



Music Informatics

Alan Smaill

Jan 15 2018

Organisation

- ▶ Lecture slots as timetabled.
- ▶ Standard exam at end of semester.
- ▶ Coursework:
 - ▶ There will also be one formative (unassessed) exercise due half way through the semester (issued 29th January, due 16th February).
 - ▶ There will be 1 assessed coursework exercise; it will be issued February 26th, due March 23rd.

- ▶ Course web page is linked from course descriptor; it is at
http://www.inf.ed.ac.uk/teaching/courses/mi/
- ▶ Slides, and other notes and links will be placed here.
- ▶ Also coursework when it is issued.
- ▶ There will be a Piazza site for the course.

We are interested in the interaction between Informatics theories and techniques and musical theory and practice.

This is not a course on music technology, signal processing, acoustics, etc – there are good courses on these elsewhere in the university!

- ▶ Background
 - ▶ Physics of music vs musical perception,
 - ▶ Music Representation
- ▶ Basic parameters
- ▶ Basic Local Analysis Algorithms
 - ▶ Beat tracking; score following
 - ▶ Tonal centre (Longuet-Higgins, Bolzano)

- ▶ Musical Grammars
- ▶ Musical Similarity
- ▶ Information Theory, Statistical methods
 - ▶ Characterising musical style
 - ▶ Music generated by statistical constraints (Xenakis)
- ▶ Machine Composition in a Given Style
- ▶ Musical Agents and Interaction

- ▶ Notions of music
- ▶ Physics of Music versus Musical Perception
- ▶ Music notations, for human-directed and machine-directed: what is a good representation for what purpose?

Here we take music as primarily **an experience of sound in time**; but this is not enough to distinguish music from sound in general. There is no general agreement as to what differentiates music from sound. A start to an answer is for example given by Scruton:

Music is an art of sound. . . . Nor is it the work of a musician to write poetry, even though poetry too is an art of sound. So what distinguishes the sound of music? The simple answer is 'organization'. But is is no answer at all if we cannot say what kind of organization we have in mind. . . .

Scruton, Aesthetics of Music, OUP, 1997, p 16

We'll look at some of the ways different sorts of music are organised in this course.

We will be interested in machine analysis and generation of such musical structures.

Other views on the nature of music:

- ▶ music is **organised sound(s)**
compare Varèse's piece 1957 piece: 'Poème électronique pour "sons organisés"' (electronic poem for organised sounds)
- ▶ Is music a language (or languages) like natural language?
There are a lot of features in common between natural language and music, and Informatics techniques developed for natural language tasks can help us in dealing with musical processing.

Music is experienced primarily as **sound**:
thus physics gives us ways to describe the sound waves
corresponding to particular musical notes, with given pitch, volume
and timbre.

It also gives us an explanation in terms of physics of how music is
generated from the human voice; from string instruments; from
wind instruments; from percussion . . .

We can for example look at the sound of a clarinet as the
superposition of sin waves with appropriate multipliers; we can

- ▶ recognise the sound waves as having this pattern, and
- ▶ generate such a sound wave artificially, to sound like a clarinet
tone.

It's tempting to think that notions of **pitch** in music correspond exactly to **frequency** in physics.

Yet there are many examples of situations where human perception of musical phenomena ends up giving the listener a different answer compared to the readings on the physicist's instrumentation. There are several good web sites that let discuss and show some of the musical “illusions”:

From Diana Deutsch:

http://philomel.com/musical_illusions/

Tritone paradox:

description: [http:](http://deutsch.ucsd.edu/psychology/pages.php?i=206)

[//deutsch.ucsd.edu/psychology/pages.php?i=206](http://deutsch.ucsd.edu/psychology/pages.php?i=206)

examples: http://philomel.com/mp3/musical_illusions/Tritone_paradox.mp3

— also see wikipedia on Shepard tones.

This course will not look at these psychoacoustic issues, but they are important. For example, the compression of sound files from CD format to mp3 uses in part standard techniques on data compression, but also known features of human sound perception.

Thus some information on a straight sound recording can be “dropped” without this affecting the perception of the average listener.

