

Balázs Horváth, Andrea Szigetvári

TIMBRE SOLFEGE

Summary: In the context of a traditional musical education, solfège training offers the possibility of developing the ability to hear pitch and rhythm analytically. This syllabus extends the notion of “ear training” to those qualities of sound that cannot be quantified in terms of pitch and rhythm. A violin, a trumpet, an electric guitar, can produce the same note at the same volume, but how is it possible that we can instantly identify which instrument produces the sound? It is because of timbre. This syllabus presents a developmental series of exercises enabling you to develop your analytical ability to understand timbre. In the 20th century – with the advent of sound recording and electroacoustic music – classifying, recognizing and remembering timbre became one of the most important areas of musical research. It is essential for sound engineers, sound designers, multimedia creators and composers to develop a highly sensitive ear not only for pitch and rhythm but also for timbre.

The exercises are built on practical interactive programmes called „Sonic Learning Applications” (SLApps). These are downloadable standalone programmes (built in MAX/MSP) that allow the student a direct learning experience and a possibility to test and assess his/her newly gained listening skills. The 15 chapters address different timbre dimensions. Each chapter includes and begins with the theoretical background. In the second part practical exercises describe how the individual SLApps works. At the end of the chapters listening strategies are proposed to help to develop analytical hearing of timbres.

Keywords: timbre, spectrum, harmonic, inharmonic, noise, sinewave, partial, loudness, pitch, amplitude envelope, transition, SLApp, interactive exercise



BUDAPESTI KOMMUNIKÁCIÓS
ÉS ÜZLETI FŐISKOLA



SZÉCHENYI TERV

Typotex Kiadó

2013

Copyright: © 2013-2018, Balázs Horváth, Andrea Szigetvári, University of Applied Sciences, Budapest

Creative Commons NonCommercial-NoDerivs 3.0 (CC BY-NC-ND 3.0)

This work can be reproduced, circulated, published and performed for non-commercial purposes without restriction by indicating the author's name, but it cannot be modified.

Reader: Stewart Collinson

ISBN 978 963 279 206 4

Prepared under the editorship [Typotex Kiadó](#)

Responsible manager: Votisky Zsuzsa

Made within the framework of the project Nr. TÁMOP-4.1.2/A/1-11/1-2011-0010 entitled „Digitális és kollaboratív művészet”.

Nemzeti Fejlesztési Ügynökség
www.ujszachenyiterv.gov.hu
06 40 638 638



A projekt az Európai Unió támogatásával, az Európai Szociális Alap társfinanszírozásával valósul meg.

Table of contents

1. Superposition of sinewaves I.	5
1.1. Theoretical background	5
1.1.1. Harmonic spectrum	5
1.1.2. Fusion	6
1.2. Practical Exercises	7
1.2.1. How SLApp01 works	7
1.2.2. Using SLApp01	9
1.2.3. Listening strategies	15
2. Superposition of sinewaves II. – Tri-stimulus.	17
2.1. Theoretical background	17
2.2. Practical Exercises	19
2.2.1. Using SLApp02-1	19
2.2.2. How SLApp02-2 works	20
2.2.3. Using SLApp02-2	21
2.2.4. Practicing strategies for tristimulus	27
3. Superposition of sinewaves III. – Transitions between waveforms	28
3.1. Theoretical background	28
3.2. Practical Exercises	30
3.2.1. How SLApp03 works	31
3.2.2. Using SLApp03	33
3.2.3. Practicing strategies for identifying different harmonic spectra (sinewave, triangle, sawtooth)	38
4. Superposition of sinewaves IV. – Inharmonic spectrum	40
4.1. Theoretical background	40
4.1.1. Sounds with nearly harmonic partials	41
4.1.2. Sounds with widely spaced (sparse) partials	41
4.1.3. Sounds with closely spaced (dense) partials	42
4.2. Practical Exercises	43
4.2.1. How SLApp04 works	43
4.2.2. Using SLApp04	45
5. Modelling musical examples – bell-like sounds	49
5.1. Theoretical background	49
5.2. Practical Exercises	50
5.2.1. How SLApp05 works	50
5.2.2. Using SLApp05	52
5.2.3. Practising strategies for hearing the fusion of partials of Risset's bell	54
6. Modelling musical examples – endless scale and glissando	55
6.1. Theoretical background	55
6.2. Practical Exercises	56
6.2.1. How SLApp06 works	57
6.2.2. Practicing strategies	58
7. Filtering white noise with low-pass and high-pass filters	59
7.1. Theoretical background	59
7.1.1. Subtractive synthesis	59
7.1.2. White noise	61
7.2. Practical Exercises	62
7.2.1. How SLApp07 works	63
7.2.2. Using SLApp07	63
7.2.4. Practicing strategies	66

8. Filtering whitenoise with band-pass filter.....	68
8.1. Theoretical background	68
8.2. Practical Exercises	70
8.2.1. How SLApp08 works	70
8.2.2. Using SLApp08	70
8.2.3. Practicing strategies.....	73
9. Filtering whitenoise with different filters	75
9.1. How SLApp09 works	75
9.2. How the timed test works	76
9.3. Using SLApp09	77
9.4. Practicing strategies	77
10. Filtering sound samples with low-pass, high-pass and band-pass filters	78
10.1. Practical Exercises	78
10.1.1 How SLApp10 works.....	79
10.1.2. Using SLApp10.....	79
10.2. Listening strategies	80
10.2.1. Spectral analysis of the original soundfiles	80
10.2.2. Practicing strategies	83
11. Filtering with an additive resonant filter	86
11.1. Theoretical background	86
11.2. Practical Exercises	86
11.2.1. How SLApp11 works.....	87
11.2.2. Using SLApp11.....	88
12. Distortion	91
12.1. Theoretical background	91
12.1.1. Clipping.....	91
12.1.2. Folding	92
12.1.3. Wrapping	92
12.2. Practical Exercises	93
12.2.1. Structure of the patch	93
12.2.2. Using SLApp12.....	94
12.2.3. Practicing strategies.....	94
13. Granulation of soundfiles	95
13.1. Theoretical background	95
13.2. Practical Exercises	96
13.2.1. How SLApp13 works.....	96
13.2.2. Using SLApp13.....	97
13.2.3. Practicing strategies	98
14. Localization – positioning sound across the left-right axes	99
14.1. Theoretical background	99
14.2. Practical Exercises	101
14.2.1. How SLApp14 works.....	102
14.2.2. Using SLApp14.....	103
15.1. Theoretical background	111
15.2. Practical Exercises	113
15.2.1. How SLApp15 works.....	113
15.2.2. Using SLApp15.....	114
15.2.3. Practicing strategies	115

The aim of the first three chapters is to help develop the analytical listening skills necessary to identify harmonic spectra built from sinewaves.

1. Superposition of sinewaves I.

In this chapter you will learn about the timbre of harmonic spectra created from the first eight partials in different combinations. Exercises presented here will contribute to the development of analytical hearing of spectra, which will help you to create timbres from sinewaves.

1.1. Theoretical background

1.1.1. Harmonic spectrum

A harmonic spectrum contains only frequency components whose frequencies are integer multiples of the fundamental frequency. Periodic waveforms always have harmonic spectra.

The series of integer ratio frequencies is also known as the harmonic overtones of the fundamental. The characteristic sounds of many musical instruments are built from overtones, as is the human voice. Since the overtones are integer multiples of the fundamental frequency, the distance between the overtones is constant. In Fig. 1.1. you can see the frequencies in Hertz (cycles per second) below the given partials.




Figure 1.1. overtone system (the frequency values are rounded to integer numbers)

But, while the neighbouring frequency values have the same difference, the distance between the **perceived** pitches gets narrower with the climb in the register. This convergence and increasingly dissonant interval is very important to hear and identify while listening to, and analyzing, harmonic spectra.




Figure 1.2. first eight partials of a harmonic sound

You can listen to these partials one by one. Concentrate on the decreasing intervallic distance between consecutive pitches.



1.01_Sound

1.1.2. Fusion

"The process by which the brain combines a previously analysed set of pure tones into a sound with only one pitch is known as fusion"¹

Fusion is a very interesting phenomenon of human audition. When sinewaves of different pitches are added together, one would suppose, the perceived result will be similar to a chord played by an instrument. This is true in many cases, but when the pitches of the partials are harmonic overtones, there is a high probability that the ear will fuse them into one single sound percept.

The phenomenon of fusion is dependent on:

- 1) the pitch ratio of the partials
A soundspectra that is harmonic and well-balanced in overtones has a clear fundamental pitch and fuses perfectly.
- 2) the amplitude of the partials
The amplitude contour of the partials should be smooth. If a partial has much higher value than the others, it could be clearly audible
- 3) the amplitude envelope of the partials
The shape of the amplitude envelopes of the partials should be similar in order to fuse them into one sound.
- 4) the overall envelope of the sound
Long sounds with slowly and smoothly changing dynamics (e.g. envelopes with trapezoidal or triangle shapes) fuse less than percussive sounds with a fast attack.
- 5) fundamental frequency
Sounds with higher fundamental frequencies fuse more easily as their partial frequencies tend to fall into common critical bands.
- 6) the overall length of the sound
Below a certain duration, the sound is heard as a click and the pitch cannot be identified. Therefore, short duration sounds tend to fuse whilst sounds of longer duration are easier to analyse by hearing.

¹ Campbell M. & Greated C. 2001 The Musician's Guide to Acoustics, Oxford. pg. 85

If a well balanced harmonic spectrum is presented, we normally hear one sound with a basic pitch and a particular colour or timbre. Experiments show that with some effort and practice we can distinguish the first eight partials in a fused harmonic spectra. Identifying partials helps us learn how an harmonic spectrum behaves.

SLApp01 contains a series of exercises where we will take the sound apart and listen to different combinations of its partials in order to learn to analyse its spectrum by ear.

SLApp01 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp01 Windows](#), [SLApp01 Mac OS X](#).

1.2. Practical Exercises

1.2.1. How SLApp01 works




Figure 1.3. layout of SLApp01

- 1 – Selection of exercises - in the pop-down menu four exercises and three tests can be selected.
- 2 – Frequency- here the fundamental frequency (first partial) is displayed. It can be changed by clicking and dragging or entering numbers.
- 3 – Length - displays the duration of the sound. By clicking and dragging or entering numbers the value can be changed.
- 4 – Amplitude envelope - in the pop-down menu three types of amplitude envelopes can be selected: trapezoid, triangle, percussive.
- 5 – Sound On-Off button.
- 6 – Volume slider
- 7 – Volume indicator
- 8 – Amplitude of partials - eight green columns represents the amplitude of the partials. In this exercise they have two states: 1 (on) or 0 (off).

- 9 – Frequency ratio of partials - the number box below each column specifies its relationship to the fundamental. If the fundamental is set to 200 Hz each frequency is its multiple ($200 \times 2 = 400$, $200 \times 3 = 600$, etc.). In this chapter sounds created only by the first eight harmonic partials are explored.
- 10 – Selection of spectra - the yellow buttons are used to select different combinations of partials specified later in each exercise.

Timed Test

In tests you can check your hearing. Timed test will allow you to assess your skills. Clicking on the Timed Test button will open a new window, where you can specify the level of the test and the total number of the sounds to want to hear.




Figure 1.4. layout of Timed Test

- 1 – Level
Select the level: easy, medium or hard, from the pop-down window.
Easy: the test sound is played twice and you have 20 seconds to answer.
Medium: the test sound is played twice and you have 10 seconds to answer.
Hard: the test sound is played only once and you have only 8 seconds to answer.
- 2 – Total number of sounds in the test
Specify the number of sounds you want to hear in your test.
- 3 – Start or stop your test
- 4 – Time remaining
Shows the time remaining to guess the current sound. The actual sound is played
- 5 – Evaluation – overall result
This window tells you (in percentage) how accurately you identified the sounds in total.
- 6 – Evaluation – one sound
This led tells you how accurately you identified the sound (green - correct, red - false)
- 7 – Current number of sound
This number shows which sound is being played currently.

1.2.2. Using SLApp01

These are series of exercises and tests, which will help you to develop your ability to hear the individual components that go to make up a timbre.

