

Advanced Functional Programming for Fun and Profit

José Pedro Magalhães

Department of Computer Science, University of Oxford
<http://dreixel.net>

August 30, 2013
Utrecht University, The Netherlands

This is a two-part talk:

- ▶ Part I: new and (even more) advanced functional programming in GHC 7.6
- ▶ Part II: Applying advanced functional programming techniques to music modelling

Part I

Advanced AFP

More features, more fun



Exciting new features in GHC 7.6:

- ▶ Data kinds
- ▶ Kind polymorphism
- ▶ Type-level literals
- ▶ Deferred type errors

We'll go through a few examples of how to put these new features to good use.

In this talk I will use:

- ▶ Blue for constructors (most of the time)
Nothing, False, Left True, 3, "abc", 'p'

In this talk I will use:

- ▶ Blue for constructors (most of the time)
`Nothing`, `False`, `Left` `True`, `3`, `"abc"`, `'p'`
- ▶ Red for types
`Int`, `String`, `Show a`, `data () = ()`

In this talk I will use:

- ▶ Blue for constructors (most of the time)
`Nothing`, `False`, `Left True`, `3`, `"abc"`, `'p'`
- ▶ Red for types
`Int`, `String`, `Show a`, `data () = ()`
- ▶ Green for kinds
`*`, `* → *`

What are kinds?



Just like types classify values...

3 :: Num a ⇒ a

'p' :: Char

Just () :: Maybe ()

"abc" :: String

What are kinds?



Just like types classify values...

3 :: Num a \Rightarrow a

'p' :: Char

Just () :: Maybe ()

"abc" :: String

... kinds classify types:

Int :: *

Char :: *

Maybe :: * \rightarrow *

[] :: * \rightarrow *

The language of kinds



However, the language of kinds, unlike that of types, is rather limited:

$$k ::= \star$$
$$| k \rightarrow k$$

In particular: no user defined kinds, no kind variables.

Diversion: the **Constraint** kind



With `-XConstraintKinds` we get one new base kind to classify constraints:

Show :: $\star \rightarrow$ **Constraint**

Functor :: $(\star \rightarrow \star) \rightarrow$ **Constraint**

Num Int :: **Constraint**

Int \sim Bool :: **Constraint**

Why do we need a better kind system? I



We often want to restrict type arguments to a particular kind:

```
data Ze
data Su n
data Vec :: * → * → * where
  Nil  :: Vec Ze a
  Cons :: a → Vec n a → Vec (Su n) a
```

Types like `Vec Int Int`, `Vec Int Bool`, and `Vec () ()` are valid (albeit uninhabited). We want to say that the first argument of `Vec` should only be `Ze` or `Su`!

Why do we need a better kind system? II



Lack of kind polymorphism leads to code duplication:

```
class Typeable (a ::  $\star$ )           where
  typeOf :: a     $\rightarrow$  TypeRep
class Typeable1 (a ::  $\star \rightarrow \star$ )  where
  typeOf1 :: a b  $\rightarrow$  TypeRep
class Typeable2 (a ::  $\star \rightarrow \star \rightarrow \star$ ) where
  typeOf2 :: a b c  $\rightarrow$  TypeRep
```

We would rather have a single, kind-polymorphic `Typeable` class!

With `-XDataKinds`, the following code is valid:

```
data Nat = Ze | Su Nat
data Vec :: Nat → * → * where
  Nil  :: Vec Ze a
  Cons :: a → Vec n a → Vec (Su n) a
```

Note the implicit promotion of the constructors `Ze` and `Su` to types `Ze` and `Su`, and of the type `Nat` to the kind `Nat`.
Types like `Vec Int Int` now trigger a kind error!

Datatype promotion II



Type families can also be indexed over promoted types:

```
type family Add (m :: Nat) (n :: Nat) :: Nat
```

```
type instance Add Ze    n = n
```

```
type instance Add (Su m) n = Su (Add m n)
```

Datatype promotion II



Type families can also be indexed over promoted types:

```
type family Add (m :: Nat) (n :: Nat) :: Nat
```

```
type instance Add Ze      n = n
```

```
type instance Add (Su m) n = Su (Add m n)
```

```
append :: Vec m a → Vec n a → Vec (Add m n) a
```

```
append Nil          v = v
```

```
append (Cons h t) v = Cons h (append t v)
```

This was all possible before, but now we can express the right kind of `Add`.

