

# Good Advice for Type-directed Programming

## Aspect-oriented Programming and Extensible Generic Functions

Geoffrey Washburn    Stephanie Weirich  
University of Pennsylvania

{geoffw,sweirich}@cis.upenn.edu

### ABSTRACT

Type-directed programming is an important idiom for software design. In type-directed programming the behavior of programs is guided by the type structure of data. It makes it possible to implement many sorts of operations, such as serialization, traversals, and queries, only once and without needing to continually revise their implementation as new data types are defined.

Type-directed programming is the basis for recent research into “scrapping” tedious boilerplate code that arises in functional programming with algebraic data types. This research has primarily focused on writing type-directed functions that are closed to extension. However, Lämmel and Peyton Jones recently developed a technique for writing openly extensible type-directed functions in Haskell by making clever use of type classes. Unfortunately, this technique has a number of limitations.

We present an alternate approach to writing openly extensible type-directed functions by using the aspect-oriented programming features provided by the language AspectML. Our solution not only avoids the limitations present in Lämmel and Peyton Jones technique, but also allows type-directed functions to be extended at any time with cases for types that were not even known at compile-time. This capability is critical to writing programs that make use of dynamic loading or runtime type generativity.

### 1. INTRODUCTION

Type-directed programming (TDP) is an invaluable idiom for solving a wide class of programming problems. The behavior of a type-directed function is guided by the type structure of its arguments. Serialization, traversals, and queries can all be implemented via TDP, with the benefit that they need only be written once and will continue to work even when new data types are added as the program evolves.

TDP has become a popular technique for eliminating boring “boilerplate” code that frequently arises in functional programming with algebraic data types [11, 12, 8, 7]. The goal in this research has been to develop easy-to-use libraries

of type-directed combinators to “scrap” this boilerplate. For example, Lämmel and Peyton Jones’s Haskell library provides a combinator called `gmapQ` that maps a query over an arbitrary value’s children returning a list of the results. Using `gmapQ` it becomes trivial to implement a type-directed function, `gsize`, to compute the size of an arbitrary value.

```
gsize :: Data a => a -> Int
gsize x = 1 + sum (gmapQ gsize x)
```

However, this line of research focused on *closed* type-directed functions. Closed type-directed functions require that all cases be defined in advance. Consequently, if the programmer wanted `gsize` to behave in a specialized fashion for one of their data types, they would need to edit the definition of `gmapQ`.

The alternative is to provide *open* type-directed functions. Lämmel and Peyton Jones developed a method for writing openly extensible type-directed functions, via clever use of Haskell type classes [13]. While their technique provides the extensibility that was previously missing, there are caveats:

- It only allows open extension at compile-time, preventing the programmer from indexing functions by dynamically created/loaded data types.
- It is not possible to openly extend type-directed functions with special cases for existential data types [14], nested data types [3], or generalized algebraic data types (GADTs) [23]. Each of these sorts of data types pose problems for constraint resolution.
- Each open type-directed function used introduces a typing constraint into the environment, which for polymorphic functions is captured as part of the function’s type. This makes it difficult to use these functions as first-class values because their types may be too constrained to be used as arguments to other functions.

Because of these restrictions, Lämmel and Peyton Jones retain their old TDP mechanism as a fallback when openly extensible type-directed functions cannot be used. Providing distinct mechanisms for open and closed TDP significantly increases the complexity of TDP for the user. Additionally, it may not be easy for a user unfamiliar with type inference for type classes to understand why an extension to an open type-directed function fails type-check.

We present an alternate approach to openly extensible TDP using the aspect-oriented programming (AOP) features provided by the language AspectML. Aspect-oriented programming languages allow programmers to specify *what* com-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2006 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

putations to perform as well as *when* to perform them. For example, AOP in AspectJ [10] makes it easy to implement a profiler that records statistics concerning the number of calls to each method. The *what* in this example is the computation that does the recording and the *when* is the instant of time just prior to execution of each method body. In AOP terminology, the specification of what to do is called *advice* and the specification of when to do it is called a *pointcut designator*. A set of pointcut designators and advice organized to perform a coherent task is called an *aspect*.

We, with Daniel S. Dantas and David Walker, developed AspectML as part of our research on extending AOP to functional languages with parametric polymorphism [5]. In a language with parametric polymorphism, it is difficult to write advice when a function may be used at many types. It is also common to simultaneously advise functions with disparate types. To address both of these issues, we found runtime type analysis necessary for writing expressive advice.

However, we began writing more complex examples in AspectML, and we found that not only was runtime type analysis necessary for expressive advice, but that advice was necessary for expressive type-directed functions. In particular, aspects enable openly extensible TDP. We use advice to specify *what* the behavior of a specific type-directed function should be for a newly defined type, or combination of types, and *when* a type-directed function should use the case provided for the new type.

Furthermore, we believe that our entirely dynamic mechanism for AOP is critical to this expressiveness. Nearly all other statically typed AOP languages have a compile-time step called “weaving” that fixes the *when*. In contrast, decisions concerning *when* are made entirely at runtime in AspectML. Consequently, not only does TDP in AspectML improve extensibility when writing programs, but also while the programs are running. For example, it is possible to safely patch AspectML programs at runtime or manipulate data types that did not exist at “compile-time”.

Our contributions in this paper include:

- examples of TDP in AspectML.
- examples of how runtime advising is essential to the expressiveness of type-directed functions.
- a comparison of our technique with other approaches to openly extensible TDP.

In the next section, we introduce AOP in AspectML via examples. In Section 3, we explain and show examples of TDP in AspectML. In Section 4, we show how AOP in AspectML can be used to openly extend type-directed functions. In Section 5, we provide a more detailed comparison with related research. Finally, in Section 6, we discuss future directions and summarize our results.

## 2. THE ASPECTML LANGUAGE

We begin by introducing the general design of AspectML. AspectML is a polymorphic functional, aspect-oriented language descended from the ML family of languages. The syntax for a fragment of AspectML is given Figure 1. It is most similar to Standard ML [19], but occasionally uses OCaml or Haskell syntax and conventions [15, 21].

We use over-bars to denote lists or sequences of syntactic objects:  $\bar{x}$  refers to a sequence  $x_1 \dots x_n$ , and  $x_i$  stands for

---

(kinds)	$k ::= *$   $k_1 \rightarrow k_2$	kind of types function kinds
(polytypes)	$s ::= \langle \bar{a}[:k] \rangle r$	
( $\rho$ -types)	$r ::= t$   $s \rightarrow s$   $pc \ pt$	monotypes higher-rank function pointcut
(pointcut type)	$pt ::= \langle \bar{a}[:k] \rangle s_1 \sim s_2$	
(monotypes)	$t ::= a$   $T$   $Unit$   $String$   $Int$   $t_1 \rightarrow t_2$   $(t_1, \dots, t_n)$   $[t]$   $t_1 \ t_2$	type variables type constructors unit strings integers functions tuples lists type application
(terms)	$e ::= x$   $C$   $()$   $" \dots "$   $i$   $e_1 \ e_2$   $let \ d \ in \ e \ end$   $(fn \ \bar{p} \Rightarrow e)$   $case \ e \ of \ \bar{p} \Rightarrow \bar{e}$   $typecase \ t \ of \ t \Rightarrow e$   $\# \bar{f}[:pt] \#$   $any$   $(e_1, \dots, e_n)$   $[e_1, \dots, e_n]$   $e:s$	term variables data constructors unit strings integers application let declarations abstraction pattern matching typecase pointcut set any pointcut tuples lists annotation wildcard
(patterns)	$p ::= \_$   $x$   $C$   $()$   $" \dots "$   $i$   $p_1 \ p_2$   $[p_1, \dots, p_n]$   $(p_1, \dots, p_n)$   $p:s$	variable binding data constructors unit strings integers application tuples lists annotation
(trigger time)	$tm ::= before$   $after$   $around$	
(declarations)	$d ::= val \ p = e$   $fun \ f \ \langle \bar{a}[:k] \rangle (p) [: s] = e$   $datatype \ T \ [: k] = C:s$   $advice \ tm \ e_1 \ \langle \bar{a}[:k] \rangle (p_1, p_1, p_3) [: s] = e_2$   $case\text{-}advice \ tm \ e_1 \ (x:t_1, p_1, p_2) [: t_2] = e_2$	value binding recursive functions data type def advice def