Excercise1 (Ex01)




Figure 1.5. Ex01

Clicking on yellow button 1 will play the fundamental frequency. Buttons 2 - 8 will add each consecutive partial. Listen and try to memorize the eight spectra.

Test01




Figure 1.6. Test01

In Test01 you can test your skills gained in Exercise1 (Ex01). The interface is similar but the amplitudes of the partials are hidden. Press the button 'NEW SOUND!', and the patch will randomly play you one of the eight spectra. As an answer you have to select the corresponding yellow button. The correct answer is indicated by the green Led in the bottom right corner. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the 'REPEAT!' button.

Try this test with different fundamental frequencies, lengths and envelopes.

After you practiced, tested yourself and feel confident in indentifying the combinations of partials, you can move on to the timed test by clicking on the button.

Excercise2 (Ex02)




Figure 1.7 Ex02

Excercise 2 (Ex02) offers seven different spectra. The partials are added to the fundamental from the top. By pressing the yellow button labelled "1+8" you will hear the fundamental and the 8th partial. Pressing button "1+7-8" you will hear the fundamental with the 7th and 8th partials and so on. Pressing the button 'all' you will hear the fundamental with all 7 overtones. This is like pressing button 8 in Ex01. Listen and try to memorize the eight spectra!

Excercise3 (Ex03)




Figure 1.8. Ex03

Excercise3 (Ex03) combines Ex01 and 02. The circular setup of the buttons illustrates a continuous timbre scale. Listen and try to memorize the 14 spectra!

Test02





Figure 1.9. Test02

In this exercise you can test your skills gained in Exercises 2 and 3. The test functions as in Test01.

Excercise4 (Ex04)*Figure 1.10. Ex04*

In Excercise4 any combination of the eight partials is possible. Because the number of these combinations is huge, using yellow button for each possible combination would be unpractical. Therefore you need to click in the space directly above the number box to trigger the fundamental and its harmonics. Try to memorize the sounds resulting from different combinations of partials.

Test03




Figure 1.11. Test03

In Test03 you need to identify different combinations of the eight partials. Pressing the button '!NEW SOUND!' will play you a random set. On the screen with red columns you should select the partials you think are present in the perceived sound. You can check your result by clicking on the 'Play my sound!' button. To compare your result with the test sound, you should click the '!REPEAT!' button and again the 'Play my sound!' button as many times as you need. To see, how close you are to the right answer, you should click on the 'Show Result!' button, which will give you the result in the right lower corner as a percentage.

Pressing the button 'Show spectrogram!' will show you the partials present in the test sound on the upper screen so that you can see any differences.

1.2.3. Listening strategies

We tend to hear harmonic timbre as a whole. The purpose of this lesson is to encourage you to develop more analytical listening in order to distinguish the partials present in the sound. It will help you to describe perceived differences between timbres.

Everybody hears the spectrum quite subjectively, so it is important to find personal strategies to identify the sonorities. To remember a sound it is important to name, identify and categorize it thru association and comparison (e.g. this sound reminds me of cicadas at the sea). This is a useful starting point, which needs to be refined with more analytical hearing. Here are three strategies to follow to increase your ability to understand what's happening in the sounds presented in this chapter:

1. Listen to the timbre as an overall sound. The fundamental sinewave (button 1) is easy to identify for its purity and emptiness. The sound becomes warmer thru adding the lower partials (2,3,4). When the higher partials are added (5 and above) depending on the fundamental frequency the sound becomes more nasal or harsher or more metallic.
2. Train your ear to hear the highest partial, because it will help you to recognise the full spectrum. Listen to the partials individually and try to remember their pitch compared to that of the fundamental. As we ascend, the pitches of the partials will follow the scale of the overtones (see fig. 1.2.).
3. Pay particular attention to the intervallic distance between the two top partials. The interval between each successive partial becomes smaller and therefore more dissonant. So partials 7 and 8, for example, will create a noticeable dissonance or harshness in the sound.

It is important to remember that the frequency of the fundamental will greatly influence the perception of the resulting sound.

- at 120 Hz and below, the partials fall into separate [critical bands](#) so it is easier to hear them separately and the intervals can be heard more clearly.
- for higher sounds (150-500 Hz) where some of the partials fall within the same critical bands you should concentrate more on hearing any dissonance present in the sound, any roughness any nasal or metallic character.

To practice the fusion of 8 sinewaves, start with the trapesoidal envelope, a length of 2000 msec or more and a relatively low fundamental frequency (80 to 250 Hz). Higher, shorter and percussive sounds tend to fuse more completely, therefore it is harder to distinguish the partials within the sound.

2. Superposition of sinewaves II. – Tri-stimulus

2.1. Theoretical background

The Tri-stimulus method is one of a number of representational approaches for understanding the perception of the harmonic spectrum of a sound. The theory was proposed by Pollard and Jansson² as an analogy of trichromatic color-vision/perception. (When light enters the eye, it falls on three different types of receptor cone cells that are tuned to the wavelengths or frequencies of red, green and blue light, the brain combines the information from each type of receptor to give rise to the perception of different wavelengths of light, or colours).

Similarly, the main properties of harmonic sound spectra can also be reduced to three parameters that correspond to specific spectral regions of the sound. Region 1 is the fundamental (1st partial). Region 2 contains the mid-range partials 2-3-4, and Region 3 contains the higher-range 5th partial and upwards. The ratio of the average loudness of each region can be represented within a triangular shaped coordinate system that is known as the **tristimulus diagram** (see Fig. 2.1.). If we add together the three regions, the result equals 1. The horizontal axis represents the mid range frequencies (m), the vertical axis, the higher range frequencies (h), the remaining part will equal the fundamental. So that, $f + m + h = 1$.

A pure sinewave with no harmonic partials is represented by a point at (0,0). Since the loudness value of the frequencies of middle and higher partials = 0, the loudness value of the fundamental = 1. A sound where the timbre is balanced is represented by the black dot in the middle of the triangle, the value of the loudness of the mid and high partial groups is equal ($m/3 + h/3 + f/3 = 1$). The yellow and red dots represent the tristimulus positions of a violin and an oboe respectively. The more balanced timbre of the violin has slightly more higher partials than mid-range partials. The oboe has predominantly more mid-range partials in its spectrum. The curved green line indicates a sound where the ratio of the loudness of mid and high range partials changes over time.

² Campbell, M.–Greated C. The Musician's Guide to Acoustics, Oxford, 2001,p. 150.




Figure 2.1. tristimulus diagram

2.2. Practical Exercises

In this chapter we will use SLApp02.1 and SLApp02.2 to further explore harmonic spectra using the tristimulus method.

SLApp02 (containing SLApp02.1 and SLApp02.2) is downloadable for Windows and Mac OS X platforms using the following links: [SLApp02 Windows](#), [SLApp02 Mac OS X](#).

2.2.1. Using SLApp02-1

With SLApp02-1 (Fig. 2.2.) you can create and listen to different tristimulus variations of a source sound created from 16 sinewaves by changing the position of the grey dot in the triangle.




Figure 2.2. Layout of SLApp02-1

2.2.2. How SApp02-2 works




Figure 2.3. Layout of SApp02-2

- 1 – Selection of exercises - in the pop-down menu four exercises and three tests can be selected.
- 2 – Frequency- here the fundamental frequency (first partial) is displayed. It can be changed by clicking and dragging or entering numbers.
- 3 – Length - displays the duration of the sound. By clicking and dragging or entering numbers the value can be changed.
- 4 – Amplitude envelope - in the pop-down menu three types of amplitude envelopes can be selected: trapesoid, triangle, percussive.
- 5 – Play the sound
- 6 – Sound On-Off button.

- 7 – Spectrum of the source sound
Sixteen green columns represent the amplitude of the partials. These can be changed by clicking and dragging. The number boxes below indicate the frequency ratio of the partials. they are fixed at the fundamental and the first 15 harmonics.
- 8 – Source sound selector
Different spectra can be selected by clicking on the buttons:
Spect01 - all partials with full amplitude
Spect02 - odd partials with full amplitude
Spect03 - odd partials with decreasing amplitudes at higher harmonics
Spect04 - random spectra
- 9 – Preset buttons
By pressing the buttons you can select and hear a spectrum created using tristimulus combination. The coloured columns represent the relative amplitudes of fundamental, mid-range, and high-range partials.
- 10 – Buttons to hide or show the spectrum of the transformed source sound (11)
- 11 – Spectrum of source sound transformed by tristimulus presets (9)

2.2.3. Using SLApp02-2

These are series of exercises and tests, which will help you to become familiar with the harmonic spectrum represented by the three tristimulus regions.

In Exercise01 only two regions of partials are used. In Exercise02 all three regions are used which offer more combinations and complex sounds.

In all presets (9) the relative levels of the tristimulus regions are set at maximum, half or zero level. This limits the possible timbral combinations allowing you to concentrate on the basics.

The upper sonogram-like display represents the spectrum of the source sound, and the result of the tristimulus transformation can be seen on the lower display. Even more combinations can be created by changing the amplitude levels of the partials in the source spectrum.

Excercise1 (Ex01)




Figure 2.4. Ex01

In Ex01 by clicking on the numbers on the left side of the 3 part scales you will hear 3 sounds representing 3 combinations of regions in sequence. The spectrogram below changes to indicate which combination is playing, helping you to identify the sounds visually. Clicking on the 'Hide spectrogram!' makes the exercise more difficult by removing it. Listen and try to memorize the different scales!

You can also click on the preset buttons to hear them individually. Listen and try to memorize the presets the different tritune combinations!

Test01




Figure 2.5. Test01

When you've spent some time listening to the scales in Ex01 you can test your listening skills here with Test01.

Press the button '!New scale!' the patch plays one of the six scales from Ex01. See if you can recognize them. To answer click on the corresponding orange number button. The correct answer is indicated by the green Led in the bottom right corner. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the '!REPEAT!' button. To see the spectra of your answer press the 'Show spectrogram!' button.

Test02




Figure 2.6. Test02

In Test02 the presets are arranged randomly not in scaled as they were in Ex01 and Test01.

Press the button '!New sound!'. To answer click on the corresponding preset button. The correct answer is indicated by the green Led to the right of the central panel. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the '!REPEAT!' button. To see the spectra of your answer press the 'Show spectrogram!' button.

Excercise2 (Ex02)




Figure 2.7. Ex02

Ex02 introduces the third region, so you can practise with a combination of all three. Sixteen different possible timbre combinations can be selected by clicking on the preset buttons. Listen and try to memorize the different combinations!

Test03




Figure 2.8. Test03

Test03 tests your ability to hear and differentiate between the combinations of three-partial regions explored in Ex02. The test interface functions exactly as in Test02.

2.2.4. Practicing strategies for tristimulus

The three regions of the tristimulus representation include three different types of sonority:

- f_0 , the fundamental sinewave has an empty, pure sound without any roughness or beating. Because of its low frequency, it will "darken" the timbre of a spectrum, it could be said to add power to the resulting timbre. Try listening to the difference of the spectrum with and without the fundamental.
- the middle region – the 2nd, 3rd, 4th partials – creates resolved, pleasant intervals (octaves and fifth). Compared with the simplicity of the fundamental sinewave this region according to Helmholtz is "rich and splendid but because of the absence of the higher partial is still soft and pleasant. Try listening to the difference of the spectrum with and without the middle region.
- in the highest region the partials are much closer to each other creating smaller intervals. Above the 6th partial the intervals between the consecutive overtones are smaller than a major second and ascending higher they become more and more dissonant. In this region the tone becomes cutting, shrill, buzzing or rough depending on the frequency of the fundamental. Try listening to the difference of the spectrum with and without the high region.

Listen to the individual spectral regions defined by the tristimulus representation, to learn to differentiate between the three kinds of percepts. After being able to identify each of them in different pitch regions, try their combinations.

3. Superposition of sinewaves III. – Transitions between waveforms

In this chapter we will explore the perceptual differences between spectra exhibited by some familiar waveforms: *sine wave*, *triangle wave*, *square wave* and *sawtooth wave*.

