

**The Book of Types**  
Type-Level Programming in Haskell

Sandy Maguire



# Contents

<b>NOTICE</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
1.1 The Value of Types . . . . .	1
<b>2 Terms, Types and Kinds</b>	<b>3</b>
2.1 The Kind System . . . . .	3
2.2 Data Kinds . . . . .	5
2.3 Type-Level Functions . . . . .	6
2.4 Promoting Built-In Types . . . . .	8
<b>3 Associated Type Families</b>	<b>11</b>
3.1 Type Proxies . . . . .	11
3.2 Generating Associated Terms . . . . .	14
3.3 Ambiguous Types and Non-Injectivity . . . . .	16
<b>4 Generics</b>	<b>19</b>
4.1 Overview . . . . .	19
4.2 Generic Representations . . . . .	20
4.3 Deriving Any Class . . . . .	21
4.4 Generic Metadata . . . . .	23
4.5 Performance . . . . .	27
4.6 Yoneda and Codensity . . . . .	28



# NOTICE

1

This document is a very early pre-release of what is planned to be a much longer book on type-level programming in Haskell. It's being shared to gauge interest in such a project, as well as to receive early feedback on its style and tone.

2

3

4

Please feel free to share it!

5

6

*Any* questions that you have while reading this book are of particularly high interest to me. Every line of prose is numbered. If you'd be willing to share these questions and the line number they struck you, I'd be forever grateful to you.

7

8

9

My email is `sandy@sandymaguire.me`, and with your help, we can make this book the best it can possibly be.

10

11

Thanks for your time,

12

Sandy Maguire

13



# Chapter 1

14

## Introduction

15

### 1.1 The Value of Types

16

Haskellers are an odd sort of folk. Most of us, I'd suspect, have spent at least one evening of our lives trying to extol the virtues of a strong typesystem to a dynamically typed colleague. They'll say things like "I like Ruby because the types don't get in my way." Though as proponents of strong typing systems, our first instinct might be to forcibly connect our head to the table, I think this is a criticism worth keeping in mind.

17  
18  
19  
20  
21  
22

We Haskellers certainly have strong opinions about the value our types. They *are* useful, and *do* carry their weight in gold when coding and when refactoring. While we can dismiss our colleague's complaints with a wave of the hand and the justification that they've never seen a "type system" more powerful than Java's, types *often do* get in the way. We've just learned to blind ourselves to these shortcomings rather than bite the bullet and consider that maybe types aren't necessarily the perfect solution to every problem.

23  
24  
25  
26  
27  
28  
29

Simon Peyton-Jones, one of the primary authors of Haskell, is quick to acknowledge the fact that there are plenty of error-free programs that are ruled out by a type system. Consider, for example, a program which has a type-error, but never actually evaluates it:

30  
31  
32  
33

```
fst ("no problems", True <> 17)
```

Evaluation of such an expression will happily produce "no problems" at runtime, despite the fact that we consider it to be "ill-typed." The usefulness of this example is admittedly low, but the point stands; types often do get in the way of perfectly reasonable programs.

34  
35  
36  
37

Sometimes this obstruction comes under the guise of "it's not clear what type this thing should have." One particularly poignant case of this is the C function `printf`.

38  
39  
40

```
int printf (const char *format, ...)
```

41 If you’ve never had the pleasure of using `printf`, it takes a “format string”  
 42 as its first parameter, parses it, and depending on its contents, will conditionally  
 43 read additional parameters. For example, the format string “`hello %s`” takes  
 44 an additional string, and will interpolate it in place of the `%s`. Likewise, the  
 45 specifier `%d` describes interpolation of a signed decimal integer.

46 The following calls to `printf` are all valid:

```
47 • printf("hello %s", "world") // hello world
48 • printf("%d + %d = %s", 1, 2, "three") // 1 + 2 = three
49 • printf("no specifiers") // no specifiers
```

50 On first inspection, it’s not at all clear what a strongly-typed `printf` should  
 51 have. Despite this, it’s inarguably a *useful* function.

52 The documentation for `printf` is quick to mention that the format string  
 53 should not be provided by the user; doing so opens vulnerabilities where attack-  
 54 ers can craft a specialized format string that exploits these specifiers and gain  
 55 access to the system. Indeed, this is often the first homework assignment in any  
 56 university-level course on software security.

57 To be clear, the vulnerabilities in `printf` occur when the format string’s  
 58 specifiers do not align with the additional arguments given to the function call.  
 59 The following calls to `printf` both cause undefined behavior (ie. crashing or  
 60 security holes):

```
61 • printf("%d") // corrupts the stack
62 • printf("%s", 1) // reads an arbitrary amount of memory
```

63 The problem here is that C’s typesystem is insufficiently expressive to de-  
 64 scribe `printf`, and thus allows type errors to make it all the way to run-  
 65 time, in the form of undefined behavior. The type of `int printf(const char`  
 66 `*format, ...)` is unsatisfying for a number of reasons: that it’s “format string”  
 67 is not really a string, that it doesn’t constrain the number of arguments it takes,  
 68 and that it doesn’t enforce anything about the types of those arguments.

69 In my mind, preventing security holes is a much more important aspect of  
 70 the value of types than “null is the billion dollar mistake” or whichever other  
 71 arguments are in vogue today. We will return to the problem of `printf` in  
 72 Chapter 3.

73 With few exceptions, the prevalent attitude of Haskellers has been to dismiss  
 74 the usefulness of ill-typed programs rather than to admit the uncomfortable  
 75 truth that our favorite language can’t do something other languages can.

76 But there is a way to have our cake and eat it too. Indeed, Haskell *is* capable  
 77 of expressing things as oddly-typed as `printf`, for those of us willing to put in  
 78 the effort to learn how. This book aims to be *the* comprehensive manual for  
 79 getting you from here to there, from a competent Haskell programmer to one  
 80 who convinces the compiler to do their work for them.



## Chapter 2

81

# Terms, Types and Kinds

82

## 2.1 The Kind System

83

It is an unfortunate fact of the world that we must walk before we can run. While most of this book will introduce new concepts as *solutions to problems*, where their uses will be motivated, we do not yet have enough vocabulary to even talk about most of the kinds of problems we'd like to be able to solve. And thus it is necessary to discuss a little about the fundamentals of type-level programming. Rest assured that we'll keep this section as brief as possible before diving into applications of the concepts we are about to learn.

84  
85  
86  
87  
88  
89  
90

In everyday Haskell programming, the fundamental building blocks are those of *terms* and *types*. Terms are the values you can manipulate, the things that exist at runtime, and the types are little more than sanity-checks: proofs to the compiler (and ourselves) that the programs we're trying to write make some amount of sense.

91  
92  
93  
94  
95

Completely analogously, the fundamental building blocks for type-level programming are *types* and *kinds*. Now it is the types we are manipulating, and the kinds become the proofs we use to keep ourselves honest.

96  
97  
98

The kind system, if you're unfamiliar with it, can be reasonably described as “the type system for types.” By that line of reasoning, then, kinds are loosely “the types of types.”

99  
100  
101

Consider the numbers 4 and 5, both of type `Int`. As far as the compiler is concerned, we could replace every instance of 4 with 5 in our program, and the whole thing would continue to compile. The program itself might do something different, but by virtue of both being of type `Int`, 4 and 5 are interchangeable to the typechecker.