---

Figure 1: AspectML syntax. Optional annotations are in math brackets.

an arbitrary member of this sequence. We use a typewriter font to indicate AspectML programs or fragments. Italicized text is used to denote meta-variables. We assume the usual conventions for variable binding and  $\alpha$ -equivalence of types and terms.

The type structure of AspectML has three layers: *polytypes*,  *$\rho$ -types*, and *monotypes*. Unless, otherwise stated when we refer to “types” in AspectML we mean monotypes. Polytypes are normally written  $\langle \bar{a}[:k] \rangle r$  where  $\bar{a}[:k]$  is a list of binding type variables with their kinds and  $r$  is a  $\rho$ -type. We often abbreviate  $\bar{a}[:k]$  as  $\bar{a}$ . When  $\bar{a}[:k]$  is empty, we abbreviate  $\langle \rangle r$  as  $r$ . This three-way partitioning of types is used as part of AspectML’s treatment of higher-rank polymorphism [22].

The category of  $\rho$ -types includes monotypes, higher-rank polymorphic functions, and `pc  $\rho t$` , the type of a pointcut, which in turn binds a list of type variables in a pair of polytypes. We explain pointcut types in more detail later.

In addition to type variables, `a`, simple base types like `Unit`, `String`, `Int`, and function types, the monotypes includes tuple and list types, `( $t_1, \dots, t_n$ )` and `[ $t$ ]` respectively, type application for higher-kinded monotypes, and type constructors `T`. Tuple types are actually syntactic sugar for an infinite number of type constructors `TUPLEn` for  $n$  greater than two. Type variables may only be instantiated with monotypes; it is not possible to construct a list of higher-rank polymorphic functions, for example.

Type constructors can be defined using `datatype` declarations. The `datatype` declaration is essentially the same as in ML and Haskell, with the exception of support for GADTs [23], and is used to define a new type constructor, `T`, via data constructors. For example, lists in AspectML are definable by giving the types of the constructors.

```
datatype List = [] : <a> List a
              | :: : <a> a -> List a -> List a
```

Note that, as in this example, it is always possible to elide the kind annotation when defining a type constructor because it can always be inferred. Furthermore, we require that the programmer provide the complete types of the data constructors because AspectML provides GADTs.

The term language of AspectML include variables `x`, term constructors `C`, unit `()`, integers `i`, strings `"..."`, tuples, syntactic sugar for lists, anonymous functions, function application and let declarations. New recursive functions may be defined in declarations. These functions may be polymorphic, and they may or may not be annotated with their argument and result types. Furthermore, if the programmer wishes to refer to type parameters within these annotations of the body of the function, they must specify a binding set of type variables `< $\bar{a};k$ >`. When the type annotations are omitted, AspectML may infer them. However, inference in AspectML is limited to monotypes if no annotations are provided. AspectML does not yet support advice for carried functions, so the syntax for recursive function definitions only allows a single argument. However, carried functions may still be written by using anonymous functions.

## 2.1 AOP in AspectML

To support AOP, AspectML provides pointcuts and advice. Advice in AspectML is second-class and includes two parts: the body, which specifies what to do, and the *pointcut designator*, which specifies when to do it. A pointcut designator has two parts, a *trigger time*, which may be `before`, `after`, or `around`, and a *pointcut* proper, which is a set of function names and an optional pointcut type annotation. The set of function names may be written out verbatim as `# $f$ #`, or, to indicate all functions, the `any` pointcut might be used.

Informally, the pointcut type, `< $\bar{a}$ >  $s_1 \rightsquigarrow s_2$` , describes the I/O behavior of a pointcut. In AspectML, pointcuts are sets of functions, and  `$s_1$`  and  `$s_2$`  approximate the domains and ranges of those functions. For example, if there are functions `f` and `g` with types `String -> String` and `String -> Unit` respectively, the pointcut `# $f,g$ #` has the pointcut type `pc (<a> String -> a)`. Because their domains are equal, the type `String` suffices. However, they have different ranges, so we use a type variable to generalize them both.

Any approximation is a valid type for a pointcut, so it would have also been fine to annotate `# $f,g$ #` with the pointcut type `pc (<a b> a -> b)`. This latter type is the most general pointcut type, and can be the type for any pointcut, and is the type of `any` pointcut.

The pointcut designator `before # $f$ #` represents the point in time immediately before executing a call to the function `f`. Likewise `after # $f$ #` represents the point in time immediately after execution. The pointcut designator `around # $f$ #` wraps around the execution of a call to the function `f`—the advice triggered by the pointcut controls whether the function call actually executes or not.

The most basic kind of advice has the form:

$$\text{advice } tm \ e_1 \ (\rho_1, \rho_2, \rho_2) = e_2$$

Here, `tm  $e_1$`  is the pointcut designator. When the pointcut designator dictates that it is time to execute the advice, the first pattern,  `$\rho_1$` , is bound either to the argument (in the case of `before` and `around` advice) or to the result of function execution (in the case of `after` advice). The second pattern,  `$\rho_2$` , is matched against the current call stack. Stack analysis is a novel feature of AspectML, but it is orthogonal to the expressiveness of TDP, so we do not discuss it in this paper. Refer to our other papers for more information on stack analysis [5]. The last pattern,  `$\rho_3$` , is matched against metadata describing the function that has been called, such as its name or location in the source text.

Since advice exchanges data with the designated control flow point, `before` and `after` advice must return a value with the same type as the first pattern  `$\rho_1$` . For `around` advice,  `$\rho_1$`  has the type of the argument of the triggering function, and the advice must return a value with the result type of the triggering function.

A common use of AOP to add tracing information to functions. These statements print out information when certain functions are called or return. For example, we can advise the program below to display messages before any function is called and after the functions `f` and `g` return. The trace of the program is shown on the right in comments.

```
(* code *)
fun f x = x + 1
fun g x = if x then f 1
          else f 0
fun h _ = False

(* Output trace *)
(*
entering g
entering f
leaving f => 2
leaving g => 2
entering h
*)
```

```
(* advice *)
advice before any (arg, _, info) =
  (print ("entering " ^ (getFunName info) ^ "\n");
   arg)

advice after #f,g# (arg, _, info) =
  (print ("leaving " ^ (getFunName info) ^
        " => " ^ (intToString arg) ^ "\n");
   arg)

val _ = h (g True)
```

Even though some of the functions in this example are monomorphic, polymorphism is essential. Because the advice can be triggered by any of these functions and they have different types, the advice must be polymorphic. Moreover, since the argument types of functions `f` and `g` have no type structure in common, the argument `arg` of the `before` advice

must be completely abstract. On the other hand, the result types of `f` and `g` are identical, so we can fix the type of `arg` to be `Int` in the after advice.