As in chapters 1 & 2, the spectra are created by adding sinewaves together (additive synthesis), but now we have many more partials. Spectra made up of 48 sinewaves are rich enough to test the differences between the above mentioned types of waveform.

3.1. Theoretical background

A *sinewave* is a smooth repetitive oscillation, which cannot be created by summing other types of waveforms. It is the simplest building block to describe and approximate any periodic waveform.


The shape of a sine wave can be described by its:

- frequency (f) that is the number of the repeating cycles in one second,
- amplitude (A) that is the peak deviation from the central middle level in each direction,
- phase (ph) that is the initial angle of a waveform at its origin.

Frequency and amplitude define the sound that we hear, phase has no effect on the auditive result. The pitch of a sine wave depends on the frequency, the loudness depends on the amplitude of the wave. (The greater the frequency, the higher the pitch; the bigger the amplitude, the louder the sound and vice versa.)

When adding together sine waves which are in harmonic relation to each other, more and more complex periodic waveforms will result (e.g. triangle or sawtooth) the spectra of which can be defined by simple rules. The timbre of these waveforms are very characteristic, that's why they were chosen as the building blocks of many synthesizers.

Sawtooth wave is a complex, non-sinusoidal waveform that contains all the integer harmonics (f , $2f$, $3f$, $4f$, etc.) with equal intensity. The sound is harsh and clear. The name of this wave was given after the shape of the sum of sine waves (see Fig. 3.1.).



http://en.wikipedia.org/wiki/File:Synthesis_sawtooth.gif

Figure 3.1. sawtooth wave

Triangle and square waves are complex waveforms in which the spectra contain only the **odd numbered partials** (f , $3f$, $5f$, $7f$, etc.). Both waveforms are named after their shape. Square wave is similar to triangle wave since it also contains the odd numbered partials. The higher harmonics of a triangle wave roll off much faster than in a square wave (proportional to the inverse square of the harmonic number as opposed to just the inverse).

The shape of a triangle wave (and the square wave) can be seen in Fig. 3.2. and 3.3.




Figure 3.2. triangle wave




Figure 3.3. square wave

To summarize:

Sinewave: the simplest form which can not be subdivided any more

Sawtooth wave: non-sinusoidal, contains all the partials: (f, 2, 3, 4, 5, 6 etc.)

Triangle wave: contains only the odd numbered partials: (f, 3, 5, 7, 9 etc.)

Squarewave: contains only the odd numbered partials: (f, 3, 5, 7, 9 etc.)

You can listen to the four basic soundtype here:

Sinewave

Sawtooth wave

Triangle wave

Squarewave



03_01_Sound

03_02_Sound

03_03_Sound

03_04_Sound

The aural difference between the spectra of the sine wave, the sawtooth wave and the triangle wave is quite obvious. It is clear that sine wave is a pure sound. The triangle wave with the odd numbered harmonics has a hollow, clarinet like sound. The sawtooth is the most complex sound being richer and nasal.

We can create interpolations between waveforms by changing the amplitudes of their partials. In this way a continuous timbre scale can be built up between the sine and the other waveforms (triangle, sawtooth, square).

3.2. Practical Exercises

SLApp03 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp03 Windows](#), [SLApp03 Mac OS X](#).

3.2.1. How SApp03 works




Figure 3.4. structure of the patch

- 1 – Fundamental frequency – keyboard
The fundamental frequency of the spectra can be selected here pressing the keys of the keyboard. The spectrum will be specified by the multiplication of this frequency by integer numbers between 1-48.
- 2 – Play the specified sound
Pressing the button the specified sound is played.
- 3 – Sound On-Off button.
- 4 – Spectrum of the sound
Fortyeight red columns represent the amplitude of the partials. These can be changed by clicking and dragging.
- 5 – Length - displays the duration of the sound. By clicking and dragging or entering numbers the value can be changed.
Amplitude envelope - in the pop-down menu three types of amplitude envelopes can be selected: trapesoid, triangle, percussive.
- 6 – Preset buttons selecting different spectra (100 Hz)
Pressing the buttons you can play the sounds from sine-wave to total spectrum at 100 Hz fundamental frequency
- 7 – Preset buttons selecting different spectra (200 Hz)
Pressing the buttons you can play the sounds from sine-wave to total spectrum at 200 Hz fundamental frequency
- 8 – Preset buttons selecting different spectra (400 Hz)

Pressing the buttons you can play the sounds from sine-wave to total spectrum at 400 Hz fundamental frequency

9 – Record the sound

Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.

10 – Selection of exercises - in the pop-down menu four exercises and four tests can be selected.

11 – Length of the individual partials

The horizontal axis corresponds to the overall duration of the sound. The lines represent the individual partials with the fundamental at the bottom and the highest partial at the top.

Click on the preset buttons above for twelve combinations of the duration of each partial and you can click and drag in the panel to change the lengths of the partials. We suggest you use this timbral dimension only after you have become familiar with the preceding exercises.

3.2.2. Using SLApp03

These are series of exercises and tests, which will help you to become familiar with the sound of the sinewave, the sawtooth and the triangle and the transitions between them.

In all examples the button marked 0 in the panel (at 8) triggers the fundamental sinewave. Buttons 1 thru 7 add more partials.

Ex01




Figure 3.5. Ex01

In Exercise1 (EX01) you can explore the transition between sinewave and sawtooth wave (all partials present). The button marked 0 in the panel (at 8) triggers the fundamental sinewave. Buttons 1 thru 7 add more partials. Listen and try to memorize the 8 spectra!

You may experiment with different fundamental frequencies, lengths and envelopes as well.

Test01




Figure 3.6. Test01 Test your ability to hear the transition between the sinewave and the sawtooth wave introduced in Ex01.

Press the button '!NEW SOUND!' to hear one of the eight spectra. To answer click on one of the numbered buttons (0-7). The correct answer is indicated by the green Led in the bottom right corner. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the '!REPEAT!' button.

Pressing 'Show spectrogram!' will display the partials and their relative amplitudes (as in EX01).

Ex02




Figure 3.7. Ex02

In Exercise2 (EX02) you can explore the transition between sinewave and triangle wave (odd partials present).

Again the button marked 0 in the panel (at 8) triggers the fundamental sinewave. Buttons 1 thru 7 add more partials. Listen and try to memorize the 8 spectra!

You may experiment with different fundamental frequencies, lengths and envelopes as well.

Test02

As Test01 just for the triangle wave.

Ex03




Figure 3.8. Ex03

Exercise03 combines the sawtooth and the triangle wave. Here you can compare the difference between the character of the sound of the three waveforms (sinewave, triangle, sawtooth). Notice that in all transitional steps (1-7) you can hear the characteristics of the sawtooth and the triangle emerging.

Practice both sawtooth and triangle wave as in Exercises01 and 02. This time pay particular attention to the difference you can hear between the waveform types.

Experiment with different fundamental frequencies, lengths and envelopes.

Test03

As Test01 just for the combination of sawtooth and triangle wave.

Ex04




Figure 3.9. Ex04

In Ex04 you have an interface combining the spectra seen in Ex03, this time with three different fundamental frequencies (100, 200, 400 Hz). The purpose of this interface is to allow you to develop the ability to hear the number and amplitude structure of the partials independent of the fundamental frequency. It is important to understand the influence of pitch register in contributing to the timbral characteristic of sounds, as it tends to mask the characteristics of the different waveforms and transitions between them.

First become familiar with the transformation from sinewave (0) to sawtooth and triangle (7) for each frequency. After that try listening "vertically" alternating different frequencies.

Test04



Figure 3.10. Test04

Test your ability to hear the transition between the sinewave and the sawtooth and/or triangle wave in different pitch registers.

Pressing "New sound!" will select sounds from any of the given fundamental frequencies and number of partials both odd and all.

The test functions exactly as in Test01.

3.2.3. Practicing strategies for identifying different harmonic spectra (sinewave, triangle, sawtooth)

To identify different harmonic spectra it is suggested you pay particular attention to:

- the different "qualities" of the timbre: learn to identify the hollow, clarinet-like quality of the triangle wave and distinguish it from the nasal oboe-like sound of sawtooth wave. The difference is a result of the absence or presence of the even numbered partials.
- the presence of higher partial will influence the brightness of the sound. Remember the higher the partials the closer to each other they are, the more dissonant they are, the harsher the resulting sound. It will result in both spectral-types (triangle and sawtooth) in a buzzing or serrated quality in the sound.

In addition to listening and exploring the qualities of sound transformed by the amplitude and presence of odd or all partials we can significantly change the quality of the sound by changing the overall length, the amplitude envelope and the duration of individual partials. The length of the partials can be changed by clicking on the presets or clicking and dragging in the panel (at 10 in 3.2.1.).

4. Superposition of sinewaves IV. – Inharmonic spectrum

In the following chapters you will learn about the timbre of inharmonic spectra created from different numbers of partials in different combinations.

4.1. Theoretical background

An inharmonic spectrum contains frequency components whose frequencies are not integer multiples of the fundamental.

The closeness of the partials, their frequency ratio, the length and the amplitude of the partials defines the fusion and the sense of the given spectra. The inharmonic spectrum could be described as the transition between harmonic sounds (see Chapters 1-3) and noise (see Chapter 7). Whilst it is obvious what pitch is heard when a harmonic sound is played, with inharmonic sounds it is not so clear.

The difference between harmonic and inharmonic sounds can be perceived as sensory dissonance. The level of inharmonicity is dependent upon the distance between the partials of a spectrum. There are different degrees of inharmonicity ranging from beating through roughness to distinguishable intervals.

The well-known phenomenon of beating occurs when two sinewaves have frequencies which are very close to each other. When the two frequencies are between 1 and a few Hertz apart, our ears hear only one frequency with a periodically changing amplitude. The speed of the beating corresponds to the difference between the frequencies of the two sinewaves. (Listen to the beating produced by two sinewaves with that are 1, 2, 4 and 10 Hz apart – 4.01_Sound.)

If the sinewaves are further from each other (approx. 10 Hz or more), we can no longer perceive the beating as separate amplitude peaks, they occur so fast, we hear them as a "roughness".

When the difference between the frequencies of the sinewaves is larger than the so called critical band, they will be heard separately and a percept of an interval or fusion will occur.



4.01_Sound

According to Beauchamp³ inharmonic sounds have three categories:

- 1) Sounds with nearly harmonic partials
- 2) Sounds with widely spaced (sparse) partials
- 3) Sounds with closely spaced (dense) partials

4.1.1. Sounds with nearly harmonic partials

The partials of this spectrum diverge slightly, but increasingly, from the harmonic partials so the higher the partial, the greater the divergence. Beating or plucking a string will produce a spectrum divergent from a clear harmonic one. This phenomenon results in a clearer, colder, more tense and brighter sound.




Figure 4.1. spectrum of nearly harmonic partials

4.1.2. Sounds with widely spaced (sparse) partials

Since the partials are dispersed and are separated from each other, there is no "roughness" between them. At the same time some partials could be in harmonic relation so they fuse better than others. This is typical for certain acoustic instruments (e.g. percussion with wooden or metal plates).

³ Beauchamp, James W.: „Analysis and Synthesis of Musical Instrumental Sounds”. In: James W. Beauchamp (ed.): *Analysis and Synthesis, and Perception of Musical Sounds, The Sound of Music*. New York: Springer Science+Business Media, 2007. 1-89.




Figure 4.2. spectrum of widely spaced (parse) partials

4.1.3. Sounds with closely spaced (dense) partials

Since the partials are very close to each other, the roughness of the spectrum is very evident. Therefore no pitch content can be heard. Most metal plate or stretched membrane instruments sound this way. Their spectra might sometimes resemble that of white-noise (see Chapter 7). The sound of a suspended cymbal is a good example of this, since the frequencies of its partials are only 20 Hz from each other. This is a tiny interval especially in the higher register.




Figure 4.3. spectrum of closely spaced (dense) partials

4.2. Practical Exercises

SLApp04 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp04 Windows](#), [SLApp04 Mac OS X](#).

