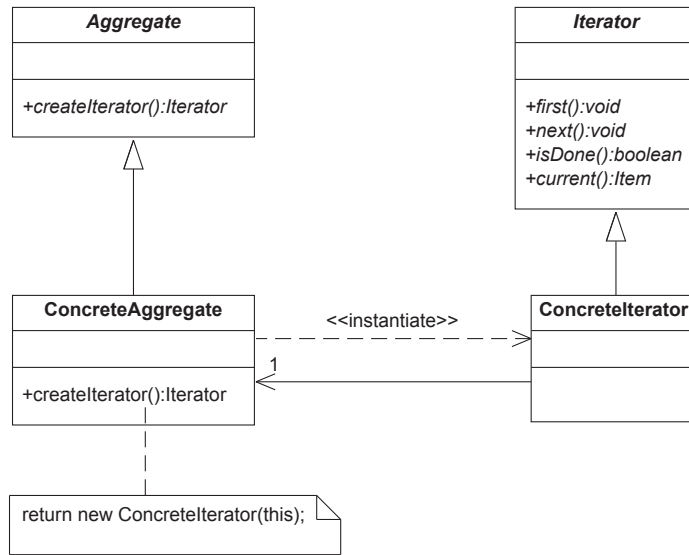


Figure 4: The class structure of the EXTERNAL ITERATOR pattern

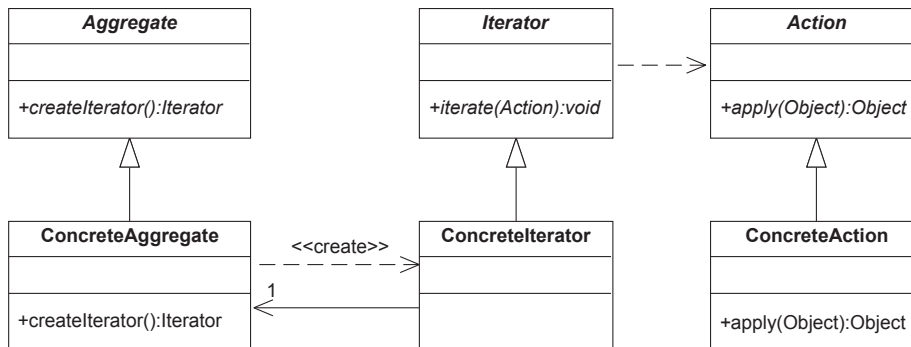


Figure 5: The class structure of the INTERNAL ITERATOR pattern

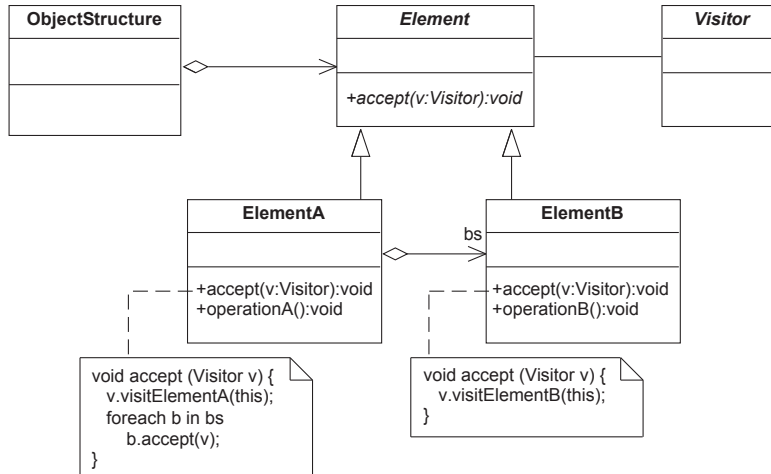


Figure 6: The class structure of the elements in the INTERNAL VISITOR pattern

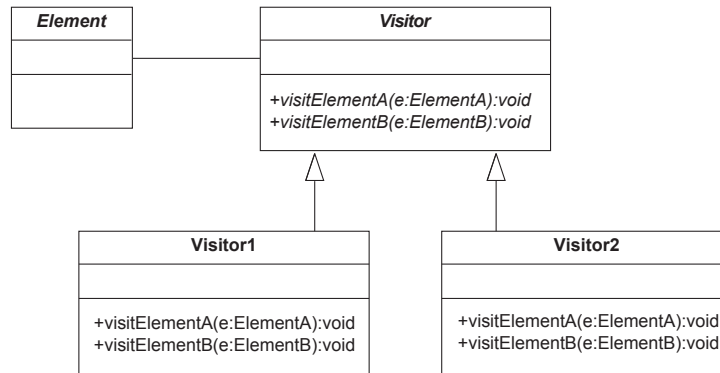


Figure 7: The class structure of the visitors in the INTERNAL VISITOR pattern

(class *Aggregate* in Figure 4) and iteration (class *Iterator*). The standard implementation is as an *external* or client-driven iterator, illustrated in Figure 4 and as embodied for example in the Java standard library.

In addition to the standard implementation, GOF also discuss *internal* or iterator-driven ITERATORS, illustrated in Figure 5. These might be modelled in Java by the following pair of interfaces:

```

public interface Action { Object apply (Object o);}
public interface Iterator { void iterate (Action a);}
  
```

An object implementing the *Action* interface provides a single method *apply*, which takes in a collection element and returns (either a new, or the same but modified) element. (The C++ STL calls such objects ‘functors’, but we avoid that term here to prevent name clashes with type functors.) A collection (implements a FACTORY METHOD to return a separate subobject that) implements the *Iterator* interface to accept an *Action*, apply it to each element in turn, and replace the original elements with the possibly new ones returned. Internal ITERATORS are less flexible than external—for example, it is more difficult to have two linked iterations over the same collection, and to terminate an iteration early—but they are correspondingly simpler to use.

### 5.3 Visitor

In the normal object-oriented paradigm, the definition of each traversal operation is spread across the whole class hierarchy of the structure being traversed—typically but not necessarily a COMPOSITE. This makes it