In general, the type of the advice argument may be the most specific polytype `s` such that all functions referenced in the pointcut are instances of `s`. Inferring `s` is not a simple unification problem; instead, it requires *anti-unification* [24]. AspectML can often use anti-unification to compute this type. It is decidable whether it is possible to reconstruct the pointcut type via anti-unification, so the implementation warns the user if they must provide the annotation.

Finally, we present an example of `around` advice. Again, `around` advice wraps around the execution of a call to the functions in its pointcut designator. The `arg` passed to the advice is the argument that would have been passed to the function had it been called. In the body of `around` advice, the `proceed` function, when applied to a value, continues that execution of the advised function with that value as the new argument.

In the following example, a cache is installed “around” the `f` function. First, a cache (`fCache`) is created for `f` with the externally-defined `cacheNew` command. Then, `around` advice is installed such that when the `f` function is called, the argument to the function is used as a key in a cache lookup (using the externally-defined `cacheGet` function). If a corresponding entry is found in the cache, the entry is returned as the result of the function. If the corresponding entry is not found, a call to `proceed` is used to invoke the original function. The result of this call is placed in the cache (using the externally-defined `cachePut` function) and is returned as the result of the `f` function.

```
val fCache : Ref List (Int,Int) = cacheNew ()

advice around ## (arg, _, _) =
  case cacheGet (fCache, arg)
  | Some res => res
  | None => let
    val res = proceed arg
    val _ = cachePut (fCache, arg, res)
  in
    res
end
```

We can transform this example into a general-purpose cache inserter by wrapping the cache creation and `around` advice code in a function that takes a first-class pointcut as its argument as described below. Finally, we do not provide their implementations here, but `cacheGet` and `cachePut` functions are polymorphic functions that can be called on caches with many types of keys. As such, the key comparisons use a polymorphic equality function that relies on the runtime type analysis described in Section 3.

One novel feature of AspectML is the ability to use pointcuts as first-class values. This facility is extremely useful for constructing generic libraries of profiling, tracing or access control advice that can be instantiated with whatever pointcuts are useful for the application. For example, recall the first example in Section 2 where we constructed a logger for the `f` and `g` functions. We can instead construct an all-purpose logger that is passed the pointcut designators of the functions we intend to log with the following code, assuming the existence of a type-directed serializer `toString`.

```
fun startLogger (toLog:pc (<a b> a ~> b)) =
  let
    advice before toLog (arg, _, info) =
      ((print ("before " ^ (getFunName info) ^ ":" ^
        (toString arg) ^ "\n")); arg)
    advice after toLog (res, _, info) =
      ((print ("after " ^ (getFunName info) ^ ":" ^
        (toString res) ^ "\n")); res)
  in
    ()
  end
```

### 3. TDP IN ASPECTML

The tracing example, in Section 2.1, could only print the name of the function called in the general case. If we would also like to print the arguments as part of the tracing aspect we must take advantage of AspectML’s primitive for runtime type analysis. For example, consider a revised tracing aspect that can print the arguments to any function that accepts integers or booleans:

```
advice before any <a>(arg : a, _, info) =
  (print ("entering " ^ (getFunName info) ^
    "with arg " ^
    (typecase a
      | Int => (int_to_string arg) ^ "\n"
      | Bool =>
        if arg then "True\n" else "False\n"
      | _ => "<unprintable>\n"));
  arg)
```

This advice is polymorphic, and the argument type `a` is bound by the annotation `<a>`. This ability to alter control-flow based on the type of the argument means that polymorphism is not parametric in AspectML—programmers can analyze the types of values at runtime. However, without this ability we cannot implement this generalized tracing aspect and many other similar examples.

Because dispatching on the type of the argument to advice is such a common idiom, AspectML provides syntactic sugar called `case-advice`, that is triggered both by the pointcut designator and the specific type of the argument. In the code below, the first piece of advice is only triggered when the function argument is an integer, the second piece of advice is only triggered when the function argument is a boolean, and the third piece of advice is triggered by any function call. (All advice that is applicable to a program point is triggered in inverse order from their definitions—newer advice has higher precedence over older advice.)

```
case-advice before any (arg : Int, _, _) =
  (print ("with arg " ^ (intToString arg) ^ "\n");
  arg)

case-advice before any (arg : Bool, _, _) =
  (print ("with arg " ^
    (if arg then "True\n" else "False\n"));
  arg)

advice before any (arg, _, info) =
  (print ("entering " ^ (getFunName info) ^ "\n");
  arg)
```

The code below and its trace demonstrates the execution of the advice. Note that even though `h`’s argument is polymorphic, because `h` is called with an `Int`, the first advice above triggers in addition to the third.

```

(* code *)          (* Output trace          *)
fun f x = x + 1     (*                          *)
fun g x = if x then (* entering g with arg True *)
  f 1              (* entering f with arg 1  *)
  else             (* entering h with arg 2  *)
  f 0
fun h _ = False
val _ = h (g True)

```

### 3.1 Writing total type-directed functions

While very important for programming in the large, data type definitions pose a problem for TDP in AspectML: A `typecase` expression can only have a finite number of cases, yet the programmer may use `datatype` to define a new type constructor at any time. For type-directed functions to be generally useful in AspectML it must be possible to write total type-directed functions that cover all possible cases. Furthermore, many type-directed functions are completely defined by the type-structure of their arguments and providing cases for each of them would introduce significant amounts of redundant code. Therefore, AspectML requires a mechanism for handling arbitrary type constructors and their associated data constructors.

We will use the term *atomic type* to refer to those monotypes in AspectML that are not defined via `datatype` and cannot be decomposed into simpler types. For example, integers, strings, and functions are atomic types in AspectML. Tuple types and type constructors are not considered atomic.

To allow programmers to write total type-directed functions, we provide a primitive `toSpine` for converting arbitrary data constructor values to *spine* abstractions, as introduced by Hinze, Löh, and Oliveira (HLO) [8, 7].

```

datatype Spine =
  | SCons : <a> a -> Spine a
  | SApp  : <a b> Spine (b -> a) -> b -> Spine a

```

Spines can be thought of “exploded” representations of a term constructor and the arguments that have been applied to it. The term constructor `SCons` is used to hold the actual constructor, and then a series of `SApp` constructors hold the arguments applied to the constructor. For example, the expression `toSpine (1::2::3::[])` produces the spine

```
SApp (SApp (SCons (op ::)) 1) (2::3::[])
```

Note that this alternate representation is only one layer deep: The two arguments of `::`, `1` and `2::3::[]`, are unchanged.

The function `toSpine` must be a primitive in AspectML. If it were implemented in terms of AspectML like

```

fun toSpine <a>(x : a) : Option (Spine a) =
  typecase a
  | (a, b) =>
    case x of (x, y) =>
      Some (SApp(SApp Tuple2 x) y)
  | (a, b, c) =>
    case x of (x, y, z) =>
      Some (SApp (SApp(SApp Tuple3 x) y) z)
  | ...