102  
103  
104  
105  
106

If kinds are the “types of types,” a reasonable question might be “what is the kind of `Int`?” As it happens, the kind of `Int` is `TYPE` (historically written as `*`). Sometimes we call types of kind `TYPE` *value types*.

107  
108  
109

`TYPE` is the kind of any type which has values that exist at runtime<sup>1</sup>.

110

---

<sup>1</sup>The pedantic reader might notice that `Void` has kind `TYPE`, but no inhabitants. Ah, but

111 Other types whose kind is `TYPE` include `Bool`, `Maybe a` and `ExceptT String`  
 112 `IO ()`.

113 However, it's often more informative to look at things which do not fall  
 114 into the same pattern. Consider `Maybe`—not `Maybe a` mind—but just the type  
 115 constructor itself: `Maybe`. `Maybe` isn't like `Int`; there are no terms whose type  
 116 is `Maybe`, nor can we replace every instance of `Bool` in our program with `Maybe`  
 117 and expect the thing to still compile. Instead what we get is a *kind error*.

118 The reason for this is not hard to see; `Maybe` is a type constructor – it has a  
 119 “type-shaped hole” that needs to be filled in before it is meaningful. As such,  
 120 we say the kind of `Maybe` is `TYPE → TYPE`. Whenever you give `Maybe` a type  
 121 of kind `TYPE`, it will give you back a type of kind `TYPE`. We call `Maybe a` a  
 122 *higher-kinded type*.

123 It's a bit of a tricky thing to put into words, but if you've written some  
 124 amount of Haskell, this should fit in with your intuitions behind how such things  
 125 work. While it's fine to talk about `Maybe String`, neither `Maybe` nor `Maybe`  
 126 `Maybe` are meaningful. That's because we're used to thinking about `TYPE` in  
 127 our everyday Haskell experiences.

128 On its own, `Maybe` has the wrong kind for us to care about, and `Maybe Maybe`  
 129 is meaningless: it's a *kind error* because `Maybe` wants a `TYPE` but we're giving  
 130 it a `TYPE → TYPE`. This is just as nonsensical as trying to pass in the function  
 131 `(+1) :: Int -> Int` as an argument to itself.

132 Several higher-kinded types you're already familiar with will be of kind `TYPE`  
 133 `→ TYPE`, including `IO` and `[ ]` (lists).

134 But what about `Either`? `Either` takes two type parameters, both of which  
 135 must be value-types. We thus say that `Either` has kind `TYPE → TYPE →`  
 136 `TYPE`: it requires two parameters of kind `TYPE` before it is fully saturated  
 137 (“out of type variables”).

138 The function arrow `(->)` also has kind `TYPE → TYPE → TYPE`, because it  
 139 requires two types to be filled in before it's a saturated value type.

140 To take things to the next level, what is the kind of the monad transformer  
 141 `MaybeT`? Well, like `Either`, it takes two type parameters, but unlike `Either`,  
 142 one of those parameters must be a `Monad` (ie. not a value type.) The eliminator  
 143 for `MaybeT` is `runMaybeT :: MaybeT m a -> m (Maybe a)`, and from here we  
 144 can work backwards.

145 We know that `Maybe` has kind `TYPE → TYPE`, so `a` must have kind `TYPE`.  
 146 Since we know the kind of `Maybe a` is `TYPE`, `m` must then have kind `TYPE →`  
 147 `TYPE`. Putting it all together, the kind of `MaybeT` must then be `(TYPE →`  
 148 `TYPE) → TYPE → TYPE`.

149 I promise, after some practice this stuff will come just as naturally to you  
 150 as the typechecking rules do.

151 However, kinds apply to everything at the type-level, not just the things we  
 152 traditionally think of as “types.” For example, the type of `show` is `Show a =>`  
 153 `a -> String`. This `Show` thing exists as part of the type signature, even though  
 154 it's not really a “type” in the traditional sense of the word. Does `Show a` have

---

in fact `undefined` is an inhabitant of `Void`!

a kind? 155

It does indeed, and that kind is `CONSTRAINT`. More generally, `CONSTRAINT` 156  
is the kind of any fully-saturated type class. 157

With knowledge of this, can you guess what kind `Show` by itself (unsaturated) 158  
has? If you said `TYPE → CONSTRAINT`, you're right! 159

What about the kind of `Functor`?<sup>2</sup> Of `Monad`?<sup>3</sup> Of `MonadTrans`?<sup>4</sup> 160

Without further language extensions, this is the extent of the expressiveness 161  
of Haskell's kind system. As you can see, it's actually quite limited – we have 162  
no notion of polymorphism, of being able to define our own kinds, or of being 163  
able to write functions. 164

Fortunately, those things are the subject matter of the remainder of this 165  
book – techniques, tools and understanding for Haskell's more esoteric language 166  
extensions. 167

## 2.2 Data Kinds 168

By enabling Haskell's `-XDataKinds` extension, we gain the ability to talk about 169  
kinds other than `TYPE` and `CONSTRAINT` (and their derivatives.) In particular, 170  
`-XDataKinds` lifts data constructors into *type constructors* and types into *kinds*. 171

As an example, given the familiar `Bool` type definition: 172

```
data Bool 173
  = True 174
  | False 175
```

we gain the following *kind* definition<sup>5</sup>: 176

```
kind Bool 177
  = 'True 178
  | 'False 179
```

Which is to say that we have now declared the types `'True` and `'False`, 180  
both of kind `BOOL`. We call `'True` and `'False` *promoted data constructors*. 181  
The leading ticks on the identifiers (the `'` in `'True`) are used to distinguish 182  
promoted data constructors from everyday type constructors, in the common 183  
case of a type with a single data constructor: 184

```
data Unit = Unit 185
```

In this example, it's very important to differentiate between the *type con-* 186  
*structor* `Unit` (of kind `TYPE`), and the *promoted data constructor* `'Unit` (of 187  
kind `UNIT`.) This is a subtle point, and can often lead to inscrutable compiler 188

<sup>2</sup>(`TYPE → TYPE`) → `CONSTRAINT`

<sup>3</sup>Likewise.

<sup>4</sup>(`TYPE → TYPE`) → `TYPE → CONSTRAINT`

<sup>5</sup>Note that this is not legal Haskell syntax.

189 errors; while it's fine to ask for values of type `Maybe Unit`, it's a *kind error* to  
 190 ask for `Maybe 'Unit`—because `'Unit` is the wrong kind!

191 Promoted data constructors are of the wrong kind to ever exist at runtime,  
 192 which raises the question “what good are they?” Without any other fancy  
 193 type-level machinery, we can use them as phantom parameters to handle state  
 194 transitions.

this is a shit example

195 Imagine a concurrent application which needs to acquire some locks, but  
 196 will deadlock if it attempts to acquire the same lock twice. It would be nice  
 197 to enforce this invariant at the type-level, so that the compiler will refuse to  
 198 compile any code which attempt to acquire the same lock twice.

```

199     data LockState
200         = Unlocked
201         | Locked
202
203     data Lock (s :: LockState) = Lock
204         { getLockId :: Int
205         }
206
207     newLock :: IO (Lock 'Unlocked)
208     newLock = ...
209
210     withLock :: Lock 'Unlocked -> (Lock 'Locked -> IO a) -> IO a
211     withLock = ...
  