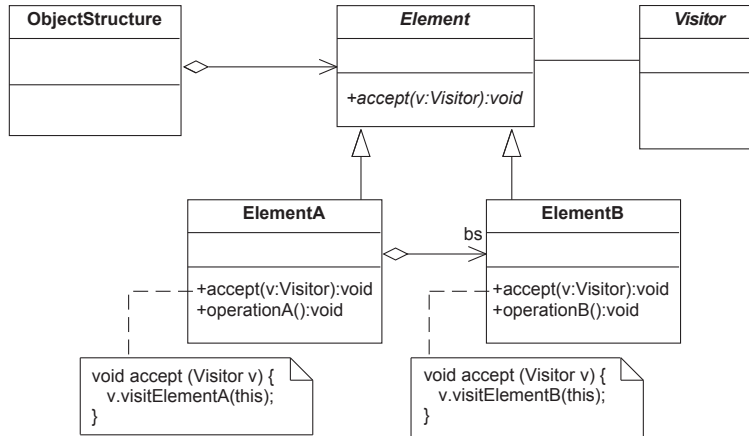


Figure 8: The class structure of the elements in the EXTERNAL VISITOR pattern

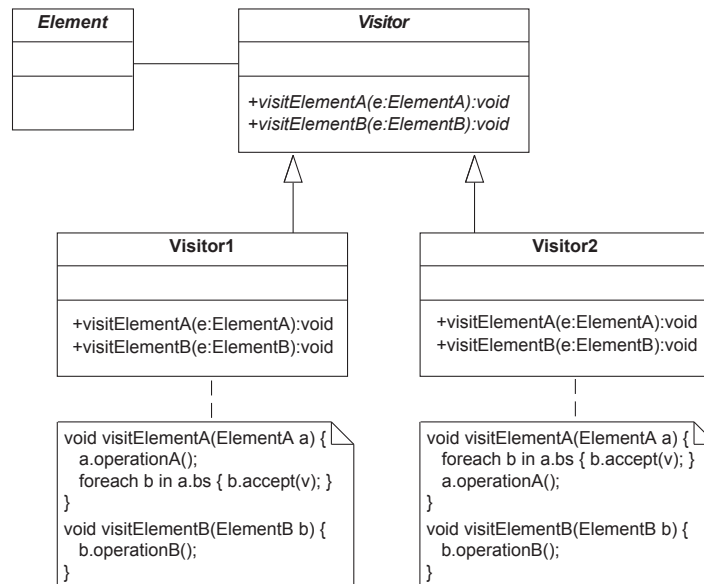


Figure 9: The class structure of the visitors in the EXTERNAL VISITOR pattern

easy to add new variants of the datatype (for example, new kinds of leaf node in the COMPOSITE), but hard to add new traversal operations.

The VISITOR pattern ‘represents an operation to be performed on the elements of an object structure’, allowing one to ‘define a new operation without changing the classes of the elements on which it operates’ [20]. This is achieved by providing a hook for associating new traversals (method *accept* in Figure 6), and an interface for those traversals to implement (interface *Visitor* in Figure 7). The effect is to simulate *double dispatch* on the types of two arguments, the element type and the operation, by two consecutive single dispatches. It is a kind of *aspect-oriented programming* [46], modularizing what would otherwise be a cross-cutting concern. It reverses the costs: it is now easy to add new traversals, but hard to add new variants. (Wadler [80] has coined the term *expression problem* for this tension between dimensions of easy extension.)

As with the distinction between internal and external iterators, there is a choice about where to put responsibility for managing a traversal. Buchlovsky and Thielecke [9] use the term ‘INTERNAL VISITOR’ for the usual presentation, with the *accept* methods of *Element* subclasses making recursive calls as shown in Figure 6. Moving that responsibility from the *accept* methods of the *Element* classes to the *visit* methods

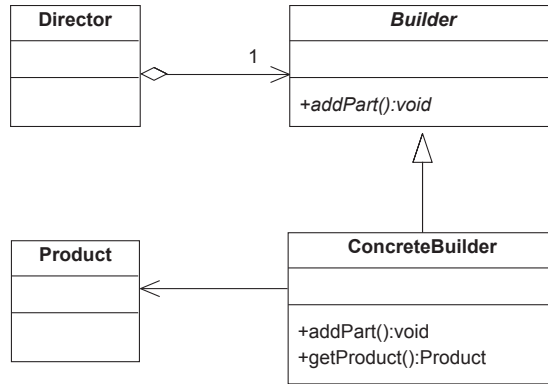


Figure 10: The class structure of the BUILDER pattern

of the *Visitor* classes, as shown in Figures 8 and 9, yields what they call an EXTERNAL VISITOR. Now the traversal algorithm is not fixed, and different visitors may vary it (for example, between preorder and postorder). One might say that this latter variation encapsulates simple case analysis or pattern matching, rather than traversals per se.

## 5.4 Builder

Finally, the BUILDER pattern ‘separates the construction of a complex object from its representation, so that the same construction process can create different representations’ [20]. As Figure 10 shows, this is done by delegating responsibility for the construction to a separate object—in fact, a STRATEGY for performing the construction.

The GOF motivating example of the BUILDER pattern involves assembling a product that is basically a simple collection; that is necessarily the case, because the operations supported by a builder object take just a part and return no result. However, GOF also suggest the possibility of building a more structured product, in which the parts are linked together. For example, to construct a tree, each operation to add a part could return a unique identifier for the part added, and take an optional identifier for the parent to which to add it; a directed acyclic graph requires a set of parents for each node, and construction in topological order; a cyclic graph requires the possibility of ‘forward references’, adding parts as children of yet-to-be-added parents.

GOF also suggest the possibility of BUILDERS that compute. Instead of constructing a large *Product* and eventually collapsing it, one can provide a separate implementation of the *Builder* interface that makes the *Product* itself the collapsed result, computing it on the fly while building.

## 5.5 An example

As an example of applying the ORIGAMI patterns, consider the little document system illustrated in Figure 11. (The complete code is given in an appendix, for reference.)

- The focus of the application is a COMPOSITE structure of documents: *Sections* have a *title* and a collection of sub-*Components*, and *Paragraphs* have a *body*.
- One can iterate over such a structure using an INTERNAL ITERATOR, which acts on every *Paragraph*. For instance, iterating with a *SpellCorrector* might correct the spelling of every paragraph body. (For brevity, we have omitted the possibility of acting on the *Sections* of a document, but it would be easy to extend the *Action* interface to allow this. We have also simplified the pattern from the presentation in Section 5.2, making the *apply* method return *void*, and so providing no way to change the identity of the document elements; more generally, *apply* could optionally return new elements.)