4.2.1. How SLApp04 works




Figure 4.4. Layout of SLApp04

- 1 – Fundamental frequency – keyboard and number box
The fundamental frequency of the spectra can be selected here pressing the keys of the keyboard and clicking and dragging the number box marked pitch. The spectrum will be specified by the multiplication of this frequency by the numbers given at Partials.
- 2 – Play the specified sound
Pressing the button the specified sound is played.
- 3 – Sound On-Off button.
- 4 – Length - displays the duration of the sound. By clicking and dragging or entering numbers the value can be changed.
Amplitude envelope - in the pop-down menu three types of amplitude envelopes can be selected: trapesoid, triangle, percussive.
- 5 – Spectrum of the sound
Fortyeight red columns represent the amplitude of the partials. These can be changed by clicking and dragging.

- 6 – Frequency ratio of the 48 partials
The frequency ratio of each partial can be specified in the 48 number boxes. The buttons to the right will load presets of different partial ratios.
(Button1 – harmonic partials used for the beating. Buttons 2-5 – different inharmonic ratios with different spectral density.)
- 7 – Preset buttons selecting different spectra with different spectral densities
(By spectral density we mean the number of partials present in the spectrum).
- 8 – Preset buttons selecting different spectra with different beating frequencies
- 9 – Record the sound
Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.
- 10 – Selection of exercises - in the pop-down menu three exercises and two tests can be selected.
- 11 – Length of the individual partials
The horizontal axis corresponds to the overall duration of the sound. The lines represent the individual partials with the fundamental at the bottom and the highest partial at the top.
Click on the preset buttons above for twelve combinations of the duration of each partial and you can click and drag in the panel to change the lengths of the partials. We suggest you use this timbral dimension only after you have become familiar with the preceding exercises.

4.2.2. Using SLApp04

Here we explore the phenomena of beating and different types of inharmonic spectra. The sounds that can be produced by inharmonic spectra are far more complex than those produced by purely harmonic ratios between the fundamental and its partials.

Ex01




Figure 4.5. Ex01

In this exercise you will explore a continuum from a pure harmonic spectrum (0) thru beating to roughness (7) as discussed earlier. The source spectrum is built from 48 partials with integer ratios. When transforming the source spectrum there are 2 sinewaves added above and below each partial. The difference between the frequency and amplitude of the partials and those of the added sinusoids will determine if we hear beating or roughness.

In the panel marked "Beating" there are two rows of buttons 1-7. In the top row the amplitude of the added sinewaves is equal to the amplitude of the source partials. In the lower row the sinewaves are 0.3 of the amplitude of the source partials. The preset

buttons 1-7 determine the frequency of the beating. In button 1 you can hear recognizable slow beating, in button 7 the beating is so fast that you hear it as beating. It is important to note that the values of the frequencies of the added sinewaves will vary randomly at each partial. The preset numbers will define the range of the randomness of the beating frequency in Hertz indicated in the panel below. Button 1 adds beating with a frequency from 1 to 6 Hz, button 2 with 5 to 10 Hz, button 3 with 8 to 14 Hz, button 4 with 13 to 18 Hz, button 5 with 17 to 24 Hz, button 6 with 23 to 33 Hz and button 7 with 32 to 51 Hz. These added sinewaves create a scale from soft beats to roughness.

You may change the fundamental frequency, the length and the envelope as well.

Test 01




Figure 4.6. Test01

Test your ability to hear and recognize the speed and amplitude of beating (roughness) introduced in Ex01.

Press the button 'NEW SOUND!' to hear one of the fifteen possibilities. To answer click on one of the numbered buttons (0-7) in the two rows. The correct answer is indicated by the green Led in the bottom right corner. If the answer is incorrect, it will be red and you can guess again.

If you need to hear the sound again, press the '!REPEAT!' button. If you need to see the beating ranges, click on "Show visuals!"

Ex02




Figure 4.7. Ex02

In this exercise now we explore inharmonic spectra with different spectral densities (ie. different number of partials of differing amplitudes present in the spectrum). The ratios of the partials to the fundamental can be seen in the central panel ranging from 1-10.72. Note that they are not integers and therefore the spectrum will be inharmonic.

There are two rows of 6 buttons. The top will trigger a fundamental frequency value of 200 Hz, the lower 100 Hz. By clicking on the buttons 1-6 you will see that each button triggers different combinations of partials of increasing density where 1 is the least dense and 6 is the most illustrating the difference between sparse and dense spectra as discussed at the beginning of the chapter (Fig. 4.2. and 4.3.)

You can also practice with different spectra by selecting from the presets on the right side of the center Partial panel. You could set your own numbers by clicking and dragging the values.

Of course as always you can change the length of the sound, its amplitude envelope and the length of its partials.

Test02

The test works the same as in Test01.

Ex03

This exercise combines the possibilities of exploring beating and spectral density together. Play with it freely, create and record your own sounds. When you feel you found an exciting sound make sure you record the settings that have produced the sound. You could do this by using a screen grab.

5. Modelling musical examples – bell-like sounds

In this chapter we will look at an analysis and a synthesis of a real world sound, that of a bell. It is a practical exercises, where you can explore how to model a bell like sound by adding inharmonic sinewaves together. You can also learn how to transform the sound radically by changing the amplitude envelopes of the individual partials.

5.1. Theoretical background

A well-known example of the fusion of inharmonic sounds is Jean-Claude Risset's synthesized bell. Basically, the amplitude envelope of the spectrum of a bell has two phases; the initial attack phase is short, the release phase is much longer. The change in both of the phases is exponential. The difference between the two phases is that in the attack the partial envelopes are synchronised, while the partials release or fade at different rates. This behaviour was imitated by Risset creating independent sinewaves of different frequency, amplitude and length.

Partial	Frequency ratios	Amplitude ratios	Length ratios
11	7.2678	1.3429	0.0762
10	6.7142	0.9714	0.1048
9	5.3571	1.3429	0.1524
8	4.8928	1.3429	0.2
7	3.5714	1.4571	0.2476
6	3.1357	1.6571	0.1048
5	2.125	2.6857	0.3238
4	1.65	1.8	0.5524
3	1.6428	1	0.8762
2	1.0053	0.6857	0.9048
1 (fundamental)	1	1	1

Figure 5.1. frequency, amplitude and length ratios of Risset's bell

The *length or duration* of the partials of a bell are different; generally, the lower ones are longer and this accounts for the sensation of a descending pitch when we listen to a bell. The *amplitude envelopes* are identical, but the *amplitude ratios* are different. The

identical amplitude envelopes help the partials to fuse into one sound percept. Normally inharmonic spectra tend not to fuse as completely as harmonic spectra, but in this case because of the synchronized fast percussive attack the partials of the inharmonic spectrum of the bell sound like arriving from the same source.

You can see the basic parameters, frequency amplitude and length ratios, in the table at Figure 5.1.

Risset not only created a synthesized imitation of a bell but he also exploited the independency of the partials to modify the sound. If the shape of the envelope is changed so that each partial has its dynamic climax point midway through its own duration, (5.2,c) the pitch of the actual partials will be heard at that point, creating an inharmonic falling melody. This also helps the listener to identify the individual partials that make up the sound of the bell. You can Risset's initial bell synthesis (both percussive and the transformation of the partial envelopes) at Sound_5.01.




Figure 5.2.a, b, c: partials and summing of bell-like sound with different envelopes



5.01_Sound

Listening to the sound examples (5.02, 5.03, 5.04) you can hear that the transformation caused by the change of form of the amplitude envelope creates similar effects in both harmonic and inharmonic spectra.

harmonic spectra

inharmonic spectra

inharmonic spectra



5.02_Sound



5.03_Sound



5.04_Sound

5.2. Practical Exercises

SLApp05 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp05 Windows](#), [SLApp05 Mac OS X](#).

5.2.1. How SLApp05 works

You will notice the numbering system has changed a little because SLApp05 is more complex having more control parameters.




Figure 5.3. structure of the patch

- 1 – Sound On-Off button.
- 2 – Play sound button.
- 3 – Selection of the exercises – in the pop-down menu two exercises can be selected.
- 4 – Length – displays the duration of the sound. By clicking and dragging or entering numbers the value can be changed.
- 5 – Frequency – the fundamental frequency can be changed here.
- 6 – Amplitude envelope
In SLApp05 for the first time one can manipulate the amplitude envelope. When you open the SLApp by clicking on the preset buttons above you can move the peak from the beginning to the end. This envelope is applied to each partial over its duration. You can click and drag on the points on the envelope curve to create your own amplitude envelopes.
- 7, 8, 9 – the length, amplitude and frequency ratios are displayed here and can be changed by clicking and dragging.
- 10 – a 7x7 matrix of presets
The 49 buttons trigger different envelope and length values. The amplitude envelope can be changed by pressing the buttons on the horizontal axis as represented along the bottom in the form percentage of time.

The individual length of the partials can be changed by pressing the buttons on the vertical axis. On the right side of the matrix is a graphical representation of the length ratios of the partials for each row.

11 – Record the sound

Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.

5.2.2. Using SLApp05

Practice the fusion of inharmonic spectra by changing the length and amplitude of partials (based on Risset's bell). The examples offer two different ways of practising based on different amplitude envelopes.

Ex01




Figure 5.4. Ex01

This exercise has envelopes with 7 different peak positions (0.1%, 2%, 5%, 8%, 11%, 14%, 20%). The peaks are played relatively early after the beginning of the sound, therefore one can hear an increasing softening of the attack.

Ex02




Figure 5.5. Interface of exercise2

This exercise has also 7 envelopes with different peak positions (0.1%, 10%, 30%, 50%, 65%, 80%, 99%). the peaks shift from the beginning (0.1%) to the end (99%) which determines wheather we hear a fused bell-like sound or a melody of the partials.

5.2.3. Practising strategies for hearing the fusion of partials of Risset's bell

There are no tests in this chapter. We encourage you to play experiment and explore.

Since there are only two timbral dimensions changing in this patch it is worth concentrating on them one by one. Learn the attack peaks and listen carefully to the softening of attack in Exercise 1 and the emerging "melodic" phenomenon in Exercise 2 as the attack time becomes longer.

The phenomenon of fusion is more obvious with short and hard attacks. Partials are heard separately when their attack phases are longer. Where the length of the individual partials differ from the others each partial is heard separately.

6. Modelling musical examples – endless scale and glissando

In the previous chapters the phenomenon of fusion has been central to our understanding of harmonic and in harmonic spectra. Here we will explore an interesting property which can be used to fool the ear completely.

In this chapter we will introduce the auditory illusion of the Shepard tone and the Shepard-Risset glissando. We think we hear endlessly rising or falling scales. It is fusion that is responsible for this illusion.

6.1. Theoretical background

The basis of this illusory phenomenon was invented in 1964 by Roger Shepard hence it is known as the Shepard tone. What we hear is a repeated set of tones organized in an ascending or descending scale. Once the ascent or descent is completed the sequence begins again. These tones are harmonic sounds made up of several partials. As they step up or down the scale Shepard changed the amplitude of the individual partials of each tone. Each partial follows an amplitude curve so that it is its loudest at the midpoint of the pitch ascent/descent, and it is inaudible at the beginning and end of the ascent/descent (see Fig. 6. 1.). Because the partials are harmonically related they fuse into a single timbre. And because each partial is louder at its midpoint, it continuously draws out attention to the middle register. By really concentrating and pulling your away from this loudest midpoint you can hear the beginning and end of the scale as it fades in and out.




Figure 6.1. visualisation of endless scale (also called Shepard's scale)

Similarly our visual perception can be tricked by the Penrose staircase (see Figure 6.2.) which appears to fold back upon itself in space. This was the inspirational source for the familiar lithograph by M. C. Escher's *Ascending and descending* (http://en.wikipedia.org/wiki/Ascending_and_Descending).




Figure 6.2. Penrose staircase

While Shepard's musical example is based on the discrete steps of a scale, Risset created a continuously sliding glissando, often called as *Shepard-Risset glissando*. The partials fuse in exactly the same way creating the illusion of an endless glissando (see Fig. 6.3.).




Figure 6.3. sonogram of endless glissando (also called Risset's endless glissando)

6.2. Practical Exercises

SLApp06 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp06 Windows](#), [SLApp06 Mac OS X](#).