```

212 In this example, we use `-XDataKinds` to lift `LockState` into `LOCKSTATE`,  
 213 which we then use as a phantom parameter to `Lock`. When we construct a  
 214 new lock via `newLock` we get back a `Lock 'Unlocked`. However, the `withLock`  
 215 function takes an unlocked lock, and gives us back a locked one, ensuring we  
 216 can't lock it again! We've successfully used the type system to prevent us from  
 217 creating a deadlock.

218 Notice that in the definition of `Lock`, we use a kind signature of `LOCKSTATE`.  
 219 This requires enabling the `-XKindSignatures` extension, which is one you will  
 220 always want enabled when doing type-level programming.

## 221 2.3 Type-Level Functions

222 Where `-XDataKinds` really begins to shine, however, is through the introduction  
 223 of *closed type families*. You can think of closed type families as *functions at the*  
 224 *type level*. Compare the function `or`, which computes the OR of two `Bools`:

```

225     or :: Bool -> Bool -> Bool
226     or True  _ = True
227     or False y = y
  
```

While we’re unfortunately unable to automatically promote term-level functions into type-level ones, we can write `or` as a closed type family (if we first remember to enable the `-XTypeFamilies` extension):

```
type family Or (x :: Bool) (y :: Bool) :: Bool where
  Or 'True y = 'True
  Or 'False y = y
```

Line for line, the similarities between `or` and `Or` are analogous. The type family `Or` requires a capital letter for the beginning of its name, because it exists at the type-level, and besides having a more verbose “kind-signature,” the two definitions proceed almost exactly in lockstep.

While the metaphor between type families and functions is enticing, it isn’t entirely *correct*. The analogues break down in several ways, but the most important one is that *type families must be saturated*. Another way of saying this is that all of a type family’s parameters must be specified simultaneously; there is no currying available.

For example, this means we’re unable to write a useful promoted `map :: (a -> b) -> [a] -> [b]`:

**Broken Code**

```
type family Map (x :: a -> b) (i :: [a]) :: [b]
  where
    Map f '[] = '[]
    Map f (x ': xs) = f x ': Map f xs

type AllTrue = Map (Or 'True)
                  '[ 'True, 'False, 'False ]
```

Attempting to compile this gives us the following error:

```
error:
  • The type family ‘Or’ should have 2 arguments, but has been
    given 1
  • In the type synonym declaration for ‘AllTrue’
```

What the error is trying to tell us is that we used the `Or` closed type-family in a non-saturated way. We only passed it one parameter instead of the two it requires, and so unfortunately GHC refuses to compile this program.

While there is nothing preventing us from writing `Map`, its usefulness is severely limited due to our inability to curry the type family we pass into it.

Before we leave this topic, look again at our definition of `Or`. Pay close attention to its “kind-signature.” We write it as `Or (x :: Bool) (y :: Bool) :: Bool`, rather than `Or x y :: Bool -> Bool -> Bool`. The kinds of type families are tricky beasts; the kind you write after the `::` is the kind of the type *returned* by the type family, *not* the kind of the type family itself.

```

260 type family Foo (x :: Bool) (y :: Bool) :: Bool
261 type family Bar x y :: Bool -> Bool -> Bool

```

262 Take a moment to think about the kinds of `Foo` and `Bar`. While `Foo` has  
 263 kind `BOOL → BOOL → BOOL`, `Bar` has kind `Type -> Type -> Bool -> Bool`  
 264 `-> Bool`.

## 265 2.4 Promoting Built-In Types

266 With `-XDataKinds` enabled, almost all<sup>6</sup> types automatically promote to kinds,  
 267 including the built-in ones. Since built-in types (strings, numbers, lists and  
 268 tuples) are special at the term level—at least in terms of syntax—we should  
 269 expect they should behave a little strange at the type level too.

270 When playing with these built-in promoted types, it’s necessary to first  
 271 import the `GHC.TypeLits` module. `GHC.TypeLits` defines the kinds themselves,  
 272 as well as all of the useful type families for manipulating them. We’ll cover this  
 273 more in detail in a second.

274 Strings promote in the way you’d expect, except that at the type level they  
 275 are *not* equivalent to lists of characters. Enabling `-XDataKinds` provides access  
 276 to the `SYMBOL` kind (defined in `GHC.TypeLits`), which corresponds to lifted  
 277 strings. Symbol type-literals can be written as `"hello" :: Symbol`, although  
 278 the kind signature is unnecessary.

279 Numbers are a little more odd, in that only the natural numbers (0, 1, 2, ...)   
 280 can be promoted. These natural numbers, naturally enough, are of kind `NAT`.

281 Lists promote exactly in the way they should, if they were defined as

```

282 data [] a
283     = []
284     | a : [a]

```

285 which is to say that when promoted, we get the following type constructors:  
 286 `' []` of kind `[A]`, and `' :` of kind `A → [A] → [A]`. When compared against the *data*  
 287 constructors of lists, `[] :: [a]` and `(:) :: a -> [a] -> [a]`, with a little  
 288 concentration, it should make some sense. Because lists’ data constructors have  
 289 symbolic names, they require the `-XTypeOperators` extension to be enabled if  
 290 you want to work with them. Don’t worry though, GHC will helpfully remind  
 291 you if you forget.

292 Note that the promoted data constructors for lists are polymorphic. We’ll  
 293 talk more about kind-level polymorphism later.

294 There is another subtle point to be noted when dealing with list-kinds. While  
 295 `[Bool]` is of kind `TYPE` and describes a term-level list of booleans, the type  
 296 `' [Bool]` is of kind `[TYPE]` and describes a type-level list with one element  
 297 (namely, the type `Bool`.)

298 Further care should be taken when constructing a promoted list; due to the  
 299 way GHC’s lexer parses character literals (`' a'`), make sure you add a space after

<sup>6</sup>Not all types. GADTs and other “tricky” data constructors fail to promote

starting a promoted list. While `'[ 'True ]` is fine, `'['True]` is unfortunately  
a parse error. 300

Finally, tuples are promoted with a leading tick. For example, `'( 5, "hello"`  
) has kind `(NAT, SYMBOL)`. The aforementioned parsing gotcha applies here 302  
as well, so be careful. 303  
304





## Chapter 3

305

# Associated Type Families

306

### 3.1 Type Proxies

307

Let's return to our earlier discussion about `printf`.

308

One of Haskell's most profound lessons is a deep appreciation for types, and with it, the understanding that `Strings` are suitable only for *unstructured* text. Our format "strings" most certainly *are* structured, and thus should not be `Strings`.

309

310

311

312

If `printf`'s format string isn't really a string, what is it?

313

314

Well, if we look only at the specifiers in the format string, they're a kind of type signature, describing not only the number of parameters, but also their types. For example, the format string `"%s%d%d"` could be interpreted in Haskell as a function that takes a string, two integers, and returns a string—the concatenation of pushing all of those parameters together. In other words, `"%s%d%d"` corresponds to the type `String -> Int -> Int -> String`.

