

# FUNCTIONAL PEARL

## *Parsing permutation phrases*

ARTHUR I. BAARS, ANDRES LÖH and S. DOAITSE SWIERSTRA

*Institute of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands*  
(e-mail: {arthurb, andres, doaitse}@cs.uu.nl)

---

### Abstract

A permutation phrase is a sequence of elements (possibly of different types) in which each element occurs exactly once and the order is irrelevant. Some of the permutable elements may be optional. We show how to extend a parser combinator library with support for parsing such free-order constructs. A user of the library can easily write parsers for permutation phrases and does not need to care about checking and reordering the recognized elements. Applications include the generation of parsers for attributes of XML tags and Haskell's record syntax.

---

### 1 Introduction

Parser combinator libraries highlight the strengths of functional programming languages: higher-order functions and the possibility to define new infix operators allow parsers to be expressed in a concise and natural notation that closely resembles the syntax of EBNF grammars. At the same time, the user has the full abstraction power of the underlying programming language at hand. Complex, often recurring patterns can be expressed by defining new combinators.

A specific parsing problem is the recognition of permutation phrases. A permutation phrase is a sequence of elements (possibly of different types) in which each element occurs exactly once and the order is irrelevant. Some of the permutable elements may be optional. Since permutation phrases are not easily expressed by a context-free grammar, the usual approach is to tackle this problem in two steps: first parse a relaxed version of the grammar, then check whether the recognized elements form a permutation of the expected elements. This method, however, has a number of disadvantages. Dealing with a permutation of typed values is quite cumbersome, and the problem is often avoided by encoding the values in a universal representation, thus adding an extra level of interpretation. Furthermore, because of the two steps involved, error messages cannot be produced until a larger part of the input has been consumed, and special care has to be taken to make them point to the right position in the code.

Permutation phrases have been proposed by Cameron (1993) as an extension to EBNF grammars, not aiming at greater expressive power, but at more clarity. Cameron also presents a pseudo-code algorithm to parse permutation phrases with

optional elements efficiently in an imperative setting. It fails, however, to address the types of the constituents.

We show how to extend any existing parser combinator library with support for parsing permutations of typed, potentially optional elements.

Possible applications include the implementation of Haskell's *read* function where it is desirable to parse the fields of a data type with labelled fields in any order, and the parsing of XML tags which have large sets of potentially optional attributes that may occur in any order. For instance, a parser for the XHTML image tag with some of its attributes can be written as follows:

```
imgtag      = token "<" * token "img" * attrs * token ">"
  where attrs = permute (Img <> field "src" uri
                        <> field "alt" string
                        <> optional (field "longdesc" uri)
                        <> optional (field "height" int)
                        <> optional (field "width" int)
                        )
```

The combinator  $\langle \rangle$  is used to separate parsers for permutable elements, and  $\langle \rangle$  can be used to apply a semantic function. Parsers for permutation phrases have to be enclosed by a call to *permute*.

Our approach makes use of two features that are not provided by all functional programming languages: existentially quantified data types are used to encode reordering information that permutes the recognized elements to a canonical order. Additionally, we utilize lazy evaluation to make the resulting implementation efficient. The administrative part of parsing permutation phrases has a quadratic time complexity in the number of permutable elements. The size of the code, however, is linear in the number of permutable elements.

We therefore choose Haskell as implementation language. Existential types are not part of the Haskell 98 standard (Peyton Jones, 2003), but are supported by several current Haskell implementations.

The paper is organized as follows: section 2 explains the parser combinators we build upon. Section 3 presents the idea of dealing with permutations in terms of permutation trees and explains how such trees are built and converted into parsers. In section 4, we take a brief look at the applications mentioned above: the parsing of data types with labelled fields and the parsing of XML attribute sets. Section 5 concludes.