If you're interested, take a look at:

<http://www.mp3-converter.com/mp3codec/waveforms.htm>

We're interested in the different ways in which music can be represented, so as to be shared, remembered and transformed, both between musicians and involving computers (and indeed virtual musicians).



The representations have different strengths and weaknesses, and we will try to analyse why this is so.

A lot of music is passed from musician to musician without being written down, or recorded except in the mind of the musicians concerned (though this is less common now in western industrialised society).



When describing instrumental music, where an instrument is not to hand, or where it is important to be clear on details, a specialised spoken vocabulary is sometimes developed.

For example, the rhythmically complex music of Indian tabla has a vocabulary of the “bols” involved:

<http://kksongs.org/talamala.html>

“WTM” is used for the family of music used classically in the west from 17th century until early 20th century, and used for most popular music today. The associated notation developed over centuries.

It is a much harder task to do optical music transcription than it is to deal with optical character recognition in English; why is this?

We can note the two-dimensionality of the syntax (because notes can be simultaneous, compared to the linear succession of words/syllables in spoken language).

Example



- ▶ x and y axis correspond to time and pitch
- ▶ non-local information (\sharp indicated at start, applies to several notes)
- ▶ relies on a lot of implicit knowledge for phrasing, tempo, dynamics ...
- ▶ written as completely regular succession of notes in time, but such a “mechanical” performance is thought of as unmusical.

Tablature uses a more procedural representation:
instead of notating what **sounds** are expected, notate what **actions**
the musician should perform (where to place fingers in strings).

The example that follows aligns this notation with standard
WTM notation.

F⁹ G^{b7} F⁹ C^{M7} F⁹

T T T T
 A A A A
 B B B B
 0 0 6 5 5 7 7 6 5 6 7 0 0 0 0 7 0 0 7 0 7

5 G^{b13} D^{oM7} F⁹ A^{M7} D^{7#9}

T T T T
 A A A A
 B B B B
 6 5 0 6 7 5 6 7 0 0 0 7 6 5 6 4

9 G^{M7} C⁹ F¹³ D^{7#9} G^{M7} C⁹

T T T T
 A A A A
 B B B B
 3 5 6 2 3 3 2 4 4 4 3 4 0

MIDI is a much-used format for interchange of music. Like the last example, it is procedural in that it is originated as actions performed on a piano-like keyboard:

keys are depressed at a given time (with a certain force), and released at a given time.

Typically the key is associated with some pitch.

A file will also have information on other aspects of the intended music (eg timbre).

There are many tools for editing, playing, communicating, sequencing, as well as MIDI capture available for many traditional instruments (relatively easy for wind instruments, percussion, . . .).

From the start, real-time message passing is envisaged for each event (note down, note release, etc). See crash course:

<http://www.skytopia.com/project/articles/midi.html>

File

Broker TCP/IP address: 127.0.0.1

1883

Connect

Disconnect

Published Messages

```
Midi/Short/NoteOn 144 67 64 0
Midi/Short/NoteOff 128 67 0 8
Midi/Short/NoteOn 144 67 64 8
Midi/Short/NoteOff 128 67 0 16
Midi/Short/NoteOn 144 69 64 16
Midi/Short/NoteOff 128 69 0 24
Midi/Short/NoteOn 144 71 64 24
Midi/Short/NoteOff 128 71 0 31
Midi/Short/NoteOn 144 67 64 31
Midi/Short/NoteOff 128 67 0 39
Midi/Short/NoteOn 144 71 64 39
Midi/Short/NoteOff 128 71 0 47
Midi/Short/NoteOn 144 69 64 47
Midi/Short/NoteOff 128 69 0 63
Midi/Short/NoteOn 144 67 64 63
Midi/Short/NoteOff 128 67 0 70
Midi/Short/NoteOn 144 67 64 70
Midi/Short/NoteOff 128 67 0 78
Midi/Short/NoteOn 144 69 64 78
Midi/Short/NoteOff 128 69 0 86
Midi/Short/NoteOn 144 71 64 86
```

Playing yankee.q.mid

Select

Play

Sound files use a common digital representation of physical sound (variation of pressure in a physical medium – air, helium, water, ...), close to the usual CD format.