315

316

317

318

319

But, a format string is not only specifiers; it can also contain arbitrary text that should be strung together between the arguments. In our earlier example, this corresponds to format strings like `"hello %s"`. The type corresponding to this function is still only `String -> String`, but its actual implementation should be `\s -> "hello " <> s`.

320

321

322

323

324

We can build the desired structure for the semantics behind these format strings by realizing that what we're trying to describe is nothing more than a sequence of `TYPE` and `SYMBOL` (promoted `Strings`), in any order. The `TYPE`s correspond to parameters to our formatting function, and the `SYMBOL`s are text to be interspersed between.

325

326

327

328

329

With this in mind, we can write the following data definition, whose only purpose is to let us keep track of our kinds:

330

331

```
data (a :: k1) :<< (b :: k2)
infixr 5 :<<
```

332

333

The `(:<<)` symbol was chosen due to the similarity it has with C++'s `<<` output stream operator, but has no other special meaning to Haskell or to us.

334

335

336 Notice here that `(<<)` doesn't have any data constructors, so we are un-  
 337 able to construct one of them at the term-level. But because we have a type  
 338 constructor, we can build types out of them.

339 In GHCi, we can use the `:kind!` command to inspect our handiwork, which  
 340 will tell us the *kind of a given type*:

```
341 > :kind! (<<)
342 (<<) :: k1 -> k2 -> Type
343 = (<<)
```

344 `(<<)`, when promoted, has kind  $K1 \rightarrow K2 \rightarrow \text{TYPE}$ , corresponding roughly  
 345 to a promoted 2-tuple. The kind of a saturated `(<<)` is `TYPE`, but this is  
 346 unimportant for our purposes; all we want is a way of defining sequences of  
 347 things of arbitrary kinds.

348 To convince ourselves that this works, we can again ask GHCi:

```
349 > :kind! "hello " :<< String :<< "!"
350 "hello " :<< String :<< "!" :: *
351 = "hello " :<< (String :<< "!" )
```

352 Notice that due to our `infixr 5 :<<` declaration, repeated applications of  
 353 `(<<)` associate as we'd expect.

354 Armed with a means of storing our format “string”, our next step is to use  
 355 it to construct the proper type signature of our formatting function. Which is  
 356 to say, given eg. a type `Int :<< ":" :<< Bool :<< "!"`, we'd like to produce  
 357 the type `Int -> Bool -> String`. This sounds like a type-level function, and  
 358 so we should immediately begin to think about type families.

359 However, instead of using *closed* type families which are useful when pro-  
 360 moting functions from the term level to the type level, we instead will use an  
 361 *associated type family*. Associated type families, as the name suggests, are as-  
 362 sociated with a type class, and provide a convenient way to bundle term-level  
 363 code with computed types. We'll talk about the bundled term-level code in a  
 364 moment, but first, we can define our associated type family:

```
365 class HasPrintf a where
366   type Printf a :: Type
```

367 Here we're saying we have a typeclass `HasPrintf a`, of which every instance  
 368 must provide an associated type `Printf a`, whose kind must be `TYPE`. `Printf`  
 369 `a` will correspond to the desired type of our formatting function, and we will  
 370 construct it in a moment.

371 Defining a type family as an associated type family allows us to take advan-  
 372 tage of Haskell's capabilities for overlapping typeclasses, which will simplify our  
 373 logic a great deal.

374 For simplicity, we will say our format types will always be of the form `a :<<`  
 375 `... :<< "symbol"`, which is to say, some they will always end with a `SYMBOL`.

Such a simplification gives us a convenient base case for the structural recursion we want to build.

If you're unfamiliar with the concept, structural recursion refers to the technique of producing something by tearing a recursive structure apart into smaller and smaller pieces, until you find a case simple enough you know how to handle. It's really just a fancy name for "divide and conquer."

In our `Printf` example, we will use three cases:

1. `instance HasPrintf (text :: Symbol)`
2. `instance HasPrintf a => HasPrintf ((text :: Symbol) :<< a)`
3. `instance HasPrintf a => HasPrintf ((param :: Type) :<< a)`

With these three cases, we can tear down any right-associative sequence of `(:<<)`s via case 2 or 3 until we run out of `(:<<)` constructors, where we are finally left with a `SYMBOL` we can handle with case 1.

Case 1 corresponds to having no more parameters, which in turn means it should expand to the desired return type of our formatting function.

```
instance HasPrintf (sym :: Symbol) where
  type Printf sym = String
```

Case 2 corresponds to having additional text we want to inject into our final formatted string, but it doesn't correspond any parameters on its own, so its `Printf` type should be the same as the one it received via structural recursion:

```
instance HasPrintf a
  => HasPrintf ((text :: Symbol) :<< a) where
  type Printf (text :<< a) = Printf a
```

We know that this is an acceptable thing to do, because `Printf a` comes associated with `HasPrintf a`, which is a constraint on our instance of `HasPrintf`.

Case 3 is the most interesting; here we want to add our `param` type as a parameter to the generated function. We can do that by defining `Printf` as an arrow type:

```
instance HasPrintf a
  => HasPrintf ((param :: Type) :<< a) where
  type Printf (param :<< a) = param -> Printf a
```

Here we're saying our formatting type requires a `param`, and then gives back our recursively-defined `Printf a` type. Strictly speaking, the `TYPE` kind signature here isn't necessary—GHC will infer it based on `param -> Printf a`—but it adds to the readability, so we keep it.

We can walk through our earlier example of `Int :<< ":" :<< Bool :<< "!"` to convince ourselves that `Printf` expands correctly. First, we see that `Int :<< ":" :<< Bool :<< "!"` is of the form `(param :: Type) :<< a` given the equations `(param ~ Int)` and `(a ~ ":" :<< Bool :<< "!")`.

415 From here, we expand the definition of `Printf (param :<< a)` into `param`  
 416 `-> Printf a`, or, substituting for our earlier type equalities: `Int -> Printf`  
 417 `(":" :<< Bool :<< "!")`.

418 We continue matching `Printf (":" :<< Bool :<< "!")` and notice now  
 419 that it matches case 2, giving us `Int -> Printf (Bool :<< "!")`. Expansion  
 420 here again follows case 3, and expands to `Int -> Bool -> Printf "!"`.

421 Finally, we have run out of `(:<<)` constructors, and so `Printf "!"` matches  
 422 case 1, where `Printf text = String`. Here our recursion ends, and we find  
 423 ourselves with the generated type `Int -> Bool -> String`, exactly the type  
 424 we were looking for.

425 Analysis of this form is painstaking and time-intensive. Instead, in the  
 426 future, we can just ask GHCi if we got it right, again with the `:kind!` command:

```
427 > :kind! Printf (Int :<< ":" :<< Bool :<< "!")
428 Printf (Int :<< ":" :<< Bool :<< "!") :: Type
429 = Int -> Bool -> String
```

430 Much easier.

## 431 3.2 Generating Associated Terms

432 Building the type `Printf a` is wonderful and all, but producing a type without  
 433 a corresponding term isn't going to do us much good. Our next step is to update  
 434 the definition of `HasPrintf` to also provide a `format` function.

```
435 class HasPrintf a where
436   type Printf a :: *
437   format :: String -> Proxy a -> Printf a
```

438 The type of `format` is a little odd, and could use an explanation. Looking at  
 439 the second parameter first, we find a term of type `Proxy a`. This `Proxy` exists  
 440 only to allow Haskell to find the correct instance of `HasPrintf` from the call-site  
 441 of `format`. You might think Haskell would be able to find an instance based on  
 442 the `a` in `Printf a`, but this isn't so for reasons we will discuss soon.