## 2 Parsing using combinator libraries

The use of a combinator library for describing parsers instead of writing them by hand or generating them from a separate formalism is a well-known technique in functional programming. As a result, there are several excellent libraries around. For this reason we just briefly present the interface we will assume in subsequent sections of this paper, but do not go into the details of the implementation. We want to stress, however, that our extension is not tied to any specific library.

```

infixl 3 <◇
infixl 4 <◇
class Parser p where
  fail    :: p a
  succeed :: a → p a
  symbol  :: Char → p Char
  (<◇)    :: p (a → b) → p a → p b
  (<◇)    :: p a → p a → p a

```

Fig. 1. Type class for parser combinators.

```

infixl 4 <◇, <◇, <◇, <◇, <◇
(<◇)    :: Parser p ⇒ (a → b) → p a → p b
f <◇ p  = succeed f <◇ p
(<◇)    :: Parser p ⇒ a → p b → p a
x <◇ p  = const x <◇ p
(<◇)    :: Parser p ⇒ p a → p b → p a
p <◇ q  = const <◇ p <◇ q
(<◇)    :: Parser p ⇒ p a → p b → p b
p <◇ q  = flip const <◇ p <◇ q
parens  :: Parser p ⇒ p a → p a
parens p = symbol '(' <◇ p <◇ symbol ')'

```

Fig. 2. Some useful parser combinators.

We make use of a simple interface (Røjemo, 1995; Swierstra & Duponcheel, 1996) that is parametrized by the result type of the parsers and assumes a list of characters as input. It can easily be implemented by straightforward list-of-successes parsers (Fokker, 1995; Wadler, 1985). Our permutation parsers have also been implemented for more advanced libraries, such as the fast, error-correcting parser combinators of Swierstra (2001) and the monadic-style (Hutton & Meijer, 1988) combinator library *Parsec* (Leijen, 2001).

The parser interface used here is given as a type class declaration in Figure 1. The function *fail* represents the parser that always fails, whereas *succeed* never consumes any input and always returns the given result value. The parser *symbol* accepts solely the given character as input. If this character is encountered, *symbol* consumes and returns this character, otherwise it fails. The <◇ operator denotes the sequential composition of two parsers, where the result of the first parser is applied to the result of the second. Finally, the operator <◇ expresses a choice between two parsers.

Many useful combinators can be built on top of these basic ones. A small selection that we use in this paper is presented in Figure 2. The most notable of the derived combinators is the application operator <◇, a parser transformer that can be used to apply a semantic function to a parse result. It is defined in terms of *succeed* and <◇.

### 3 Permutation parsers

We compute the parser for a permutation phrase not directly, but from an intermediate tree. We first introduce an auxiliary data type to represent such

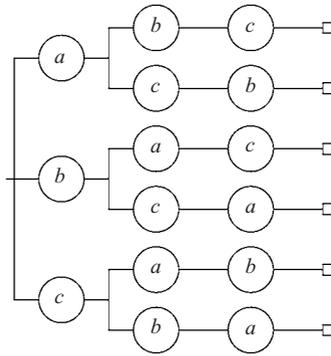


Fig. 3. A permutation tree containing three elements.

permutation trees, and show how they can be converted to parsers. Later, we introduce special combinators that provide a convenient notation for constructing permutation trees.

### 3.1 Permutation trees

A permutation phrase of a set of elements can be expanded into an EBNF definition by summing up all possible permutations of the elements. Consider for example the permutation phrase of three elements  $a$ ,  $b$ , and  $c$ . Using Cameron's notation for permutation phrases, we can write it as:

$$s ::= \langle\langle a \parallel b \parallel c \rangle\rangle$$

Expanding and subsequently left-factorizing this permutation phrase gives us the following EBNF production rule:

$$\begin{array}{l}
 s ::= a b c \mid a c b \\
 \quad \mid b a c \mid b c a \\
 \quad \mid c a b \mid c b a
 \end{array}
 \rightsquigarrow
 \begin{array}{l}
 s ::= a (b c \mid c b) \\
 \quad \mid b (a c \mid c a) \\
 \quad \mid c (a b \mid b a)
 \end{array}$$

The left-factorized production rule can be represented by a tree as illustrated in Figure 3. Each path from the root to a leaf in the tree represents a particular permutation. Permutations with a common prefix share the same subtree, hence the number of choices in each node is limited by the number of permutable elements. Such a permutation tree is very suitable as intermediate data structure for a permutation parser. Parsing a permutation phrase boils down to checking whether the input matches one of the paths in the permutation tree.