```

an infinite number of cases would be necessary, as there are an infinite number of tuple types. However, Writing the inverse function, `fromSpine`, is straightforward.

```

fun fromSpine <a>(x : Spine a) : a =
  case x
  | SCons con => con
  | SApp spine arg => (fromSpine spine) arg

```

---

```

val escape : String -> String
val toString : <a> a -> String

fun toString <a>(x : a) =
  let
    fun spineHelper <a>(x : Spine a) =
      case x
      | SCons c =>
        (case (conToString c)
         | Some s => "(" ^ s ^ ")")
         | None => abort "impossible")
      | SApp spine arg =>
        "(" ^ (spineHelper spine) ^ " " ^
          (toString arg) ^ ")"
  in
    typecase a
    | Int => intToString x
    | String => "\"" ^ (escape x) ^ "\""
    | b -> c => "<fn>"
    | _ => (case (toSpine x)
             | Some spine =>
               spineHelper spine
             | None =>
               abort "impossible")
  end

```

Figure 2: `toString` in AspectML

---

Although the primitive `toSpine` function has type `<a> a -> Option (Spine a)` and can therefore be used with any input, it only returns a spine when provided with a non-atomic type. This is because spines simply do not make sense for atomic types; for example,  $\lambda$ -abstractions have no data constructor that “builds” functions.

Finally, AspectML provides primitive implementations of equality and string conversion for integers, strings, and unapplied data constructors. Unapplied constructors are data constructors that have not been applied to any arguments. For example, `::` versus `1::[]`. The functions `eqCon` and `conToString` have the somewhat unusual types: `<a> a -> Option Bool` and `<a> a -> Option String` respectively. The reason they return `Options` is much like the reason `toSpine` returns an `Option`—the polymorphic type allows them to be applied to things other than unapplied data constructors. Finally to address the problem that there are infinitely many tuple type constructors we introduce a primitive `tupleLength` of type `<a> a -> Option Int` that if given an unapplied tuple type constructor will return an `Option` containing its length. For all other values it will return `None`.

Putting everything together, we can now write a type-directed pretty printer, `toString`, as shown in Figure 2. If `toString` is applied to an integer it uses the primitive `intToString` function to convert it to a string. If it is applied to a string it calls the `escape` function to escape any special characters and then wraps the string in quotes. For function types, `toString` simply returns the constant string “<fn>” because there is no way to inspect the structure of functions. Finally, for all other types `toString` will attempt to convert the value to a spine and then call `spineHelper`. Even though we know that `toSpine` will always succeed, because the `typecase` covers all atomic types, we must still perform the check to make sure a `Spine` was actually constructed.

The function `spineHelper` walks down the length of the spine and makes recursive calls to `toString` on the arguments and upon reaching the head converts the unapplied

---

```

val gmapT : <a> a -> (<b> b -> b) -> a
val gmapTspine : <a> Spine a ->
  (<b> b -> b) -> Spine a

fun gmapT <a>(x:a) = fn (f:<b> b -> b) =>
  (case (toSpine x)
   | Some y => fromSpine (gmapTspine y f)
   | None => x))

and gmapTspine <a>(x:Spine a) = fn (f:<b> b -> b) =>
  case x
  | SCons _ => x
  | SApp spn arg =>
    SApp (gmapTspine spn f) (f arg)

```

Figure 3: **gmapT** in AspectML

constructor to a string using `conToString`. As with `toSpine`, it is again necessary to check the output of `conToString` because the type system does not know that the argument of `SCons` is guaranteed to be an unapplied constructor.

Furthermore, why did we choose to write `spineHelper` rather than just adding a branch to the `typecase` expression for the type constructor `Spine`, and then calling `toString` recursively directly? There are two problems that arise with such an implementation. Firstly, the behavior of `toString` actually be incorrect in the instance that the argument is a `Spine`. For example, `toString (toSpine[1])` would produce the string "Cons 1 Nil" rather than "(SApp (SApp (SCons Cons) 1) Nil)".

Secondly, such an implementation may only be given “well-formed” `Spines`, that is, spines that have unapplied data constructors at their head. If we recall the definition of the `Spine` data type

```

data Spine =
  | SCons : <a> a -> Spine a
  | SApp  : <a> b Spine (b -> a) -> b -> Spine a

```

we see that there is nothing preventing a client from calling `toString` on spines that do not have unapplied constructors in them, like `(SCons 5)`. The implementation in Figure 2 avoids this problem by converting all non-atomic types into a spine before converting them to strings. Because `toSpine` is guaranteed to construct well-formed `Spines`, it is guaranteed to never fail.

## 3.2 Scrapping your boilerplate

It is now straightforward to use the TDP infrastructure of AspectML to implement the various combinators Lämmel and Peyton Jones developed for generic programming.

In Figure 3 is an implementation of `gmapT`, a “single-layer” generic mapping function. Given any value and a polymorphic function, it applies the function to any children the value has, if any. It is implemented in terms of a helper function for spines, `gmapTspine`. This helper is fairly straightforward. Because the `Spine` data type makes the arguments to a constructor explicit via the `SApp` constructor, it only need walk down the spine and apply the function `f` to each argument.

Finally, `gmapT` ties everything together by attempting to convert the argument to a `Spine`. If `toSpine` fails, the argument must be atomic and has no children to map over. If `toSpine` succeeds, `gmapTspine` is applied to the spine and `fromSpine` is used to convert back to the original type.

---

```

val == : <a> a -> a -> Bool

fun == <a>(x:a) = fn (y:a) =>
  let
    fun spineHelper <a>(x:Spine a, y:Spine a) =
      case (x, y)
      | (SCons c1, SCons c2) =>
        (case (eqCon c1 c2)
         | Some z => z
         | None => abort "impossible")
      | (SApp sp1 arg1, SApp sp2 arg2) =>
        (case (cast sp2, cast arg2)
         | (Some sp2', Some arg2') =>
           (sp1 == sp2') andalso (arg1 == arg2')
         | _ => False)
    in
      typecase a
      | Int => eqInt x y
      | String => eqString x y
      | _ =>
        (case (toSpine x, toSpine y)
         | (Some x', Some y') => spineHelper (x', y')
         | (None, None) => abort "equality undefined")
    end

```

Figure 4: **Polymorphic structural equality** in AspectML

To effectively use `gmapT` we need a way to construct a polymorphic function of type `<b> b -> b` that is not either the identity or the diverging function. Lämmel and Peyton Jones initially implemented this using a combinator `extT` that lifts a monomorphic function of type `t -> t`, for any type `t`, to be `<b> b -> b`. They do this using a primitive casting operation, but it is simple to implement in AspectML.

```

fun extT <a>(f:a -> a) = fn <b> (x:b) =>
  typecase a of b => f x
  | _ => x

```

It is then straightforward to use `gmapT` to write a generic traversal combinator

```

fun everywhere <a>(f:<b> b -> b) = fn (x:a) =>
  f (gmapT x (everywhere f))

```

While very useful, `toString`, `gmapT` and `everywhere` may only traverse over a single value at a time. Many type-directed functions need to be able to operate on multiple values simultaneously. The generic structural equality, shown in Figure 4, is the prototypical example. For atomic values that are integers or strings, the primitive equality functions are used. However, if both values have a spine representation, the spines are compared with the function `spineHelper`. If both spines are unapplied data constructors, `spineHelper` uses the primitive equality, `eqCon`. Otherwise, if both spines are application spines, the arguments and remainder of the spines are compared recursively.

