

FUNCTIONAL PEARL

Type-safe pattern combinators

MORTEN RHIGER*

The IT University of Copenhagen, Copenhagen, Denmark
 (e-mail: mir@itu.dk)

Abstract

Macros still haven't made their way into typed higher-order programming languages such as Haskell and Standard ML. Therefore, to extend the expressiveness of Haskell or Standard ML, one must express new linguistic features in terms of functions that fit within the static type systems of these languages. This is particularly challenging when introducing features that span across multiple types and that bind variables. We address this challenge by developing, in a step by step manner, mechanisms for encoding patterns and pattern matching in Haskell in a type-safe way.

1 Untyped patterns

Let us tacitly ignore the existence of Haskell's pattern-matching facility, and let us attempt to extend Haskell with our own implementation of patterns and pattern matching:

```

data Pat = Pcst Int | Ppair Pat Pat | Pvar String
data Val = Vcst Int | Vpair Val Val

match :: Pat → Val → [(String, Val)]
match (Pvar x) v      = [(x, v)]
match (Pcst i) (Vcst j) = if i ≡ j then [] else error "match"
match (Ppair p q) (Vpair v w) = match p v ++ match q w
match _ _ = error "type error"

```

This solution defines data types of patterns (consisting of integer constants, pairs, and pattern variables) and of values (consisting of integer constants and pairs) and a function that matches a pattern against a value and that produces a list of bindings on success. The solution is *first order* and *untyped*: All type checks involving patterns are performed at runtime, by inspecting patterns and values. For example while an application of Haskell's built-in pattern-matching facility such as

```

case ((1, 2), (3, 4)) of
  (5, x) → 2 + x

```

*At the time of writing, the author is on leave from Roskilde University, Roskilde, Denmark.

results in compile-time type errors (in this particular case since the types of the pattern 5 and the value (1,2) mismatch and since the variable x is used as an integer but bound to a pair), the “type errors” in its encoded counterpart

```

case (match (Ppair (Pcst 5) (Pvar "x"))
          (Vpair (Vpair (Vcst 1) (Vcst 2))
                (Vpair (Vcst 3) (Vcst 4)))) of
  [("x", Vcst v)] → 2 + v
  _                → error "type error"

```

pass undetected through Haskell’s type checker.

The goal of this paper is to implement a library of type-safe patterns that can be extended with new and (hopefully) interesting patterns. Unlike the first-order solution above, our implementation does not require values to be encoded in terms of a data type, and it does not demand the use of Haskell’s built-in pattern matching to decode the values of bound variables.

The means are functional programming techniques: We demonstrate that polymorphic higher-order functions alone enable us to develop the mechanisms required to encode type-safe patterns in Haskell and Standard ML in a way that goes beyond their built-in pattern-matching facilities.

Example 1. Using the operations *match*, *pair*, *cst*, *var* and \rightarrow from the library implemented in Section 4 of this paper, the example presented above is encoded as follows:

```

match ((1, 2), (3, 4)) $
  pair (cst 5) var → λx → 2 + x

```

As required, this expression does not type check in Haskell.

Example 2. Again using the library from Section 4, the following example does type check and yields 7:

```

match (5, (3, 4)) $
  pair (cst 5) (pair var var) → λx y → x + y

```

2 Patterns exposing their type

We first introduce machinery allowing us to expose to Haskell’s type system the types of patterns and the structure of bindings produced on successful matches.

2.1 Exposing types

We want the *value* of the first argument (the pattern) to *match* to determine the *type* of the second argument (the value). For example, if the first argument is *Ppair (Pint 1) (Pint 2)*, then the second argument should be of type *(Int, Int)*.