If the grammar (and thus the permutation tree) is ambiguous, large parts of the tree might need to be evaluated before it can be decided which path must be followed. Therefore, ambiguous grammars should (as always) be avoided. However, if the ambiguity in the grammar stems from optional elements in the permutation phrase, the permutation tree can be modified in a simple way to resolve the ambiguity.

In section 3.2, we develop an implementation for permutations without optional elements, and in section 3.3 extend this solution to cover optionality.

### 3.2 Permutation trees without optional elements

We introduce a data type *Perms* for permutation trees, which is parametrized by a type constructor  $p$  (e.g. the parser type) and a result type  $a$ .

```
data Perms p a = Empty a
                | Choice [Branch p a]
data Branch p a =  $\exists x$ . Br (p x) (Perms p (x  $\rightarrow$  a))
```

The *Empty* constructor represents a leaf of the tree, the *Choice* constructor stores a branching node. A single branch is constructed using *Br* and consists of an element, represented as a parser, plus a subtree. As the types of the elements may differ between branches, we hide their types by existentially quantifying the  $x$  in the definition of *Br*. The subtrees have an element type that is different from the type of the original tree, making *Perms* a non-regular data type. Each subtree must contain information how to construct a value of the result type  $a$  from a value of the element's type. This is achieved by storing at the end of each path a function that effectively reorders the elements on the path.

To show that reordering is determined by the type of the components, we will henceforth write the type (or the element type for type constructors) of a variable as an index to its name.

The idea that each path in the tree represents the parser for one of the possible permutations is reflected by the following simple conversion function from permutation trees to parsers. A leaf, i.e. an *Empty*, is interpreted as an always succeeding parser; a *Choice* constructor is converted into a choice between parsers. For each branch, we have to supply the value resulting from the parser to the reordering function resulting from the subtree.

```
permute          :: Parser p  $\Rightarrow$  Perms p a  $\rightarrow$  p a
permute (Empty va) = succeed va
permute (Choice bsa) = choice (map branch bsa)
branch          :: Parser p  $\Rightarrow$  Branch p a  $\rightarrow$  p a
branch (Br px tx $\rightarrow$ a) = ( $\lambda x$  f  $\rightarrow$  f x)  $\triangleleft$  px  $\triangleleft$  permute tx $\rightarrow$ a
choice         :: Parser p  $\Rightarrow$  [p a]  $\rightarrow$  p a
choice         = foldr ( $\triangleleft$ ) fail
```

Lazy evaluation plays an important role here, in that it ensures that the permutation tree is never computed completely. In fact, for a permutation of  $n$  elements, just the  $n$  tree elements immediately below the root of the tree are required to decide which subtree will be used to parse the rest of the permutation. From then on, only that subtree (a permutation tree of size  $n - 1$ ) is relevant. Iterating this argument leads to a complexity of only  $O(n^2)$ .

There are two potential problems here. First, the underlying parser combinator library might try to optimize parsing behaviour by evaluating different possible paths.

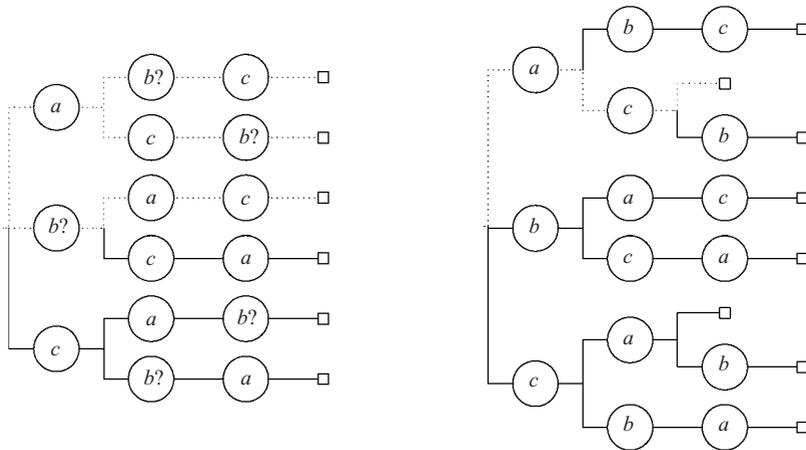


Fig. 4. The ambiguous and the adapted permutation tree for optional *b*.

