



Centrum voor Wiskunde en Informatica

REPORTRAPPORT

SEN

Software Engineering



Software ENgineering

Strategic polymorphism requires just two combinators!

Ralf Lämmel, Joost Visser

REPORT SEN-E0324 DECEMBER 23, 2003

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).

CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2003, Stichting Centrum voor Wiskunde en Informatica

P.O. Box 94079, 1090 GB Amsterdam (NL)

Kruislaan 413, 1098 SJ Amsterdam (NL)

Telephone +31 20 592 9333

Telefax +31 20 592 4199

ISSN 1386-369X

Strategic polymorphism requires just two combinators!

ABSTRACT

In previous work, we introduced the notion of functional strategies: first-class generic functions that can traverse terms of any type while mixing uniform and type-specific behaviour. Functional strategies transpose the notion of term rewriting strategies (with coverage of traversal) to the functional programming paradigm. Meanwhile, a number of Haskell-based models and combinator suites were proposed to support generic programming with functional strategies. In the present paper, we provide a compact and matured reconstruction of functional strategies. We capture strategic polymorphism by just two primitive combinators. This is done without commitment to a specific functional language. We analyse the design space for implementational models of functional strategies. For completeness, we also provide an operational reference model for implementing functional strategies (in Haskell). We demonstrate the generality of our approach by reconstructing representative fragments of the Strafunski library for functional strategies.

1998 ACM Computing Classification System: D.3.3

Keywords and Phrases: Functional Strategies; Strategic Polymorphism; Strafunski

Strategic polymorphism requires just two combinators!

Ralf Lämmel^{1,2} and Joost Visser^{1,3}

¹ CWI, Kruislaan 413, NL-1098 SJ Amsterdam

² Vrije Universiteit, De Boelelaan 1081a, NL-1081 HV Amsterdam

³ Software Improvement Group, Kruislaan 419, NL-1098 SJ Amsterdam

Email: (Ralf.Laemmel|Joost.Visser)@cwi.nl

Abstract

In previous work, we introduced the notion of *functional strategies*: first-class generic functions that can traverse terms of any type while mixing uniform and type-specific behaviour. Functional strategies transpose the notion of term rewriting strategies (with coverage of traversal) to the functional programming paradigm. Meanwhile, a number of Haskell-based models and combinator suites were proposed to support generic programming with functional strategies.

In the present paper, we provide a compact and matured reconstruction of functional strategies. We capture strategic polymorphism by just two primitive combinators. This is done without commitment to a specific functional language. We analyse the design space for implementational models of functional strategies. For completeness, we also provide an operational reference model for implementing functional strategies (in Haskell). We demonstrate the generality of our approach by reconstructing representative fragments of the *Strafunski* library for functional strategies.

1 Introduction

In [25], we introduced the notion of *functional strategies* for which we assume the following matured definition throughout this paper:

Definition 1 *Functional strategies are functions that*

1. *are generic,*
2. *can mix uniform and type-specific behaviour,*
3. *can traverse terms, and*
4. *are first-class citizens.*

Functional strategies go beyond parametrically polymorphic functions because of the abilities to traverse terms and to dispatch to type-specific behaviour. We call this extra polymorphism simply ‘strategic polymorphism’. In the present paper, we will capture strategic polymorphism by just two fundamental function combinators. Most of our presentation will avoid commitment to a specific functional language, but we ultimately define a Haskell-based reference model.

Functional strategies were derived from the notion of (typed) term rewriting strategies [8, 4, 31, 7, 5, 22]. In fact, the notion of (traversal) strategies can be amalgamated with different programming paradigms. We use the term *strategic programming* for generic programming with strategies in whatever language. It is at the heart of strategic programming that traversal schemes are programmer-definable. The native application domain of strategic programming is language processing, in particular, the implementation of functionality for program transformations and analyses; see [25, 21, 26] for a few typical applications.

With strategies, one can operate on large syntaxes or formats in a scalable and flexible manner. Scalability is implied by genericity, and flexibility by a combinator style that enables the definition of appropriate traversal schemes. In fact, the ability to traverse terms while mixing uniform and type-specific behaviour is beneficial in almost every non-trivial software application. This is demonstrated in [23], where an approach to *general purpose generic programming* is based on original expressiveness for ‘strategic polymorphism’.

Our efforts to amalgamate strategic and functional programming are scoped by the *Strafunski* project.¹ Distributions of the Haskell-centred *Strafunski* bundle for generic programming and language processing include generative tool support that, given the programmer-supplied datatypes, generates the code needed for strategic programming. In previous work, we came up with different models of functional strategies [25, 20]. These models differ regarding the selection of primitive combinators and their types, but they all share Haskell as the base language.

In the present paper, we capture strategic polymorphism by just two combinators:

- *adhoc* for type-based function dispatch, and
- *hfolds* for folding over constructor applications.

The combinator couple *adhoc* and *hfolds* was identified in [20], and a very similar couple was employed in [23].² In the present paper, we use these two combinators for a very compact and matured reconstruction of functional strategies. Our goal here is to avoid an invasive commitment to Haskell. Also, we want to clearly maintain the link to strategic programming as it was initiated in the context of term rewriting. The above two combinators can be mapped to Def. 1 as follows. The combinator *adhoc* is crucial for mixing uniform and type-specific behaviour. It can be considered as a disciplined form of type case. The combinator *hfolds* is the mother of all (one-layer, i.e., non-recursive) traversal. It folds over the immediate subterms of a constructor application — very much in the style of a list fold. This turns traversal schemes into programmable entities where ordinary recursive function definition suffices to complete folding into recursive traversal.

The paper is structured as follows. In Sec. 2, we approach to the essential expressiveness for strategic polymorphism via a motivating example. In Sec. 3, we define the two key combinators for strategic polymorphism. In Sec. 4, we discuss models of functional strategies. Sec. 3 and Sec. 4 are largely language-independent. In Sec. 5, we provide a reference model to extend Haskell with our two combinators. In Sec. 6, we demonstrate the power that results from our language extension by reconstructing representative parts of *Strafunski*’s library. In Sec. 7, the paper is concluded.

2 Strategic polymorphism — a motivating example

We choose a simple but challenging example that demonstrates all the characteristics of functional strategies. We will define a combinator *query* with an argument *f*, such that *query f* performs a top-down, left-to-right traversal to find a subterm that can be processed by *f* so that a value is extracted from the subterm. A suitable subterm has to meet two criteria. Firstly, its type must coincide with the domain of *f*. Secondly, *f* applied to the subterm should not fail, that is, it should not return *Nothing*.

This is the Haskell type of the *query* combinator:

```

query :: ∀α β u. (Term α, Term β)
  ⇒ (α → Maybe u)  -- Recogniser / extractor
  → β                -- Input term
  → Maybe u          -- Found entity (if any)

```

¹*Strafunski* home page: <http://www.cs.vu.nl/Strafunski/> — *Stra* refers to strategies, *fun* refers to functional programming, and their harmonious composition is a homage to the music of Igor Stravinsky.