We expose this relationship by representing patterns of type α as functions whose argument has type α , as follows (the intricacies of pattern variables having deliberately been postponed to the discussion below):

```

var x      = λv → ...
cst v'     = λv → if v ≡ v' then [] else error "match"
pair p q   = λv → p (fst v) ++ q (snd v)
match p v  = p v

```

This implementation is *typed*: All type errors that result from a mismatch between the pattern and the value are detected at compile time. Furthermore, run-time values need no longer be encoded and decoded using a data type. An additional advantage is that the combinator *cst*, unlike its first-order counterpart *Pcst*, is not restricted to integers but may be applied to values of any type that exhibit equality. As an example,

```
match (pair (cst 1) (cst True))
```

is a function that expects an argument of type $(Int, Bool)$ and matches this argument against the pattern $(1, True)$.

The technique applied here has been used by Danvy (1998) to implement *printf* in Standard ML, by Fridlender and Indrika (2000) to implement a generic *zipWith* in Haskell, and by Filinski (1999), Yang (2004), and the author (Rhiger, 1999) to implement type-directed partial evaluation in Standard ML. It is described in detail by Yang (2004).

2.2 Exposing bindings

In the first-order untyped implementation of patterns, the right-hand side of a clause would be represented by a function accepting a list of bindings and fetching the values of bound variables from this list as follows:

```
λenv → ...lookup "x" env ...
```

We must avoid such representations of pattern variables as strings, since they hinder reasoning about the number and the types of pattern variables within Haskell's type system.

For the time being, we will be content with fetching values of bound variables from a structure that is isomorphic to the pattern in question rather than from a list-like structure. For this purpose, we introduce a pattern *var* that marks a place in the pattern at which a variable occurs, without requiring us to name that variable. This idea results in a notion of *nameless* patterns, implemented as follows:

```

var        = λv → v
cst v'     = λv → if v ≡ v' then () else error "match"
pair p q   = λv → (p (fst v), q (snd v))
match p v  = p v

```

Again, this implementation is *typed*: All type errors that result from mismatches between the definition and the use of variables are detected at compile time. For example we have the following typings (with the left-hand side of the function arrow being the type of the value to match against and the right-hand side representing the bindings produced on a successful match):

```
match (pair var (pair (cst 2) var)) :: (α, (Int, β)) → (α, ((, β))
```

Therefore, the Haskell expression

```
case (1, (2, 3)) of (x, (2, y)) → x + y
```

is represented by

```
case match (pair var (pair (cst 2) var)) (1, (2, 3)) of
  (x, ((, y)) → x + y
```

And indeed, we're not required to encode values in a data type and there's no dynamic look-up of pattern variables.

However, as satisfying as this may seem, we're not quite done yet: The implementation still requires us to fetch the values of pattern variables inside a structure that resembles the original value that the pattern was matched against. Therefore, we must encode the structure of a pattern twice, first using the pattern combinators and then using Haskell's built-in patterns.

3 Patterns producing flat bindings

To improve the implementation presented above, we employ a *flattened* representation of bindings as *heterogeneous sequences*. The structure of such bindings is independent of the structure of the corresponding pattern, just as in the first-order untyped implementation of patterns. But unlike the first-order implementation, the heterogeneous sequences correctly expose their length and the types of their elements within Haskell's type system.

Assuming the existence of operations to build empty sequences (*nil*) and singleton sequences (*one v*) and to append two sequences (*m # n*), the implementation of patterns is defined as follows:

```
var      = λv → one v                -- bind v
cst v'   = λv → if v ≡ v' then nil else error "match" -- bind nothing or fail
pair p q = λv → p (fst v) # q (snd v) -- append bindings
```

We present two implementations of heterogeneous sequences below, one that results in a *direct-style program* that uses an accumulator and one that results in a style of programming resembling *continuation-passing style*.

3.1 Direct style, uncurried

We first represent sequences as nested pairs terminated by an empty tuple. Such a flattened representation encodes traditional lists by replacing `[]` by `()` and `(x : xs)` by `(x, xs)`. It allows elements of different types in one sequence.

More precisely, a sequence is represented as a function that accepts a tail and that "conses" (using a pair) its elements onto that tail, as follows:

```
nil      = λac → ac                -- cons no element
one v    = λac → (v, ac)           -- cons v
m # n    = λac → m (n ac)         -- let n cons first, then m
```

It is straightforward to verify that *nil* is both a left unit and a right unit of *#* and that *#* is associative. In other words, *#* and *nil* form a *monoid* over heterogeneous sequences.