Because spines are existential data types, it is necessary to use a `cast` in the case for `SApp` in order to ensure that the arguments and constituent spines have the same types. Otherwise it would not be possible to call `==` or `spineHelper` recursively, as they both require two arguments of the same type. Writing `cast` is trivial in AspectML, because it is possible to use abstracted types within type patterns

```

datatype Queue = Q : <a> List a -> List a -> Queue a

val emptyQueue : <a> Queue a
val emptyQueue = Q [] []

val enqueue : <a> a -> Queue a -> Queue a
fun enqueue e = fn (Q l1 l2) => Q l1 (e::l2)

val dequeue : <a> Queue a -> Option (a, Queue a)
fun dequeue q =
  case q
  | Q [] [] => None
  | Q [] l2 => dequeue (Q (rev l2) [])
  | Q (h::t) l2 => Some (h, Q t l2)

```

Figure 5: Amortized constant time functional queues

```

fun cast <a b>(arg : a) : Option b =
  typecase a of b => Some arg
  | _ => None

```

So far we have only transliterated some standard examples of TDP into AspectML. We believe that our realization of these functions is quite elegant, but the most significant benefit comes from the dynamic open extension allowed by AspectML. We will examine this capability in the next.

#### 4. OPEN EXTENSIBILITY

By default, `toString`, `gmapT`, `everywhere`, and `==` work with all AspectML monotypes, but the programmer has no control over the behavior for specific types. This is a problem because the programmer might define a new data type for which the default behavior is incorrect. This problem is easily solved in AspectML using advice.

Starting out simple, while `toString` from Figure 2 can produce string representation for all types, its string representations of tuples is not what we would expect. The application `toString (1, True)` will generate the string `"((Tuple2 1) True)"`. We can improve the output by writing advice for `toString` to provide a specialized case for tuples.

```

advice around #toString# <a>(x : a, _, _) =
  let
    fun helper <a>(spine : Spine a) =
      case spine
      | SCons _ => ""
      | SApp spine arg =>
        (helper spine) ^ ", " ^ (toString arg)
  in
    case (tupleLength x)
    | Some i =>
      (case (toSpine x)
      | Some spine => "(" ^ (helper spine) ^ ")"
      | None => abort "impossible")
    | None => proceed x
  end

```

This advice triggers on any invocation of `toString`, and checks whether the argument is a tuple using the `tupleLength` primitive we discussed in Section 3.1. If it is not a tuple, it uses `proceed` to return control to the `toString` function. Otherwise, it converts the tuple to a `Spine` and then walks down the spine producing a comma separated string representation.

While it is nice that we can use advice to specialize the behavior of `toString`, the modification was simply for aes-

```

datatype Exists = Ex : <a> a -> Exists

val pickle : <a> a -> String
val unpickle : String -> Exists

val openFile : String -> Handle
val writeFile : Handle -> String -> Unit
val readFile : Handle -> String
val closeFile : Handle -> Unit

```

Figure 6: Simple primitives for dynamic loading

thetic purposes. There are cases where the default behavior of a type-directed function is incorrect. Say the programmer defines a new data type and the default traversal behavior provided by `everywhere` is not correct; traversal orders is especially relevant in the presence of effects the traversal order could be important. For example, consider the common functional implementation of queues via two lists in Figure 5. Elements are dequeued from the first list and enqueued to the second; if the dequeue list becomes empty the enqueue list is reversed to become the new dequeue list. By default, `everywhere` will traverse the dequeue list in the expected head to tail order (FIFO) but the default traversal will walk over the enqueue list in the reverse of what a user might expect.

One solution would be for the author of `Queue` to edit `everywhere` or `gmapT` to include a special case for queues. But, this solution is undesirable because their source code might not be available, and continually adding special cases would clutter them with orthogonal concerns.

If the author of the `Queue` data type knows that `everywhere` is implemented in terms of `gmapT`, can simply write the following advice for `gmapT` to ensure that the elements are traversed in FIFO order.

```

case-advice around #gmapT# (x:Queue a, _, _) =
  fn (f:<b> b -> b) =>
    case x
    | Q l1 l2 => Q (map f (l1 @ (rev l2))) []

```

This case-advice intercepts a call to `gmapT` when the first argument is type `Queue a` for any `a`, otherwise `gmapT` is executed.

The author or the `Queue` data type will encounter another problem when using the type-directed structural equality function described in Section 3.2, because it distinguishes between too many values. For example, `==` does not equate the queues `Q [1] []` and `Q [] [1]` even though both of these queues are extensionally equivalent. In fact, this is an example of how structural equality breaks representation independence.

Fortunately it is again quite simple to use advice to extend `==` so that it correctly compares `Queues`.

```

case-advice around # == # (x:Queue a, _, _) =
  (fn (y : Queue a) =>
  case (x, y)
  | (Q l1 l2, Q l3 l4) =>
    (l1 @ (rev l2)) == (l3 @ (rev l4))

```

This advice triggers when `==` is called on values of type `Queue a` and converts the constitutes lists to a canonical form before comparing them.

The examples of open extension we have shown so far

---

```

datatype Extension = Ext : Exists ->
  (Unit -> Unit) ->
  Extension

datatype Exists = Ex : <a> a -> Exists

val pickle : <a> a -> String
val unpickle : String -> Exists

val openFile : String -> Handle
val writeFile : Handle -> String -> Unit
val readFile : Handle -> String
val closeFile : Handle -> Unit

advice after #unpickle# (ex : Exists, _, _) =
  case ex
  | Ex <a> x =>
    (typecase a
     | Extension =>
       (case x of (Ext ex f) => (f ()); ex)
     | _ => ex)

```

**Figure 7: Extended dynamic loading**

are quite useful, but are essentially the same as what is possible using Lämmel and Peyton Jone’s technique, modulo its restrictions. Our approach however, extends to uses that are simply not possible with static approaches. For example, consider programming an extension to AspectML to support dynamic loading, using the primitives defined in Figure 6. The functions `openFile`, `writeFile`, `readFile`, and `closeFile` are straightforward functions for file I/O.

The two primitives `pickle` and `unpickle` are a bit more complicated. The `pickle` primitive is similar to the function `toString` we defined, but includes support for meaningful serialization of functions. Its inverse `unpickle` returns an existential datatype, `Exists`. This is because we cannot know the type of the value that will be produced in advance. However, thanks to TDP it is possible to still do useful things with the data, even if we have no idea what it might be. For example, our type-directed functions `gmapT`, `everywhere`, and `==` can still be used on the value the existential package is hiding.

It is even possible for a program to manipulate data types that were not defined when it was compiled. Specifically, we might have two separate programs each with some data types unique to them. For example, imagine one of our programs used the queues described in Figure 5, but another did not. This first program would be able to `pickle` a queue, write it to a file, and then the second program could then read the queue out of the file and manipulate it as it would any other existentially packaged value.

However, this second host will encounter the exact same problems we discussed earlier with `gmapT` and `==` behaving incorrectly on functional queues. If it was only possible to extend type-directed functions at compile-time there would be nothing that could be done. However, because we might use advice to dynamically extend type-directed functions, we opt to package data written to disk with the necessary extensions. We show a fairly naïve implementation of this in Figures 7 and 8.