²There are tiny technical differences. In [20], first-class polymorphism is employed as opposed to proper rank-2 function types in the present paper and in [23]. Also, in [23], type cast is favoured as opposed to function dispatch in the present paper and in [20]. Furthermore, in [23], a left-associative fold combinator is used as opposed to the right-associative combinator in the present paper and in [20].

There are three universally quantified type parameters:³

- α denotes the type of ‘relevant’ subterms from which a value is extracted.
- β denotes the type of the term that is eventually passed to *query*.
- u denotes the type of the extracted value.

The first argument of *query* is a function of type $\alpha \rightarrow \text{Maybe } u$ that is meant to interrogate subterms of type α to extract a value of type u . The second argument is the input term of type β . The result type of *query* is the extracted value u (if any), wrapped in the partiality monad *Maybe*. There are two class constraints *Term* α and *Term* β which point out that α and β are place-holders for term types. The Haskell type class *Term* hosts expressiveness for ‘strategic polymorphism’, that is, basically the two combinators *adhoc* and *hfoldr* as we will say later.

Before we describe the actual definition of *query*, let us first point out that this sort of function is useful for all kinds of scenarios in program transformation and analysis:

- Query the type of an abstraction. For instance, from a Java class declaration one may want to retrieve the signature of a method with a given name.
- Query the definition of an abstraction. For instance, from an XML DTD, one may want to retrieve the content specification of a declared element type.
- Query the focused entity in the course of refactoring. For example, from a Cobol program, one may want to retrieve a focused group of data description entries.

One can easily think of numerous variations on *query* that are equally useful, e.g., implementing bottom-up search, or transforming terms rather than querying them, and so on.

Let’s turn to the Haskell definition of the *query* combinator:

```

query f x
= case adhocMTU (const Nothing) f x of  -- Attempt f
    Just u → Just u                               -- Done
    Nothing → oneMTU (query f) x      -- Recurse

```

We boxed the two combinators that involve strategic polymorphism. (*adhocMTU* is a type-specialised variant of *adhoc*, and *oneMTU* is defined in terms of *hfoldr* as we will show later.) For clarity, their names end on “*MTU*” to remind us of the strategy type at hand: *Monadic Type-Unifying*. Monadic style is in place here because of the partiality of querying. By ‘type unification’ we mean that querying returns a value of a specific type, regardless of the type of the input term. The combinator *adhocMTU* is used to attempt an application of f to the input term x . The composed function behaves like the polymorphic function *const Nothing* passed as the first argument to *adhocMTU* except for the type α handled by the function f . If the attempt to apply f results in a value u (first branch of **case**), the search is done, and u is returned. Otherwise (second branch of **case**), the combinator *oneMTU* is used to call *query* recursively on the children of x , i.e., its immediate subterms. The *oneMTU* combinator is meant to attempt application of its argument to each child in left-to-right order, and returns the result of the first application that succeeds. We will later see that *oneMTU* is just one example of many ‘one-layer’ traversal combinators — they can all be defined in terms of the fundamental combinator *hfoldr*.

To summarise, the *query* combinator illustrates all characteristics of strategies (recall 1.–4. in Def. 1):

1. The aspect that the *query* combinator is *generic*, i.e., applicable to *all* term types, is expressed by the universally quantified β in its type.

³For the purpose of a homogeneous notation, we always use *explicit* (top-level) universal quantification in all Haskell type signatures. Note: the type of *query* is just a plain rank-1 type, that is, all type variables are indeed quantified at the top-level. We will later also employ rank-2 types, and thereby go beyond Haskell 98. Rank-2 types make *explicit* quantification mandatory. So we decide to switch to explicit quantification all-over the place.

2. The aspect that the *query* combinator *mixes uniform and type-specific behaviour* is reflected by the application of the combinator *ad hocMTU*.
3. The aspect that the *query* combinator *traverses* the input term is reflected by the recursive definition of *query* in terms of the traversal combinator *oneMTU*.
4. The *first-class* status of functional strategies is illustrated by their featuring as combinator arguments (cf. the first argument of *ad hocMTU* and *oneMTU*).

We refer the reader to [25, 21, 26] for many more examples of strategic polymorphism. The present example has been carefully chosen to cover all aspects of strategic polymorphism, and to be representative for strategic programming.

3 Just two combinators for strategic polymorphism

In the sequel, we define the two combinators *ad hoc* and *hfoldr* for strategic polymorphism. The definition is language-independent. We use a semi-formal style for the semantics of the combinators, and we use rank-2 types to assign types to the combinators. The interested reader is referred to [22] for a formal definition of typed rewriting strategies in a basically first-order and many-sorted term-rewriting setting. The below definition makes heavy use of higher-orderness and polymorphism. This is the key to capturing strategic expressiveness in just two dedicated combinators as opposed to the several combinators in previous work.

The *ad hoc* combinator for type-based function dispatch

We want to give a single definition of *ad hoc*, which is valid for all different types of functional strategies (think of querying vs. transformation). So we assume the general type scheme $\overline{\forall}\alpha. \alpha \rightarrow c \alpha$ for functional strategies. Here, c is a type constructor that derives the co-domain of the function type from the domain α . If we instantiate the co-domain constructor c with the identity type constructor (i.e., $I \alpha = \alpha$), we obtain the type of type-preserving strategies. The constant type constructor (i.e., $C u \alpha = u$) handles the type-unifying scheme (i.e., the result type is always u regardless of the input type α). We use an over-lined version $\overline{\forall}$ of the universal quantifier \forall to distinguish ‘strategic polymorphism’ from ordinary parametric polymorphism.

$$\begin{array}{ll}
 \textit{ad hoc} & : \quad \forall c. & \text{-- Co-domain constructor} \\
 & \quad \overline{\forall}\beta. & \text{-- Term type of specific functionality} \\
 & \quad (\overline{\forall}\alpha. \alpha \rightarrow c \alpha) & \text{-- Generic default} \\
 & \rightarrow (\beta \rightarrow c \beta) & \text{-- Type-specific ad-hoc case} \\
 & \rightarrow (\overline{\forall}\gamma. \gamma \rightarrow c \gamma) & \text{-- Constructed strategy}
 \end{array} \tag{i}$$

$$\textit{ad hoc } p m x = \begin{cases} m x, & \text{if } \textit{typeOf}(x) = \textit{domOf}(m) \\ p x, & \text{otherwise} \end{cases}$$

The placing of the various $\overline{\forall}$ quantifiers in the (rank-2) type of *ad hoc* emphasises that the combinator constructs a polymorphic function from a polymorphic function argument p and a monomorphic function argument m .⁴ In the definition, the expressions *typeOf*(x) and *domOf*(m) denote the *specific* type of x and the *specific* domain of m , i.e., the respective types to which γ and β are instantiated *at run-time*. So, following this definition, *ad hoc* $p m x$ dispatches to the monomorphic m if m is applicable to the x at hand, but otherwise it resorts to the polymorphic, generally applicable p .