Example 3. The following expression constructs a sequence containing an integer, a boolean, and a string:

```
one 5 # nil # one True # one "abc" # nil
```

The type of this expression, $\alpha \rightarrow (Int, (Bool, (String, \alpha)))$, indeed reflects the number and the types of the elements of the corresponding sequence, as required. Passing this expression the empty tuple $()$ yields $(5, (True, ("abc", ())))$.

When encoding heterogeneous sequences in direct style like this, the well-typed Haskell expression

```
case ((1, 2), 3) of ((x, 2), y) → x + y
```

can be represented by the well-typed

```
case (match (pair (pair var (cst 2)) var) ((1, 2), 3)) of
  (x, (y, ())) → x + y
```

3.2 Continuation-passing style, curried

Even though the structure of bindings is now independent of the structure of the corresponding pattern, we still rely on Haskell's built-in pattern matching to fetch the values of bound pattern variables. (Alternatively, we may use the functions *fst* and *snd* to fetch the values from the sequence, but such a solution would be cumbersome.)

To overcome this problem and to make bindings digestible, we *curry* the functions that process them. To enable such a change, we adopt a representation of sequences as functions that pass all the elements of a sequence to a given curried continuation. The operations on sequences are defined as follows:

```
nil      = λk → k           -- pass no values to k
one v    = λk → k v        -- pass v to k
m # n    = λk → n (m k)    -- let m pass values first, then n
```

It is again straightforward to verify that *#* and *nil* form a monoid over heterogeneous sequences. Notice that according to this implementation, any non-functional value is a “continuation” of zero arguments: In general, the representation of a sequence of length *n* passes its elements to a function of *n* arguments. When *n* = 0, such a “function” can be a value of any type. As we shall see in the next section, this generalisation allows us to represent a right-hand side of a pattern-matching clause as a value of base type when the corresponding pattern is “ground” (that is, one that does not contain variables).

Example 4. The following expression constructs a sequence containing an integer, a boolean, and a string:

$$\text{one } 5 \# \text{ nil} \# \text{ one True} \# \text{ one "abc"} \# \text{ nil}$$

The type of this expression, $(Int \rightarrow Bool \rightarrow String \rightarrow \alpha) \rightarrow \alpha$, again reflects the number and the types of the elements of the corresponding sequence. Passing this expression the continuation $\lambda x y z \rightarrow (x, y, z)$ yields $(5, True, "abc")$.

When encoding heterogeneous sequences in continuation-passing style like this, the well-typed Haskell expression

$$\text{case } ((1, 2), 3) \text{ of } ((x, 2), y) \rightarrow x + y$$

can be represented by the well-typed

$$\begin{aligned} & \text{match } (\text{pair } (\text{pair var } (\text{cst } 2)) \text{ var}) \\ & \quad ((1, 2), 3) \\ & \quad (\lambda x y \rightarrow x + y) \end{aligned}$$

Thus, we have replaced the need for Haskell's built-in pattern matching to fetch the values of pattern variables by the application of a curried function.

Danvy's statically typed implementation of *printf* in Standard ML (Danvy, 1998) used curried continuations in a similar fashion. But where Danvy's formatting combinators *consume* a heterogeneous sequence of values, our pattern-matching combinators *produce* such a heterogeneous sequence of values.

4 A library of typed patterns

With the basic mechanisms at hand, let us design a library that takes us beyond the capabilities of Haskell's built-in pattern matching. To this end, we extend the language of patterns presented so far with *disjunctions* and *conjunctions* (also known as *or* patterns and *and* patterns), *predicates*, and a *set-like view* on lists.

To support disjunctions, we add a failure continuation to the encoding of sequences. However, this addition rules out the implementation above in which currying happens *on the fly*. (In a curried implementation, the failure continuation is expected to consume bindings. This contradicts the intuition that the failure continuation is applied only when a pattern does not match and hence when no bindings are available.)