Promoted lists and tuples



Haskell lists are natively promoted, so we can encode heterogeneous lists as follows:

```
data HList :: [*] → * where  
  HNil   :: HList []  
  HCons  :: a → HList t → HList (a : t)
```

As an example, here is a heterogeneous collection:

```
hetList :: HList [Int, Bool]  
hetList = HCons 3 (HCons False HNil)
```

Tuples are also promoted, e.g. $(*, * \rightarrow *, \text{Constraint})$.

Kind polymorphism reduces code duplication:

```
data EqT a b where
```

```
  Refl :: EqT a a
```

Previously the kind of `EqT` would default to `* → * → *`. With `-XPolyKinds` it doesn't, so the following types are all valid: `EqT a Int`, `EqT f Maybe`, `EqT t Either`.

Kind-polymorphic `Typeable`



Now we can define a single kind-polymorphic `Typeable` class:

```
data Proxy (t :: k) = Proxy  
class Typeable (t :: k) where  
  typeRep :: Proxy t → TypeRep
```

Note that `Proxy` is kind polymorphic!

Kind-polymorphic `Typeable`



Now we can define a single kind-polymorphic `Typeable` class:

```
data Proxy (t :: k) = Proxy  
class Typeable (t :: k) where  
  typeRep :: Proxy t → TypeRep
```

Note that `Proxy` is kind polymorphic!

We can now give `Typeable` instances for types of various kinds:

```
instance Typeable Char where ...  
instance Typeable [] where ...  
instance Typeable Either where ...
```

Thanks to Iavor Diatchki's work, we now have efficient type-level naturals:

$0, 1, 2, \dots :: \text{Nat}$

Note the colours!

These type-level naturals come with associated operations:

$(\leq) :: \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Constraint}$

$(+)$ $:: \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat}$

$(*)$ $:: \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat}$

$(^)$ $:: \text{Nat} \rightarrow \text{Nat} \rightarrow \text{Nat}$

How do we manipulate values representing type-level naturals? There is a family of singleton types, parameterised by literals:

```
newtype Sing :: a → ★
```

How do we manipulate values representing type-level naturals? There is a family of singleton types, parameterised by literals:

```
newtype Sing :: a → ★
```

From types to values:

```
fromSing :: Sing a → SingRep a
```

```
type family SingRep a
```

```
type instance SingRep (a :: Nat)    = Integer
```

```
type instance SingRep (a :: Symbol) = String
```

Note that we can have type-level literals other than naturals, and `SingRep` is a *kind-indexed* family!

Revisiting vectors, now with type-level naturals:

```
data Vec :: Nat → * → * where  
  Nil  :: Vec 0 a  
  Cons :: a → Vec n a → Vec (n + 1) a
```

Vector concatenation uses type-level natural number addition:

```
append :: Vec m a → Vec n a → Vec (m + n) a  
append Nil      ys = ys  
append (Cons x xs) ys = Cons x (append xs ys)
```


Why are type-level naturals hard to implement?



```
append :: Vec m a → Vec n a → Vec (m + n) a
append Nil          ys = ys
append (Cons x xs) ys = Cons x (append xs ys)
```

Function `append` requires GHC to prove equalities between natural number expressions:

- ▶ Could not deduce $(n \sim (0 + n))$ from the context $(m \sim 0)$ bound by a pattern with constructor `Nil` :: $\forall a. \text{Vec } 0 \ a$
- ▶ Could not deduce $((m + n) \sim ((n' + n) + 1))$ from the context $(m \sim (n' + 1))$ bound by a pattern with constructor `Cons` :: $\forall a (n :: \text{Nat}). a \rightarrow \text{Vec } n \ a \rightarrow \text{Vec } (n + 1) \ a$

We need an equation solver!

The illogical next step



What is the next thing that you want, when you have data kinds, polymorphic kinds, and type-level literals?

The illogical next step



What is the next thing that you want, when you have data kinds, polymorphic kinds, and type-level literals?

Naturally, to turn off type checking! :-)

Why would you want to do that?



For instance:

- ▶ Prototyping
- ▶ Large refactoring
- ▶ IDE

Example 1



With the flag `-fdefer-type-errors`, this example:

```
p, q :: Int
p = 1
q = '1'
main = print p
```

Compiles with warning: “couldn't match expected type `Int` with actual type `Char` in an equation for `q`: `q = '1'`”. Runs and returns 1!