In strategic programming, the common usage of *ad hoc* is to derive *generic* ‘rewrite steps’ from *type-specific* ones [24]. The monomorphic ingredient m is then a function which rewrites terms of a specific

⁴Reminder: $\forall/\overline{\forall}$ quantifiers expand to the right as far as possible. Lifting the $\overline{\forall}\gamma$ to the top of the function type would be acceptable, because $\overline{\forall}\gamma$ is placed on the right of the outermost \rightarrow , and hence, lifting leads to an equivalent type. By contrast, moving the inner $\overline{\forall}\alpha$ to the top would lead to a too liberal function that also accepted monomorphic arguments for p . Dually, pushing the $\overline{\forall}\beta$ to the second argument (where it is used) would lead to a too restrictive function that insisted on polymorphic ad-hoc cases for m instead of type-specific functionality.

type, potentially based on pattern matching, whereas the polymorphic ingredient p provides a trivial generic default. As a simple example, consider the following polymorphic function $negbool$ which behaves like the identity function by default but applies Boolean negation “ \neg ” when faced with a Boolean:

$$negbool = \text{ad hoc id } (\neg)$$

For this type-preserving strategy, the co-domain type constructor c is instantiated to the identity type constructor I . In our *query* sample, *ad hoc* was used to construct a type-unifying generic rewrite step “*ad hoc (const Nothing) f*”. Here, the co-domain constructor c is instantiated to the constant type constructor C . The generic default *const Nothing* models the failure of the *query*. The function f interrogates terms of a certain type to extract a *Maybe* value.

The *hfoldr* combinator for folding over constructor applications

Let us recall the basic idiom of list traversal. Without loss of generality, we consider *right-associative* list traversal. Folding a list $[x_1, x_2, \dots, x_n]$ according to ingredients f and z for the non-empty and the empty list form is defined as usual:

$$\text{foldr } f z [x_1, x_2, \dots, x_n] = f x_1 (f x_2 (\dots (f x_n z) \dots))$$

Folding over constructor applications is similar: instead of folding over the elements of a homogeneous list, we fold over the children of a term, i.e., its immediate subterms. Since the children of a term are potentially of different types, we need a *heterogeneous* fold. So we use the name *hfoldr*. Without loss of generality, we assume *curried* constructor applications. Folding a term $C x_n \dots x_2 x_1$ is now defined as follows:

$$\begin{array}{ll} \text{hfoldr} & : \quad \forall c. & \text{-- Co-domain constructor} \\ & (\forall \alpha \forall \beta. \alpha \rightarrow c (\alpha \rightarrow \beta) \rightarrow c \beta) & \text{-- 'cons' case} \\ & \rightarrow (\forall \gamma. \gamma \rightarrow c \gamma) & \text{-- 'nil' case} \\ & \rightarrow (\forall \delta. \delta \rightarrow c \delta) & \text{-- Constructed strategy} \end{array} \quad [\text{ii}]$$

$$\text{hfoldr } f z (C x_n \dots x_2 x_1) = f x_1 (f x_2 (\dots (f x_n (z C)) \dots))$$

The indices in $C x_n \dots x_2 x_1$ clarify that we treat the children in the curried expression as a ‘snoc list’ with the rightmost child as ‘head’. Also note that the empty constructor application C is passed to z so that the constructor can contribute to the result of folding. The type of *hfoldr* is somewhat involved, especially regarding the argument f . Here, α denotes the type of the current head x_i , i.e., the next subterm, and β denotes the type of a fragment of $C x_n \dots x_i$. The type $c (\alpha \rightarrow \beta)$ denotes the recursively processed tail, and it reflects that this tail lacks the i^{th} child of type α . The type of *hfoldr*’s argument z of *hfoldr* is less of a headache. Since z is meant to process constructors, the γ in its type denotes the type of the constructor C , applied to no children.

This folding operation is now sufficient to define arbitrary *one-layer* traversal combinators which in turn can be completed into recursive traversal schemes in different ways. In the sample section, we assumed a one-layer traversal combinator *oneMTU* which can now be defined concisely as follows:

$$\text{oneMTU } s = \text{hfoldr } (\lambda h t \rightarrow t \text{ 'mplus' } s h) (\text{const } mzero)$$

Here we assume an extended monad with operations *mplus* for a kind of choice, and *mzero* for failure. The *Maybe* monad is a typical representative of this class of monads. The definition states that the argument strategy s is applied to the head h of the constructor application, and *mplus* is used to combine the recursively processed tail t and the processed head. Folding starts from *mzero*. In Sec. 6, we will provide definitions of more one-layer traversal combinators, and we derive several typical recursive traversal combinators — just in the same way as *query* was derive from *oneMTU* by means of ordinary recursive function definition.

4 Implementational models — a detailed analysis

The two combinators *adhoc* and *hfoldr*, which capture strategic polymorphism, can be modelled in several ways. In this section, we will analyse the dimensions of this design space for implementation. This also allow us to refer to the large body of related work on generic programming. The exploration will avoid commitment to a specific functional language. This will clarify that different typed functional languages can be made fit for strategic programming, e.g., Clean, Haskell, and SML.

Models at a glance

As with any generic functionality, there are three ways to enable functional strategies in a given language:

built-in The combinators *adhoc* and *hfoldr* are implemented as language primitives.

defined They are defined in terms of already available expressiveness.

per-type They are defined using a term interface which is implemented per datatype.

Let us make some important side remarks regarding these overall approaches. The **built-in** option is the preferred one but it requires changing the language and its implementations. The **defined** option may fail to be faithful or may require inconvenient encodings, depending on the given expressiveness (as demonstrated below). To be practical and scalable, the **per-type** option needs generative tool support to implement the term interface per datatype. If the generative component becomes an integral part of the language implementation (as is the case for other generic functionality, e.g., equality predicates in SML and Haskell 98 [29]), the **per-type** option evolves into the **built-in** option.

Whichever option is chosen, another seven dimensions span a design space:

type reflection How to query type information and how to perform coercion?

term reflection How to generically destruct and construct terms?

application How to *apply* a strategy to an actual term?

quantification How to separate strategic and parametric polymorphism?

ranking How to rank the types of strategy combinators?

reduction How to deal with eager vs. lazy reduction and with effects?

modularisation How to maintain separate compilation?