To address this issue, we employ an uncurried encoding of bindings as nested pairs during the construction of patterns, as in the direct-style implementation above. We then uncurry a given curried (success) continuation before passing it to a pattern. To this end, we use the following typed and generic implementation of uncurrying of functions:

$$\begin{aligned} \text{zero } f &= \lambda() \rightarrow f \\ \text{succ } n f &= \lambda(x, xs) \rightarrow n (f x) xs \\ \text{uncurry } n f &= n f \end{aligned}$$

The functions *zero* and *succ* define a representation of numerals with the property that an n -ary curried function can be uncurried by the representation of n .¹ Fridlender and Indrika (2000) also use a higher-order representation of numerals in their generic and well-typed implementation of an n -ary *zipWith*.

Example 5. To uncurry a function of two arguments, we apply the following representation of the numeral 2 to the function:

$$\text{succ } (\text{succ } \text{zero}) :: (\alpha \rightarrow \beta \rightarrow \gamma) \rightarrow (\alpha, (\beta, ())) \rightarrow \gamma$$

4.1 Basic patterns

We represent heterogeneous sequences in continuation-passing style, using two continuations (a success continuation k_s and a failure continuation k_f) and an accumulator. We also add two constructions, *fail* and *catch*, that introduce an exception mechanism into the language of sequences:

$$\begin{aligned} \text{nil} &= \lambda k_s k_f ac \rightarrow k_s ac \\ \text{one } v &= \lambda k_s k_f ac \rightarrow k_s (v, ac) \\ m \# n &= \lambda k_s k_f ac \rightarrow n (m k_s k_f) k_f ac \\ \text{fail} &= \lambda k_s k_f ac \rightarrow k_f () \\ m \text{ 'catch' } n &= \lambda k_s k_f ac \rightarrow m k_s (\lambda () \rightarrow n k_s k_f ac) ac \end{aligned}$$

Both ($\#$, *nil*) and (*catch*, *fail*) form monoids over heterogeneous sequences.

A pattern consists of a pair of functions. The first increments a given numeral (in the representation presented above) by the number of variables of the pattern. Eventually, this component is used to curry success continuations. The second component implements the matching itself. The basic patterns are defined as follows:

$$\begin{aligned} \text{var} &= (\text{succ}, \lambda v \rightarrow \text{one } v) \\ \text{cst } v' &= (\text{id}, \lambda v \rightarrow \mathbf{if } v \equiv v' \mathbf{ then nil else fail}) \\ \text{pair } p \ q &= (\text{curry}_{pq}, \lambda v \rightarrow \text{match}_p (\text{fst } v) \# \text{match}_q (\text{snd } v)) \\ &\quad \mathbf{where } (\text{curry}_p, \text{match}_p) = p \\ &\quad \quad (\text{curry}_q, \text{match}_q) = q \\ &\quad \quad \text{curry}_{pq} = \text{curry}_p \cdot \text{curry}_q \end{aligned}$$

4.2 Matching

Before we continue, let us introduce a more realistic main function. To this end, we introduce *clauses* (pairs of patterns and right-hand sides) using the infix operator \rightarrow and we separate two such clauses by $\|$. It is \rightarrow that curries continuations properly. These operations are defined as follows:

¹ Our use of the term “uncurried” is unconventional: Rather than taking n -tuples as arguments, our uncurried functions take nested pairs of depth n as arguments.

$$\begin{aligned}
 p \rightarrow k &= \lambda v \ k_f \rightarrow \text{match}_p \ v \ (\text{curry}_p \ \text{zero} \ k) \ k_f \ () \\
 &\quad \mathbf{where} \ (\text{curry}_p, \text{match}_p) = p \\
 c_1 \parallel c_2 &= \lambda v \ k_f \rightarrow c_1 \ v \ (\lambda() \rightarrow c_2 \ v \ k_f)
 \end{aligned}$$