This is problematic because permutation trees are so large that the precomputation is clearly undesirable. If the library has such features, they should be locally disabled for permutation trees. Second, repeated parsing of different permutations might cause multiple paths of the tree to be evaluated. In practice, however, the number of different permutations that actually occur in the input is small compared to the number of possible permutations.

### 3.3 Permutation trees with optional elements

Optional elements can be represented by parsers that can recognize the empty string and return a default value for this element. However, if *permute* is called on a permutation tree that contains such parsers, the resulting parser is ambiguous.

Consider the left tree in Figure 4, which contains all permutations of *a*, an optional *b* and *c*. Suppose we want to recognize *ac*. This can be done in three different ways since the empty *b* can be recognized before *a*, after *a* or after *c*. The three possible parses are shown as dotted paths in the figure. Fortunately, it is irrelevant for the result of a parse where exactly the empty *b* is derived, since order is not important. This allows us to use a strategy similar to the one proposed by Cameron (1993): parse nonempty constituents as they are seen and allow the parser to stop if all remaining elements are optional. When the parser stops the default values are returned for all optional elements that have not been recognized. The right tree in Figure 4 depicts this strategy for our three element example. The additional leaves mark the positions where the parser is allowed to stop. The string *ac* can now be parsed in only one way.

To implement this strategy we need to be able to determine whether a parser can derive the empty string and split it into its default value and its non-empty part, i.e. a parser that behaves the same except that it does not recognize the empty string. Both parts are represented as *Maybe* values: the first component is *Nothing* if and only if the parser cannot recognize the empty string. If the second component

is *Nothing*, the parser does not consume any input. Hence, it is either a *succeed* (i.e. it just carries semantics) or a *fail*. The splitting of parsers is represented by the *ParserSplit* class that is an extension of the normal *Parser* class. If the underlying parser combinator library cannot be easily adapted to cover this extension, one can alternatively introduce additional combinators, similar to  $\langle \triangleright$  and  $\langle \triangleright \triangleright$ , and let the user mark the optional elements explicitly (Leijen, 2001).

```
class Parser p  $\Rightarrow$  ParserSplit p where
  split :: p a  $\rightarrow$  (Maybe a, Maybe (p a))
```

In the solution that does not deal with optional elements a parser for a permutation follows a path from the root of a permutation tree to a leaf, i.e. an *Empty* node. In the presence of optional elements, however, a parser may stop in any node that stores only optional elements. We adapt the *Perms* data type to incorporate this additional information. If all elements stored in a tree are optional then their default values are stored in *defaults*, otherwise *defaults* is *Nothing*. The parser stored in each *Branch* is not allowed to derive the empty string. We can express the former *Empty* constructor as a function now.

```
data Perms p a = Choice { defaults :: (Maybe a), branches :: [Branch p a] }
  empty x       = Choice (Just x) []
```

The function *permute* is straightforwardly adapted to the generalized data type:

```
permute                :: Parser p  $\Rightarrow$  Perms p a  $\rightarrow$  p a
permute (Choice da bsa) = exit da
                        $\langle \triangleright$  choice (map branch bsa)
exit                   :: Parser p  $\Rightarrow$  Maybe a  $\rightarrow$  p a
exit (Just va)        = succeed va
exit Nothing           = fail
```

Note the similarity to the old definition. The only difference is the possibility to exit early where a default value is present in the tree.

### 3.4 Building a permutation tree

Permutation trees are created by adding the elements of the permutation one by one to an initially empty tree. The function *add* takes a pair consisting of an optional default value and a parser that does not recognize the empty string, and adds it to an existing tree.

```
add                :: (Maybe a, p a)  $\rightarrow$  Perms p (a  $\rightarrow$  b)  $\rightarrow$  Perms p b
add (da, pa) ta $\rightarrow$ b = case ta $\rightarrow$ b of
  Choice da $\rightarrow$ b bsa $\rightarrow$ b  $\rightarrow$ 
    Choice (ap da $\rightarrow$ b da)
      (Br pa ta $\rightarrow$ b : map ins bsa $\rightarrow$ b)
where ins (Br px tx $\rightarrow$ a $\rightarrow$ b) = Br px (add (da, pa) (mapPerms flip tx $\rightarrow$ a $\rightarrow$ b))
```