The type-reflection dimension

Type-based function dispatch (*adhoc*) assumes some *type information* at run-time. In addition, *coercion* of type-specific cases is needed to apply them to terms encountered at run-time. Dynamic typing [1, 2] provides expressiveness to **define** *adhoc*. Then, the terms that are processed by strategies had to be of type `Dynamic`, and *adhoc* is implemented by ‘dynamic type case’ for type matching and coercion. (There are few language implementations with built-in support for dynamic typing but Yale Haskell used to support it, and it was recently added to Clean. **Per-type** support for dynamic typing has been suggested in various ways.) In fact, dynamic typing is more powerful than needed for *adhoc* because it involves the special type `Dynamic`. A lightweight dynamic-typing approach is to maintain a universe in which all datatypes are embedded [34]. This approach is particularly suited for **per-type** generative support. Embedding can be performed in two ways: either via a constructor per type [5], or on the basis of a universal term representation that includes a type *representation* [25]. As an alternative to dynamic typing, one may also consider intensional polymorphism [12, 33] as a means to **define** *adhoc*. This is a major language extension. (It is not available in widespread language implementations.) More seriously, this approach is not applicable because all work in the area of intensional type analysis favours *structural* type analysis while our kind of dispatching requires *nominal* type analysis as argued in [11, 23]. This closes the case to **define** *adhoc* in terms of other expressiveness. As for a **built-in** *adhoc*, the following approaches are at our disposal. Access to run-time type information can be based on a term representation with type tags [11, 22].

We should note that such tags are in conflict with type erasure. Also, they slightly enlarge the run-time representation of terms, and in turn slow down term manipulation. Alternatively, term constructors can be used to retrieve type information via a mapping from constructors to type tags. Yet another approach is to rely on carrying dictionaries in the sense of type classes [28] instead of carrying types in the terms themselves. Yet another approach is to rely on run-time type information complemented by unsafe type coercion as discussed in [23].

The term-reflection dimension

The poor man’s way to observe term structure is based on a universal term representation. Here, explosion and implosion functions mediate between the programmer-supplied datatypes and the universal term representation. This expressiveness is usually provided `per-type` [5, 25]. Folding over constructor applications boils then down to ordinary folding over homogeneous lists of term representations. Note that it is imperative to effectively hide the representation type in order to guarantee ‘implosion safety’ [25]. Then, this approach shines because of its simplicity. Note that implosion and explosion are likely to lead to a performance degradation but the choice of a suitable representation type can limit the depth of term conversion to a traversal’s extent. A built-in definition of `hfoldr` would entirely avoid this rather indirect style of operating on terms. In fact, built-in support is straightforward: Def. [ii] can be defined on any run-time term representation as is. As our Haskell-specific reference model of the upcoming section will demonstrate, term reflection can also be supported elegantly `per-type`. That is, `hfoldr` directly operates on terms on the basis of a `per-type` implementation of Def. [ii]. The most prominent expressiveness to attempt a definition of `hfoldr` is presumably polytypism [18, 16, 13, 14, 15, 3] as implemented in PolyP and Generic Haskell. The combinator `hfoldr` can indeed be expressed by structural induction on the type for its traversed argument. However, nominal run-time type case as needed for `adhoc` is beyond the scope of polytypism. The various constructs to customise polytypic definitions⁵ are compile-time means as opposed to a combinator for run-time type-based dispatch. One may also consider recent proposals for generalised pattern-match constructs [17, 9] to define `hfoldr`. Pattern matching is then not restricted to a single type, and a pattern-match case does not insist on a specific constructor. Again, these approaches do not offer type case.

The application dimension

Ideally, strategies are plain functions on the programmer’s term types. In fact, this characterises a challenging corner in our multi-dimensional design space. There are the following reasons why strategy application might deviate from function application. (a) Strategies might operate on `Dynami c` or a representation type in the interest of type and/or term reflection. (b) Strategies might be wrapped inside datatype constructors for reasons of opaqueness [25], or for reasons of rank-2 polymorphism [20]. (c) Strategies might operate on datatypes constructed from term types for reasons of a uniform definition [20, 23]: recall the co-domain constructor `c` in Def. [i] and Def. [ii]. This constructor will be normally a proper datatype because type-level lambdas are hardly supported in functional programming. Strategy application differs for (a)–(c). As for (a), to deal with `Dynami c` or a representation type, the ordinary terms need to be converted before and after strategy application. For convenience, this can be encapsulated via an overloaded application operator [25]. Then the programmer does not need to provide type tags. As for (b), a corresponding application operator is trivially defined by unwrapping. Basic strategy combinators also need to perform wrapping and unwrapping all-over the place. As for (c), we are saved by the fact that usually only a small number of co-domain constructors are used. Hence, one can specialise the types of `adhoc` and `hfoldr` for these few cases instead of postponing the type adjustment until strategy application. The built-in approach can hide this problem via a closed-world assumption regarding co-domain constructors.

⁵Cf. ad-hoc definitions for Generic Haskell as of [13], type-specific instances for derivable type classes [15, 3] just as for ordinary Haskell type or constructor classes [32], copy lines and constructor cases [6] added to Generic Haskell as of [14].

The quantification dimension

It is clear that strategies go beyond parametrically polymorphic functions. To reflect this fact, we used $\bar{\forall}$ instead of \forall in Def. [i] and Def. [ii]. In an actual language design, we can extend the interpretation of the universal quantifier, i.e., we equate our $\bar{\forall}$ with \forall . This is common practice for intensional polymorphism. We introduced $\bar{\forall}$ for the sake of a clear separation of parametric and strategic polymorphism. Having in mind a **per-type** approach, we can also introduce explicit type constraints in the sense of ad-hoc polymorphism [32], as supported by Haskell and Clean, i.e., we equate $\bar{\forall}\alpha. \dots$ with $\forall\alpha. Term\ \alpha \Rightarrow \dots$. Thus, the class constraint $Term\ \alpha$ points out where we go beyond parametric polymorphism. A built-in approach does not rely on type classes or class constraints.

The ranking dimension

To enable the combinator style of strategic programming, it is indispensable that strategy combinators consume polymorphic arguments. The most basic example is the *hfoldr* combinator which must insist on a polymorphic first argument because it is applied to children of potentially different term types. Normally, the need for polymorphic function arguments necessitates second-order polymorphism [10, 30] as opposed to ‘simple’ polymorphism with top-level quantification. One can attempt to organise generic traversal in terms of rank-1 expressiveness [17] but this will rule out the key idioms of strategic programming — in particular one-layer traversal. As an aside, second-order polymorphism is generally avoidable if strategies are weakly typed as functions on a representation type. Some form of rank-2 types (or even higher ranks) are supported in several functional language implementations. A well-understood form is first-class polymorphism [19] as employed for modelling functional strategies in [20]. First-class polymorphism means to wrap up polymorphic functions as constructor components. This necessitates unwrapping prior to function application. In [23], we employ rank-2 types (as supported in the current GHC implementation of Haskell) for generic traversal combinators. Regarding the earlier discussion of expressiveness to define our combinators, we should now add that ‘second order’ is also indispensable if existing forms of polymorphism were considered. In particular, polytypism in Generic Haskell [13, 14] is not second order because polytypic functions cannot involve polytypic function arguments.