The following main operation then matches a given value against a group of clauses:²

$$\text{match } v \ cs = cs \ v \ (\lambda() \rightarrow \text{error } \text{"match"})$$

Example 6. With the correct precedences, these definitions allow us to write expressions such as the following. (Here \$ is an infix apply operator of low precedence commonly used as an alternative to parenthesising a function's argument. In the rest of this article, we read it as “with”.)

$$\begin{aligned}
 \text{match } n \ \$ \\
 \quad \text{cst } 0 &\rightarrow \text{"zero"} \\
 \quad \parallel \text{cst } 1 &\rightarrow \text{"one"} \\
 \quad \parallel \text{var} &\rightarrow \lambda i \rightarrow \text{error } (\text{"not a binary digit: " } \# \text{ show } i)
 \end{aligned}$$

Since patterns are first-class values, *pattern abstraction* (Fähndrich & Boyland, 1997) is at our hands to extend the library of patterns. As the first example, we can introduce anonymous pattern-matching functions as follows, using the infix operator \mapsto :

$$p \mapsto k = \lambda v \rightarrow \text{match } v \ \$ \ p \rightarrow k$$

Example 7. The following expression sums each pair of numbers in a list:

$$\text{map } (\text{pair } \text{var } \text{var} \mapsto \lambda x \ y \rightarrow x + y) :: [(Int, Int)] \rightarrow [Int]$$

4.3 Disjunctions

The disjunctive aspect is represented by the pattern combinators *none* and \vee , defined below. The pattern *none* does not match any value.

$$\begin{aligned}
 \text{none} &= (id, \quad \lambda v \rightarrow \text{fail}) \\
 p \vee q &= (\text{curry}_{pq}, \lambda v \rightarrow \text{match}_p \ v \ \text{'catch' } \text{match}_q \ v) \\
 &\quad \mathbf{where} \ (\text{curry}_p, \text{match}_p) = p \\
 &\quad \quad (\text{curry}_q, \text{match}_q) = q \\
 \text{curry}_{pq} &= \mathbf{if } \text{True} \ \mathbf{then} \ \text{curry}_p \ \mathbf{else} \ \text{curry}_q
 \end{aligned}$$

The combinators \vee and *none* form a monoid over patterns. The conditional in the implementation of disjunction serves to unify the types of the currying component associated with the two patterns. As a result, the two branches are required to produce bindings of equal type.

² To give it a stronger resemblance with Haskell's **case-of-** construct, this final version of *match* takes the value to match against as its first argument.

Example 8. The pattern given below matches pairs whose first *or* second component is 1. For successful matches, it binds a variable to the other component:

$$\text{pair } (\text{cst } 1) \text{ var } \vee \text{pair var } (\text{cst } 1)$$

4.4 Conjunctions

The conjunctive aspect is represented by the pattern combinators *any* and \wedge , as shown below. The pattern *any* matches any value. It corresponds to the wildcard `_` in Haskell.

Conjunctions are implemented as follows:

$$\begin{aligned} \text{any} &= (\text{id}, \lambda v \rightarrow \text{nil}) \\ p \wedge q &= (\text{curry}_p \cdot \text{curry}_q, \lambda v \rightarrow \text{match}_p v \# \text{match}_q v) \\ &\quad \mathbf{where} \ (\text{curry}_p, \text{match}_p) = p \\ &\quad \quad (\text{curry}_q, \text{match}_q) = q \end{aligned}$$

These combinators also form a monoid over patterns. Conjunction generalises layered patterns: A pattern like $\text{var} \wedge p$ matches any value that p matches and binds the value to a variable.