443 The first parameter, the `String` is an implementation detail, and will act as  
 444 an accumulator where we can keep track of all of the formatting done by earlier  
 445 steps in the recursion.

446 We can update our instance definitions of the three cases so they correctly  
 447 implement `format`. In the first case, we have no work to do, so the only thing  
 448 necessary is to return the accumulator and append the final text to it.

```
449 instance KnownSymbol text => HasPrintf (text :: Symbol) where
450   type Printf text = String
451   format s _ = s <> symbolVal (Proxy @text)
```

`symbolVal` is a function that converts a `SYMBOL` into a `String`. For example, it will turn the type `"hello" :: Symbol` into the *term* `"hello" :: String`. `symbolVal`'s is `KnownSymbol sym => Proxy sym -> String`, and all this `KnownSymbol` stuff is simply a proof that GHC knows what `SYMBOL` we're talking about; it will automatically generate the `KnownSymbol` instance for us, so it's nothing we need to worry about.

Case 2 is very similar; here we want to update our accumulator with the `symbolVal` of `text`, but also structurally recursively call `format`. This requires conjuring up a `Proxy a`, which we can do via `-XTypeApplications`:

```
instance (HasPrintf a, KnownSymbol text)
  => HasPrintf ((text :: Symbol) :<< a) where
  type Printf (text :<< a) = Printf a
  format s _ = format (s <> symbolVal (Proxy @text))
                    (Proxy @a)
```

All that's left is case 3, which if you've been paying attention to the other cases, should look familiar:

```
instance (HasPrintf a, Show p)
  => HasPrintf ((param :: Type) :<< a) where
  type Printf (param :<< a) = param -> Printf a
  format s _ p = format (s <> show p) (Proxy @a)
```

Notice the `p` parameter to our `format` function here—this corresponds to the `param` parameter in case 3's `Printf` instance. For any specifier, we use its `Show` instance to convert the parameter into a string, and append it to our accumulator.

With all three of our cases covered, we appear to be finished. We can define a helper function to hide the accumulator from the user, since it's purely an implementation detail:

```
printf :: HasPrintf a => Proxy a -> Printf a
printf = format ""
```

Firing up GHCi allows us to try it:

```
> printf (Proxy @"test")
"test"
> printf (Proxy @(Int :<< " + " :<< Int :<< " = 3")) 1 2
"1 + 2 = 3"
> printf (Proxy @(String :<< " world!")) "hello"
"\hello\ world!"
```

It works pretty well for our first attempt, all things considered. One noticeable flaw is that `Strings` gain an extra set of quotes due to being `shown`. We can fix this infelicity by providing a special instance of `HasPrintf` just for `Strings`:

```

492 instance {-# OVERLAPPING #-} HasPrintf a
493     => HasPrintf (String :<< a) where
494     type Printf (String :<< a) = String -> Printf a
495     format s _ param = format (s <> param) (Proxy @a)

```

496 Writing this instance will require the `-XFlexibleInstances` extension, since  
497 the instance head is no longer just a single type constructor and type variables.  
498 We mark the instance with the `{-# OVERLAPPING #-}` pragma because we'd  
499 like to select this instance instead of case 3 when the parameter is a `String`.

```

500 > printf (Proxy @(String :<< " world!")) "hello"
501 "hello world!"

```

502 Marvelous.

503 What we've accomplished here is a type-safe version of `printf`, but recog-  
504 nizing that C++'s "format string" is better thought of as a "structured type  
505 signature." Using type-level programming, we were able to convert such a thing  
506 into a function with the correct type, that implements nontrivial logic.

### 507 3.3 Ambiguous Types and Non-Injectivity

508 We return now to the question of what's going with that pesky `Proxy` thing in  
509 the earlier examples. The purpose of `Proxy a` in the type signature of `format` is  
510 to gently guide GHC towards finding the instance of `HasPrintf` that we want.

511 But `format :: String -> Proxy a -> Printf a` already has an `a` visible  
512 in its type signature, in the form of `Printf a`. Why isn't this enough to drive  
513 instance resolution?

514 To see why this must be the case, let's look at a simpler example of a closed  
515 type family:

```

516 type family AlwaysUnit a where
517     AlwaysUnit a = ()

```

518 The `AlwaysUnit` family maps every type to `()`. The usefulness of such a  
519 thing is limited, but it serves to illustrate our problem.

**Broken Code**

```

class TypeName a where
  typeName :: AlwaysUnit a -> String

instance TypeName String where
  typeName _ = "String"

instance TypeName Bool where
  typeName _ = "Bool"

name :: String
name = typeName ()

```



The problem, of course is that both `AlwaysUnit String ~ ()` and `AlwaysUnit Bool ~ ()`, which means that given `AlwaysUnit a ~ ()`, we can't learn anything about `a`. More specifically, the problem is that `AlwaysUnit` doesn't have an inverse; there's no `Inverse` type family such that `Inverse (AlwaysUnit a) ~ a`. In mathematics, this property is known as *non-injectivity*.

Consider an analogous example from cryptography; just because you know the hash of someone's password is `1234567890abcdef` doesn't mean you know what the password is; any good hashing function, like `AlwaysUnit`, is *one way*. Just because we can go one way doesn't mean we can go backwards.

In the case of `name` above, when GHC sees `typeName ()`, it doesn't have enough information to go backwards and figure out what `a` is supposed to be, because all it know is `AlwaysUnit a ~ ()`. But again, this is true for all types `a`, so it is not a particularly helpful fact to know.

Going back to `Printf a`, we can see that both `Printf "hello" ~ String` and `Printf "goodbye" ~ String`, so `Printf` can't possibly be injective.

The solution to non-injectivity is to give GHC some other way of determining the otherwise ambiguous type. This can be done like in our examples by adding a `Proxy a` parameter whose only purpose is to drive inference, or it can be accomplished by enabling `-XAllowAmbiguousTypes` at the definition site, and using `-XTypeApplications` at the call-site to fill in the ambiguous parameter manually. We will discuss these two extensions more in a later chapter.





# Chapter 4

541

# Generics

542

## 4.1 Overview

543

Haskell offers two kinds of polymorphism: parametric polymorphism, which has one definition for every possible type (think `head :: [a] -> a`); and ad-hoc polymorphism, where every type can do its own thing (`mempty :: a`). But for our purposes, there's also a third category to keep in mind: *boilerplate polymorphism*.

544

545

546

547

548

Boilerplate polymorphism doesn't have any formal definition, but it's the kind of thing you recognize when you see it. It's ad-hoc polymorphism that doesn't require any sort of thought in order to write. `Eq`, `Show` and `Functor` instances are good examples of boilerplate polymorphism—there's nothing interesting about writing these instances. The tedium of writing boilerplate polymorphism is somewhat assuaged by the compiler's willingness to write some of them for us.