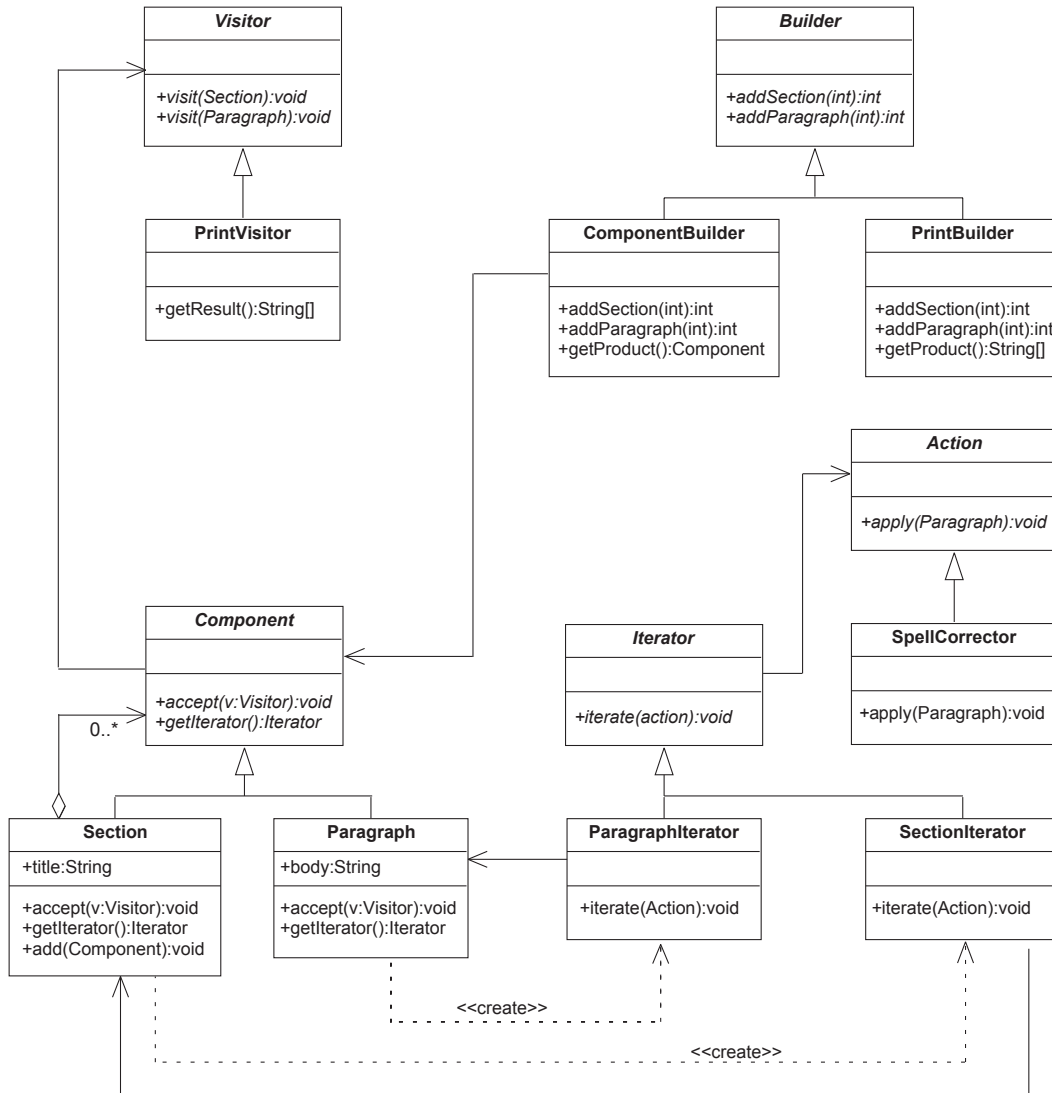


Figure 11: An application of the ORIGAMI patterns

- One can also traverse the document structure with a VISITOR, for example to compute some summary of the document. For instance, a *PrintVisitor* might yield a string array with the section titles and paragraph bodies in order.
- Finally, one can construct such a document using a BUILDER. We have used the structured variant, adding *Sections* and *Paragraphs* as children of existing *Components* via unique integer identifiers. A *ComponentBuilder* constructs a *Component* as expected, whereas a *PrintBuilder* incorporates the printing behaviour of the *PrintVisitor* incrementally, actually constructing a string array instead.

This one application is a paradigmatic example of each of the four ORIGAMI patterns. We therefore claim that any alternative representation of the patterns cleanly capturing this application structure is a faithful rendition of (at least the code parts of) those patterns. In Section 6 below, we provide just such a representation, in terms of the higher-order datatype-generic programs from Section 4. Section 6.5 justifies our claim of a faithful rendition by capturing the structure of the document application in this alternative representation.

## 6 Patterns as HODGPs

We now revisit the ORIGAMI patterns from Section 5, showing that each of the four patterns can be captured as a higher-order datatype-generic program (HODGP). However, we consider them in a slightly different order; it turns out that the datatype-generic representation of the ITERATOR pattern builds on that of VISITOR.

### 6.1 Composite in HODGP

COMPOSITES are recursive data structures; in the OO setting, they are packaged together with some operations, but in a functional setting the operations are represented separately. So actually, these correspond not to programs, but to types. Recursive data structures come essentially for free in functional programming languages.

```
data Fix s a = In{ out :: s a (Fix s a) }
```

What is datatype-generic about this definition is that it is parametrized by the shape  $s$  of the data structure; thus, one recursive datatype serves to capture *all* (regular) recursive data structures, whatever their shape.

### 6.2 Visitor in HODGP

The VISITOR pattern collects fragments of each traversal into one place, and provides a hook for performing such traversals. The resulting style matches the normal functional programming paradigm, in which traversals are entirely separate from the data structures traversed. No explicit hook is needed; the connection between traversal and data is made within the traversal by dispatching on the data, either by pattern matching or (equivalently) by applying a destructor. What was a double dispatch in the OO setting becomes in HODGP the choice of a function to apply, followed by a case analysis on the variant of the data structure. A common case of such traversals, albeit not the most general, is the fold operator introduced above.

```
fold :: Bifunctor s => (s a b -> b) -> Fix s a -> b
fold f = f . bimap id (fold f) . out
```

This too is datatype-generic, parametrized by the shape  $s$ : the same function *fold* suffices to traverse any shape of COMPOSITE structure.

For a detailed study of an implementation of the VISITOR pattern that is generic both in the datatype and along various ‘strategy’ dimensions such as between internal and external control, see [61].

### 6.3 Iterator in HODGP

EXTERNAL ITERATORS give sequential access to the elements of collection. The functional approach would be to provide a view of the collection as a list of elements, at least for read-only access. Seen this way, the ITERATOR pattern can be implemented using the VISITOR pattern, traversing using a body *combiner* that combines the element lists from substructures into one overall element list.

```
contents :: Bifunctor s => (s a (List a) -> List a) -> Fix s a -> List a
contents combiner = fold combiner
```

With lazy evaluation, the list of elements can be generated incrementally on demand, rather than eagerly in advance: ‘lazy evaluation means that lists and iterators over lists are identified’ [81].

In the formulation above, the *combiner* argument has to be provided to the *contents* operation. Passing different *combiners* allows the same COMPOSITE to yield its elements in different orders; for example, a tree-shaped container could support both preorder and postorder traversal. On the other hand, it is clumsy always to have to specify the *combiner*. One could specify it once and for all, in the class *Bifunctor*, in effect making it another datatype-generic operation parametrized by the shape *s*. In the languages of tomorrow, one might expect that at least one, obvious implementation of *combiner* could be inferred automatically.

Of course, some aspects of external ITERATORS can already be expressed linguistically; the interface *java.util.Iterator* has been available for years in the Java API, the iterator concept has been explicit in the C++ Standard Template Library for even longer, and recent versions of Java and C# even provide language support (the **foreach** statement in C#) for iterating over the elements yielded by such an operator. Thus, element consumers can be written datatype-generically today. But still, one has to implement the *Iterator* anew for each datatype defined; element producers are still datatype-specific.

An internal ITERATOR is basically a map operation, iterating over a collection and yielding another collection of the same shape but with different or modified elements; it therefore supports write access to the collection as well as read access. In HODGP, we can give a *single generic* definition of this.

```
map :: Bifunctor s => (a -> b) -> Fix s a -> Fix s b
map f = In · bimap f (map f) · out
```

This is in contrast with the object-oriented approach, in which *Iterator* implementations are datatype-specific. Note also that the HODGP version is more general than the OO version, because the elements in the return collection may be of a different type.

Although the internal ITERATOR explains both read and write access to a collection, it does not explain imperative access, with impure aspects such as side-effects, I/O and so on. Moreover, it does not subsume the HODGP external ITERATOR, because it does not allow *accumulation* of some measure of the elements (for example, to compute the size of the collection in passing). Recent work on *idiomatic traversals* [51, 32] overcomes both of these shortcomings: idiomatic traversals support imperative features and mapping and accumulating aspects simultaneously, using *idioms* or *applicative functors*, a slight generalization of monads [79]. One small extra piece of ad-hockery is required: a mechanism for pulling an idiomatic effect out of the shape of a data structure.

```
class Bifunctor s => Bitraversable s where
  bidist :: Idiom m => s (m a) (m b) -> m (s a b)
```

Given this tool, a datatype-generic *traverse* operator turns out to be an instance of *fold*:

```
instance Bitraversable s => Traversable (Fix s) where
  traverse f = fold (fmap In · bidist · bimap f id)
```

Applications of *traverse* include maps, accumulations and imperative iterations over collections [32].