Example II



```
p, q :: Int
p = 1
q = '1'
main = print q
```

Fails at runtime with: “couldn't match expected type `Int` with actual type `Char` in an equation for `q`: `q = '1'`”.

Example III



```
t1 :: Int
t1 = '1'

t2 :: a → String
t2 = show

data T a where
  T1 :: Int → T Int
  T2 :: a → T a

t3 :: T a
t3 = T1 0

main = print 1
```

Runs fine!

GHC's core language uses *coercions* to (safely) cast terms:

data $T\ a = T_1\ (a \sim Int)\ Int \mid T_2\ a$

$unT :: T\ a \rightarrow a$

$unT\ (T_1\ c\ n) = n \triangleright (\text{sym } c)$

$unT\ (T_2\ x) = x$

$\triangleright :: b \rightarrow (b \sim a) \rightarrow a$

Evidence, or values of type (\sim) , is automatically generated by GHC during type checking. Deferring type errors simply means generating runtime errors as evidence!

It's not dynamic typing!



Note that deferring type errors doesn't mean any form of checks are performed at runtime. Consider this example:

```
f ::  $\forall a. a \rightarrow a \rightarrow a$   
f x y = x  $\wedge$  y  
main = print (f True False)
```

It still fails at runtime!

A better kind system gives us:

- ▶ Increase type safety
- ▶ Increase expressivity
- ▶ Reduce code duplication
- ▶ Allow for writing clearer code

A better kind system gives us:

- ▶ Increase type safety
- ▶ Increase expressivity
- ▶ Reduce code duplication
- ▶ Allow for writing clearer code

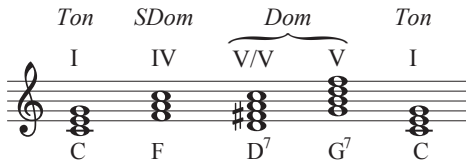
And if we get tired of it we can always defer errors to runtime!

Part II

Haskell, music, fun... and maybe even profit!

- ▶ Modelling musical harmony using Haskell
- ▶ Avoiding programming
- ▶ Applications of a model of harmony:
 - ▶ Musical analysis
 - ▶ Finding cover songs
 - ▶ Generating chords for melodies
 - ▶ Correcting errors in chord extraction from audio sources

What is harmony?



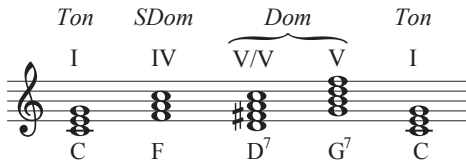
Ton *S**Dom* *Dom* *Ton*

I IV V/V V I

C F D⁷ G⁷ C

- ▶ Harmony arises when at least two notes sound at the same time
- ▶ Harmony induces tension and release patterns, that can be described by music theory and music cognition
- ▶ The internal structure of the chord has a large influence on the consonance or dissonance of a chord
- ▶ The surrounding context also has a large influence

What is harmony?



Ton *SDom* *Dom* *Ton*

I IV V/V V I

C F D⁷ G⁷ C


- ▶ Harmony arises when at least two notes sound at the same time
- ▶ Harmony induces tension and release patterns, that can be described by music theory and music cognition
- ▶ The internal structure of the chord has a large influence on the consonance or dissonance of a chord
- ▶ The surrounding context also has a large influence

Demo: how harmony affects melody

An example harmonic analysis


Ton *SDom* *Dom* *Ton*

I IV V/V V I



C F D⁷ G⁷ C

Detailed description: A musical staff in treble clef showing five chords. Above the staff are functional labels: 'Ton' above 'I', 'SDom' above 'IV', 'Dom' above 'V/V' and 'V', and 'Ton' above 'I'. The chords are represented by vertical lines with dots for notes. Below the staff are the chord names: C, F, D⁷, G⁷, and C.



Why are harmony models useful?



Having a model for musical harmony allows us to automatically determine the functional meaning of chords in the tonal context. The model determines which chords “fit” on a particular moment in a song.

Why are harmony models useful?