In Figure 7, we extended the dynamic loading primitives with a notion of “extension”. The new data type `Extension` packages up an existential along with a function from `Unit ->`

---

```

val myQueue = enqueue 2 (enqueue 1 emptyQueue)

val f =
  (fn () =>
   let
     case-advice around #toString# (x:Queue a,_,_) =
       case x
       | Q l1 l2 =>
         "[" ^
           concat " " (map toString (l1 @ (rev l2))) ^
           "]"

     case-advice around #gmapT# (x:Queue a,_,_) =
       fn (f:<b> b -> b) =>
         case x
         | Q l1 l2 => Q (map f (l1 @ (rev l2))) []

     case-advice around # == # (x:Queue a,_,_) =
       fn (y : Queue a) =>
         case (x, y)
         | (Q l1 l2, Q l3 l4) =>
           (l1 @ (rev l2)) == (l3 @ (rev l4))

   in
     ()
   end)

val file = openFile "testing.data"
val () = writeFile file
  (pickle (Extension (Ex myQueue) f))
val () = closeFile file

```

**Figure 8: Program compiled with functional queues**

`Unit` that can be used to perform any necessary configuration required for the packaged value. We then provide advice for the `unpickle` primitive to checks whether it has created an `Extension` value. If so, it invokes the function contained in the `Extension` and returns included the existential package as the result.

Our first program, that was compiled with functional queues can now store them to disk as shown in Figure 8. Here, it creates an configuration function that installs the relevant extensions to `toString`, `gmapT`, and `==` for functional queues. It then packages the function with a queue and writes it to disk. Now any other program built with the “extensible” version of the dynamic loading primitives in Figure 7 can read the queue from the disk, and at the same time receive appropriate advice for manipulating the data in a generic fashion. For example, it can call `toString` on the value it obtains from the file `testing.data` and obtain `"(Ex [|1 2|])"`.

Furthermore, data type definitions in AspectML are generative, so that every time they are evaluated a new distinct data type is created. Therefore, type-directed functions will in practice encounter types that were not defined at compile-time anywhere. In the following example, every time `nonce` is called not only does it give a distinct value, but a distinct type!

```

val getTime : Unit -> Time

val nonce : Unit -> Exists
fun nonce () =
  let
    datatype Nonce = Nonce : Time -> Nonce
  in
    Ex (Nonce (getTime ()))
  end

```

Without the dynamic extensibility provided by aspects there would be no way to handle these kinds of examples.

## 5. RELATED-WORK

In this section we will examine the wide spectrum of research that this paper is built upon.

### 5.1 Extensible type-directed programming

*Scrapping your boilerplate with class.* Our examples in Sections 3 and 4 were adapted from the work of Lämmel and Peyton Jones on “scrapping your boilerplate” [11, 12, 13]. In their latest paper, Lämmel and Peyton Jones have attempted to address the problem of providing openly extensible type-directed functions. This is done by associating each type-directed function with a type class. However, this technique has many limitations. Most importantly, the Haskell type class mechanism requires that all instances for these classes be defined at compile-time. Consequently, it is impossible to use their library with types that are not known until runtime. This is necessary for any applications that make significant use of dynamic loading or generative types, such as the examples in the previous section.

The second major limitation of their technique is that it is not possible to write openly extensible type-directed functions that have specialized behavior on several significant classes of data types: existential data types, nested data types, and GADTs. It is not possible to write open type-directed functions for existential data types because their use of the type class mechanism. Consider the following example where we attempt to define a specialized case for `Size`, an open type-directed function to generically compute the size of a value.

```
class Data Size a => Size a where
  gsize :: a -> Int

data Exists :: * where Ex :: a -> Exists

instance Size Exists where
  -- Does not type check
  gsize (Ex (ex :: b)) = 1 + gsize ex
```

Here, in order to define an instance of `Size` for `Exists`, we need to call `gsize` on the existentially packaged value `ex`. However, calling `gsize` requires that there exist an instance of `Size` for type `b`. However, because `b` could be anything there is no way of determining whether such an instance exists. The only solution is to require that the `Exists` data type package the necessary type class dictionary itself.

```
data Exists :: * where
  Ex :: Size a => a -> Exists

instance Size Exists where
  -- Type checks
  gsize (Ex (ex :: b)) = 1 + gsize ex
```

However, this is simply not scalable as every time we would like to use an open type-directed function on values of type `Exists` it will be necessary to add another constraint to the data type definition.

Nested data types, such as Okasaki’s square matrices [20], cause other difficulties with type class constraints.

```
data Sq a = Zero a | Succ Sq (a, a)
```

In order to solve type-classes constraints for Lämmel and Peyton Jones’s implementation, a form of coinductive constraint solving is necessary. Therefore, a solver attempting determine whether the type class constraint for `Size [a]` is satisfiable will go through a constraint back-chaining process that looks like the following.

```
Size [a]
  ↓
Size [a], Data Sized [a]
  ↓
Size [a], Data Sized [a], Sat (Sized [a])
```

First, `Size [a]` requires the instance `Data Sized [a]`. And `Data Sized [a]` needs `Sat (Sized [a])`. Finally, because `Size [a]` implies the existence of `Sat (Sized [a])` and at this point the solver reaches as fixed-point.

With nested data types a fixed-point can never be reached. For example, an instance for `Size (Sq a)` will require an instance of `Data Sized (Sq a)` which needs both `Sat (Sized a)` and `Sat (Sized (Sq (a,a)))`. While `Sat (Sized (Sq (a, a)))` can be obtained from `Size (Sq (a, a))`, it in turn needs `Data Sized (Sq (a,a))` which in requires `Sat (Sized (Sq ((a,a), (a,a))))`, and so forth ad nauseam.

A similar problem arises with GADTs. Consider the following contrived GADT.

```
datatype Foo :: * -> * where
  Bar :: a -> Foo a
  Qux :: Foo Int -> Foo Bool

instance Size a => Size (Foo a) where
  gsize (Bar x) = gsize x
  gsize (Qux fi) = 1 + (gsize fi)
```

Here, attempting to call `gsize` on `fi` will require an instance of `Size (Foo Int)` which will require an instance of `Data Sized (Foo Int)` which will in turn need the instance `Sat (Sized (Foo Int))`. However, `Sat (Sized (Foo Int))` would require an instance `Size (Foo Int)` to be derivable, but there is no such instance.

Finally, because each open type-directed function used introduces a constraint into a function’s type it becomes difficult to use them as first-class functions. This is because they may become more constrained than other functions are expecting their arguments to be. For example,

```
shrink :: Shrink a => a -> [a]

shrinkList :: (Eq a, Shrink a) => [a] -> [a]
shrinkList xs = nub (concat (map shrink xs))

gmapQ :: Data ctx a => Proxy ctx ->
  (forall b. Data ctx b => b -> r) ->
  a -> [r]

-- Ill-typed
shrinkTwice x = gmapQ shrinkProxy
  (shrinkList . shrink) x
```

Here `shrinkTwice` is ill-typed because `gmapQ` is expecting a polymorphic function with the constraint `Data ctx b` for some `ctx` and some `b`. However, the composition of `shrinkList` and `shrink` has the constraint `(Eq a, Shrink a)` which resolves to `(Eq a, Data Shrink a)` which is too constrained by the need for the `Eq` type class to be a valid argument.

This problem with constraints can be alleviated by either making `shrinkList` an openly extensible type-directed

function or by introducing new dummy type classes that encapsulate constraints. The former leads to large amounts of boilerplate for otherwise simple helper functions and the latter leads to a combinatorial explosion as the number of type-directed functions grow.

*Scrapping your boilerplate with spines.* Our use of spines to implement type-directed functions was heavily influenced by the closely related work by Hinze, Löh, and Oliveira on understanding Lämmel and Peyton Jones’s techniques in terms of spines. In their work, they opt to implement `toSpine` directly using type representations built from GADTs. They assume the existence of some mechanical method of producing all the necessary representations and cases for `toSpine`. However, in order to provide open extension, they use type classes in a fashion inspired by Lämmel and Peyton Jones. Therefore, their approach to open extension has many of the same problems with type-classes.