549

550

551

552

553

554

555

Consider the `Eq` typeclass; while every type needs its own implementation of `(==)`, these implementations are always of the form:

556

557

```
instance Eq Foo where
  F0      == F0      = True
  F1 a1   == F1 a2   = a1 == a2
  F2 b1 c1 == F1 b2 c2 = b1 == b2 && c1 == c2
  _       == _       = False
```

558

559

560

561

562

Instances of `Eq` are always the same; the same data constructors are equal if and only if all of their components are equal.

563

564

Boilerplate polymorphism is mindless work to write, but needs to be done. In the case of some of the standard Haskell typeclasses, GHC is capable of writing these instances for you via `deriving`. Unfortunately, for custom typeclasses we're on our own, without any direct support from the compiler.

565

566

567

568

As terrible as this situation appears, all hope is not lost. Using `GHC.Generics`, we're capable of writing our own machinery for helping GHC derive our typeclasses, all in regular Haskell code.

569

570

571

## 572 4.2 Generic Representations

573 Before diving into the generic machinery, it will be instructive to spend some  
 574 time thinking about the *functoriality* of datatypes. As it happens, all datatypes  
 575 have a “canonical” representation as a *sum-of-products*. In Haskell, the canon-  
 576 ical sum type is `Either`, and the canonical product is the 2-tuple `(,)`. As it  
 577 happens, every `data` and `newtype` definition is equivalent to some composition  
 578 of sums and products.

579 Take for example, `Maybe a`, which you’ll recall has this definition:

```
580 data Maybe a
581     = Just a
582     | Nothing
```

583 `Maybe a`, in its canonical sum-of-products form is `Either a ()`, which we  
 584 can prove is the same thing via an isomorphism:

```
585 toCanonical :: Maybe a -> Either a ()
586 toCanonical (Just a) = Left a
587 toCanonical Nothing = Right ()
588
589 unCanonical :: Either a () -> Maybe a
590 unCanonical Left a   = Just a
591 unCanonical (Right ()) = Nothing
```

592 Since we can convert to and from the canonical representation without losing  
 593 any information in either direction, we know that these types must be equivalent.  
 594 But the question remains: who cares?

595 Well, the point is that if we have a small number of primitive building blocks,  
 596 we can write our generic code to operate over those primitives. If we then also  
 597 had the ability to convert to and from our canonical representations at will, we’d  
 598 be able to take regular Haskell data, convert it to its canonical representation,  
 599 manipulate it, and then convert it back.

600 How can such a thing be possible? The secret is in the `-XDeriveGeneric`  
 601 extension, which will automatically derive an instance of `Generic` for you:

```
602 class Generic a where
603     type Rep a :: Type
604     from :: a -> Rep a x
605     to   :: Rep a x -> a
```

606 `Rep a` is the canonical sum-of-products representation of our type `a`, while  
 607 `from` and `to` correspond to our above isomorphism between the two types.

608 Let’s look at `Rep Bool` for inspiration about what this thing might look like.

```
609 > :kind! Rep Bool
610 Rep Bool :: * -> *
```

```

= D1 ('MetaData "Bool" "GHC.Types" "ghc-prim" 'False) 611
    ( C1 ('MetaCons "False" 'PrefixI 'False) U1        612
      :+: C1 ('MetaCons "True" 'PrefixI 'False) U1     613
    )                                                  614

```

Quite a mouthful, but at its heart the interesting parts of this are the `(:++:)` and `U1` types. These correspond to the *canonical sum* and *canonical unit*, respectively. When viewed beside the definition of `Bool`, some similarities appear:

```

data Bool 619
  = False 620
  | True   621

```

While `(:++:)` and `U1` give us the *shape* of `Bool`, the `D1` and `C1` give us metadata about `Bool`'s definition itself. `D1` describes the definition of `Bool`, including its name, module, package and whether or not the type is a `newtype`. `C1` is used to describe data constructors, with its name, fixity definition, and whether or not it contains record selectors. All of this metadata is provided via a data kind, so it exists in the type-level and can be manipulated accordingly.

### 4.3 Deriving Any Class

Armed with the knowledge of `Rep`, we can write an illustrative example of generically deriving `Eq`. The approach is threefold:

1. Define a typeclass to act as a “carrier.”
2. Provide inductive instances of the class for the generic constructors.
3. Finally, write a helper function to map between the `Rep`.

We begin by defining our typeclass. A good convention is add a `G` prefix to a generic typeclass. If you want to derive `Eq` generically, call your carrier typeclass `GEq`.

```

class GEq a where 637
  geq :: a x -> a x -> Bool 638

```

Our `GEq` class has a single method, `geq`, whose signature closely matches `(==) :: a -> a -> Bool`.

Notice that the type parameter `a` to `GEq` has kind `TYPE -> TYPE`. This is a quirk of `GHC.Generics`, and allows the same `Rep` machinery when dealing with higher-kinded classes. When writing classes for types of kind `TYPE`, we will always saturate `a` with a dummy type `x`.

A good approach when writing generic instances is to work “inside-out.” Start with the innermost constructors (`K1`, `U1` and `V1`), as these are the base cases of the structural induction.

648 In our case, `U1` is the simplest, so we will start there. Recall that `U1` repre-  
 649 sents a data constructor with no parameters, in which case it's just `()` with a  
 650 different name. Since `()` is always equal to itself, so too should `U1` be.

```
651 instance GEq U1 where
652   geq U1 U1 = True
```

653 Similarly for `V1` which corresponds to types that can't be constructed. It  
 654 might seem silly to provide an `Eq` instance for such types, but it costs us nothing.  
 655 Consider instances over `V1` as being vacuous; if you *could* give me a value of `V1`,  
 656 I claim that I could give you back a function comparing it for equality. Since  
 657 you *can't* actually construct a `V1`, then my claim can never be tested, and so we  
 658 might as well consider it true.

659 Strictly speaking, `V1` instances often aren't necessary, but we might as well  
 660 provide one if we can.

```
661 instance GEq V1 where
662   geq _ _ = True
```

663 The one other case we need to consider is what should happen for real  
 664 parameters to data constructors? Such things are denoted via `K1`, and in this  
 665 case, we want to fall back on an `Eq` (*not* `GEq`!) instance to compare the two.  
 666 The analogous non-generic behavior for this is how the `Eq` instance for `Maybe`  
 667 `a` is `Eq a => Eq (Maybe a)`; most datatypes simply want to lift equality over  
 668 their constituent fields.

```
669 instance Eq a => GEq (K1 _i a) where
670   geq (K1 a) (K1 b) = a == b
```

671 But why should we use an `Eq` constraint rather than `GEq`? Well we're using  
 672 `GEq` to help derive `Eq`, which implies `Eq` is the actual type we care about. If  
 673 we were to use a `GEq` constraint, we'd remove the ability for anyone to write a  
 674 non-generic instance of `Eq`!