### 6.4 Builder in HODGP

The standard protocol for the BUILDER pattern involves a *Director* sending *Parts* one by one to a *Builder* for it to assemble, and then retrieving from the *Builder* a *Product*. Thus, the product is assembled in a

step-by-step fashion, but is unavailable until assembly is complete. With lazy evaluation, we can in some circumstances construct the *Product* incrementally: we can yield access to the root of the product structure while continuing to assemble its substructures. In the case that the data structure is assembled in a regular fashion, this corresponds in the HODGP style to an unfold operation.

$$\begin{aligned} \text{unfold} &:: \text{Bifunctor } s \Rightarrow (b \rightarrow s \ a \ b) \rightarrow b \rightarrow \text{Fix } s \ a \\ \text{unfold } f &= \text{In} \cdot \text{bimap } \text{id} \ (\text{unfold } f) \cdot f \end{aligned}$$

When the data structure is assembled irregularly, a build operator has to be used instead.

$$\begin{aligned} \text{build} &:: \text{Bifunctor } s \Rightarrow (\forall b. (s \ a \ b \rightarrow b) \rightarrow b) \rightarrow \text{Fix } s \ a \\ \text{build } f &= f \ \text{In} \end{aligned}$$

These are both datatype-generic programs, parametrized by the shape of product to be built. In contrast, the GOF BUILDER pattern states the general scheme, but requires code specific for each *Builder* interface and each *ConcreteBuilder* implementation.

Turning to GOF’s computing builders, with lazy evaluation there is not so pressing a need to fuse building with postprocessing. If the structure of the consumer computation matches that of the producer—in particular, if the consumer is a fold and the producer a build or an unfold—then consumption can be interleaved with production, and the whole product never need be in existence at once.

Nevertheless, naive interleaving of production and consumption of parts of the product still involves the creation and immediate disposal of those parts. Even the individual parts need never be constructed; often, they can be deforested [78], with the attributes of a part being fed straight into the consumption process. When the producer is an unfold, the composition of producer and consumer is (under certain mild strictness conditions [52]) a hylomorphism.

$$\begin{aligned} \text{hylo} &:: \text{Bifunctor } s \Rightarrow (b \rightarrow s \ a \ b) \rightarrow (s \ a \ c \rightarrow c) \rightarrow b \rightarrow c \\ \text{hylo } f \ g &= g \cdot \text{bimap } \text{id} \ (\text{hylo } f \ g) \cdot f \end{aligned}$$

More generally, but less conveniently for reasoning, the producer is a build, and the composition simply replaces the constructors in the builder by the body of the fold.

$$\begin{aligned} \text{foldBuild} &:: \text{Bifunctor } s \Rightarrow (\forall b. (s \ a \ b \rightarrow b) \rightarrow b) \rightarrow (s \ a \ b \rightarrow b) \rightarrow b \\ \text{foldBuild } f \ g &= f \ g \end{aligned}$$

Once again, both definitions are datatype-generic; both take as arguments a producer  $f$  and a consumer  $g$ , both with types parametrized by the shape  $s$  of the product to be built. Note especially that in both cases, the fusion requires no creativity; in contrast, GOF’s computing builders can take considerable insight and ingenuity to program (as we shall see in the appendix).

## 6.5 The example, revisited

To justify our claim that the higher-order datatype-generic representation of the ORIGAMI patterns is a faithful rendition, we use it to re-express the document application discussed in Section 5.5 and illustrated in Figure 11. (It is instructive to compare these 40-odd lines of Haskell code with the equivalent Java code in the appendix.)

- The COMPOSITE structure has the following shape.

```
data DocF a b = Para a | Sec String [b]
type Doc = Fix DocF String

instance Bifunctor DocF where
    bimap f g (Para s) = Para (f s)
    bimap f g (Sec s xs) = Sec s (map g xs)
```

We have chosen to consider paragraph bodies as the ‘contents’ of the data structure, but section titles as part of the ‘shape’; that decision could be varied.

- We used an INTERNAL ITERATOR to implement the *SpellCorrector*; this would be modelled now as an instance of *map*.

```
correct :: String → String -- definition omitted
corrector :: Doc → Doc
corrector = map correct
```

- The use of VISITOR to print the contents of a document is a paradigmatic instance of a *fold*.

```
printDoc :: Doc → [String]
printDoc = fold combine

combine :: DocF String [String] → [String]
combine (Para s) = [s]
combine (Sec s xs) = s : concat xs
```

- Finally, in place of the BUILDER pattern, we can use *unfold* for constructing documents, at least when doing so in a structured fashion. For example, consider the following simple representation of XML trees.

```
data XML = Text String | Entity Tag Attrs [XML]
type Tag = String
type Attrs = [(String, String)]
```

From such an XML tree we can construct a document, with *Text* elements as paragraphs and *Entity*s as sections with appropriate titles.

```
fromXML :: XML → Doc
fromXML = unfold step

step :: XML → DocF String XML
step (Text s) = Para s
step (Entity t kvs xs) = Sec (title t kvs) xs

title :: Tag → Attrs → String
title t [] = t
title t kvs = t ++ paren (join (map attr kvs)) where
  join [s] = s
  join (s : ss) = s ++ ", " ++ join ss
  attr (k, v) = k ++ "=" ++ v ++ "'"
  paren s = " (" ++ s ++ ")"
```

Printing of a document constructed from an XML tree is the composition of a fold with an unfold, and so a hylomorphism:

```
printXML :: XML → [String]
printXML = hylom step combine
```

- For constructing documents in a less structured fashion, we have to resort to the more general and more complicated *build* operator. For example, here is a builder for a simple document of one section with two sub-paragraphs.

```
buildDoc :: (DocF String b → b) → b
buildDoc f = f (Sec "Heading" [f (Para "p1"), f (Para "p2")])
```

We can actually construct the document from this builder, simply by passing it to the operator *build*, which plugs the holes with document constructors.

```
myDoc :: Doc
myDoc = build buildDoc
```

If we want to traverse the resulting document, for example to print it, we can do so directly without having to construct the document in the first place; we do so by plugging the holes instead with the body of the *printDoc* fold.

```
printMyDoc :: [String]
printMyDoc = buildDoc combine
```

## 7 Discussion

We have shown that two advanced language features—*higher-order functions* and *datatype genericity*—suffice (in the presence of other standard features such as datatypes and interfaces) to capture as reusable code a number of the familiar GOF design patterns; specifically, the patterns we have considered are COMPOSITE, ITERATOR, VISITOR and BUILDER, which together we call the ORIGAMI patterns. We also believe that these or equivalent features are necessary for this purpose, since the design patterns are parametrized by actions and by the shape of datatypes.

Our intentions in doing this work are not so much to criticize the existing informal presentations of these four and other patterns; indeed, as we explain below, the informal presentations contribute much useful information beyond the code. Rather, we aim to promote the uptake of higher-order and datatype-generic techniques, and to encourage their incorporation in mainstream programming languages. In this regard, we are following in the footsteps of Norvig [56], who wrote that 16 of the 23 GOF patterns are ‘invisible or simple’ in Lisp, and others who argue that design patterns amount to admissions of inexpressiveness in programming languages. However, in contrast to Norvig and the others favouring dynamic languages [69], our presentation provides genericity while preserving strong static typing.