6.2.1. How SApp06 works




Figure 6.4. layout of SApp06

- 1 – Direction of scale/glissando
Select ascending or descending scale/glissando.
- 2 – Length of period
The period is the duration of one cycle of ascent/descent
- 3 – Lowest frequency of the spectrum – can be set between 20 and 200 Hz.
- 4 – Interval between partials
The interval between the frequencies of partials in semitones from 1 to 15.29
- 5 – Number of partials
Number of partials present (1 to 10).
- 6 – Volume
- 7 – Sound On-Off button.
- 8 – Spectrogram/Sonogram – toggle between spectrogram and sonogram display
- 9 – Spectrogram/Sonogram display
- 10 – Partial amplitude envelope presets and display
Five amplitude envelope presets can be selected for the partials.
- 11 – Amplitude envelope of fused sound (only for Shepard scale and not the glissando).
Percussive, end-peak and trapesoid envelopes can be selected.
- 12 – Time interval
Interval between tones in Shepard scale (only works in periodoc mode). It can be set between 100 and 1000 msec.

- 13 – Periodic/Continuous toggle
You can select between continuous (Risset's endless glissando) and periodic (Shepard's scale) sound.
- 14 - Preset button triggering different parameter combinations influencing Shepard scale and Risset glissando

6.2.2. Practicing strategies

There are no special exercises or tests in this SLApp. When you first launch it, you will hear an ascending Shepard scale. Notice that there are eight partials each an octave (12 semitones) apart. Try changing the different parameters to see how they affect the illusion of endless ascent. Change to preset 2 (at 14) to hear what happens if we have only 2 partials present and the lowest frequency of the spectrum is 200 Hz. Listen to the beginning and the end of the sequence. Notice that the fusion and so the illusion is lost.

To hear the same effect with Risset glissando press presets 3 then 4.

Experiment and try altering different parameters of the SLApp. Pay particular attention to the partial amplitude envelope settings and how these are crucial in influencing the fusion. Also the harmonic relationship of the partials will greatly influence whether or not we hear the partials fused in the glissando sequence. Find out the least number of partials necessary for the illusion.

7. Filtering white noise with low-pass and high-pass filters

In the following chapters you will learn about synthesis based on filtering, known as **subtractive synthesis**. This type of synthesis is the opposite of additive synthesis used in the first chapters. In additive synthesis, complex sound is created by adding simple waveforms together. Conversely, in subtractive synthesis, the starting point is a relatively complex waveform from which parts of the sound are removed by a process known as filtering. White noise (Fig.7.1, Sound7.1) is a useful starting point because it contains all frequencies, which we are going to filter in the following three SLApps.




Figure 7.1. [sonogram](#) of white noise



7.01_Sound

7.1. Theoretical background

7.1.1. Subtractive synthesis

A synthesis technique in which parts of the spectrum are subtracted by filtering. A filter is a device that performs some sort of transformation on the spectrum of a signal.

The most common types of filters are:

1) Low-pass filter (Fig. 7. 2.): passes the partials below the given cutoff frequency (or cuts the partials above it).




Figure 7.2. sonogram of low-pass filtered white noise

2) High-pass filter (Fig. 7.3.): passes the partials above a given cutoff frequency (or cuts the partials below it).




Figure 7.3. sonogram of high-pass filtered white noise

3) Band-pass filter (7.4.): passes the partials around the given centre frequency in the range of a given bandwidth.




Figure 7.4. sonogram of band-pass filtered white noise

To summarize: to filter a sound, we need to specify the following parameters:

- the cutoff frequency,
- the center frequency:
- the bandwidth (the size of range of frequency to be filtered:),

Our present SLApp is limited to low-pass and high-pass filters as an introduction to the concepts of subtractive synthesis.

7.1.2. White noise

Audible white noise is a random signal that contains sinewaves of all frequencies with constant density between 20 Hz and 20 KHz. The name 'whitenoise' originates from its visual equivalent, where white light contains all visible frequencies; red, orange, yellow, etc.

The sonogram of white noise can be seen in [Figure 7.1.](#) and heard in [7.01 Sound.](#)

You can see from Fig. 7.5. that the waveform of white noise is made of random amplitude values with no regular repetition. The sonogram shows that white noise is a saturated and well-balanced spectrum containing all frequencies. It is possible to think of white noise as a block of marble from which a sculpture can be carved. Although, one can filter or subtract parts from any sound spectrum, the most productive filtering will occur with a source sound rich in partials. Since white noise contains all possible frequencies, we chose it as starting point for our subtractive synthesis which can be thought of as sculpting in sound.

Filtering will completely change the timbre of white noise.




Figure 7.5. waveform of white noise

7.2. Practical Exercises

SLApp07 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp07 Windows](#), [SLApp07 Mac OS X](#).

7.2.1. How SApp07 works




Figure 7.6. layout of SApp07

- 1 – Selection of exercises - in the pop-down menu three exercises and three tests can be selected.
- 2 – Original: play white noise without filtering
- 3 – Filter presets – the preset buttons will select low-pass or high-pass filters at different cutoff frequencies. HP means high-pass filter, LP means low-pass filter. The numbers on the buttons indicate the cutoff frequency.
- 4 – The test interface
- 5 – Amplitude envelope – you can change the amplitude envelope with the 2 preset buttons (trapezoid, percussive) or you can create your own by clicking and dragging on the envelope panel
- 6 – Duration – you can change the duration of a sound by clicking and dragging.
- 7 – Volume slider and indicator
- 8 – Sound On-Off button.

7.2.2. Using SApp07

In SApp07 you will discover that the process of filtering white noise with simplest low-pass and high-pass filters will produce a surprising variety of sounds. Some of the sounds will sound very similar to ones we hear in the environment. These similarities can be helpful in memorizing the some of the filter parameters.

Ex01




Figure 7.7. Ex01

Listen to the different parameter settings of the high-pass filter. The numbers you see on the preset buttons tell you the cutoff frequency in Hertz. Notice that the cutoff frequency increases by a multiple of 2, this is the same as an octave interval.

Try to memorize the sound produced by each cutoff frequency.

Always start practicing with trapesoidal amplitude envelope and sound durations of 3 seconds. Then you can move on to shorter sounds and percussive envelope too.

Test01




Figure 7.8. Test01

Here you can test your ability to recognize the quality of sound produced by filtering white noise at different cutoff frequencies.

Press the button '!NEW SOUND!' to hear one of the seven presets. To answer click on one of the cutoff frequency buttons. The correct answer is indicated by the green Led. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the '!REPEAT!' button.

When you feel confident in recognising the sounds produced by the high-pass filters you can move on to the Timed Test.

Ex02




Figure 7.9. Ex02

Listen to the different parameter settings of the low-pass filter. The numbers you see on the preset buttons tell you the cutoff frequency in Hertz. Again the cutoff frequency increases by a multiple of 2, this is the same as an octave interval.

Try to memorize the sound produced by each cutoff frequency.

Test02

As Test01 just for the low-pass filter.

Ex03




Figure 7.10. Ex03

Now listen to both filter-types. Try to memorize the filter types and the cutoff frequencies whilst alternating between the two kinds of filter.

Test03

As Test01 just for both high-pass and low-pass filter..

7.2.4. Practicing strategies

To practice low-pass and high-pass filtering first concentrate on high-pass filters first then move onto low-pass ones then combine them.

When practicing the filters, keep referring back to the original white noise so that you learn to hear the difference between them. Try to find strategies of what to concentrate on. There are some tips to follow:

1) you may find vowel-like sounds produced by some of the cutoff frequencies. It happens more often with low-pass filtered sounds. Specific [formants](#) might be heard around certain cutoff frequencies. Remember that not everybody will hear the same vowels.

2. Also some consonants can be heard around the higher frequencies. The high-pass filter at 4000 Hz gives a percept of a sizzling "sh" consonant, at 8000 Hz a clear "s" can be heard. The same cutoff frequencies thru the low-pass filter will take away those consonants, but you will hear them reappear in the upper low-pass filters (4000, 8000 Hz).

Therefore high-pass filtered sounds and those low-pass filters that have higher frequencies (cutoff freq.: 4000 Hz or 8000 Hz) might be identified with consonants of our speech.

2) The very high and very low partials are easily recognizable. Frequencies around 125 Hz have a kind of gurgling sound. The higher partials around 4000 Hz are similar to sizzling, and around 8000 Hz, to hissing sounds. If you can not hear hissing in a sound

you know that frequencies of 8000 Hz and above are not present (therefore the right answer is LP 8000)

3) it is worth listening the cutoff frequency edge. After a while you can learn to hear the partials just above (with high-pass filter) or just below (with low-pass filter) the cutoff frequency

4) a good strategy is to listen to a high-pass and a low-pass filter of the same cutoff frequency repeatedly so you hear everything above or below a certain frequency. Also refer back to the full spectrum white noise. Because high-pass and low-pass filtering are subtractive processes, the filtered frequencies obviously will not be present in the filtered sound. It is important to get a sense of what is missing from your filtered sound because it will help you recognize and to analyze accurately the frequency components present in any given sound.

Working with sound at this level we do not have the reference points of pitch, melody, or rhythm, but we can use our subjective reference system. It is important that you develop your own system of categories, connotations and responses. Share your experiences because it will help create a common language for the description of timbre-based sound manipulations. Below you will find one person's subjective response to the effects of high- and low-pass filters at different cutoff frequencies. Some observations you may agree with, some you may not.

High-pass filter

- 125 Hz: the very low gurgling sounds are missing. This spectrum has a fricative 'fffff' sound
- 250 Hz: this spectrum sounds more powerful than at 125 Hz because of the missing partials around 100-200 Hz which will soften the spectrum.
- 500 Hz: the 'fffff' sound becomes stronger and higher than before and the /u/ (see [IPA](#)) vowel can be heard softly.
- 1000 Hz: a bit of sizzling can be heard with the vowel /a:/ (see [IPA](#))
- 2000 Hz: this sound is quite sharp, like the 'sh' consonant and has the /e:/ (see IPA) vowel in it.
- 4000 Hz: sizzling sound with the vowel /i/ (see IPA)
- 8000 Hz: hissing sibilant sound.

Low-pass filter

- 8000 Hz: hard 'f' sound without the very high hissing
- 4000 Hz: duller sound with 'f', 'sh' and the vowel /ε/ (see IPA)
- 2000 Hz: duller sound with the vowel /a:/ (see IPA) vowel in it like someone who cannot whistle. From here on consonants cannot be heard.
- 1000 Hz: a blunt but open sound, 'aah', with the vowel of /ɒ/ (see IPA)
- 500 Hz: blunt sound with 'oo' vowels /u/ (see IPA) and sounds under water.
- 250 Hz: this is like a dull roar under the sea.
- 125 Hz: very low gurgling subterranean sounds without consonants.

8. Filtering whitenoise with band-pass filter

8.1. Theoretical background

A band-pass filter will pass frequencies in a certain range and will reject other frequencies outside that range (Fig. 8.1.)

The band-pass filter has two main parameters

- center frequency of the filtering (f_0 in Fig. 8.1.)
- bandwidth - the size of the range to be filtered (B in Fig. 8.1.)




Figure 8.1. band-pass filter

In this chapter we will explore the band-pass filter. The filters we will use have specific ranges: the intervals of a third and of an octave. An interval is a musical term meaning the distance between two pitches. When talking about frequency we need to think of an interval as the ratio between two frequencies. This means that, for example, successive octaves result in an exponential increase of frequency, even though the human ear perceives this as a linear increase in pitch. For example, any two notes an octave apart have a frequency ratio of 2:1 (Fig. 8.2.) Notice the increasing distance between each successive octave on the frequency (horizontal) axis. The [lowest A0 on the piano](#) is 27.5 Hertz, its octave is 55 Hertz. The next A2 is 110 Hz, its octave being the 220 Hz.




Figure 8.2. octave interval in Hertz

In this chapter we will explore two fixed bandwidths, the perfect octave (2:1) and the just third (4:5). The reason for this is to hear the difference between a wide bandwidth (the octave) and a much narrower bandwidth (the third).