675 With our base cases complete, we're ready to lift them over products and  
 676 sums. We can lift equality over products component-wise, and over sums on  
 677 whether or not they are the same data constructor.

```
678 instance (GEq a, GEq b) => GEq (a ::: b) where
679   geq (a1 ::: b1) (a2 ::: b2) = geq a1 a2 && geq b1 b2
680
681 instance (GEq a, GEq b) => GEq (a :+: b) where
682   geq (L1 a1) (L1 a2) = geq a1 a2
683   geq (R1 b1) (R1 b2) = geq b1 b2
684   geq _ _ = False
```

685 Finally, we want to lift all of our `GEq` instances through the `Rep`'s metadata  
 686 constructors, since the names of things aren't relevant for defining `Eq`. Fortu-  
 687 nately, all of the various types of metadata provided by `GHC.Generics` are all  
 688 type synonyms of `M1`:

```
instance GEq a => GEq (M1 _x _y a) where           689
    geq (M1 a1) (M1 a2) = geq a1 a2                690
```

Using `-XDefaultSignatures`, we're also capable of getting `-XDeriveAnyClass` 691  
to work for us! 692

```
class MyEq a where                                 693
    eq :: a -> a -> Bool                            694
    default eq :: (Generic a, GEq (Rep a)) => a -> a -> Bool 695
    eq a b = geq (from a) (from b)                 696

data Foo                                           697
    = F0                                           698
    | F1 String                                     699
    deriving (Generic, MyEq) . . . . . ❶          700
```

Notice how at ❶, we have *not* derived an instance of `Eq`. We can fire up the 702  
REPL to see how we did: 703

```
> eq F0 F0                                         704
True                                               705

> eq (F1 "foo") (F1 "foo")                         706
True                                               707

> eq F0 (F1 "hello")                               708
False                                              709

> eq (F1 "foo") (F1 "bar")                         710
False                                              711
```

Just as we'd expect! Awesome! 715

## 4.4 Generic Metadata 716

For reasons beyond the author's comprehension, rather than imbue JavaScript 717  
with a reasonable typesystem, its proponents have instead blessed us with JSON 718  
Schema. For those of you lucky enough to be unfamiliar with it, JSON Schema 719  
is, in its own words “a vocabulary that allows you to annotate and validate 720  
JSON documents.” It's sort of like a typesystem, but ad-hoc and not very 721  
powerful. 722

For example, the following Haskell type: 723

```
data Person = Person                               724
    { name  :: String                               725
    , age   :: Int                                  726
    , phone :: Maybe String                          727
    }                                               728
```

729 would be described in JSON Schema as:

```
730 { "title": "Person"
731   , "type": "object"
732   , "properties":
733     { "name": { "type": "string" }
734     , "age": { "type": "integer" }
735     , "phone": { "type": "string" }
736     }
737   , "required": ["name" , "age"]
738 }
```

739 I'm sorry to have had to introduce you to this abomination, but it is what  
740 it is. We can't do anything about JSON Schema, but at least it provides a  
741 motivating example of generating code generically, lest anyone need to write  
742 such a monstrosity themselves.

743 We begin as always with the definition of our typeclass. Obviously such  
744 a thing needs to produce a `Value` (aeson's representation of a JSON value).  
745 Less clear is that in order to produce the `required` property, we'll also need  
746 to propagate information about required-ness upwards. As such, we decide our  
747 `GSchema` typeclass will look like this:

```
748 class GSchema (a :: Type -> Type) where
749   gschema :: (Value, [Text])
```

750 Notice that `gschema` doesn't have any way of binding the `a` type parameter.  
751 While we *could* use a `Proxy` to drive the instance lookups the way we did for  
752 `HasPrintf`, a cleaner interface is to enable `-XAllowAmbiguousTypes` and later  
753 use `-XTypeApplications` to fill in the desired type variable. We will see this  
754 usage in a moment.

755 For our purposes, we will assume we only want to generate JSON Schema  
756 for Haskell records – anonymous products are out, and it's not clear how to  
757 represent sums anyway.

758 To begin with, we'll write a helper function to generate the property objects  
759 for us—things of the form `{"foo": {"type": "integer"}}`. We'll write this  
760 function to take its name (the "foo") from a `KnownSymbol`, because that's the  
761 form we'll receive it in from `GHC.Generics`.

use a type fam instead of  
Typeable

```
762 makePropertyObj
763   :: forall name. KnownSymbol name
764   => String --^ The JSON type.
765   -> Value
766 makePropertyObj ty = object
767   [ pack (symbolVal $ Proxy @name) .= object
768     [ "type" .= String (pack ty)
769     ]
770   ]
```

In order to get access to the record name, it's insufficient to simply define an instance of `GSchema` for `K1`, since by the time we get to `K1` we've lost access to the metadata. Instead, we can do (type-level) pattern matching on `M1 S meta (K1 _ a)`; the `S` type is used as a parameter to `M1` to describe *record selector metadata*.

If the type we're eventually wrapping is a `Maybe Int`, we'd like to say that its type in JSON Schema is an `integer`, and that it isn't required. Since `Maybe a` is a more specific type than `a`, it seems like a good place to start. We'll mark it as an `OVERLAPPING` instance:

```
instance {-# OVERLAPPING #-} (Typeable a, KnownSymbol nm)
  => GSchema (M1 S ('MetaSel ('Just nm) _1 b_c _d)
              (K1 _e (Maybe a))) where
  gschema =
    ( makePropertyObj @nm (show . typeRep $ Proxy @a)
      , mempty
    )
```

This instance will give us back our property object, some JSON of the form `{"foo": {"type": "integer"}}`, and an empty list corresponding to having not yet discovered any required fields. We leave the responsibility of combining these JSON objects to other instances of `GSchema`.

The case for `K1 a` is very similar, except that we also emit the name as a required field:

```
instance (Typeable a, KnownSymbol nm)
  => GSchema (M1 S ('MetaSel ('Just nm) _1 b_c _d)
              (K1 _e a)) where
  gschema =
    ( makePropertyObj @nm (show . typeRep $ Proxy @a)
      , pure . symbolVal $ Proxy @nm
    )
```

We can define products `(:*)` for our types as simply merging our property objects together via their underlying semigroup:

```
addObjects
  :: Semigroup a
  => (Value, a)
  -> (Value, a)
  -> (Value, a)
addObjects (Object a, x) (Object b, y)
  = (Object $ a <> b, x <> y)

instance (GSchema f, GSchema g) => GSchema (f :*: g) where
  gschema = addObjects (gschema @f) (gschema @g)
```

open type family to generate js names

814 It's unclear how to represent sum-types in JSON schema, so we won't provide  
 815 an instance of `(:+:)`. Furthermore, it doesn't seem valuable to marshal unit  
 816 and void types to JavaScript, so we won't provide instances for `U1` or `V1` either.

817 All that remains is the remaining `M1` instances; `M1 D` is used for type infor-  
 818 mation, which doesn't matter to us, and `M1 C` which gives us data constructor  
 819 information. It is from `M1 C` that we will get our object's `"title"`, and set up  
 820 the top-level information about our object:

```
821 instance GSchema a => GSchema (M1 D _1 a) where
822     gschema = gschema @a
823
824 instance (GSchema a, KnownSymbol nm)
825     => GSchema (M1 C ('MetaCons nm _1 _2) a) where
826     gschema =
827     let (sch, req) = gschema @a
828     in ( object
829         [ "title" .= (String . pack . symbolVal $ Proxy @nm)
830         , "type"  .= String "object"
831         , "properties" .= sch
832         ]
833     , req
834     )
```

835 Here we use `_1` and `_2` as “wildcard types.” These are not special in any way,  
 836 but Haskell doesn't allow us to write `_` as a type variable. Most of the time, most  
 837 the data in the `M1` constructor won't be relevant to you, and so you'll probably  
 838 end up using a lot of these wildcard types.