Having a model for musical harmony allows us to automatically determine the functional meaning of chords in the tonal context. The model determines which chords “fit” on a particular moment in a song. This is useful for:

- ▶ Musical information retrieval (find songs similar to a given song)
- ▶ Audio and score recognition (improving recognition by knowing which chords are more likely to appear)
- ▶ Music generation (create sequences of chords that conform to the model)

Representing harmony as a datatype, naively



```
data Piece    = Piece [Phrase]
data Phrase  = PT    Ton    | PD    Dom
data Ton     = TIMaj DegNum
data Dom     = DVMaj DegNum | DSDom SDom Dom
data SDom    = SIVMaj DegNum
data DegNum  = I | II | III ...
```

Representing harmony as a datatype, naively



```
data Piece    = Piece [Phrase]
data Phrase  = PT    Ton    | PD    Dom
data Ton     = TIMaj DegNum
data Dom     = DVMaj DegNum | DSDom SDom Dom
data SDom    = SIVMaj DegNum
data DegNum  = I | II | III ...
```

Problem: the term `Piece [PT (TIMaj II)]` typechecks, but represents an invalid harmonic structure. `TIMaj` should only take `I` as argument. How can we enforce this? (And why do we want it?)

Representing harmony as a datatype I



```
data Piece = Piece [Phrase]
data Phrase = PT Ton | PD Dom
data Ton = TIMaj (Deg I)
data Dom = DVMaj (Deg V) | DSDom SDom Dom
data SDom = SIVMaj (Deg IV)

data Deg  $\delta$  = Deg DegNum
```

We do not export the `Deg` constructor. So how do we build `Degs`?

Representing harmony as a datatype II



```
deg :: ToDegree  $\delta$   $\Rightarrow$   $\delta$   $\rightarrow$  Deg  $\delta$   
deg d = Deg (toDegree d)
```

We need degrees at the type level:

```
data I; data II; ... data VII;
```

And type-to-value conversions for degrees:

```
class ToDegree  $\delta$  where toDegree ::  $\delta$   $\rightarrow$  DegNum  
instance ToDegree I where toDegree _ = I  
instance ToDegree II where toDegree _ = II  
... -- instances for all degrees
```

Representing harmony as a datatype III



Now the term `Piece [PT (TIMaj (deg (\perp :: II)))]` does not typecheck, as we wanted:

```
*ghci> Piece [PT (TIMaj (deg (undefined :: II)))]  
    Couldn't match expected type 'I'  
    with actual type 'II'
```

Representing harmony as a datatype IV



We have a well-typed harmony model in Haskell that accepts only “correct” harmonic sequences.

Here we have only seen a very basic subset of our model. In reality, it consists of 14 datatypes with a total of 46 constructors.

Also, while developing the model we have to change it often, trying to find a balance between complexity and expressiveness.

Parsing chord sequences I



We now consider the problem of parsing (textual) chord sequences into our datatype representing musical harmony:

```
class Parse  $\alpha$  where  
  parse :: ParserMusic  $\alpha$ 
```

Most instances are trivial:

instance Parse Phrase where

parse = PT <\$> parse

<|> PD <\$> parse

instance Parse Ton where

parse = T_{IMaj} <\$> parse

instance Parse Dom where

parse = D_{VMaj} <\$> parse

<|> D_{SDom} <\$> parse <*> parse

...

Most instances are trivial:

instance Parse Phrase **where**

parse = PT <\$> parse

<|> PD <\$> parse

instance Parse Ton **where**

parse = T_{IMaj} <\$> parse

instance Parse Dom **where**

parse = D_{VMaj} <\$> parse

<|> D_{SDom} <\$> parse <*> parse

...

So trivial that we do not write them; we use a *generic parser*.

For degrees we need an adhoc parser:

```
instance (ToDegree  $\delta$ )  $\Rightarrow$  Parse (Deg  $\delta$ ) where  
  parse = pChord (toDegree ( $\perp$  ::  $\delta$ ))
```

```
pChord :: DegNum  $\rightarrow$  ParserMusic (Deg  $\delta$ )  
pChord = ...
```

Again we use the type-to-value conversion class `ToDegree`.

Increasing complexity



The model we presented so far is very simple and cannot account for:

Repeated chords $\underline{G_{Maj}} \underline{G_{Maj}} C_{Maj}$

Diatonic secondary dominants $\underline{E_{min}} \underline{A_{min}} \underline{D_{min}} G^7 C_{Maj}$

Chromatic secondary dominants $\underline{A^7} \underline{D^7} G^7 C_{Maj}$

Minor key dominant borrowing $\underline{G_{min}} C_{Maj}$

Tritone substitutions $\underline{A^b7} G^7 C_{Maj}$

...

