

Functional Modelling of Musical Harmony

and its applications

José Pedro Magalhães

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

October 22, 2012
ICT.OPEN, Rotterdam, The Netherlands

This talk



- ▶ Modelling musical harmony using Haskell

This talk



- ▶ Modelling musical harmony using Haskell
- ▶ Avoiding programming

This talk



- ▶ Modelling musical harmony using Haskell
- ▶ Avoiding programming
- ▶ Applications of the harmony model:

This talk



- ▶ Modelling musical harmony using Haskell
- ▶ Avoiding programming
- ▶ Applications of the harmony model:
 - ▶ Musical analysis

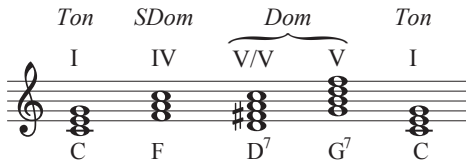
This talk



- ▶ Modelling musical harmony using Haskell
- ▶ Avoiding programming
- ▶ Applications of the harmony model:
 - ▶ Musical analysis
 - ▶ Finding cover songs

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

What is harmony?



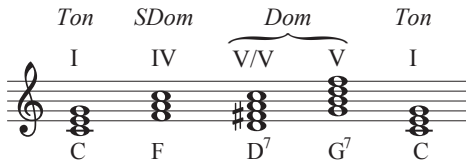
Ton *SDom* *Dom* *Ton*

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

What is harmony?



The diagram illustrates a sequence of chords on a treble clef staff. Above the staff, functional labels are placed: *Ton* above C, *SDom* above F, *Dom* above D⁷ and G⁷ (grouped by a bracket), and *Ton* above the final C. Below the staff, Roman numerals are placed: I above C, IV above F, V/V above D⁷, V above G⁷, and I above the final C. The chords themselves are represented by groups of notes on the staff: C (C4, E4, G4), F (C4, F4, A4), D⁷ (D4, F4, A4, C5), G⁷ (B3, D4, F4, G4), and C (C4, E4, G4).

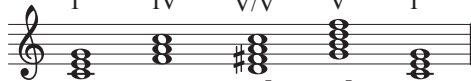
- ▶ Harmony arises when at least two notes sound at the same time
- ▶ Harmony induces tension and release patterns, that can be described by music theory and music cognition
- ▶ The internal structure of the chord has a large influence on the consonance or dissonance of a chord
- ▶ The surrounding context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

Demo: how harmony affects melody

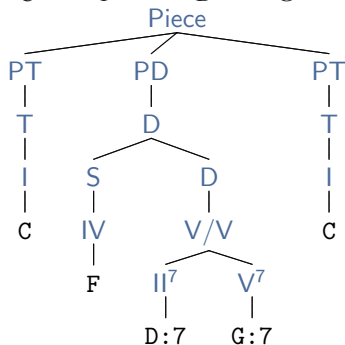
An example harmonic analysis

Ton *SDom* *Dom* *Ton*

I IV V/V V I



C F D⁷ G⁷ C



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)

Why Haskell?



Haskell is a strongly-typed pure functional programming language:

Strongly-typed All values are classified by their type, and types are known at compile time (statically). This gives us strong guarantees about our code, avoiding many common mistakes.

Pure There are no side-effects, so Haskell functions are like mathematical functions.

Functional A Haskell program is an expression, not a sequence of statements. Functions are first class citizens, and explicit state is avoided.

Representing harmony as a datatype, naively



data Piece = Piece [Phrase]

data Phrase = PT Ton | PD Dom

data Ton = T_{IMaj} DegNum

data Dom = D_{VMaj} DegNum | D_{SDom} SDom Dom

data SDom = S_{IVMaj} 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
```

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.

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

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

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

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 begins! We adapt the `Deg` type again:

```
type DegSD δ = BDegSD δ
```

```
data BDegSD δ where
```

```
  Cons :: BDegSD (PerfV δ) → Deg δ → BDegSD δ
```

```
  Base :: Deg δ → BDegSD δ
```

```
data Deg δ = Deg ... -- as before
```

```
type family PerfV δ -- computes the perfect fifth
```

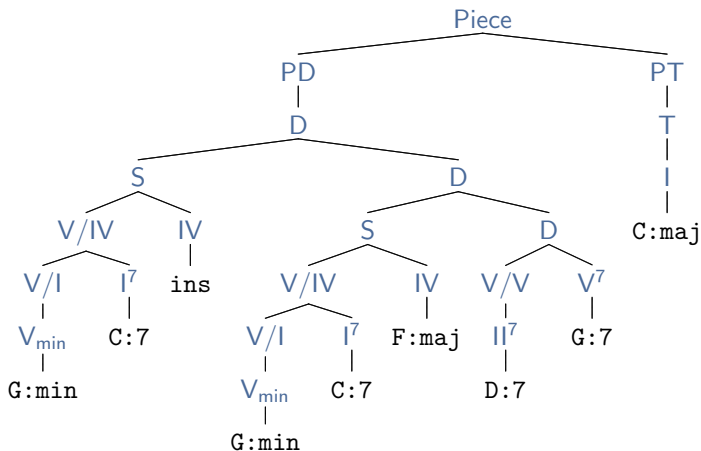
```
type instance PerfV I = V
```

```
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: chord recognition



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.

Application: chord recognition



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>

What we have so far:

- ▶ A model for musical harmony as a Haskell 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
- ▶ Recognising harmony from audio sources

Play with it!



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

`http://chordify.net`