The reduction dimension

Functional strategy combinators were inspired by term rewriting strategies [31, 22] which are eager. As for ordinary function application, the eager functional programmer can effectively postpone strategy application by the lazy “*if*” in the definition of new strategy combinators. When adding the *hfoldr* combinator in an eager framework, Def. [ii] must be implemented with some care to prevent premature evaluation of applications to subterms. This is not a problem in a straightforward inductive implementation (as opposed to the maybe too eager reading of Def. [ii]). There is also no problem with using impure effects such as in SML rather than monadic effects. This is again demonstrated by strategic term rewriting à la Stratego because Stratego supports some effects, e.g., for hygenic name generation or I/O.

The modularisation dimension

Strategic programs do not require compile-time specialisation of generic functionality as opposed to polytypic programs. This is because strategic programs operate on the programmer-supplied datatypes only via *adhoc* and *hfoldr*. These combinators, in turn, either operate on suitable run-time term representations as built-ins, or they are overloaded **per-type**. Hence, separate compilation is maintainable. Some techniques for type or term reflection might however imply a closed-world assumption. Firstly, the definition of a universe with embedding constructors per type is not extensible unless we assume extensible datatypes. In [20], we deal with the same problem: type case is encoded in a way that relies on a class member per term type. Separate compilation is also sacrificed when one provides **per-type** functionality by a ‘monster switch’, that is, by a central authority which would be meant to cover all types as opposed to proper overloading with support for separate compilation.

5 A reference model for Haskell

We will now define a Haskell-based reference model for the implementation of the combinators *adhoc* and *hfoldr*. In view of the previous section, we can provide the following characterisation:

- The reference model relies on **per-type** functionality.
- Strategies directly operate on the terms of the programmer-supplied datatypes.
- Strategy application is plain function application.

To start with, we define the type scheme for generic functions that model strategies:

```
type Parametric  $\alpha$   $c$  =  $\alpha \rightarrow c$   $\alpha$ 
type Strategic  $c$       =  $\forall \alpha. Term \alpha \Rightarrow Parametric \alpha c$ 
```

These two type synonyms make a distinction between unconstrained, i.e., parametrically polymorphic functions and constrained, i.e., strategically polymorphic functions. The class constraint points out where we go beyond parametric polymorphism. Roughly, the *Term* class hosts our two combinators, but the details follow below. Let us first list the Haskell types of our two combinators:

```
adhoc  ::  $\forall \alpha. Term \alpha \Rightarrow Strategic c \rightarrow Parametric \alpha c \rightarrow Strategic c$ 
hfoldr ::  $\forall c. HCons c \rightarrow HNil c \rightarrow Strategic c$ 
```

```
type HCons  $c$  =  $\forall \alpha \beta. Term \alpha \Rightarrow \alpha \rightarrow c (\alpha \rightarrow \beta) \rightarrow c \beta$ 
type HNil  $c$   =  $\forall \gamma. \gamma \rightarrow c \gamma$ 
```

The two type synonyms above are defined for convenience to make the type of *hfoldr* more comprehensible. In the types of the combinators, we make use of rank-2 polymorphism as provided by the GHC implementation of Haskell. It remains to define the combinators, to provide the complete declaration of the *Term* class, and to describe the derivation of the *Term* instances. We start with the implementation of Def. [i] for *adhoc*. To this end, we assume an operation *typeOf* which maps ‘typeable’ values of any term type to a type representation. Note that this is the kind of mapping that we proposed earlier in order to avoid carrying type information in the terms themselves. The operation *typeOf* is placed in a *Typeable* class as follows:

```
class Typeable  $\alpha$  where
  typeOf ::  $\alpha \rightarrow TypeRep$ 
```

The type *TypeRep* models type representations. Furthermore, we assume an operation *unsafeCoerce* to cast a value of any type to another type. Then, Def. [i] can be rephrased in Haskell as follows:

```
adhoc  $p$   $m$   $x$  = if (typeOf  $x$ )  $\equiv$  (domOf  $m$ ) then (unsafeCoerce  $m$ )  $x$  else  $p$   $x$ 
domOf ( $f :: a \rightarrow b$ ) = typeOf ( $\perp :: a$ )
```

It is important to notice that the argument of *typeOf* is only used to carry type information via overloading. The assumed two features are readily available in Haskell implementations. Of course, a proper language extension, which offers *adhoc* as a built-in, does not need to expose either *typeOf* or *unsafeCoerce*. These two operations are folklore in the Haskell community because they form the foundation of the *Dynamic* library, which has been a standard part of Haskell distributions for several years. The folk’s wisdom to derive *Typeable* and to perform type-safe cast on top of it with the help of *unsafeCoerce* is found in [23]. One can also use other, maybe safer, but also more involved approaches to dynamic typing. Regardless of the specific approach, the important thing to remember is that our implementational model of *adhoc* only involves term types but no universe such as `Dynamic`. This means that *adhoc* is very simple in nature, and *adhoc* is not inherently dynamically typed.

From the above development it is clear that *adhoc* is indeed defined in terms of **per-type** functionality for accessing run-time type representations based on overloading. Then, the class *Term* that captures strategic polymorphism has to be constrained by the *Typeable* class so that every term type is also known to be typeable. Thus, we have:

```
class Typeable  $\alpha \Rightarrow Term \alpha$  where
  -- Completed below
```

Def. [ii] for the *hfldr* combinator immediately necessitates **per-type** support because Haskell does not offer any expressiveness to generically observe the structure of terms. We place *hfldr* itself as a member in the *Term* class. In fact, the very comprehensible type of *hfldr* from above is not immediately suited for an overloaded class member. So we place a primed member in the class with an equivalent type:

```
class Typeable  $\alpha \Rightarrow Term \alpha$  where
  hfldr' ::  $\forall c. HCons\ c \rightarrow HNil\ c \rightarrow Parametric\ \alpha\ c$ 
```

We define *hfldr* simply in terms of *hfldr'*:

```
hfldr ::  $\forall c. HCons\ c \rightarrow HNil\ c \rightarrow Strategic\ c$ 
hfldr f z = hfldr' f z
```

The type of the primed member uses *implicit* quantification over the class parameter α . The non-primed version uses *explicit* quantification hidden in *Strategic*. The *Term* class exhibits an intriguing feature: it is defined recursively in the sense that the *Term* class itself is used to constrain the signature of its member *hfldr'* hidden in *HCons*. This can be viewed as a sign of the first-class status of strategies.

The classes *Term* and *Typeable* can be now added to the Haskell 98 [29] language definition in the same way as the standard classes *Eq*, *Ord*, *Show*, and *Read*. Just as the Haskell language definition contains a specification of derived instances for these built-in classes, so do we need to complete our extension by specifying the derived *Term* instances. The derivation of the member *hfldr'* is completely straightforward. We need to provide one equation per constructor based on the scheme for heterogeneous, right-associative fold according to Def. [ii]. That is, given a constructor *C* of type τ with *n* arguments, we need a pattern-match case defined as follows:

```
instance Term  $\tau$  where
  hfldr f z (C  $x_n \dots x_2 x_1$ ) = f  $x_1$  (f  $x_2$  ( $\dots$  (f  $x_n$  (z C))  $\dots$ ))
  -- Continue for the other constructors
```

This simple scheme is applicable to mutually recursive, parameterised (perhaps over higher-kinded type variables) datatypes. Not even non-uniform recursion is a problem. For all basic datatypes such as *Int*, we assume instances that apply the nil case. This is also a sensible choice for function types as there is no way to traverse them but it should be safe to encounter them in the course of traversal.