Having a default value means that the parser described by the permutation tree can accept the empty string. Surely, for the constructed tree, that is only possible if both

the original tree has a default value and  $d_a$  is not *Nothing*. Then, the new default value can be built from the two using function application. The function  $ap$  does exactly this – it is function application lifted to *Maybe* types.

$$\begin{aligned} ap & && :: \text{Maybe } (a \rightarrow b) \rightarrow \text{Maybe } a \rightarrow \text{Maybe } b \\ ap \text{ (Just } f \text{) (Just } x \text{)} &= \text{Just } (f \ x) \\ ap \text{ } \_ \text{ } \_ &= \text{Nothing} \end{aligned}$$

We can add a new element  $(d_a, p_a)$  to a permutation tree by inserting it in all possible positions to every permutation that is already in the tree. The function  $add$  explicitly constructs the tree that represents the permutation in which  $p_a$  is the top element. Additionally, for each branch of the original tree, the top element is left unchanged, and  $(d_a, p_a)$  is inserted everywhere (by a recursive call to  $add$ ) in the subtree. Because the new element and the top element of the branch are now swapped, the function resulting from the subtree of the branch gets its arguments passed in the wrong order, which is repaired by applying  $flip$  to that function.

The function  $mapPerms$  is a mapping function on permutation trees. In a branch,  $f_{a \rightarrow b}$  is composed with the function that results from the subtree.

$$\begin{aligned} mapPerms & && :: (a \rightarrow b) \rightarrow \text{Perms } p \ a \rightarrow \text{Perms } p \ b \\ mapPerms \ f_{a \rightarrow b} \text{ (Choice } d_a \ bs_a \text{)} &= \text{Choice } (fmap \ f_{a \rightarrow b} \ d_a) \\ & && \quad (map \ (mapBranch \ f_{a \rightarrow b}) \ bs_a) \\ mapBranch & && :: (a \rightarrow b) \rightarrow \text{Branch } p \ a \rightarrow \text{Branch } p \ b \\ mapBranch \ f_{a \rightarrow b} \text{ (Br } p_x \ t_{x \rightarrow a} \text{)} &= \text{Br } p_x \ (mapPerms \ (f_{a \rightarrow b} \circ) \ t_{x \rightarrow a}) \end{aligned}$$

We now define two operators for building permutation trees. The first is an operator that extends a permutation tree with a new element.

$$\begin{aligned} (\diamond) & && :: \text{ParserSplit } p \Rightarrow \text{Perms } p \ (a \rightarrow b) \rightarrow p \ a \rightarrow \text{Perms } p \ b \\ t \diamond p &= \\ \text{case split } p \text{ of} & && \\ \text{(Just } e \ , \text{Just } ne \text{)} &\rightarrow \text{add } (\text{Just } e \ , \text{ne}) \ t && \text{-- optional element} \\ \text{(Nothing, Just } ne \text{)} &\rightarrow \text{add } (\text{Nothing, } ne) \ t && \text{-- required element} \\ \text{(Just } e \ , \text{Nothing)} &\rightarrow \text{mapPerms } (\lambda f \rightarrow f \ e) \ t && \text{-- pure semantics} \\ \text{(Nothing, Nothing)} &\rightarrow \text{Choice } \text{Nothing } [] && \text{-- fail} \end{aligned}$$

The second provides an empty permutation tree with initial semantics. It has a similar functionality as the  $\diamond$  for normal parsers.

$$\begin{aligned} (f \diamond) & && :: \text{ParserSplit } p \Rightarrow (a \rightarrow b) \rightarrow p \ a \rightarrow \text{Perms } p \ b \\ f \diamond p &= \text{empty } f \diamond p \end{aligned}$$

An example with three permutable elements, corresponding to the tree in Figure 3, can now be realized by:

$$\text{permute } ((,) \diamond \text{int} \diamond \text{char} \diamond \text{bool})$$

where  $int$ ,  $char$ , and  $bool$  are parsers for literals of type *Int*, *Char*, and *Bool*, respectively. Then all permutations of an integer, a character and a boolean are accepted, and the results of a successful parse are combined using the triple constructor  $(,,)$ , thus yielding a value of type  $(Int, Char, Bool)$ .