Handling repeated chords



Easy: adapt the `Deg` type and `pChord`:

```
data Deg  $\delta$  = Deg { chordRole :: DegNum  
                    , numReps  :: Int }
```

```
pChord :: DegNum  $\rightarrow$  ParserMusic (Deg  $\delta$ )  
pChord = ... -- adapted accordingly
```

Secondary dominants—datatype representation



Here the fun really begins! We adapt the `Deg` type again:

```
type DegSD  $\delta$  = BDegSD  $\delta$ 
```

```
data BDegSD  $\delta$  where
```

```
  Cons :: BDegSD (PerfV  $\delta$ )  $\rightarrow$  Deg  $\delta$   $\rightarrow$  BDegSD  $\delta$ 
```

```
  Base :: Deg  $\delta$   $\rightarrow$  BDegSD  $\delta$ 
```

```
data Deg  $\delta$  = Deg ... -- as before
```

```
type family PerfV  $\delta$  -- computes the perfect fifth
```

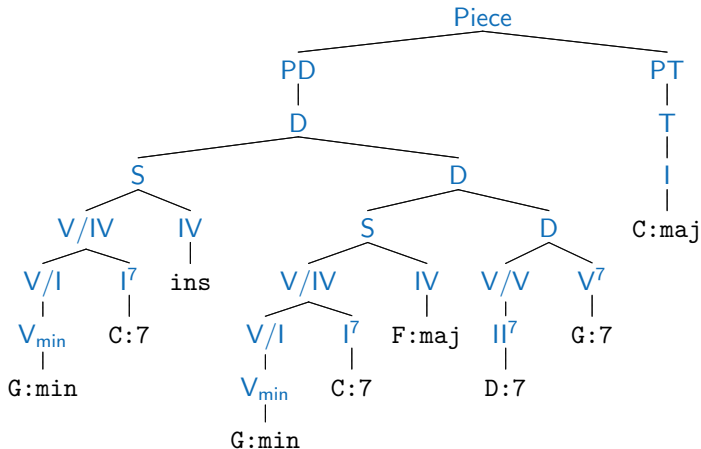
```
type instance PerfV I = V
```

```
type instance PerfV V = II
```

```
...
```

Application: harmony analysis

Parsing the sequence G_{\min} C^7 G_{\min} C^7 F_{Maj} D^7 G^7 C_{Maj} :



- ▶ A practical application of a harmony model is to estimate harmonic similarity between songs
- ▶ The more similar the trees, the more similar the harmony
- ▶ We don't want to write a diff algorithm for our complicated model; we get it automatically by using a *generic diff*
- ▶ The generic diff is a type-safe tree-diff algorithm, part of a student's MSc work at Utrecht University
- ▶ Generic, thus working for any model, and independent of changes to the model

Application: automatic harmonisation of melodies



Another practical application of a harmony model is to help selecting good harmonisations (chord sequences) for a given melody:

The image displays a musical score for a single system. The top staff is a treble clef with a common time signature (C). The melody consists of the following notes: G4 (quarter), A4 (quarter), G4 (quarter), F4 (quarter), E4 (quarter), D4 (quarter), C4 (half). The bottom staff is a bass clef with a common time signature (C). The accompaniment consists of the following chords: G2-A2-B2 (quarter), G2-A2-B2 (quarter), G2-B2 (quarter), G2-A2-B2 (quarter), G2-A2-B2 (quarter), G2-A2-B2 (quarter), G2-A2-B2 (quarter), G2-A2-B2 (quarter), G2 (half). Below the bass staff, Roman numeral chord symbols are placed: V, III, I, III, II, IV, III, IV, V.

We generate candidate chord sequences, parse them with the harmony model, and select the one with the least errors.

Application: chord recognition



Yet another practical application of a harmony model is to improve chord recognition from audio sources.

Chord candidates	0.92 C	0.96 Em	
	0.94 Gm	0.97 C	
	1.00 C	1.00 G	1.00 Em
Beat number	1	2	3

How to pick the right chord from the chord candidate list? Ask the harmony model which one fits best.

Demo: <http://chordify.net>

Musical modelling with Haskell:

- ▶ A model for musical harmony as a GADT
- ▶ The datatypes can change, the code does not have to:
 - ▶ Generic parser
 - ▶ Generic pretty-printer
 - ▶ Generic diff
- ▶ Multiple models, lots of code reuse
- ▶ When chords do not fit the model: error correction
- ▶ Harmonising melodies
- ▶ Recognising harmony from audio sources

Play with it!



`http://hackage.haskell.org/package/HarmTrace`

chordify[®]

`http://chordify.net`