Bandpass-filtering white noise will result in very different sounds depending on the center frequency and the bandwidth. In some cases we hear clear pitches, in others, more pitches or even cluster-like chords are heard, and in other cases we simply hear noisebands with no pitch content. You will find that the octave (8) band-pass filter is much noisier than the third (3) filter. But the noisiness of both filters increases, the higher the center frequency.

It is important to note that even though we are discussing pitched sound versus noisiness as though they were two separate things, you will hear many states inbetween. There is no clear boundary that separates pitch and noise. We could say, band-pass filtering will add pitch to white noise: it augments the noisiness instead of replacing it. We still do not know why it is that at certain bandwidths and certain center frequencies we hear individual or grouped pitches or what causes certain pitches (the center frequency or the upper or lower cutoff frequencies or even others) to be heard. Much more research is needed to answer these questions.

8.2. Practical Exercises

SLApp08 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp08 Windows](#), [SLApp08 Mac OS X](#).

8.2.1. How SLApp08 works




Figure 8.3. layout of SLApp08

- 1 – Selection of exercises - in the pop-down menu three exercises and three tests can be selected.
- 2 – Original: play white noise without filtering
- 3 – Filter presets – the preset buttons will select low-pass or high-pass filters at different cutoff frequencies. HP means high-pass filter, LP means low-pass filter. The numbers on the buttons indicate the cutoff frequency.
- 4 – The test interface
- 5 – Amplitude envelope – you can change the amplitude envelope with the 2 preset buttons (trapezoid, percussive) or you can create your own by clicking and dragging on the envelope panel
- 6 – Duration – you can change the duration of a sound by clicking and dragging.
- 7 – Volume slider and indicator
- 8 – Sound On-Off button.

8.2.2. Using SLApp08

In SLApp08 you will explore filtering white noise with two the band-pass filters each with a fixed bandwidth.

Ex01




Figure 8.4. Ex01

Listen to the band-pass filter with the octave (8) interval starting with a center frequency of 125 Hz thru to 8000 Hz. As with the previous filters each button is an octave above the last providing clearly identifiable changes in the filtered sound.

Test01




Figure 8.5. Test01

Here you can test your ability to recognize the quality of sound produced by band-pass filtering white noise at different center frequencies with octave (8) bandwidth.

Press the button '!NEW SOUND!' to hear one of the seven presets. To answer click on one of the cutoff frequency buttons. The correct answer is indicated by the green Led. If

the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the 'REPEAT!' button.

When you feel confident in recognising the sounds produced by the band-pass filters you can move on to the Timed Test.

Ex02




Figure 8.6. Ex02

Listen to the band-pass filter with the third (3) interval starting with a center frequency of 125 Hz through to 8000 Hz. As with the previous filters each button is an octave above the last providing clearly identifiable changes in the filtered sound.

Test02

As Test01 just for the band-pass (3) filter.

Ex03




Figure 8.7. Ex03

Now listen to both filters (octave and third). Try to memorize the different filter settings.

Test03

As Test01 just for both band-pass filters.

8.2.3. Practicing strategies

Listen to the series of filtered sounds with the octave bandwidth in ascending order and do the same with the third bandwidth. Then listen to the two filters at the same center frequencies, for example compare the sound of 8-500 and 3-500 listening carefully to the differences between them.

When practicing the filters, return and listen to the original white noise regularly so that you hear and learn the difference between them. Try to find strategies of what to concentrate on. Here are some useful tips:

1) Try concentrating on the different sounds produced by the two filters at the same center frequencies. Notice that with a narrower bandwidth (third) the pitches are easy to identify especially in the lower registers. In this region with this narrow bandwidth there is clearly more pitch than noise. In some cases we hear only one pitch with complex beating.

3-125, 3-250, 3-500 and 3-1000 produce a wind-like "floating" sound. 3-2000 produces a scream-like sound with a higher vowel. 3-4000 and 3-8000 have a whistling quality and are also similar to the sound of crickets on a summer's night.

2) When the sound of the octave band-filtered noise is pitched, we never hear only one pitch, rather a group of pitches.

8-125 has a floating, gurgling sound. It's apparent pitch(es) seems to be lower than that of the 3-125 sound.

8-250 has a floating sound as if two or three frequency nodes can be heard. One node is a third higher than that of the 3-250 sound – that maybe the top edge of the filtered sound. The other nodes are a third (and maybe a fifth) below.

8-500 is similar to 8-250 but with one more element. This spectrum sounds like a chord of 3-4 pitches close to each other.

8-1000 is kind of the continuation of the two previous sounds. This spectrum is built up of 4-6 nodes and more dissonant than 8-500.

8-2000 - the top edge of the filtered spectra can be heard clearly and colors the whole sound.

8-4000 seems strangely to be lower than 3-4000 and is heard at the centre of the spectra surrounded by a noise field.

8-8000 is heard as a total noise. The typical consonant sound 'fffff' is heard in a high register without any salient pitches.

Always start practicing with trapesoidal amplitude envelope and sound durations of 3 seconds. Then you can move on to shorter sounds and percussive envelopes too.

9. Filtering whitenoise with different filters

SLApp09 combines all the filter types introduced and explored in Chapters 7 and 8: high-pass, low-pass and third and octave band-pass filtering of white noise. Use this patch to learn the similarities and differences between them.

SLApp09 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp09 Windows](#), [SLApp09 Mac OS X](#).

9.1. How SLApp09 works




Figure 9.1. Layout of SLApp09

- 1 – Original: play white noise without filtering
- 2 – Filter presets – the preset buttons will select low-pass, high-pass and band-pass filters at different frequencies and bandwidths.
- 3 – The test interface

- 4 – Amplitude envelope – you can change the amplitude envelope with the 2 preset buttons (trapesoid, percussive) or you can create your own by clicking and dragging on the envelope panel
- 5 – Duration – you can change the duration of a sound by clicink and dragging.
- 6 – Volume slider and indicator
- 7 – Sound On-Off button.

9.2. How the timed test works




Figure 9.2. structure of the timed test

- 1 – Level
Select the level: easy, medium or hard, from the pop-down window.
Easy: the test sound is played twice and you have 20 seconds to guess the filter-type.
Medium: the test sound is played twice and you have 10 seconds to guess the filter-type.
Hard: the test sound is played only once and you have only 8 seconds to guess the filter-type.
- 2 – Type of test
You can choose between high-pass, low-pass or both, band-pass filter 8, band-pass filter 3 or both band-pass filters, or all of the filter types used in the last three chapters.
- 3 – Total number of sounds in the test
Specify the number of sounds you want to hear in your test.
- 4 – Start or stop your test

- 5 – Current number of sound
This number shows which sound is being played currently.
- 6 – Time remaning
Shows the time remaining to guess the current sound. The actual sound is played once or twice within that time.
- 7 – Evaluation – filter type
This window tells you (in percentage) how accurately you identified the type (HP, LP, octave or third interval) of filter.
- 8 – Evaluation – filter frequency
This window tells you how accurately you identified the filter frequency.
- 9 – Evaluation – overall result
This window tells you how accurately you identified the filtering in total.

9.3. Using SLApp09

SLApp09 has only one interface combining Exercise and Test window. Here you can listen to all the filter setting explored in Chapter 7-8.

In this window you can also test your recognition of all the practised filter types.

In addition when you feel confiden in recognizing all the different filter settings you can try the Timed Test (8).

9.4. Practicing strategies

Since we gave you some strategies of how and what to listen to in the different types of filterings in Chapter 7 and 8, we will not explain them again. However, it is worth mentioning that combining the filter types may create new comparisons. Low-pass filtered noise with the cutoff frequencies of 125 and 250 Hz might sound similar to band-pass filtered noise with these center frequencies. High-pass filtered noise with the cutoff frequencies of 4000 and 8000 Hz might sound similar to band-pass filtered noise with these center frequencies.

Therefore one strategy is to click on the preset buttons in rows and then try clicking on them in columns. Work thru the SLApp horizontally and vertically.

10. Filtering sound samples with low-pass, high-pass and band-pass filters

In SLApp10 we use the filters you are familiar with from Chapters 7-9 with recorded sound samples instead of white noise.

There are differences between the perception of white noise and prerecorded soundfiles taken from the real world in two respects:

1. the perception of white noise does not change over time, all frequencies are always present at the same amplitude. On the other hand most sounds we hear in the real world change constantly, the amplitude and frequency of their partials is in continuous flux.
2. the spectrum of white noise has a uniform spectral density, while real world sounds have characteristic patterns in their spectra which helps us to differentiate between them.

Now we try to apply what we have learnt about filtering to these real world sounds. Because they have very characteristic qualities as a result of their highly structured spectral content, filtering will have a different effect on these sounds. What you've learnt about filtering using white noise as a source will now be colored by the spectra of prerecorded sounds.

The soundfiles we will filter were chosen on the basis of their different characteristics: cymbal, voice, machine-gun fire, piano and viola da gamba. They have very different spectra and rhythms. As you will hear, the variety of rhythm, pitch and timbre changes will contribute to the outcome of the filtering process.

With soundfiles, although the filter will work the same way as with white noise, the resulting sound will still retain some of the characteristics of the soundfile, as you will find when you begin to listen to the filtered soundfiles.

The reason why all filter parameter settings will work with white noise is that it contains all frequencies. With soundfiles the result will depend to a large extent on the spectrum of the soundfile itself. Imagine filtering one sinewave at 2000 Hz with a high-pass filter. Depending on where you set the cutoff frequency you will either hear the sinewave or nothing. It makes no difference if the cutoff frequency is at 500 Hz, a 1000 Hz or 1900 Hz, you will still hear the sinewave. In filtering soundfiles the process is much more interactive than filtering white noise. You always have to bear in mind the spectrum of the soundfile when choosing filter parameters.

10.1. Practical Exercises

SLApp10 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp10 Windows](#), [SLApp10 Mac OS X](#).

10.1.1 How SLApp10 works




Figure 10.1. layout of SLApp10

- 1 – Selection of the sound source (original soundfile)
In the pop-down menu, five soundfiles (Piatti, Voice, Machine_gun, Piano) can be selected. These serve as sound sources for filtering.
- 2 – Original: play soundfile without filtering
- 3 – Filter presets – the preset buttons will select low-pass, high-pass and band-pass filters at different frequencies and bandwidths.
- 4 – The test interface
- 5 – Volume slider and indicator
- 6 – Sound On-Off button.

10.1.2. Using SLApp10

SLApp10 has only one interface combining Exercise and Test window. So listen to each of the soundfiles through each of the filters.

To test yourself first choose one of the soundfiles. Then press the button 'NEW SOUND!' to hear one of the 28 presets. To answer click on one of preset buttons. The correct answer is indicated by the green Led. If the answer is incorrect, it will be red and you can guess again. If you need to hear the sound again, press the 'REPEAT!' button.

When you feel confident in recognising the sounds produced by the filters you can move on to the Timed Test (see description in Chapter 9 at 9.2.).

10.2. Listening strategies

10.2.1. Spectral analysis of the original soundfiles

White noise - for reference




Figure 10.2. sonogram white noise

Machine gun




Figure 10.3. sonogram of machine gun



10.01_Sound

Listen and filter this soundfile first. Compare the filtered sound of the machine gun with that of filtered white noise.

As you can see from the sonogram this sound is less dense than that of white noise (Fig. 10.2). You can clearly see it has a regular rhythm and is not continuous. Also there are variations vertically. However, the spectrum and the sound of the machine gun still

resembles white noise. Each shot has a dense spectrum, the partials being close to each other.

You see some darker areas in the lower and middle part of the spectrum. These are coming not from the shot itself but from the metallic body of the gun as it resonates.

The action of filtering this sound will be similar to that of filtering white noise but with a clearly audible rhythmical shape with those additional resonances.

Piatti – a pair of cymbals




Figure 10.4. sonogram of piatti



10.02_Sound

The the sound of a cymbal is a perfect example of an inharmonic spectrum. It also has a dense spectrum, but with more recognizable frequency structure.