- ▶ CD version has typically two channels of 44,100 samples per second, 16 bits per sample;
- ▶ good audio quality, expensive in space use;
- ▶ very hard to manipulate directly, eg to bring up level of a musical voice over backing instruments.

MP3 is audio part an MPEG standard (MPEG-1 Audio layer 3). It is used for a simplified representation from the wav format, in part by omitting information that is perceptually irrelevant. mp3 file is typically 10% of the size of corresponding wav file. There is a good account at the hydrogen-audio wiki:

```
http://wiki.hydrogenaudio.org/index.php?  
title=MP3
```

Successor formats such as mp4 give more flexibility, supporting eg streaming over the internet, multiplexing of multiple video and audio streams in one file,

Mark up languages (eg html) for text are important to give an abstract structure to text for the purposes of supporting different display formats, and indications of syntactic and semantic roles within the text.

XML (eXtensible Markup Language) gives a tighter markup on documents; it is intended to support internet communication, and there are many tools and editors developed in this context.

Wikipedia's list of XML-based markup languages has about 150 entries. There are several proposals for music versions of this. Note that this gives a representation of (a conventional WTM) score, rather than a more direct representation of musical structure. The most widely used seems to be MusicXML; info at <http://www.recordare.com/xml.html>. Below is a snippet from a marked-up piano and voice piece.

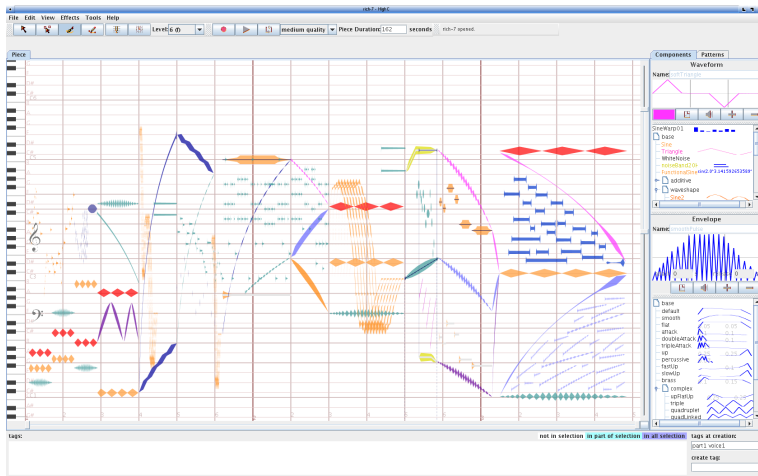
```
<measure number="4" width="546">  
  <note default-x="19">  
    <pitch>  
      <step>B</step>  
      <alter>-1</alter>  
      <octave>4</octave>  
    </pitch>  
    <duration>48</duration>  
    <voice>1</voice>  
    <type>quarter</type>  
    <stem default-y="-55.5">down</stem>  
    <lyric default-y="-82" number="1">  
      <syllabic>end</syllabic>  
      <text>a!</text>  
    </lyric>
```


The composer and architect Iannis Xenakis developed in late 1970s a graphical composition tool, which is accessible to non-experts, that allows music to be put together, both at the level of creating a sound palette, and distributing sounds in time and pitch from this palette by drawing on a screen. This was influenced by his experience of software tools for architectural design, which have been widely used for a long time.

See <http://membres.multimania.fr/musicand/INSTRUMENT/DIGITAL/UPIC/UPIC.htm>

A modern re-engineering of this is now available, called HighC:

See <http://highc.org/>



It is worth considering where recent languages like Pure Data fit in the landscape we are looking at:

`http://puredata.info/`

- ▶ It is a (visual) programming language, aimed at manipulating rather than representing music.
- ▶ It builds on many of the features mentioned above, in ways that are intended to facilitate many sorts of musical interactions.
- ▶ Like UPIC and highC, there is an emphasis on **visual** correlations to heard music.

There are many other representations out there, both human oriented and machine oriented.

Think of: ancient Chinese notation, music box cylinders, piano rolls, medieval western notation, magnetic tape, vinyl disc, IRCAM soundfile format, . . .

- ▶ Physical and perceptual accounts of music
- ▶ Different representations support different manipulations and different goals.