839 All that's left is to put a nice veneer in front of `gschema`, and to build our  
 840 `"required"` field.

```
841 schema :: forall a. (GSchema (Rep a), Generic a) => Value
842 schema =
843     let (v, reqs) = gschema @(Rep a)
844     in fst $ addObjects (v, ())
845         ( object
846           [ "required" .=
847             Array (fromList $ String . pack <$> reqs)
848           ]
849         , ()
850         )
```

851 And we're done. We've successfully used the metadata in `GHC.Generics`  
 852 to automatically marshal a description of our Haskell datatypes into JSON  
 853 Schema. We didn't need to resort to using code generation—which would have  
 854 complicated our compilation pipeline—and we've written nothing but everyday  
 855 Haskell in order to accomplish it.



## 4.5 Performance

Before diving much further into `GHC.Generics` and all of the wonderful things we can do with them, it's worthwhile to take a minute to slow down. Sure, describing all of this behavior in Haskell is great, but do we pay an overhead at runtime to get it? If so, it might not be a very good investment; writing things by hand is annoying and tedious, but at least we have some understanding of what's going on under the hood. With `GHC.Generics`, it's certainly less clear.

There is good and bad news here. The good news is that usually adding `INLINE` pragmas to each of your class' methods is enough to optimize away all usage of `GHC.Generics`. The bad news is that this is *usually* enough to optimize them away. Since there is no separate compilation step when working with `GHC.Generics`, it's quite a lot of work to actually determine whether or not your generic code is being optimized away.

Enter the `inspection-testing`[1] library. `inspection-testing` provides a plugin to GHC which allows us to make assertions about our generated code. We can use it to ensure GHC optimizes away all of our usages of `GHC.Generics`, and instead spits out the same code that we would have written by hand.

We can use `inspection-testing` like so:

1. Enable the `{-# OPTIONS_GHC -O -fplugin Test.Inspection.Plugin #-}` pragma.
2. Enable `-XTemplateHaskell`.
3. Import `Test.Inspection`.
4. Write some code that exercises the generic code path. Call it `foo`, for example.
5. Add `inspect $ hasNoGenerics 'foo` to your top level module.

And that's it. Now when you compile your module, GHC will yell at you and refuse to compile if your generic code has any runtime overhead. Unfortunately for us, `inspection-testing` isn't magic and can't guarantee our implementation is as good as a hand-written example, but at least it can prove the generic representations don't exist at runtime.

In order to prove two implementations (eg. one written generically and one written by hand) *are* equal, you can use `inspection-testing`'s `(===)` combinator, which causes a compile-time error if the actual generate code isn't identical. This is often impractical to do for complicate usages of `GHC.Generics`, but it's comforting to know that it's possible in principle.

However, there is a particularly egregious case that GHC is unable to optimize. It's described colloquially as "functions that are too polymorphic." But what does it mean to be *too polymorphic*?

This class of problems sets in when GHC requires knowing about the functor/applicative/monad laws in order to perform the inlining, but the type itself

real example here

897 is polymorphic. That is to say, a generic function that produces a `forall m`.  
 898 `m a` will perform poorly, but `Maybe a` is fine. A good rule of thumb is that if  
 899 you have a polymorphic higher-kinded type, your performance is going to go  
 900 into the toolies.

## 901 4.6 Yoneda and Codensity

902 The good news is that reclaiming our performance from the clutches of too-  
 903 polymorphic generics isn't very hard. The secret is to rewrite our types from  
 904 the form `forall f. Functor f => f a` into `forall f. Yoneda f a`, and  
 905 from the form `forall f. Monad m => m a` (and `Applicatives`) into `forall`  
 906 `m. Codensity m a`. Both come from the `kan-extensions`[2] We'll talk more  
 907 about this transformation in a moment.

908 In essence, the trick here is to write our "too polymorphic" code in a form  
 909 that amenable to GHC's inlining abilities, and then transform it back into the  
 910 desired form at the very end. `Yoneda` and `Codensity` are tools that can help  
 911 with this transformation.

912 Consider the definition of `Yoneda`:

```
913 newtype Yoneda f a = Yoneda
914   { runYoneda :: forall b. (a -> b) -> f b
915   }
```

916 When we ask GHCi about the type of `runYoneda`, we see something inter-  
 917 esting:

```
918 > :t runYoneda
919 runYoneda :: Yoneda f a -> (a -> b) -> f b
```

920 `runYoneda` looks a lot like `flip fmap :: f a -> (a -> b) -> f b`, doesn't  
 921 it? This is not an accident. The `Functor` instance for `Yoneda` is particularly  
 922 enlightening:

```
923 instance Functor (Yoneda f) where
924   fmap f (Yoneda y) = Yoneda $ f . y
```

925 Note the lack of a `Functor f` constraint on this instance! `Yoneda f` is a  
 926 `Functor` *even when f isn't*. In essence, `Yoneda f` gives us a free instance of  
 927 `Functor` for any type of kind `TYPE -> TYPE`. We call `Yoneda` the *free Functor*.  
 928 There's lots of interesting category theory behind all of this, but it's not impor-  
 929 tant to us.

930 But how does `Yoneda` work? Keeping in mind the functor law that `fmap f .`  
 931 `fmap g = fmap (f . g)`, the implementation of `Yoneda's Functor` instance  
 932 looks very similar. All `Yoneda` is doing is accumulating all of the functions we'd  
 933 like to `fmap`, so that it can perform them all at once.

day currying?

As interesting as all of this is, the question remains: how does `Yoneda` help  
GHC optimize our programs? Recall that GHC’s failure to inline “too polymor-  
phic” functions is due to it being unable to perform the functor/etc laws while  
inlining polymorphic code. But since `Yoneda f` is a functor even when `f` isn’t,  
`Yoneda`’s `Functor` instance can’t possibly depend on `f`. That means `Yoneda f`  
is never “too polymorphic,” and as a result, acts as a fantastic carrier for our  
optimization tricks.

Finally, the functions `liftYoneda :: Functor f => f a -> Yoneda f a`  
and `lowerYoneda :: Yoneda f a -> f a` witness that `Yoneda f a` is isomor-  
phic to `f a`. Whenever your generic code needs to do something in `f`, it should  
use `liftYoneda`, and the final interface to your generic code should make a call  
to `lowerYoneda` to hide it as an implementation detail.

`Codensity` is to `Monad` as `Yoneda` is to `Functor`. This means `Codensity`  
is the *free* `Monad`, and can likewise be used to optimize “too polymorphic”  
monadic code.

Both types have other interesting uses—notably, `Codensity` is invaluable  
in flattening JavaScript-esque “callback pyramids of doom”—though we will  
not dwell on them any further, as such things do not fall within the scope of  
this book. The motivated reader is encouraged to familiarize themselves with the  
`kan-extensions` package.



# Bibliography

954

[1] Joachim Breitner. <https://github.com/nomeata/inspection-testing>

955

[2] Ed Kmett. <https://github.com/ekmett/kan-extensions>

956