At its very beginning when it is struck, as you can see at the left edge of the sonogram, the whole range of frequencies is produced, but you can see, as the sound progresses from left to right, clear darker bands which you hear as pitches. Notice how the higher, noisier frequencies fade quickly. As the frequencies get lower their decay is slower. Notice how the very lowest frequencies last the longest.

Filtering will very clearly reveal the different frequency regions in the sound of the cymbal.

Voice – speech in Hungarian




Figure 10.5. sonogram of human voice (Hungarian)



10.03_Sound

In this sonogram analysis you can see the human speaking voice. The rhythm of the words is clearly visible as are the different spectral territories of the voice.

The sibilant consonants (such as s, c and f) cover a wide, dense (noisy) spectrum with most of their energy being in the high frequency range. The vowels are harmonic sounds. You can see this in the groups of dark equidistant lines. The curves represent the melody of the speech. The darker areas represent the [formants](#) which determine the vowel we hear (such as a, o, e, i)

Because the human speaking voice is spectrally rich, containing both noise and harmonic sound, the filtering possibilities are great.

Piano – acoustic piano




Figure 10.6. sonogram of piano



10.04_Sound

Notice how the sonogram of the piano is much more coherent than the previous examples. The spectrum of the piano is essentially harmonic. The parallel horizontal lines clearly represent the partials making up a range of pitches. What you see are six chords being played on the piano keyboard. Notice how the beginning of each chord is slightly darker and noiser, this is caused by the hammers hitting the strings. The sound is at its loudest here. Similarly, as we saw on the cymbal sonogram, the higher partials of the piano fade quicker than the lower.

Filtering will allow you to explore inside each chord creating different pitch structures. Some parameter settings will soften, others brighten the sound. You will be able to turn the piano into a tambourine.

10.2.2. Practicing strategies

Here are some listening notes made whilst exploring SLApp10. Some of the comments are analytical and some more subjective. This is how we develop our understanding of complex and varied sound spectra.

Begin to listen to the sounds from the world around you and analyze and think about them in the same way.

Machine gun

High-pass filter - the source is still recognizable. It rather emphasizes the lower region of the actual sound.

Low-pass filter changes the timbre similarly to that of white noise. The important difference is that the attack of each shot is mainly made up of the higher frequencies. Therefore with LP1000 and below the rhythmic clarity disappears and the sound becomes more blurred. This also happens with octave and third interval filters at a 1000 Hz and below.

Piatti

High-pass filter:

- HP125, 250, 500 and 1000 the special low resonancies disappear without destroying the timbral identity of the cymbal.
- HP4000 and above you just hear the noisy components of the sound, it is similar to filtered white noise at the same frequencies.

Low-pass filter:

- LP500 and above the timbral characteristics (low resonancies) are present, the higher consonant-like partials (h, f, s) disappear. It sounds much more like a gong.
- LP125 and 250 low resonating sinewave-like partials are revealed (which never hear in the whole spectrum of the cymbal).

Band-pass filter:

- BP125 both octave and third filtering will reveal a low beating sound. Notice that the beating was not present at LP125. This is because the partials are very close together.
- higher bands (above 1000 Hz) work very similarly to filtered white noise: the third interval will create more pitched sounds whilst the octave filter reveals dense clusters and noisy sounds

Voice

High-pass filter:

as we progress thru the high-pass filters from low to high we lose the intelligibility of the speech. Between 250 and a 1000 it sounds as if the voice is coming from smaller and smaller radios and then a telephone. At HP2000 the vowels get distorted as if the speaker talks through gritted teeth. At 4000 and 8000 all we hear are higher frequency sibilant sounds but with a recognizable rhythm of speech.

Low-pass filter:

- LP125-250 low rhythm, no vowels, no consonants
- LP500 - reconizable as speech but not intelligible. The low vowels such as (a, o, u) start to emerge.
- LP1000 - we hear all the vowels, but no consonants as if the speaker had no teeth
- LP4000 and above - speech is intelligible but some higher frequencies are missing, LP8000 is almost like the original.

Band-pass filters:

- 8-125 and 250 are similar to LP125 and 250.
- 8-500 - 2000 - small radios, badly tuned. Consonants emerge at around 1000-2000. Speech becomes intelligible at 2000.
- 8-4000 - the gritted teeth sound again
- 8-8000 - is all sibilant
- 3-125 - 250 - muffled rhythmic sound
- 3-500 - highly pitched not recognizable as male voice
- 3-1000-8000 - highly pitched becoming noisy and not really intelligible however the low fundamental frequency (around 100 Hz) of the speech is recognizable although it's been filtered out. This is the phenomenon of the [missing fundamental](#).

Piano

High-pass filter:

- HP250-500: sounds thinner as though from a small radio
- HP1000 and above: the filtered sounds become metallic and the melody becomes lost
- HP8000: no recognizable piano sound, it is almost like a rhythm played on a small cymbal.

Low-pass filter:

- LP125: now the rhythm is played on a huge gong.
- LP250-500: pitch variations just start to emerge, but the piano is hardly recognizable.
- LP1000-4000: both piano and chord progression are recognizable just lacking in high frequencies to different degrees as though the piano were in another room.

Band-pass filter:

- 8-125-250: similar to LP125 and 250 but with narrower pitch content. Not recognizable as piano, no melody.
- 8-500-1000: recognizable piano played somewhere down a long corridor, the sound getting thinner.
- 8-2000-8000: the sound becomes purely rhythmic and more metallic.
- 3-125: gong-like, low sound, where the attacks of the chords are completely softened
- 3-250: attacks are little harder
- 3-500: pitch content emerges with coloration of the 'o' vowel
- 3-1000: pitch content start to disappear, the sound has a nasal quality
- 3-2000: more metallic, just rhythm
- 3-4000: rhythm being played on the bell of a cymbal
- 3-8000: rhythm played on the edge of a small cymbal

11. Filtering with an additive resonant filter

11.1. Theoretical background

All objects have a natural resonant frequency. Tapping a wine glass for example will produce long resonating, high sounds. Tapping a lump of dough will resonate much less.

Resonance is common throughout physics, mechanics, electricity and magnetism. It is the tendency to vibrate or oscillate with higher amplitudes at certain frequencies. Resonance is the property which the human voice and all musical instruments exploit it. Here it is desired. In other situations it is unwanted and even disastrous as in the collapse of the Tacoma bridge (see video 11.1).

In our chapter we will model resonance with a help of a filter. You will have experienced using band-pass filters that can emphasize certain groups of frequencies, most noticeably the third interval band-pass filters. With a resonant filter we narrow this interval even further and through introducing a feedback loop, create a ringing sound. As we narrow the bandwidth of the filter the ringing or the resonance will become longer.

Our basic resonant filter needs three parameters:

- center frequency
- bandwidth: the range of frequencies affected by the filter (in Hz) above and below the centre frequency.
- amplitude of the filtered frequency band.

In our present test we use white noise as a source spectrum with 9 resonant filters. Here we just use a percussive amplitude envelope, which can not be changed. We are familiar with white noise from Chapter 7-9 where we used one filter at a time. Here in this SLApp we will use 9 filters to subtract frequencies from white noise, and add their outputs together. So in this way we combine additive and subtractive synthesis.

11.2. Practical Exercises

SLApp11 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp11 Windows](#), [SLApp11 Mac OS X](#).

11.2.1. How SApp11 works




Figure 11.1. layout of SApp11

- 1 – Play the sound
You can change all three the parameters of each filters. Click on the "Play" button to hear your sound.
- 2 – Fundamental frequency
- 3 – Panel of resonant filter parameters
Each filter has 3 parameters. The top number is the frequency ratio of the fundamental frequency (2). If the fundamental frequency is set to 500 Hz and Filter No2 is set 2, the frequency of Filter No2 is at 1000 Hz.
The second number in the coloumn is the gain or loudness.
The third parameter is the bandwidth (in Hertz).
- 4 – Duration of triggered sounds
- 5 – Amplitude envelope of triggered sounds.
In the pop-down menu two types of amplitude envelopes can be selected: percussive and triangle.
- 6 – Selection of the exercisces
In the pop-down menu two exercisces and two tests can be selected.
- 7 – Record the sound
Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.
- 8 – Volume slider and indicator
- 9 – Sound On-Off button
- 10 – Matrix of preset buttons
Clicking along the rows you select different combinations of filter gains, clicking up or down the coloumn you change the bandwidth. As you click on the presets you will see the values change in the panel of resonant filter parameters (3).
- 11 – Amplitudes of filtered bands

Nine sliders represent the amplitude of the filtered bands. These can be changed by clicking and dragging. As you change the sliders you will see the values change in the panel of resonant filter parameters.

11.2.2. Using SLApp11

In SLApp11 you will discover that we will completely transform white noise by using resonant filters. To create complex, interesting sounds you need to change a large number of parameters simultaneously. In a way the purpose of this SLApp is to illustrate this.

Ex01




Figure 11.2. Ex01

In Ex01 resonant filterbands are added together somewhat similarly to Chapter01, where we explored how to identify individual sinusoidal components of a complex spectrum. Here the bandwidth of the components – the nine filtered bands – can be also modified. In Ex01 clicking on button one will play the lowest band, it is white noise filtered around the fundamental frequency. Buttons 2 - 9 will add each consecutive partial. Clicking up or down the columns you change the bandwidth. Listen and try to memorize 9x5=45 spectra.

Explore different frequency ratios saved at preset panel above the frequency ratios. Experience also with percussive and triangle amplitude envelopes and different lengths.

Test01




Figure 11.3. Test01

Here you can test your hearing of spectra composed from resonant bands of different bandwidths. Press '!NEW SOUND!' then guess the corresponding filter. In case your answer was correct you will see X in the box. (You can always repeat the original sound with the '!REPEAT!' button.) If the test is too difficult you can make visible the filter parameters pressing 'Show visuals!' button. Using the 'Hide visuals!' button you can hear the sounds without seeing their parameters. After you practiced, tested yourself and feel confident in indentifying the combinations of filtered bands, you can move on to the timed test by clicking on the button.

Ex02




Figure 11.4. Ex02

In this exercise we will explore how to create resonant areas in the spectrum. Two neighbors of the nine filtered bands are more resonant (having narrower bandwidth) and louder than the others. Clicking along the rows you select the positions of the two strongest filtered bands, clicking up or down the columns you change the bandwidth of those two bands. Listen and try to memorize $8 \times 3 = 24$ spectra.

Explore different frequency ratios saved at preset panel above the frequency ratios. Experience also with percussive and triangle amplitude envelopes and different lengths.

Test02

As Test01 just for the position and bandwidth of the two strongest filtered bands.

12. Distortion

12.1. Theoretical background

Here we discuss distortion, a technique to modify existing partials and create new components in the spectrum. Under normal circumstances care is taken to prevent it coloring sound (e.g. studio recording, PA systems). Though of course it is intrinsic to the sound of the electric guitar. Distortion has become part of the vocabulary of electronic music as a creative tool.

Distortion is created when parts of the amplitude of a wave exceed the amplitude threshold. Therefore we can call distortion a waveshaping technique.

Three types of distortion of sinewaves are illustrated in the diagrams in Figs. 12.1., 12.2., 12.3. If we distort a pure sinewave the result will be a harmonic sound with multiple integers of the input sine-wave. Distorting a complex sound produces sums and differences of the components. This phenomenon is called intermodulation distortion.

12.1.1. Clipping

Clipping is the most common distortion method also known as overdrive. It occurs when the preset threshold is exceeded (indicated by the red line in the diagram) and is replaced with threshold itself. Literally, the top and bottom parts of the waveform are "clipped". This adds higher partial to our sinewave.




Figure 12.1. sine-wave distorted with clipping

You can listen to a sinewave distorted with clipping.



12.01_Sound

12.1.2. Folding

Folding is an overdrive method of distortion where amplitudes exceeding the preset threshold will be mirrored below the threshold line. The resulting waveform will be coarser, because folding is a stronger distortion method than clipping.

You can see the spectrogram of a sine-wave distorted with folding.




Figure 12.2. sine-wave distorted with folding

You can listen to a sinewave distorted with folding.