In their more recent work, Hinze and Löh have shown how to extend spines so that it is possible to implement type-directed functions that can not only process arbitrary data types, but built values of arbitrary type. Additionally, they show how to perform type-directed operations on higher-kinds using what they call *lifted spines*. AspectML does not presently have support for writing type-directed functions that construct arbitrary data types. We could easily adapt the modifications to spines that Hinze and Löh developed, but we will most likely pursue an approach based upon Weirich’s concurrent research [32]. It would be straightforward to extend AspectML with support for *lifted spines*, but we think it would be better for us to focus on developing a general solution for writing functions with polykinded types [6].

*A core calculus for open and closed TDP.* Along with Dimitrios Vytiniotis, we developed a core calculus,  $\lambda_{\mathcal{L}}$ , that could be used to implement type-directed functions that could be either closed or open to extension [27]. While it was possible to extend type-directed functions in  $\lambda_{\mathcal{L}}$ , the programmer must manage the extensions herself, and manually close off the recursion before the functions could be used. However, as  $\lambda_{\mathcal{L}}$  was designed as a core calculus rather than a high-level programming language, it is not expected to be easy to write software in  $\lambda_{\mathcal{L}}$ .

A novelty of  $\lambda_{\mathcal{L}}$  is its use of *label-sets* to statically ensure that a type-directed function was never given an argument for which it lacked a case. However, *label-sets* are very heavy-weight and lack sufficient precision to describe whether some inputs are acceptable. In AspectML, the programmer must rely on the initial implementation of a type-directed function covers all the necessary types. Lämmel and Peyton Jones use of type classes does not have this problem as constraint solving can be relied upon to know that the input will have instances for all the appropriate type classes. However, checking that a type-directed function covers all cases statically, requires knowing all types that may be used statically.

## 5.2 Generic programming

The research on Generic Haskell [17] is related to our work on AspectML in many ways. However, Generic Haskell has a number of limitations in comparison to AspectML. Firstly, in Generic Haskell all uses of type information are resolved at compile-time, making it impossible to extend to TDP with types that are created at runtime or loaded

dynamically. Secondly, Generic Haskell does not provide any direct mechanism for open extension. The latest version of Generic Haskell has added syntactic sugar easily extend a type-directed function with new cases. However, these new cases do not change the behavior of the original function, as advice does.

Achten and Hinze examined an extension to the Clean language [1] to allowing dynamic types and generic programming to interact cleanly [2]. This extension is quite flexible, even providing polykinded types. However, the generic functions in this language extension did not provide any kind of open extension. Cheney and Hinze developed a library with very similar capabilities in Haskell in terms of type representations [4]. However, it is not possible to provide special cases for user defined types in their library without defining new type representations, and modifying the infrastructure.

## 5.3 Aspect-oriented programming

There has been an increasing amount of research into AOP in the presence of parametric polymorphism. Our language, AspectML, directly builds upon the framework proposed by Walker, Zdancewic, and Ligatti [28], but extends it with polymorphic versions of functions, labels, label sets, stacks, pattern matching, advice and the auxiliary mechanisms to define the meaning of each of these constructs. We also define “around” advice and a novel type inference algorithm that is conservative over Hindley-Milner inference, which were missing from WZL’s work.

Concurrently with the development of AspectML, Matsuhara, Tatsuzawa, and Yonezawa [18] developed an aspect-oriented version of core OCaml they call Aspectual Caml. Their implementation effort is impressive and deals with several features we have not considered, including curried functions and extensible data types. There are similarities between AspectML and Aspectual Caml, but there are also many differences: Pointcut designators in AspectML can only reference names that are in scope, Aspectual Caml does not check pointcut designators for well-formedness, pointcuts in Aspectual Caml are second-class, and there is no formal description of the Aspectual Caml type system or operational semantics.

Later, Meng Wang, Kung Chen, and Siau-Cheng Khoo began examining language design problems in combining aspects with a polymorphic functional language [29]. Their design makes fundamentally different assumptions about aspects that lead to significant differences in expressiveness: Their advice is scoped such that it is not possible to install advice that will affect functions that have already been defined, their advice is named, like Aspectual Caml their pointcuts are second-class, finally their design does not provide a mechanism for examining the call-stack or obtaining information about the specific function being advised.

Jagadeesan, Jeffrey, and Riely have shown how to extend Generic Featherweight Java with aspects [9]. However, in their work they did not emphasize the importance of type analysis in the presence of parametric polymorphism, perhaps partly because of the predominance of subtype polymorphism in idiomatic Java programming. Their aspects however do implicitly perform an operation equivalent to `subtypeof` in Type-directed Java [34]. Finally, their language extension does not provide any mechanism for installing new advice at runtime.

The only other statically-typed implementation of aspects

that provides dynamically installable advice was Lippert’s modification of AspectJ [16] so that installing and uninstalling plug-ins, for the Eclipse development environment, that provide advice gives the expected behavior.

## 5.4 The expression problem

Our use of advice in this paper can be seen as a solution to the “expression problem”, but at the level of types. Coined by Phil Wadler, the expression problem arises when data types and operations on these types must be extended simultaneously. We can view types as an openly extensible data type, and then use advice to simultaneously extend type-directed functions with the necessary cases.

One of Masuhara, Tatsuzawa, and Yonezawa’s primary goals in developing Aspectual Caml was as a solution to the “expression problem”. Aspectual Caml’s ability to openly extend algebraic data types with additional data constructors required some method of extending functions over these data types to support the new cases. Much like we use advice to openly extend functions with specialized cases for user defined types, they used advice to extend functions with the needed cases for newly defined data constructors.

## 6. CONCLUSIONS AND FUTURE WORK

We have shown that runtime advising is critical to the expressive power of type-directed functions, just as runtime type analysis is key to the expressive power of AOP in polymorphic functional languages. Furthermore, our fully dynamic approach to advice allows the benefits of open extensibility to be realized even for types not known at compile-time. However, our solution is hardly the final word and there are many modifications or extensions we would like to explore to improve upon this work.

One problem that we did not address is that our type-directed functions cannot restrict the domain of types they may operate upon. For example, in AspectML functions cannot be compared for equality. Therefore, polymorphic structural equality in AspectML cannot provide a sensible notion of equality on functions types. However, its type, `<a> a -> a -> Bool`, does not provide any indication that it will abort execution when provided two functions. As mentioned, our language  $\lambda_C$  attempted to solve this problem using label-sets, but this solution seems too heavyweight for practical programming. Furthermore, label-sets are still too imprecise as they can be used to prevent using polymorphic structural equality on values of function type, but will also prevent the comparison of reference cells containing functions, which do have a notion of equality in AspectML. We plan to explore whether separating the kind `*` into two or more subkinds could provide an acceptable middle-ground. A more expressive solution might be to allow the programmer to use regular patterns restricting to specify the accepted domain.

Another direction we would like to explore is the capability to temporarily advise the behavior of functions in AspectML. Currently, once advice is installed, there is no way of removing it. Therefore, it is not possible to temporarily alter the behavior of a type-directed function, even if the desired scope of the extension is limited. Therefore, we plan to consider adding an `advise-with` form that allows programmers to write advice that only applies within the scope of a specific expression. A closely related extension would be to extend AspectML with named advice. We compile AspectML into a core calculus that already has first-class advice, so

there seems to be little reason not to expose the additional expressive power this provides.