We do not claim to have captured all 23 of the GOF patterns, or for that matter any deuterocanonical ones either. In particular, we do not see yet how to capture *creational* design patterns as higher-order datatype-generic programs. This is perhaps because our approach is to model object-oriented ideas in a functional framework, and that framework has no direct analogue of object creation. However, we hope and expect that the languages of tomorrow will provide higher-order datatype-generic features in a more traditional framework, and then we may be able to make better progress. Indeed, Alexandrescu’s *type list* implementation of a GENERIC ABSTRACT FACTORY [2] is essentially a datatype-generic metaprogram written using C++ templates.

We also appreciate that there is much more to design patterns than their extensional characteristics, which can be expressed as class and sequence diagrams and captured as programs or programming constructs. Also important are their intensional characteristics: motivation for their use, paradigmatic examples, trade-offs in their application, and other aspects of the ‘story’ behind the pattern. Our presentation impinges only on the limited extensional aspects of those patterns we treat; the intensional characteristics are not amenable to ‘implementation’.

## 8 Related work

This paper is based on ideas from the Algebra of Programming (‘Squiggol’) community, and especially the work of Backhouse and Malcolm [50, 5], Bird and de Moor [7, 8], Fokkinga, Meijer and Paterson [19, 52], Jeurig and Hinze [41, 35, 37], and Hughes [40]. For their inspiration, I am indebted. For further details on the datatype-generic style presented here, see [22, 23] and the above references.

Jay has an alternative approach to datatype-generic programming, which he calls *shape polymorphism* [43, 44]. He and Palsberg have also done some work on a generic representation of the VISITOR pattern [62], but this relies heavily on reflection rather than his work on shape.



For other recent discussions of the meeting between functional and object-oriented views of genericity, see [17, 21].

## 9 Conclusions

Design patterns are traditionally expressed informally, using prose, pictures and prototypes. We have argued that, given the right language features, certain patterns at least could be expressed more usefully as reusable library code. The language features required, in addition to those provided by mainstream languages, are *higher-order functions* and *datatype genericity*; for some aspects, *lazy evaluation* also turns out to be helpful. These features are familiar in the world of functional programming; we hope to see them soon in more mainstream programming languages.

## 10 Acknowledgements

This paper elaborates on arguments developed in a course presented while the author was a Visiting Erskine Fellow at the University of Canterbury in New Zealand in 2005, and explored further at tutorials at ECOOP [24] and OOPSLA [25]. The contribution of participants at those venues and at less formal presentations of the same ideas is gratefully acknowledged, as is that of several anonymous referees (and one brave soul who signed his review). An earlier version appeared as a workshop paper [26], and some of the material was also incorporated into lecture notes from a Spring School [27]. The work was carried out as part of the EPSRC-funded *Datatype-Generic Programming* project at Oxford and Nottingham; we thank members of the project for advice and encouragement, and especially Bruno Oliveira for his significant contribution.

## References

- [1] M. Aigner and G. M. Ziegler. *Proofs from The Book*. Springer-Verlag, third edition, 2004.
- [2] A. Alexandrescu. *Modern C++ Design*. Addison-Wesley, 2001.
- [3] R. Backhouse and T. Sheard, editors. *Proceedings of the Workshop on Generic Programming*. Unpublished, Marstrand, Sweden, June 1998. <http://www.cs.uu.nl/people/johanj/wgp98.html>.
- [4] R. C. Backhouse and B. A. Carré. Regular algebra applied to path-finding problems. *Journal of the Institute of Mathematics and Applications*, 15:161–186, 1975.
- [5] R. C. Backhouse, P. Jansson, J. Jeuring, and L. G. L. T. Meertens. Generic programming: An introduction. In *Advanced Functional Programming*, volume 1608 of *Lecture Notes in Computer Science*, pages 28–115, 1998.
- [6] R. Bird. *Introduction to Functional Programming Using Haskell*. Prentice-Hall, 1998.
- [7] R. Bird and O. de Moor. *The Algebra of Programming*. Prentice-Hall, 1996.
- [8] R. Bird, O. de Moor, and P. Hoogendijk. Generic functional programming with types and relations. *Journal of Functional Programming*, 6(1):1–28, 1996.
- [9] P. Buchlovsky and H. Thielecke. A type-theoretic reconstruction of the Visitor pattern. *Electronic Notes in Theoretical Computer Science*, 155, 2005. 21st Conference on Mathematical Foundations of Programming Semantics.
- [10] L. Cardelli and P. Wegner. On understanding types, data abstraction and polymorphism. *ACM Comput. Surv.*, 17(4):471–522, Dec. 1985.
- [11] J. Cheney and R. Hinze. A lightweight implementation of generics and dynamics. In *Haskell Workshop*, pages 90–104, 2002.

- [12] K. Claessen and J. Hughes. Specification-based testing with QuickCheck. In J. Gibbons and O. de Moor, editors, *The Fun of Programming*, pages 17–40. Palgrave, 2003.
- [13] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, 1990.
- [14] K. Czarnecki and U. Eisenecker. *Generative Programming: Methods, Tools and Applications*. Addison-Wesley, 2000.
- [15] O. Danvy and A. Filinski. Representing control: A study of the CPS transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, 1992.
- [16] J. C. Dehnert and A. Stepanov. Fundamentals of generic programming. In M. Jazayeri, R. Loos, and D. Musser, editors, *Report of the Dagstuhl Seminar on Generic Programming*, volume 1766 of *Lecture Notes in Computer Science*, pages 1–11. Springer-Verlag, Heidelberg, Germany, 2000.
- [17] G. Dos Reis and J. Järvi. What is generic programming? In *Library-Centric Software Design*, 2005.
- [18] J. Edmonds. Matroids and the Greedy Algorithm. *Mathematical Programming*, 1:125–136, 1971.
- [19] M. M. Fokkinga and E. Meijer. Program calculation properties of continuous algebras. Technical Report CS-R9104, CWI, Amsterdam, Jan. 1991.
- [20] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
- [21] R. Garcia, J. Järvi, A. Lumsdaine, J. G. Siek, and J. Willcock. A comparative study of language support for generic programming. In *Object-Oriented Programming, Systems, Languages, and Applications*, Oct. 2003.
- [22] J. Gibbons. Calculating functional programs. In R. Backhouse, R. Crole, and J. Gibbons, editors, *Algebraic and Coalgebraic Methods in the Mathematics of Program Construction*, volume 2297 of *Lecture Notes in Computer Science*, pages 148–203. Springer-Verlag, 2002.
- [23] J. Gibbons. Origami programming. In J. Gibbons and O. de Moor, editors, *The Fun of Programming*, pages 41–60. Palgrave, 2003.
- [24] J. Gibbons. Design patterns as higher-order datatype-generic programs. <http://2005.ecoop.org/8.html>, June 2005. Tutorial presented at ECOOP.
- [25] J. Gibbons. Design patterns as higher-order datatype-generic programs. <http://www.oopsla.org/2005/ShowEvent.do?id=121>, Oct. 2005. Tutorial presented at OOPSLA.
- [26] J. Gibbons. Design patterns as higher-order datatype-generic programs. In R. Hinze, editor, *Workshop on Generic Programming*, Sept. 2006.
- [27] J. Gibbons. Datatype-generic programming. In *Spring School on Datatype-Generic Programming*, volume 4719 of *Lecture Notes in Computer Science*. Springer-Verlag, 2007.
- [28] J. Gibbons, R. Backhouse, B. Oliveira, and F. Reig. Datatype-generic programming project. <http://web.comlab.ox.ac.uk/oucl/research/pdt/ap/dgp/>, Oct. 2003.
- [29] J. Gibbons, G. Hutton, and T. Altenkirch. When is a function a fold or an unfold? *Electronic Notes in Theoretical Computer Science*, 44(1), Apr. 2001. Proceedings of Coalgebraic Methods in Computer Science.
- [30] J. Gibbons and G. Jones. The under-appreciated unfold. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*, pages 273–279, Baltimore, Maryland, Sept. 1998.
- [31] J. Gibbons, D. Lester, and R. Bird. Enumerating the rationals. *Journal of Functional Programming*, 16(4), 2006.