### 3.5 Separators

Permutable elements are often separated by symbols that do not carry meaning – typically commas or semicolons. Consider extending the three-element example to the Haskell tuple syntax: not just the elements, but also the parentheses and the commas should be parsed. Since there is one separation symbol less than there are permutable elements, our current variant of *permute* cannot handle this problem.

Therefore we define *permuteSep* as a generalization of *permute* that accepts an additional parser for the separator as an argument. The semantics of the separators are ignored for the result.

$$\begin{aligned} \text{permuteSep} & \quad :: \text{Parser } p \Rightarrow p \ b \rightarrow \text{Perms } p \ a \rightarrow p \ a \\ \text{permuteSep } s \ \text{perm} & = \text{permuteSep}' (\text{succeed } ()) \ s \ \text{perm} \end{aligned}$$

The function *permuteSep'* now converts a permutation tree into a parser in almost the same way as the former *permute*, except that before each permutable element a separator is parsed. To prevent that a separator is expected before the first permutable element, we make use of the following simple trick. The *permuteSep'* function expects two extra arguments: the first one will be parsed immediately before the first element, and the second will be used subsequently. Using *succeed* () as first extra argument in *permuteSep* leads to the desired result.

$$\begin{aligned} \text{permuteSep}' & \quad :: \text{Parser } p \Rightarrow p \ c \rightarrow p \ b \rightarrow \text{Perms } p \ a \rightarrow p \ a \\ \text{permuteSep}' \ sf \ s \ (\text{Choice } d_a \ bs_a) & = \text{exit } d_a \\ & \quad \triangleleft sf \ \bowtie \ \text{choice } (\text{map } (\text{branchSep } s) \ bs_a) \\ \text{branchSep} & \quad :: \text{Parser } p \Rightarrow p \ b \rightarrow \text{Branch } p \ a \rightarrow p \ a \\ \text{branchSep } s \ (\text{Br } p_x \ t_{x \rightarrow a}) & = \text{flip } (\$) \ \triangleleft p_x \ \triangleleft \text{permuteSep}' \ s \ s \ t_{x \rightarrow a} \end{aligned}$$

The *permute* function can now be implemented in terms of *permuteSep*.

$$\begin{aligned} \text{permute} & \quad :: \text{Parser } p \Rightarrow \text{Perms } p \ a \rightarrow p \ a \\ \text{permute} & = \text{permuteSep } (\text{succeed } ()) \end{aligned}$$

To return to the small example, triples of an integer, a character, and a boolean – in any order – are parsed by:

$$\text{parens } (\text{permuteSep } (\text{symbol } ' , ')) \ ((,) \ \triangleleft \text{int} \ \triangleleft \text{char} \ \triangleleft \text{bool})$$

## 4 Applications

### 4.1 XML attributes

We now demonstrate the use of the permutation parsers by showing how to parse XML tags with attributes. For simplicity, we just consider one tag (the *img* tag of XHTML) and only deal with a subset of the attributes allowed. In a Haskell

program, this tag might be represented by the following data type.

```

data XHTML = Img { src      :: URI
                    , alt      :: String
                    , longdesc  :: Maybe URI
                    , height    :: Maybe Int
                    , width     :: Maybe Int
                    }
  | ...

```

Our variant of the `img` tag has five attributes of three different types. We use Haskell's record syntax to keep track of the names. The first two attributes are mandatory whereas the others are optional. We choose the *Maybe* variant of their types to reflect this optionality. Our parser should be able to parse the attributes in any order, where any of the optional arguments may be omitted. For the parsing process, we ignore whitespace and assume that there is a parser

$$\text{token} :: \text{Parser } p \Rightarrow \text{String} \rightarrow p \text{ String}$$

that consumes just the given token and fails on any other input.