The infix operator $?$ matches a ground pattern against a value and yields a boolean value representing the status of the match:

$$p \ ? \ v = \text{match } v \ \$ \ p \ \rightarrow \ \text{True} \ || \ \text{any} \ \rightarrow \ \text{False}$$

Example 9. The following expression removes from a list all pairs whose first component is non-zero:

$$\text{filter } (\text{pair } (\text{cst } 0) \ \text{any} \ ?) :: [(Int, \alpha)] \rightarrow [(Int, \alpha)]$$

4.5 Predicates

It is sometimes useful to employ a predicate during matching. The following operation yields a pattern that matches only the values for which a given predicate p returns *True*:

$$\text{is } p = (\text{id}, \lambda v \rightarrow \mathbf{if } p \ v \ \mathbf{then } \text{nil} \ \mathbf{else } \text{fail})$$

Example 10. The following function implements integer exponentiation, using the Russian peasant algorithm:

$$\begin{aligned} \text{power } x \ n = & \\ \text{match } n \ \$ & \\ \text{cst } 0 & \ \rightarrow 1 \\ || \ \text{is even} & \ \rightarrow \text{square } (\text{power } x \ (n \ \text{div } 2)) \\ || \ \text{is odd} & \ \rightarrow x \times \text{power } x \ (n - 1) \\ \mathbf{where} \ \text{square } x &= x \times x \end{aligned}$$

Example 11. The following expression removes from a list all those pairs whose first component is odd:

$$\text{filter } (\text{pair } (\text{is even}) \text{ any } ?) :: [(Int, \alpha)] \rightarrow [(Int, \alpha)]$$

It is tempting to write

$$\text{filter } (\text{is even } ?) :: [Int] \rightarrow [Int]$$

but one can prove that the “left section”

$$(\text{is } p \ ?)$$

(which is a partial application of the infix operator $?$ to its first argument $\text{is } p$) is equivalent to p , for any function p of type $\alpha \rightarrow \text{Bool}$.

4.6 Sets

It may also be useful to view a list as if it were a set of values. The pattern *has* p scans a list for the first element matched by the pattern p . If no element matches, then the entire pattern fails:

$$\begin{aligned} \text{has } p &= (\text{curry } p, \text{foldr } \text{catch fail} \cdot \text{map } \text{match } p) \\ &\quad \text{where } (\text{curry } p, \text{match } p) = p \end{aligned}$$

Example 12. The function *fetch*, of type $\text{String} \rightarrow [(String, \alpha)] \rightarrow \alpha$, implements a look-up function:

$$\text{fetch } v = \text{has } (\text{pair } (\text{cst } v) \text{ var}) \mapsto \lambda x \rightarrow x$$

Example 13. The function *first*, of type $(\alpha \rightarrow \text{Bool}) \rightarrow [\alpha] \rightarrow \alpha$, implements a function that takes a list and returns the first element of the list that satisfies the predicate f :

$$\text{first } f = \text{has } (\text{var } \wedge \text{is } f) \mapsto \lambda x \rightarrow x$$

We can take Example 12 a step further by implementing look-up as a pattern rather than as a function. The resulting pattern binds a variable:

$$\text{get } v = \text{has } (\text{pair } (\text{cst } v) \text{ var})$$

Example 14. We can not only implement *fetch* using the pattern combinator *get*, but also more complex matches. The code snippet below produces the value associated with either “a” or “A” in the list of pairs *env*. It complains if either both or none of the bindings exist (but not if two occurrences of “a” exists in the list):

$$\begin{aligned} \text{match } \text{env } \$ \\ \quad \text{get "a" } \wedge \text{get "A" } &\rightarrow (\lambda_ _ \rightarrow \text{error "ambiguity"}) \\ \quad \|\ \text{get "a" } \vee \text{get "A" } &\rightarrow (\lambda a \rightarrow a) \\ \quad \|\ \text{any} &\rightarrow \text{error "unbound"} \end{aligned}$$

4.7 Data abstraction

Since pattern combinators are first-class values, an abstract data type can be associated with a set of patterns without exposing its implementation. In the sense of Wadler (1987), we can define patterns that provide a concrete *view* of an abstract data type. For example the function below uses pattern matching to define f^n , the n -times composition of a function f with itself. However, rather than fixing the representation of the integer n , it allows any representation that can be viewed as natural numbers built using a zero and a successor. This is accomplished by parameterising the implementation over two patterns combinators, *zero* and *succ*:

```
makeCompose zero succ = iterate
  where iterate n f x =
        match n $
          zero      →→ x
        || succ var →→ λn' → iterate n' f (f x)
```