- [32] J. Gibbons and B. C. d. S. Oliveira. The essence of the Iterator pattern. *Journal of Functional Programming*, 19(3-4):307–402, 2009.
- [33] A. Gill, J. Launchbury, and S. Peyton Jones. A short cut to deforestation. In *Functional Programming Languages and Computer Architecture*, 1993.
- [34] C. Hall, K. Hammond, S. Peyton Jones, and P. Wadler. Type classes in Haskell. *ACM Transactions on Programming Languages and Systems*, 18(2):19–138, 1996.
- [35] R. Hinze. Polytypic values possess polykinded types. In R. C. Backhouse and J. N. Oliveira, editors, *Mathematics of Program Construction*, volume 1837 of *Lecture Notes in Computer Science*, pages 2–27. Springer, 2000.
- [36] R. Hinze. Generics for the masses. *Journal of Functional Programming*, 16(4-5):451–483, 2006.
- [37] R. Hinze and J. Jeuring. Generic Haskell: Practice and theory. In R. Backhouse and J. Gibbons, editors, *Summer School on Generic Programming*, volume 2793 of *Lecture Notes in Computer Science*, pages 1–56. Springer-Verlag, 2003.
- [38] R. Hinze and S. Peyton Jones. Derivable type classes. In G. Hutton, editor, *Haskell Workshop*, volume 41.1 of *Electronic Notes in Theoretical Computer Science*. Elsevier Science, Aug. 2000.
- [39] C. A. R. Hoare. Notes on data structuring. In O.-J. Dahl, E. W. Dijkstra, and C. A. R. Hoare, editors, *Structured Programming*, pages 83–174. Academic Press, 1972.
- [40] J. Hughes. Why functional programming matters. *Computer Journal*, 32(2):98–107, Apr. 1989.
- [41] P. Jansson and J. Jeuring. PolyP – a polytypic programming language extension. In *Principles of Programming Languages*, pages 470–482, 1997.
- [42] J. Järvi, M. Marcus, and J. N. Smith. Programming with C++ concepts. *Science of Computer Programming*, 75(7):596–614, 2010.
- [43] C. B. Jay. A semantics for shape. *Science of Computer Programming*, 25:251–283, 1995.
- [44] C. B. Jay, G. Bellè, and E. Moggi. Functorial ML. *Journal of Functional Programming*, 8(6):573–619, 1998.
- [45] G. Kiczales, J. des Rivières, and D. G. Bobrow. *The Art of the Metaobject Protocol*. MIT Press, 1991.
- [46] G. Kiczales, J. Lamping, A. Menhdhekar, C. Maeda, C. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-oriented programming. In M. Akşit and S. Matsuoka, editors, *European Conference on Object-Oriented Programming*, volume 1241 of *Lecture Notes in Computer Science*, pages 220–242. Springer-Verlag, Berlin, Heidelberg, and New York, 1997.
- [47] B. Korte, L. Lovász, and R. Schrader. *Greedoids*. Springer-Verlag, 1991.
- [48] B. Liskov. A history of CLU. In T. J. Bergin, Jr. and R. G. Gibson, Jr., editors, *History of Programming Languages—II*, pages 471–510, 1996.
- [49] A. Löh. *Exploring Generic Haskell*. PhD thesis, Utrecht University, 2004.
- [50] G. Malcolm. Data structures and program transformation. *Science of Computer Programming*, 14:255–279, 1990.
- [51] C. McBride and R. Paterson. Applicative programming with effects. *Journal of Functional Programming*, 18(1):1–13, 2008.
- [52] E. Meijer, M. Fokkinga, and R. Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In J. Hughes, editor, *Functional Programming Languages and Computer Architecture*, volume 523 of *Lecture Notes in Computer Science*, pages 124–144. Springer-Verlag, 1991.

- [53] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML*. MIT Press, revised edition, 1997.
- [54] MITRE Corp. *Ada Reference Manual*. ISO, 2000. IEC 8652.
- [55] D. R. Musser and A. A. Stepanov. *The Ada Generic Library linear list processing packages*. Springer-Verlag New York, Inc., New York, NY, USA, 1989.
- [56] P. Norvig. Design patterns in dynamic programming. In *Object World*, Boston, MA, May 1996. Tutorial slides at <http://norvig.com/design-patterns/>.
- [57] M. Odersky, L. Spoon, and B. Venners. *Programming in Scala*. Artima, 2008.
- [58] B. Oliveira, A. Moors, and M. Odersky. Type classes as objects and implicits. In *OOPSLA*, pages 341–360, 2010.
- [59] B. C. d. S. Oliveira and J. Gibbons. TypeCase: A design pattern for type-indexed functions. In D. Leijen, editor, *Haskell Workshop*, 2005.
- [60] B. C. d. S. Oliveira and J. Gibbons. Scala for generic programmers. *Journal of Functional Programming*, 20(3,4):303–352, 2010.
- [61] B. C. d. S. Oliveira, M. Wang, and J. Gibbons. The Visitor pattern as a reusable, generic, type-safe component. In *OOPSLA*, Oct. 2008.
- [62] J. Palsberg and C. B. Jay. The essence of the Visitor pattern. In *22nd Annual International Computer Software and Applications Conference*, pages 9–15, Vienna, Austria, August 1998.
- [63] A. Pardo. Fusion of recursive programs with computation effects. *Theoretical Comput. Sci.*, 260:165–207, 2001.
- [64] A. Pardo. Combining datatypes and effects. In *Advanced Functional Programming*, volume 3622 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [65] L. C. Paulson. *ML for the Working Programmer*. Cambridge University Press, second edition, 1996.
- [66] S. Peyton Jones. *The Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [67] F. Ruehr. *Analytical and Structural Polymorphism Expressed Using Patterns over Types*. PhD thesis, University of Michigan, 1992.
- [68] J. Siek, L.-Q. Lee, and A. Lumsdaine. *The Boost Graph Library*. Addison-Wesley, 2002.
- [69] G. T. Sullivan. Advanced programming language features for executable design patterns: Better patterns through reflection. Artificial Intelligence Laboratory Memo AIM-2002-005, Artificial Intelligence Lab, MIT, Mar. 2002.
- [70] W. Taha. A gentle introduction to multi-stage programming. In C. Lengauer, D. Batory, C. Consel, and M. Odersky, editors, *Domain-Specific Program Generation*, number 3016 in *Lecture Notes in Computer Science*, pages 30–50. Springer-Verlag, 2004.
- [71] The Programatica Team. Programatica tools for certifiable, auditable development of high-assurance systems in Haskell. In *High Confidence Software and Systems Conference*. National Security Agency, April 2003.
- [72] E. Unruh. Prime number computation. Technical Report X3J16-94-0075/ISO WG21-462, ANSI, 1994.
- [73] T. Uustalu and V. Vene. Primitive (co)recursion and course-of-value (co)iteration, categorically. *Informatica*, 10(1):5–26, 1999.