For completeness, we have investigated the option to define our combinators by means of *derivable type classes* as they were proposed for Haskell and Clean [15, 3]. Derivable type classes are precisely meant for the definition of classes for which the instances follow a common scheme. To this end, polytypic patterns for sums and products [18, 16] are included in the pattern-match syntax for class member definition. The *adhoc* combinator (and hence, the *Typeable* class) does not take advantage of derivable type classes because it necessitates a nominal approach rather than structural induction. However, the *hfldr* combinator (and hence, the *Term* class) can be defined as follows:

```
class Typeable  $\alpha \Rightarrow Term \alpha$  where
  hfldrTP ::  $HCons\ I \rightarrow HNil\ I \rightarrow Parametric\ \alpha\ I$ 
  hfldrTP { Unit } f z Unit      = z Unit
  hfldrTP {  $\alpha : + : \beta$  } f z (Inl x) = Inl (hfldrTP f z x)
  hfldrTP {  $\alpha : + : \beta$  } f z (Inr x) = Inr (hfldrTP f z x)
  hfldrTP {  $\alpha : * : \beta$  } f z (x : * : y) = f x ( $\lambda x' \rightarrow x' : * : hfldrTP f z y$ )
```

There are equations for *Unit*, $\alpha : + : \beta$ (i.e., sums), and $\alpha : * : \beta$ (i.e., right-associative products). In fact, we only show the special case for the type-preserving scheme where the co-domain constructor is instantiated to the identity type constructor *I*.

6 Reconstruction of *Strafunski*

The two simple combinators *adhoc* and *hfoldr* are sufficient to obtain the full power of strategic programming. We will demonstrate this by the reconstruction of essential parts of the strategic programming library of *Strafunski*. In fact, all previously published examples of functional strategy combinators [25, 21, 26] can be reconstructed with the identified primitives.

Normally we distinguish two broad categories of strategies, namely *type-preserving* vs. *type-unifying* ones. In case of the former, input and output term are of the same type. In case of the latter, the output type is fixed regardless of the input type. Another dimension of categorisation arises from the issue of possibly ‘effective’ traversal. In previous work, all our strategy combinators adhered to monadic style to be prepared for effects such as partiality, environment propagation, and I/O. If no such effect is present, the trivial identity monad can be used to ‘recover’ from monadic style. In the present paper, we explicitly distinguish *monadic* and *non-monadic* strategies. This allows a strategic programmer to resort to the simpler non-monadic types whenever this is sufficient. The variation in the two aforementioned dimensions — type-preservation vs. type-unification, and non-monadic vs. monadic strategies — can be captured by the following Haskell type synonyms:

```

type TP      =  $\forall \alpha. \text{Term } \alpha \Rightarrow \alpha \rightarrow \alpha$ 
type TU u    =  $\forall \alpha. \text{Term } \alpha \Rightarrow \alpha \rightarrow u$ 
type MTP m   =  $\forall \alpha. \text{Term } \alpha \Rightarrow \alpha \rightarrow m \alpha$ 
type MTU u m =  $\forall \alpha. \text{Term } \alpha \Rightarrow \alpha \rightarrow m u$ 

```

Note that these strategy types are just instances of the more abstract type *Strategic* which we defined in the previous section. This can be demonstrated as follows:

```

type TP      = Strategic I           type I  $\alpha$       =  $\alpha$ 
type TU u    = Strategic (C u)      type C u  $\alpha$     = u
type MTP m   = Strategic (MI m)     type MI m  $\alpha$   = m  $\alpha$ 
type MTU u m = Strategic (MC u m)   type MC u m  $\alpha$  = m u

```

The synonyms *I*, *C*, *MI*, and *MC* are Haskell implementations of non-monadic and monadic versions of the identity and constant type constructors. The four broad categories of strategy types are reconstructed by passing one of these type constructors to *Strategic*. We assume specialised versions of *adhoc* and *hfoldr* for each of these categories. We postfix the combinators by the corresponding category. These specialised combinators shall be used in strategic code if a specific category is intended.⁶ The specialised combinators for the category *TP*, for example, are declared as follows:

```

adhocTP ::  $\forall \alpha. \text{Term } \alpha \Rightarrow TP \rightarrow \text{Parametric } \alpha \ I \rightarrow TP$ 
hfoldrTP ::  $HCons \ I \rightarrow HNil \ I \rightarrow TP$ 

```

These specialisations allow us to maintain “strategy application = function application” in absence of type-level lambdas in Haskell. Recall that the types of *adhoc* and *hfoldr* involve a parameter *c* for co-domain construction. Any instantiation of the combinator types had to use *datatypes* or *newtypes* for *c* as opposed to type synonyms. This is no problem for the category *MTP* where a monad instantiates *c* but the other categories were defined above via type synonyms for identity and constant type construction. We can easily match up the types of *adhoc* and *hfoldr* with the simple types favoured by the programmer. The trick is to wrap and unwrap extra ‘coaching’ constructors inside the definitions of the specialisations. We illustrate this technique for *adhocTP*:

```

newtype I'  $\alpha$  = I'  $\alpha$            -- The newtype variation on I
unI' (I' x)    = x                -- Unwrapping function
adhocTP f g   = unI'  $\circ$  adhoc (I'  $\circ$  f) (I'  $\circ$  g) -- Specialisation of adhoc for TP

```

⁶So we should have used *hfoldrMTU* instead of *hfoldr* in the definition of *oneMTU* in Sec. 3.

At this point, we have defined all helper types and specialised combinators so that strategic programming can commence. We can define numerous one-layer traversal combinators in terms of *hfldr*, e.g.:

Attempt to process one child; try from left-to-right

$$\begin{aligned} \text{oneMTU} &:: \forall u m. \text{MonadPlus } m \Rightarrow \text{MTU } u m \rightarrow \text{MTU } u m \\ \text{oneMTU } s &= \text{hfldrMTU } (\lambda h t \rightarrow t \text{ 'mplus' } s h) (\text{const } mzero) \end{aligned}$$

Reduce all children via a monoid

$$\begin{aligned} \text{allTU} &:: \forall u. \text{Monoid } u \Rightarrow \text{TU } u \rightarrow \text{TU } u \\ \text{allTU } s &= \text{hfldrTU } (\lambda h t \rightarrow t \text{ 'mappend' } s h) (\text{const } mempty) \end{aligned}$$

Map over the children

$$\begin{aligned} \text{allTP} &:: \text{TP} \rightarrow \text{TP} \\ \text{allTP } s &= \text{hfldrTP } (\lambda h t \rightarrow t (s h)) \text{ id} \end{aligned}$$