As witnessed by the following definitions of actual pattern combinators, we can view both built-in integers and lists as natural numbers:

```
zeroint  = (id,    λv → if v ≡ 0 then nil else fail)
succint p = (curryp, λv → if v ≡ 0 then fail else matchp (v - 1))
           where (curryp, matchp) = p
zerolist = (id,    λv → if null v then nil else fail)
succlist p = (curryp, λv → if null v then fail else matchp (tail v))
           where (curryp, matchp) = p
```

We instantiate the general function as follows:

```
composeint = makeCompose zeroint succint
composelist = makeCompose zerolist succlist
```

Using these specialised functions, both of the following expressions yield 8:

```
composeint 3 (×2) 1
composelist [10, 20, 30] (×2) 1
```

4.8 On type safety

The type safety of the pattern combinators thus implemented follows by construction: There are no dynamic type checks in their implementation, and they are well typed according to Haskell's type checker.

4.9 On efficiency

Matching is linear in the size of patterns that do not involve the combinator *has*. (Matching the pattern *has p* against a list amounts to matching *p* against each element in the list.) Experiments performed using Hugs suggest that matching of pattern combinators require two to four times as many reductions as matching using

Haskell's built-in patterns. This overhead is primarily induced by managing continuations and by currying. Standard partial-evaluation techniques (β reduction and η expansion) can completely compile away both the manipulation of continuations and currying, hence eliminating the overhead.

5 Conclusion

New languages are often constructed by piling new features on top of an existing language's definition and by integrating these features in the existing language's implementation. However, it is a sign of expressiveness if new features can be implemented *within* an existing language without changing its definition.

Short of macros, functional languages such as Haskell and Standard ML require new features to be expressed in terms of typed higher-order functions. We have demonstrated how to extend—or, in Guy Steele's terminology, to “grow” (Steele, Jr., 1999)—Haskell with our own statically typed implementation of pattern matching, and we have shown how to extend this framework with patterns not currently supported by Haskell.

References

- Danvy, O. (1998) Functional unparsing. *J. Funct. Prog.* **8**(6), 621–625.
- Fähndrich, M. & Boyland, J. (1997) Statically checkable pattern abstractions. *Pages 75–84 of: Tofte, Mads (ed), Proceedings of the 1997 ACM SIGPLAN international conference on functional programming (ICFP 1997)*. Amsterdam, The Netherlands: ACM Press.
- Filinski, A. (1999) A semantic account of type-directed partial evaluation. In: *International conference on principles and practice of declarative programming (PPDP 1999)*, Nadathur, G. (ed). Lecture Notes in Computer Science, vol. 1702. Paris, France: Springer-Verlag, pp. 378–395.
- Fridlender, D. & Indrika, M. (2000) Do we need dependent types? *J. Funct. Prog.* **10**(4), 409–415.
- Rhiger, M. (1999) Deriving a statically typed type-directed partial evaluator. In: *Proceedings of the ACM SIGPLAN workshop on partial evaluation and semantics-based program manipulation (PEPM 1999)*, Danvy, O. (ed). BRICS Note Series, no. NS-99-1, Department of Computer Science, University of Aarhus, Denmark, pp. 25–29.
- Steele, Jr., G. L. (1999) Growing a language. *Higher-Order Symb. Comput.* **12**(3), 221–236.
- Wadler, P. (1987) Views: A way for pattern matching to cohabit with data abstraction. In: *Proceedings of the fourteenth annual ACM symposium on principles of programming languages (POPL 1987)*, O'Donnell, M. J. (ed). Munich, West Germany: ACM Press, pp. 307–313.
- Yang, Z. (2004) Encoding types in ML-like languages. *Theoret. Comput. Sci.* **315**(1), 151–190.