- [74] T. Veldhuizen. *Active Libraries and Universal Languages*. PhD thesis, Computer Science, Indiana University, 2004.
- [75] V. Vene and T. Uustalu. Functional programming with apomorphisms (corecursion). *Proceedings of the Estonian Academy of Sciences: Physics, Mathematics*, 47(3):147–161, 1998. 9th Nordic Workshop on Programming Theory.
- [76] D. Vytiniotis, G. Washburn, and S. Weirich. An open and shut typecase. In *International Conference on Functional Programming*, 2004.
- [77] P. Wadler. Theorems for free! In *Functional Programming and Computer Architecture*, 1989.
- [78] P. Wadler. Deforestation: Transforming programs to eliminate trees. *Theoretical Computer Science*, 73:231–248, 1990.
- [79] P. Wadler. Monads for functional programming. In M. Broy, editor, *Program Design Calculi: Proceedings of the Marktoberdorf Summer School*, 1992.
- [80] P. Wadler. The expression problem. Posting to java-genericity mailing list, 12th Nov 1998.
- [81] P. Wadler. How to solve the reuse problem? Functional programming. In *Internal Conference on Software Reuse*, pages 371–372. IEEE, 1998.

## 11 Appendix: Java programs

Section 6.5 provides a nearly complete implementation of the document application in a higher-order datatype-generic style; all that is missing is a definition for the spelling corrector *correct*. In contrast, Section 5.5 and Figure 11 present only the outline of a Java implementation of the same application. For completeness, this appendix presents the Java code. (Note that we have used Java generics throughout, in order to avoid type casts. Nevertheless, we wish to emphasize that the datatype-genericity discussed in this paper is a different kind of genericity to that provided in Java or C# generics.)

### 11.1 Components

The *Component* interface forms the root of the COMPOSITE hierarchy.

```
public interface Component {
    void accept (Visitor v);
    Iterator getIterator ();
}
```

The class *Section* plays the *Composite* role in the COMPOSITE pattern—a section has a *Vector* of child *Components*, and is itself a *Component*.

```
import java.util.Vector;
import java.util.Enumeration;

public class Section implements Component {
    protected Vector<Component> children;
    protected String title;
    public Section (String title) {
        children = new Vector<Component> ();
        this.title = title;
    }
    public String getTitle () {
        return title;
    }
}
```

```

public void addComponent (Component c) {
    children.addElement (c);
}
public Enumeration<Component> getChildren () {
    return children.elements ();
}
public Iterator getIterator () {
    return new SectionIterator (this);
}
public void accept (Visitor v) {
    v.visitSection (this);
}
}

```

The class *Paragraph* plays the *Leaf* role in the COMPOSITE pattern—a paragraph has a body, which is a *String*, and no children.

```

public class Paragraph implements Component {
    protected String body;
    public Paragraph (String body) {
        setBody (body);
    }
    public void setBody (String s) {
        body = s;
    }
    public String getBody () {
        return body;
    }
    public Iterator getIterator () {
        return new ParagraphIterator (this);
    }
    public void accept (Visitor v) {
        v.visitParagraph (this);
    }
}

```

## 11.2 Iterators

The *Iterator* interface provides INTERNAL ITERATOR behaviour over a document.

```

public interface Iterator {
    void iterate (Action a);
}

```

As discussed in Section 5.5, we have chosen to model the *Actions* of the INTERNAL ITERATOR as applying only to the paragraphs in a document, not the sections; that could be changed. Also, *Action.apply* returns **void** for simplicity: an action can update the state of a paragraph, but it cannot change the identity of that paragraph, nor delete it.

```

public interface Action {
    void apply (Paragraph p);
}

```

The natural implementation of an *Iterator* for *Sections* is to iterate over the child components of that section.

```

import java.util.Enumeration ;
public class SectionIterator implements Iterator {
    protected Section s ;
    public SectionIterator (Section s) {
        this.s = s ;
    }
    public void iterate (Action a) {
        for (Enumeration<Component> e = s.getChildren () ; e.hasMoreElements ();) {
            e.nextElement ().getIterator ().iterate (a) ;
        }
    }
}

```

The *Iterator* for a *Paragraph* applies the *Action* provided to the paragraph.

```

public class ParagraphIterator implements Iterator {
    protected Paragraph p ;
    public ParagraphIterator (Paragraph p) {
        this.p = p ;
    }
    public void iterate (Action a) {
        a.apply (p) ;
    }
}

```

For example, correcting the spelling of each paragraph body could be implemented using the INTERNAL ITERATOR pattern.

```

public class SpellCorrector implements Action {
    public void apply (Paragraph p) {
        p.setBody (correct (p.getBody ())) ;
    }
    public String correct (String s) {
        return s.toLowerCase () ;
    }
}

```

(For simplicity, we have provided only a dummy implementation of spelling correction; a real implementation would be beyond the scope of this example.)

### 11.3 Visitors

The VISITOR pattern represents traversals over the composite structure—in this case, a *Visitor* provides methods to visit a *Paragraph* and a *Section*.

```

public interface Visitor {
    void visitParagraph (Paragraph p) ;
    void visitSection (Section s) ;
}

```

(As with our instantiation of the INTERNAL ITERATOR pattern, we have restricted attention to visitors that operate by means of side-effects, rather than directly returning a result; there are alternative designs.)

For example, a *PrintVisitor* implements the *Visitor* interface to assemble a string array, with one string for each section title and for each paragraph body, all indented appropriately to reflect the hierarchical structure.

```

import java.util.Enumeration ;
import java.util.Vector ;

public class PrintVisitor implements Visitor {
    protected String indent = "" ;
    protected Vector<String> lines = new Vector<String> () ;
    public String [] getResult () {
        String [] ss = new String [0] ;
        ss = lines.toArray (ss) ;
        return ss ;
    }
    public void visitParagraph (Paragraph p) {
        lines.addElement (indent + p.getBody ()) ;
    }
    public void visitSection (Section s) {
        String currentIndent = indent ;
        lines.addElement (indent + s.getTitle ()) ;
        for (Enumeration<Component> e = s.getChildren () ; e.hasMoreElements ()) {
            indent = currentIndent + " " ;
            e.nextElement ().accept (this) ;
        }
        indent = currentIndent ;
    }
}

```

## 11.4 Builders

The *Builder* interface represents the structured variant of the BUILDER pattern, in which each *part* is added as a child of an existing part in the product being assembled. The *addPart* methods each return a part identifier (an integer), to be used to identify the parent of a subsequent part. The process is intended to be initialized by specifying a negative number as the ‘parent’ of the first part to be added, which will then form the root of the document hierarchy.

```

public interface Builder {
    int addParagraph (String body, int parent) throws InvalidBuilderId ;
    int addSection (String title, int parent)    throws InvalidBuilderId ;
}