Using the *permute* combinator, writing the parser for the `img` tag is easy:

```

imgtag      :: ParserSplit p => p XHTML
imgtag      = token "<" > token "img" > attrs < token ">"
  where attrs = permute $ Img <>
    <> field "src" uri
    <> field "alt" string
    <> optional (field "longdesc" uri)
    <> optional (field "height" int)
    <> optional (field "width" int)

optional    :: Parser p => p a -> p (Maybe a)
optional p  = Just <> p <> succeed Nothing

```

The order in which we denote the attributes determines the order in which the results are returned. Therefore, we can apply the *Img* constructor to form a value of the *XHTML* data type. The helper function *field* is used to parse a single attribute.

$$\begin{aligned} \text{field} &:: \text{Parser } p \Rightarrow \text{String} \rightarrow p a \rightarrow p a \\ \text{field } s \text{ } p &= \text{token } s \text{ } \> \text{symbol '=' } \> p \end{aligned}$$

## 4.2 Haskell's record syntax

Haskell allows data types to contain labelled fields. If one wants to construct a value of that data type, one can make use of these names. The advantage is that the user does not need to remember the order in which the fields of the constructor have been defined. Furthermore, all fields are considered as optional. If a field is not explicitly set to a value, it is silently assumed to be  $\perp$ .

Whereas compilers support order-free syntax (the record fields in a program can be ordered arbitrarily), the *read* function expects the fields in the same order as in the data type declaration. The resulting asymmetry is unfortunate. Using the

*permuteSep* combinator, it is an easy task to write more flexible parsers for data types with labelled fields.

```

justOrNothing  :: Parser p => p a -> p (Maybe a)
justOrNothing p = Just << token "Just" << p
               << Nothing << token "Nothing"

img            :: ParserSplit p => p XHTML
img           = token "Img" >> token "{" >> fields << token "}"
  where fields = permuteSep (symbol ',') $
    Img <<< recordfield "src"    uri
      <<< recordfield "alt"    string
      <<< recordfield "longdesc" (justOrNothing uri)
      <<< recordfield "height"  (justOrNothing int)
      <<< recordfield "width"   (justOrNothing int)

```

We use *recordfield* here to parse a single optional record field, returning  $\perp$  if it is not present.

```

recordfield    :: Parser p => String -> p a -> p a
recordfield f p = field f p << succeed  $\perp$ 

```

## 5 Conclusion

We have shown how to extend a parser combinator library with the functionality to parse free-order constructs. It can be placed on top of any combinator library that implements the *Parser* interface. A user of the library can easily write parsers for free-order constructs and does not need to care about checking and reordering the parsed elements. Due to the use of existentially quantified types the implementation of reordering is type safe and hidden from the user.

The underlying parser combinators can be used to handle errors, such as missing or duplicate elements, since the extension inherits their error-reporting or error-repairing properties.

We have shown how our extension can be used to parse XML attributes and Haskell records. Other interesting examples mentioned by Cameron (1993) include citation fields in  $\text{\LaTeX}$  bibliographies and attribute specifiers in C declarations. Cameron's pseudo-code algorithm uses a similar strategy. It does not show, however, how to maintain type safety by undoing the change in semantics resulting from reordering, nor can it deal with the presence of separators between free-order constituents.

## References

- Cameron, R. D. (1993) Extending context-free grammars with permutation phrases. *ACM Lett. Program. Lang. Syst.* **2**(4), 85–94.
- Fokker, J. (1995) Functional parsers. *Advanced Functional Programming, First International Spring School: LNCS 925*, pp. 1–23. Springer-Verlag.
- Hutton, G. and Meijer, H. (1988) Monadic parser combinators. *J. Funct. Program.* **8**(4), 437–444.

- Leijen, D. (2001) *Parsec, a fast combinator parser*. <http://www.cs.uu.nl/~daan/parsec.html>.
- Peyton Jones, S. (2003) *Haskell 98 Language and Libraries*. Cambridge University Press. <http://www.haskell.org/report>.
- Røjemo, N. (1995) *Garbage collection and memory efficiency in lazy functional languages*. PhD thesis, Chalmers University of Technology.
- Swierstra, D. (2001) Combinator parsers: From toys to tools. In: Hutton, G. (editor), *Electronic Notes in Theoretical Computer Science*, vol. 41. Elsevier.
- Swierstra, D. and Duponcheel, L. (1996) Deterministic, error correcting combinator parsers. *Advanced Functional Programming, Second International Spring School: LNCS 1129*, pp. 184–207. Springer-Verlag.
- Wadler, P. (1985) How to replace failure with a list of successes. *Functional Programming Languages and Computer Architecture: LNCS 201*, pp. 113–128. Springer-Verlag.