Another limitation we did not discuss was that AspectML does not offer full reflexivity [26]. Because `typecase` may only analyze monotypes, and type variables may only be instantiated with monotypes, it is not possible to use type-directed operations on pointcuts and higher-rank polymorphic functions. We do not have enough experience to gauge whether this is a significant limitation in practical programming. However, we have studied the problem in other settings and believe that extending AspectML with higher-order type analysis would be feasible [31, 33].

Concurrently with this paper, Weirich developed a new approach to TDP in Haskell [32]. We expect that we will revise the primitives we have chosen for handling user defined data types and tuples in AspectML to take advantage of what her study has learned about the use of representation types.

Despite all of the benefits of TDP, another problem we did not address was that type-directed functions can break representation independence and potentially even the invariants of abstract data types. We have shown that it is possible to prove a generalization of Reynolds’s parametricity [25], theorem by tracking the information content of types and terms [30]. In the near future, we hope to extend AspectML with an information-flow type and kind system to evaluate how well this solution works in practice.

Finally, advice is an extremely imperative language feature. It does not seem plausible to us that, for example, wrapping advice inside a monad would allow the the benefits we described to be realized in a purely functional setting. It would be worthwhile for us to explore how to reconcile aspects and purely functional languages.

## 7. REFERENCES

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin. Dynamic typing in a statically typed language. *ACM Trans. Program. Lang. Syst.*, 13(2):237–268, 1991.
- [2] P. Achten and R. Hinze. Combining generics and dynamics. Technical Report NIII-R0206, Nijmegen Institute for Computing and Information, 2002.
- [3] R. Bird and L. Meertens. Nested datatypes. In J. Jeuring, editor, *Proceedings 4th Int. Conf. on Mathematics of Program Construction, MPC’98, Marstrand, Sweden, 15–17 June 1998*, volume 1422, pages 52–67. Springer-Verlag, Berlin, 1998.
- [4] J. Cheney and R. Hinze. A lightweight implementation of generics and dynamics. In *Haskell ’02: Proceedings of the 2002 ACM SIGPLAN workshop on Haskell*, pages 90–104, New York, NY, USA, 2002. ACM Press.
- [5] D. S. Dantas, D. Walker, G. Washburn, and S. Weirich. PolyAML: A polymorphic aspect-oriented functional programming language (extended version). Technical Report MS-CIS-05-07, University of Pennsylvania, May 2005.
- [6] R. Hinze. Polytypic values possess polykinded types. In *MPC ’00: Proceedings of the 5th International Conference on Mathematics of Program Construction*, pages 2–27, London, UK, 2000. Springer-Verlag.
- [7] R. Hinze and A. Löh. “Scrap your boilerplate” revolutions. In *MPC*, 2006.
- [8] R. Hinze, A. Löh, and B. C. D. S. Oliveira. “Scrap your boilerplate” reloaded. In *FLOPS*, pages 13–29, 2006.

- [9] R. Jagadeesan, A. Jeffrey, and J. Riely. Typed parametric polymorphism for aspects. 2006. To appear.
- [10] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. Griswold. An overview of AspectJ. In *Proceedings of the 15th European Conference on Object-Oriented Programming*. Springer-Verlag, 2001.
- [11] R. Lämmel and S. Peyton Jones. Scrap your boilerplate: a practical design pattern for generic programming. *ACM SIGPLAN Notices*, 38(3):26–37, Mar. 2003. Proceedings of the ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI 2003).
- [12] R. Lämmel and S. Peyton Jones. Scrap more boilerplate: reflection, zips, and generalised casts. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP 2004)*, pages 244–255. ACM Press, 2004.
- [13] R. Lämmel and S. Peyton Jones. Scrap your boilerplate with class: extensible generic functions. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming (ICFP 2005)*, pages 204–215. ACM Press, Sept. 2005.
- [14] K. Läufer and M. Odersky. An extension of ML with first-class abstract types. In *Proceedings of the SIGPLAN Workshop on ML and its Applications*, pages 78–91, June 1992.
- [15] X. Leroy. The Objective Caml system: Documentation and user’s manual, 2000. With Damien Doligez, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. Available from <http://caml.inria.fr>.
- [16] M. Lippert. Ajeer: an aspectj-enabled eclipse runtime. In *OOPSLA ’04: Companion to the 19th annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications*, pages 180–181, New York, NY, USA, 2004. ACM Press.
- [17] A. Löb, D. Clarke, and J. Jeuring. Dependency-style generic haskell. In *ICFP ’03: Proceedings of the eighth ACM SIGPLAN international conference on Functional programming*, pages 141–152, New York, NY, USA, 2003. ACM Press.
- [18] H. Masuhara, H. Tatsuzawa, and A. Yonezawa. Aspectual caml: an aspect-oriented functional language. In *Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming*, pages 320–330, New York, NY, USA, Sept. 2005. ACM Press.
- [19] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, Cambridge, Massachusetts, 1997.
- [20] C. Okasaki. From fast exponentiation to square matrices: an adventure in types. In *ICFP ’99: Proceedings of the fourth ACM SIGPLAN international conference on Functional programming*, pages 28–35, New York, NY, USA, 1999. ACM Press.
- [21] S. Peyton Jones. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [22] S. Peyton Jones, D. Vytiniotis, S. Weirich, and M. Shields. Practical type inference for arbitrary-rank types. Accepted for publication to the Journal of Functional Programming, 2005.
- [23] S. Peyton Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for GADTs. To appear in ICFP 2006, Apr. 2006.
- [24] G. D. Plotkin. A note on inductive generalization. In *Machine Intelligence*, volume 5, pages 153–163. Edinburgh University Press, 1970.
- [25] J. C. Reynolds. Types, abstraction, and parametric polymorphism. *Information processing*, pages 513–523, 1983.
- [26] V. Trifonov, B. Saha, and Z. Shao. Fully reflexive intensional type analysis. In *ICFP ’00: Proceedings of the fifth ACM SIGPLAN international conference on Functional programming*, pages 82–93, New York, NY, USA, 2000. ACM Press.
- [27] D. Vytiniotis, G. Washburn, and S. Weirich. An open and shut typecase. In *ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI)*, Long Beach, CA, USA, Jan. 2005. 15 pages.
- [28] D. Walker, S. Zdancewic, and J. Ligatti. A theory of aspects. In *Proceedings of the 8th ACM SIGPLAN International Conference on Functional Programming*, New York, NY, USA, Aug. 2003. ACM Press.
- [29] M. Wang, K. Chen, and S.-C. Khoo. On the pursuit of staticness and coherence. In *Proceedings of the 5th Workshop on Foundations of Aspect-Oriented Languages*, Mar. 2006.
- [30] G. Washburn and S. Weirich. Generalizing parametricity using information flow. In *IEEE Symposium on Logic in Computer Science (LICS)*, Chicago, IL, USA, June 2005.
- [31] S. Weirich. Higher-order intensional type analysis. In D. L. Métayer, editor, *11th European Symposium on Programming (ESOP)*, pages 98–114, Grenoble, France, Apr. 2002.
- [32] S. Weirich. RepLib: A library for derivable type classes, June 2006. Submitted. Available from <http://www.cis.upenn.edu/~sweirich/RepLib/>.
- [33] S. Weirich. Type-safe run-time polytypic programming. *Journal of Functional Programming*, 2006. 30 pages. To appear.
- [34] S. Weirich and L. Huang. A design for type-directed Java. In V. Bono, editor, *Workshop on Object-Oriented Developments*, ENTCS, 2004. 20 pages.