Map over the children; monadic variation

$$\begin{aligned} \text{allMTP} &:: \forall m. \text{Monad } m \Rightarrow \text{MTP } m \rightarrow \text{MTP } m \\ \text{allMTP } s &= \text{hfldrMTP } (\lambda h t \rightarrow t \gg= \lambda t' \rightarrow s h \gg= \text{return } \circ t') \text{ return} \end{aligned}$$

These examples reconstruct combinators as introduced in [25]. The first combinator, *oneMTU*, was already defined in Sec. 3; here, we also provide its type, and we use the specialised combinator *hfldrMTU* to reflect the kind of strategy at hand. There is a class constraint *MonadPlus m* because *oneMTU* finds the suitable child via try and failure. The second combinator, *allTU*, uses monoid operations to combine the results of applying the type-unifying argument strategy to all children. The third combinator, *allTP*, is a variation on the folklore list *map*. The type-preserving argument strategy *s* is mapped over the children of a term. In the case of a non-empty constructor application, *s* is applied to the head *h*, and then the result is passed to the recursively processed tail. The empty constructor application is simply preserved via the identity function *id*. The last combinator, *allMTP*, is a monadic variation on *allTP*. It uses monadic bind ($\gg=$) to sequence the applications of the type-preserving argument strategy to all children, and reconstructs the term with the processed children.

Recursive traversal combinators can now be fabricated. The following portfolio provides *traversal schemes* that are parameterised by strategies for node-processing:

Full traversal of all nodes in top-down manner

$$\begin{aligned} \text{full_tdTP} &:: \text{TP} \rightarrow \text{TP} \\ \text{full_tdTP } s x &= \text{allTP } (\text{full_tdTP } s) (s x) \end{aligned}$$

Full traversal of all nodes in bottom-up manner

$$\begin{aligned} \text{full_buTP} &:: \text{TP} \rightarrow \text{TP} \\ \text{full_buTP } s x &= s (\text{allTP } (\text{full_buTP } s) x) \end{aligned}$$

Top-down traversal with cut after success

$$\begin{aligned} \text{stop_tdMTP} &:: \forall m. \text{MonadPlus } m \Rightarrow \text{MTP } m \rightarrow \text{MTP } m \\ \text{stop_tdMTP } s x &= s x \text{ 'mplus' } \text{allMTP } (\text{stop_tdMTP } s) x \end{aligned}$$

Accumulating a list query all-over the place

$$\begin{aligned} \text{collect} &:: \forall u. \text{TU } [u] \rightarrow \text{TU } [u] \\ \text{collect } s x &= s x \text{ 'mappend' } \text{allTU } (\text{collect } s) x \end{aligned}$$

Selection from the first node that admits success

$$\begin{aligned} \text{select} &:: \forall u m. \text{MonadPlus } m \Rightarrow \text{MTU } u m \rightarrow \text{MTU } u m \\ \text{select } s x &= s x \text{ 'mplus' } \text{oneMTU } (\text{select } s) x \end{aligned}$$

The first two combinators model non-monadic, type-preserving, and full traversal, i.e., they visit every node in the input term and the result type coincides precisely with the type of the input type. The monadic, type-preserving combinator *stop_tdMTP* performs a *partial* traversal, since it does not descend below nodes where its argument strategy is applied successfully. The type-unifying combinator *collect* performs a full traversal while intermediate list results are concatenated. The monadic, type-unifying combinator *select* performs a partial traversal following the same scheme as *query* from the sample section. However, the argument for recognition and extraction is not a function on a certain term type, but a strategy. This generality is appropriate whenever ‘relevant’ subterms can be of different types. All these combinators are reconstructions of combinators that are available in the strategy library distributed with *Strafunski*.

7 Concluding remarks

We have realized a compact amalgamation of term rewriting strategies and functional programming with emphasis on traversal strategies. Functional strategic programming features first-class generic functions that traverse terms of any type while mixing uniform and type-specific behaviour. Our reconstruction of functional strategies is based on just two combinators. Our first combinator is *ad hoc* for type-based function dispatch. Our second combinator is *hfoldr* for folding over constructor applications. We have given concise definitions of these two basic combinators. We have demonstrated how they are used to define one-layer traversal combinators and recursive traversal combinators as used in strategic programming. The abilities to traverse terms and to mix uniform and type-specific behaviour provide, in our experience, the key to practical application of functional programming to program analysis and transformation problems.

We have discussed implementational models of functional strategies without commitment to a specific functional language. We have, for example, argued that adding built-in support for *ad hoc* and *hfoldr* to functional language implementations requires only modest modifications, if second-order polymorphism is available. More specifically, using a reference model which requires generative support *per-type*, we have shown how the Haskell language needs to be extended to include our two combinators. Fully operational support for functional strategic programming is available in the form of the *Strafunski* bundle for generic programming and language processing in Haskell. The bundle can be configured to use one out of several alternative models. Generative support is based on the DrIFT preprocessing technology for Haskell. *Strafunski* has been applied for Java refactoring, Cobol reverse engineering, grammar engineering, Haskell program analysis and transformation, XML document transformation, and others.

Thus, our approach is lightweight, highly expressive, well-founded, and has already proven its value in important application domains. Language users take advantage of the generality and simplicity provided by our combinator style of generic programming. Language implementors take advantage of the fact that strategic programming does not require any new language constructs, but only two simple combinators which are easily defined *per-type* or as built-ins. There is no need for compile-time specialisation, and separate compilation is easily maintained.

Acknowledgement We are grateful for discussions with David Basin, Robert Ennals, Kevin Hammond, Ralf Hinze, Barry Jay, Johan Jeuring, Claude Kirchner, Paul Klint, Jan Kort, Andres Löh, Simon Peyton Jones, Claus Reinke, Simon Thompson, Phil Trinder, Jurgen Vinju, and Stephanie Weirich.

References

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic Typing in a Statically Typed Language. *ACM Transactions on Programming Languages and Systems*, 13(2):237–268, Apr. 1991.
- [2] M. Abadi, L. Cardelli, B. Pierce, and D. Rémy. Dynamic Typing in Polymorphic Languages. In *Proc. of the 1992 ACM Workshop on ML and its Applications*, pages 92–103, San Francisco, June 1992. ACM Press.