```

Here is a simple *Exception* class for when things go wrong—for example, when attempting to create a second root, when passing in an unrecognized parent identifier, or when attempting to add a child to a *Paragraph* node.

```

public class InvalidBuilderId extends Exception {
    public InvalidBuilderId (String reason) {
        super (reason) ;
    }
}

```

The class *ComponentBuilder* provides the most obvious implementation of the *Builder* interface; its *getProduct* method yields a *Component* as the final product. The assembly process makes use of a *HashMap* from node identifiers to document components.

```

import java.util.AbstractMap ;
import java.util.HashMap ;

```



```

public class ComponentBuilder implements Builder {
    protected int nextId = 0;
    protected AbstractMap<Integer, Component> cs = new HashMap<Integer, Component> ();
    public int addParagraph (String body, int pId) throws InvalidBuilderId {
        return addComponent (new Paragraph (body), pId);
    }
    public int addSection (String title, int pId) throws InvalidBuilderId {
        return addComponent (new Section (title), pId);
    }
    public Component getProduct () {
        return cs.get (0);
    }
    protected int addComponent (Component c, int pId) throws InvalidBuilderId {
        if (pId < 0) { // root component
            if (cs.isEmpty ()) {
                cs.put (nextId, c);
                return nextId++;
            }
            else
                throw new InvalidBuilderId ("Duplicate root");
        } else { // non-root
            Component parent = (Component) cs.get (pId);
            if (parent == null) {
                throw new InvalidBuilderId ("Non-existent parent");
            } else {
                if (parent instanceof Paragraph) {
                    throw new InvalidBuilderId ("Adding child to paragraph");
                } else {
                    Section s = (Section) parent;
                    s.addComponent (c);
                    cs.put (nextId, c);
                    return nextId++;
                }
            }
        }
    }
}

```

The *PrintBuilder* implementation of the *Builder* interface is the only class with a non-obvious implementation. It constructs on the fly the printed representation (a *String* []) of a *Component*. In order to do so, it needs to retain some of the tree structure. This is done by maintaining, for each *Component* stored, the unique identifier of its right-most child (or its own identifier, if it has no children). This is stored in the *last* field of the corresponding *Record* in the vector *records*. This vector itself is stored in the order the lines will be returned, that is, a preorder traversal. When adding a new *Component*, it should be placed after the rightmost descendent of its immediate parent, and this is located by following the path of *last* references. Note that devising this implementation requires considerable insight and ingenuity; in contrast, the HODGP COMPUTING BUILDER discussed in Section 6.4 arises completely mechanically.

```

import java.util.Vector;
public class PrintBuilder implements Builder {
    protected class Record {
        public int id;
        public int last;
        public String line;
    }
}

```

```

    public String indent ;
    public Record (int id, int last, String line, String indent) {
        this.id = id ;
        this.last = last ;
        this.line = line ;
        this.indent = indent ;
    }
}
protected Vector<Record> records = new Vector<Record> () ;
protected Record recordAt (int i) {
    return records.elementAt (i) ;
}
protected int find (int id, int start) {
    while (start < records.size () && recordAt (start).id != id)
        start++ ;
    if (start < records.size ())
        return start ;
    else
        return - 1 ;
}
protected int nextId = 0 ;
protected SpellCorrector c = new SpellCorrector () ;
public int addParagraph (String body, int pid) throws InvalidBuilderId {
    return addComponent (c.correct (body), pid) ;
}
public int addSection (String title, int pid) throws InvalidBuilderId {
    return addComponent (title, pid) ;
}
public String [] getProduct () {
    String [] ss = new String [records.size ()] ;
    for (int i = 0 ; i < ss.length ; i++)
        ss [i] = recordAt (i).indent + recordAt (i).line ;
    return ss ;
}
protected int addComponent (String s, int pId) throws InvalidBuilderId {
    if (pId < 0) { // root component
        if (records.isEmpty ()) {
            records.addElement (new Record (nextId, nextId, s, "")) ;
            return nextId++ ;
        }
        else
            throw new InvalidBuilderId ("Duplicate root") ;
    }
    else { // non-root
        int x = find (pId, 0) ;
        Record r = recordAt (x) ;
        String indent = r.indent ;
        if (x == - 1) {
            throw new InvalidBuilderId ("Non-existent parent") ;
        }
        else {
            int y = x ; // ids [x] = ids [y] = pid
            while (r.id != r.last) {
                y = x ;
            }
        }
    }
}

```

```

        x = find (r.last, x);
        r = recordAt (x);
    } // lasts [y] = lasts [x] = ids [x]
    records.insertElementAt (new Record (nextId, nextId, s, indent + "  "), x + 1);
    recordAt (y).last = nextId; // lasts [y] = lasts [x + 1] = nextId
    return nextId++;
}
}
}
}
}

```

## 11.5 A program

Finally, here is a simple application, providing a *main* method. Depending on whether the argument on the command line is "building" or "computing", the application either uses the ordinary ('building') implementation of the *Builder* interface to construct a document, then prints it with a *Visitor*, or it uses the computing implementation of the *Builder* to construct the printed representation directly—the output is identical in either case.

```

public abstract class Main {
    protected static void build (Builder b) {
        try {
            int rootId = b.addSection ("Doc", -1);
            int sectId = b.addSection ("Sec 1", rootId);
            int subsId = b.addSection ("Subsec 1.1", sectId);
            int id = b.addParagraph ("Para 1.1.1", subsId);
            id = b.addParagraph ("Para 1.1.2", subsId);
            subsId = b.addSection ("Subsec 1.2", sectId);
            id = b.addParagraph ("Para 1.2.1", subsId);
            id = b.addParagraph ("Para 1.2.2", subsId);
            sectId = b.addSection ("Sec 2", rootId);
            subsId = b.addSection ("Subsec 2.1", sectId);
            id = b.addParagraph ("Para 2.1.1", subsId);
            id = b.addParagraph ("Para 2.1.2", subsId);
            subsId = b.addSection ("Subsec 2.2", sectId);
            id = b.addParagraph ("Para 2.2.1", subsId);
            id = b.addParagraph ("Para 2.2.2", subsId);
        } catch (InvalidBuilderId e) {
            System.out.println ("Exception: " + e);
        }
    }
}

protected static void usage () {
    System.err.println ("java Main {building|computing}");
}

public static void main (String [] args) {
    String [] lines = new String [0];
    if (args == null ∨ args.length != 1) {
        usage (); System.exit (1);
    }
    if (args [0].equals ("building")) { // build then compute
        ComponentBuilder b = new ComponentBuilder ();
        build (b);
        Component root = b.getProduct ();
    }
}

```

```

    root.getIterator ().iterate (new SpellCorrector ());
    PrintVisitor pv = new PrintVisitor ();
    root.accept (pv);
    lines = pv.getResult ();
} else if (args [0].equals ("computing")) { // computing builder
    PrintBuilder b = new PrintBuilder ();
    build (b);
    lines = b.getProduct ();
} else {
    usage (); System.exit (1);
}
for (int i = 0; i < lines.length; i++)
    System.out.println (lines [i]);
}
}

```