- [3] A. Alimarine and M. Plasmeijer. A Generic Programming Extension for Clean. In T. Arts and M. Mohnen, editors, *Implementation of Functional Languages, 13th International Workshop, IFL 2002 Stockholm, Sweden, September 24-26, 2001, Selected Papers*, volume 2312 of *LNCS*, pages 168–185. Springer-Verlag, 2002.
- [4] P. Borovansky, C. Kirchner, and H. Kirchner. Controlling Rewriting by Rewriting. In Meseguer [27].
- [5] P. Borovanský, C. Kirchner, H. Kirchner, and C. Ringeissen. Rewriting with strategies in ELAN: a functional semantics. *International Journal of Foundations of Computer Science*, 12(1):69–95, 2001.
- [6] D. Clarke and A. Löh. Generic Haskell, Specifically. In J. Gibbons and J. Jeuring, editors, *Proc. of the IFIP TC2 Working Conference on Generic Programming*. Kluwer Academic Publishers, 2002. To appear.
- [7] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and J. Quesada. The Maude System—System Description. In P. Narendran and M. Rusinowitch, editors, *Proceedings of the 10th International Conference on Rewriting Techniques and Applications (RTA-99)*, volume 1631 of *LNCS*, pages 240–243, Trento, Italy, July 1999. Springer-Verlag.
- [8] M. Clavel, S. Eker, P. Lincoln, and J. Meseguer. Principles of Maude. In Meseguer [27].
- [9] E. Dolstra and E. Visser. First-class rules and generic traversal. Technical report, Universiteit Utrecht, UU-CS-2001-38, 2001.
- [10] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures dans l’arithétique d’ordre supérieur*. PhD thesis, Université Paris VII, 1972.
- [11] N. Glew. Type dispatch for named hierarchical types. In *Proc. of the Fourth ACM SIGPLAN International Conference on Functional Programming (ICFP-99)*, volume 34.9 of *ACM Sigplan Notices*, pages 172–182, N.Y., Sept. 27–29 1999. ACM Press.
- [12] R. Harper and G. Morrisett. Compiling polymorphism using intensional type analysis. In *Conference record of POPL ’95: the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 130–141. ACM Press, 1995.
- [13] R. Hinze. A New Approach to Generic Functional Programming. In T. W. Reps, editor, *Conference record of POPL ’00: the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 119–132. ACM Press, Jan. 2000.
- [14] R. Hinze. Polytypic values possess polykinded types. In R. Backhouse and J. Oliveira, editors, *Proc. of the Fifth International Conference on Mathematics of Program Construction (MPC 2000), July 3-5, 2000*, volume 1837 of *LNCS*, pages 2–27. Springer-Verlag, July 2000.
- [15] R. Hinze and S. Peyton Jones. Derivable type classes. In G. Hutton, editor, *ENTCS*, volume 41. Elsevier Science Publishers, 2001.
- [16] P. Jansson and J. Jeuring. PolyP - a Polytypic Programming Language Extension. In *Conference record POPL ’97: the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 470–482. ACM Press, 1997.
- [17] C. Jay. Distinguishing data structures and functions: the constructor calculus and functorial types. In S. Abramsky, editor, *Typed Lambda Calculi and Applications: 5th International Conference TLCA 2001, Kraków, Poland, May 2001, Proceedings*, volume 2044 of *LNCS*, pages 217–239. Springer-Verlag, 2001.
- [18] J. Jeuring and P. Jansson. Polytypic programming. In J. Launchbury, E. Meijer, and T. Sheard, editors, *2nd Int. School on Advanced Functional Programming, Olympia, WA, USA, 26–30 Aug 1996*, volume 1129 of *LNCS*, pages 68–114. Springer-Verlag, Berlin, 1996.

- [19] M. Jones. First-class polymorphism with type inference. In *Conference record of POPL '97: the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 483–496, Paris, France, 15–17 Jan. 1997. ACM Press.
- [20] R. Lämmel. The Sketch of a Polymorphic Symphony. In B. Gramlich and S. Lucas, editors, *Proc. of International Workshop on Reduction Strategies in Rewriting and Programming (WRS 2002)*, volume 70 of *ENTCS*. Elsevier Science, 2002. 21 pages.
- [21] R. Lämmel. Towards Generic Refactoring. In *Proc. of Third ACM SIGPLAN Workshop on Rule-Based Programming RULE'02*, Pittsburgh, USA, Oct.5 2002. ACM Press. 14 pages.
- [22] R. Lämmel. Typed Generic Traversal With Term Rewriting Strategies. *Journal of Logic and Algebraic Programming*, 54:1–64, September 2002.
- [23] R. Lämmel and S. Peyton Jones. Scrap your boilerplate: a practical design pattern for generic programming. In A. SIGPLAN, editor, *Proc. of the ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI 2003)*. ACM Press, 2003. To appear in ACM SIGPLAN Notices.
- [24] R. Lämmel and J. Visser. Design Patterns for Functional Strategic Programming. In *Proc. of Third ACM SIGPLAN Workshop on Rule-Based Programming RULE'02*, Pittsburgh, USA, Oct.5 2002. ACM Press. 14 pages.
- [25] R. Lämmel and J. Visser. Typed Combinators for Generic Traversal. In S. Krishnamurthi and C. Ramakrishnan, editors, *Proc. of PADL 2002*, volume 2257 of *LNCS*, Portland, OR, USA, Jan. 2002. Springer-Verlag. 18 pages.
- [26] R. Lämmel and J. Visser. A Strafunski Application Letter. In V. Dahl and P. Wadler, editors, *Proc. of PADL 2003*, LNCS, New Orleans, LA, USA, Jan. 2003. Springer-Verlag. To appear; 18 pages.
- [27] J. Meseguer, editor. *Proceedings of the 1st International Workshop on Rewriting Logic and its Applications, RWLW'96, (Asilomar, Pacific Grove, CA, USA)*, volume 4 of *ENTCS*, Sept. 1996.
- [28] J. Peterson and M. Jones. Implementing type classes. *ACM SIGPLAN Notices*, 28(6):227–236, June 1993.
- [29] S. Peyton Jones and J. Hughes (eds.). Haskell 98: A non-strict, purely functional language. <http://www.haskell.org/onlinereport/>, 1999–2002.
- [30] J. Reynolds. Towards a Theory of Type Structures. In *Programming Symposium (Colloque sur la Programmation, Paris)*, volume 19 of *LNCS*, pages 408–425. Springer-Verlag, 1974.
- [31] E. Visser, Z. Benaïssa, and A. Tolmach. Building Program Optimizers with Rewriting Strategies. In *Proceedings of the third ACM SIGPLAN International Conference on Functional Programming (ICFP'98)*, pages 13–26. ACM Press, September 1998.
- [32] P. Wadler and S. Blott. How to make *ad-hoc* polymorphism less *ad hoc*. In *Conference record of POPL '89: the 16th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Austin, TX, USA*, pages 60–76. ACM Press, Jan. 1989.
- [33] S. Weirich. Higher-order intensional type analysis. In D. L. Métyer, editor, *Programming Languages and Systems: 11th European Symposium on Programming, ESOP 2002, Grenoble, France*, number 2305 in *LNCS*, pages 98–114, 2002.
- [34] Z. Yang. Encoding types in ML-Like languages. In M. Berman and S. Berman, editors, *Proceedings of the third ACM SIGPLAN International Conference on Functional Programming (ICFP-98)*, volume 34, 1 of *ACM SIGPLAN Notices*, pages 289–300, New York, Sept. 27–29 1998. ACM Press.