12.02_Sound

12.1.3. Wrapping

Wrapping is another overdrive method of distortion where amplitudes exceeding the preset threshold will be shifted up or down to guarantee that the signal wouldn't exceed the threshold. The threshold value is subtracted or added to those parts of the signal which exceed its value. This results in a jagged, serrated waveform which sounds much more distorted than clipping and folding.




Figure 12.3. sine-wave distorted with wrapping

You can listen to a sinewave distorted with wrapping.



12.03_Sound

12.2. Practical Exercises

12.2.1. Structure of the patch




Figure 12.4. structure of the patch

Since the basic structure of the patch is identical to that of Chapter 11, here we describe only the parts of the patch that differ. This patch is built up with resonance filters (see details in Chapter 11).

- 1 – Filter parameters of resonant filters No1-No9 (similar to resonant filter presented in Chapter 11)
Each filter has 3 parameters. The top number is the frequency. For filter No1 you set the frequency in Hertz. For filters No2-9 the number is the ratio of that frequency. If filter No1 is set to 500 and Filter No2 is set 2, the frequency of Filter No2 is at 1000 Hz.
The second number in the column is the gain or loudness.
The third parameter is the bandwidth (in Hertz).
Filters No2-7 also have two small toggles for the the bandwidth. If these toggles are checked (so you see an 'X' in the box) when you next change the values of Filter No1, these values will change also at the selected Filters.
- 2 – Random ranges specified for the filter parameters
By dragging or entering numbers the ranges of the filter paramters can be set. The triggered sounds (or sound sequences) will have random durations, frequencies and bandwidths within the specified ranges.
- 3 – Panel for triggering distorted sounds or sequences
The buttons trigger filterd sounds processed with the different distortion-types and threshold amplitudes. The three rows represent the three distortion types, clip, fold and wrap. The seven columns represent seven distortion thresholds.
- 4 – Playback mode

In the menu you can select from two possibilities: playback of one sound or of a sequence of sounds. When clicking on the buttons representing distortion thresholds, you will hear one sound or a sequence.

- 5 – Panel of the parameters of sound sequences
The number of consecutive sounds in a sequence and their minimum and maximum tempo (in BPM) can be specified here.
- 6 – Record the sound or sequence
Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.
- 7 – Sound On-Off button.
- 8 – Volume slider and indicator

12.2.2. Using SLApp12

SLApp12 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp12 Windows](#), [SLApp12 Mac OS X](#).

There are no individual examples in this SLApp since each waveform reacts to distortion differently. Therefore tendencies to be structured into scales were not possible. You can practice the three distortion-types with different threshold levels comparing different filtered sounds saved in the preset panel (2). Listening to the distortion parameters with different sound sequences gives an idea about using distortion in musical situations.

Try the three types of distortion; clipping, folding and wrapping for each threshold level. The threshold levels are set between 0 and 1.0 where 1.0 is the peak level of the input sound. The lower the value (0.7, 0.55, 0.4, 0.2, and 0.05 etc.) of the threshold the more distorted the original input sound becomes.

12.2.3. Practicing strategies

Practicing with the preset parameters you should hear increasing distortion resulting from clipping, folding and wrapping. It is typical that the latter one becomes very noisy. Low threshold warping can even destroy the input sound while the clipping still creates a harmonic sound.

It is worth listening for the difference between the same distortion applied to different fundamental frequencies. Higher frequency input sounds are more susceptible to distortion. This means that the same distortion threshold level applied to a higher frequency input sound will result in more distortion than it would if applied to a lower frequency input sound.

Notice that the end of each sound is similar to its original spectra. This is because of the shape of the percussive amplitude envelope, high at the beginning, low towards the end.

After studying the preset values, you can change them (length and fundamental frequency) and listen to different possibilities. Changing the general bandwidth and/or the values of the resonance filter you can create completely different sounds. Do experiment with this patch!

13. Granulation of soundfiles

13.1. Theoretical background

Granular synthesis is a technique, where the sound is built from short fragments of sounds called grains. The grains can be synthesized or cut from a prerecorded soundfile.

In this we look at using soundfiles as our source sound. Once we chop sounds into tiny grains we can treat them in many different ways to produce sounds far removed from the original source.

One grain is represented by its frequency, duration and amplitude envelope.

On one hand, granular synthesis is very similar to additive synthesis since building up a sound event can be achieved by putting grains with different durations after each other. Also it requires a large amount of data and complicated controlling. Therefore it can be quite difficult to create even a 2000 msec long sound of tiny grains. On the other hand it is an effective synthesis technique capable of creating new timbres and sound events if we can control it by predetermined parameters.

The main parameters of the synthesis technique are:

- the waveform of the grain: this will be totally determined by the source soundfile
- length of the grain (in ms): the length of grain will determine whether we hear the timbre of the source sound or the process of granulation in the form of noisiness.

if the grain is:

- shorter than 2 ms: we can not hear the original sound, we hear noisy clicks
- at about 5 ms: the percept of pitch starts to emerge
- above 25 ms: we hear the pitch and timbre of the source sound.

Between these grainlengths there are different transitional phases which contain combinations of noise and the pitch and timbre recognizable from the source.

- grainspeed: the rate at which one grain follows another. Certain grainspeeds will influence the sound according to their frequencies. Very high speeds will add a pitch to the overall sound. Slowing down the grainspeed significantly (e.g. 8 Hz and below) will transform the sound from a continuous stream to a succession of separate entities.
- position in the soundfile (ms): the grain can be taken from anywhere in the source sound. If the grain is taken from one point, this part of the soundfile will be just repeated resulting in a frozen kind of sound. Taking each successive grain

from random points will result in a rapidly changing sound provided, the source sound changes over time.

- transposition (there are two types)
 1. general transposition where all grains are transposed up or down by an equal amount
 2. random transposition: where all grains are transposed randomly within a specified range

There are other parameters, but we will not use them in the current chapter.

13.2. Practical Exercises

SLApp13 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp13 Windows](#), [SLApp13 Mac OS X](#).

13.2.1. How SLApp13 works




Figure 13.1. layout of SLApp13 (Ex01)

- 1 – Selection of excercises: one exercise, one test
- 2 – Selection of soundfiles
select sound files for granulation including sounds of fire, didgeridoo, chain, opera gong, voice
- 3 – Play original soundfile for comparison with granulation
- 4 – Volume slider and indicator
- 5 – Sound On-Off button.
- 6 – Recording.
- 7 – Thirty granulation presets.

In Ex01 the presets are visualized with a graphic, in Test01 they are represented with normal buttons. The figures on the buttons in Test01 are transposition and grain length parameters.

13.2.2. Using SLApp13

In SLApp13 we apply granular synthesis to different sound sources. We will concentrate on the effect of changing two parameters: grainlength and maximum random transposition.

Ex01

In this exercise the thirty presets are arranged in a matrix of six columns and five rows. Rows represent different levels of random transposition, where the maximum amount of transposition is specified. In the top row there is no transposition (0). In the bottom row there is a maximum of 2 octaves of transposition.

If you are looking at the grid as we move from left to right we have increasing grainlength with 5 ms on the left and 200 ms on the right. The top row has 0 transposition applied the bottom row has a maximum of two octaves applied. So, pressing on the preset in the top left corner will chop the soundfile into 5 ms grains and apply no transposition. Pressing the button in the bottom right hand corner will chop the soundfile into 200 msec grainlength and transpose it to a maximum of two octaves.

Test01




Figure 13.2. Test01

Here you can test your ability to recognize the transformations of sound which occur thru granular synthesis. Notice in this test the interface changes. The visualisation of the grains is replaced by the actual parameters so 0.00/5 is no transposition and 5 ms grainlength exactly as in Ex01 interface.

Press the button 'NEW SOUND!' to hear one of the thirty presets. To answer click on one of preset buttons. The correct answer is indicated by the green Led. If the answer is

incorrect, it will be red and you can guess again. If you need to hear the sound again, press the 'REPEAT!' button.

Try testing yourself with different soundfiles.

When you feel confident in recognising the sounds produced by the different granulation parameters you can move on to the Timed Test (see Chapter9 at 9.2).

13.2.3. Practicing strategies

Select the Voice (2) and play it at two presets, top left and bottom right. What do you notice about the two sounds you hear? Although they are produced from the same soundfile, they sound radically different. At 5 ms and no transposition the sound is not recognizable as speech. You will hear a clear pitch not present in the original sound. This 'artefact' is created because of the very fast rate. You hear grains of 5 ms duration being played one after the other in rapid succession. It's this that give the pitch. This is present in all presets of 5 ms grainlength. At 200 ms and 2 octave maximum transposition we can clearly hear the human voice because of the longer grainlength. We can also hear the effects of the transposition where each successive grain is randomly transposed from non to two octaves. So sometimes we hear fragments of the voice at its original pitch.

Now listen to the preset in the top righthand corner, still 200 ms grainlength but this time no transposition. You can clearly hear the voice at its original pitch, the grains are long enough to allow for recognition of parts of words.

Now play the preset in the bottom left corner. Again the 5 second grainlength imposes a pitch, but we also hear more pitches caused by the transposition.

Now explore the presets. Click along the rows so the transposition stays the same and the grainlength changes. Then click up and down the columns so the transposition changes whilst the grainlength stays the same. Do this for each soundfile and notice how the nature of the source sound will interact with the granulation.

14. Localization – positioning sound across the left-right axes

Sound waves spread out in every direction from their point of origin. We are able to sense the position of a sound's source because we have two ears. The capability of two ears to locate sound sources within an acoustic space is known as spatial or binaural localization. A pair of normally functioning ears can easily localize a source to within a few degrees in the horizontal plane.

Through binaural hearing, humans are able to identify the position of a sound in three dimensions: Left-Right, Front-Rear, Up-Down. As well as perceiving a sound's location, we can also gather indispensable information about the nature (size and quality) of the space itself.

The present chapter concentrates only on the Left-Right axis in order to develop understanding and skills involved in localization.

14.1. Theoretical background

The position of a sound source on the Left-Right axis can be identified by its relation to two points: our two ears. There are three acoustic cues received by the ears which help us locate a sound:

1) Interaural intensity differences

If the sound source is precisely in front of the listener (S1 at Fig. 14.1), the direct sound will arrive at each ear at exactly the same time with the same amplitude. When it comes from one side (S2 at Fig. 14.1) the soundwaves will arrive at the ears with different intensity. Since one ear (the left in our present example) is obscured by the head, the waves will arrive at the ear only as reflected sounds. After losing most of its energy it will be quieter. This interaural intensity difference is used by the brain to locate the sound source.

The interaural intensity difference is more significant in the case of middle and higher frequencies (i.e. above 700 Hz). This is because their wavelengths are shorter. The longer wavelengths of lower frequencies are not effected by the "shadow" of the head therefore we locate lower frequency sounds more with the help of the interaural arrival-time difference.




Figure 14.1. sound sources in different horizontal positions.

2) Interaural arrival-time differences

In Figure 14.1. it can be seen that the distance between S2 (sound source 2) and the left ear is longer than the distance to the right one. The difference in distances is used for localization in the case of lower frequencies because the sound will arrive later at the ear furthest away from the source. It is important to note that it is the combination of interaural arrival-time and interaural intensity differences that allows us to localise sound across the whole frequency range.

3) The effects of the pinnae (outer ears)

Because our ears are situated more or less symmetrically on either side of the head, they can localize sound in the horizontal plane in front of the face. The sounds located 130° from the central axis (i.e. partly behind the ear) are localized by the pinna. Because of the detailed structure of the ridges of the pinna, the waves are reflected and delayed before entering the ear. The delayed sounds are then combined with the direct sound, producing different colorations, which further helps the brain in locating the sound source.

The spatial position of a sound behind or above the listener is identified through its timbral change. The timbre is modified by the head, the nose, the shape of the ears and the hair of the listener.

The brain determines the horizontal nearness and vertical location of sound through the differences in intensity and timbre arriving at each ear, although in this chapter we will only focus on horizontal, left-right localization.

Stereophonic sound systems were developed to create the illusion of directionality and audible perspective to emulate natural hearing. Originally the term referred to surround systems as well, but nowadays it means 2 channel arrangements.

The stereo amplification system is a configuration of 2 loudspeakers, where the ideal listening position for a stereo audio signal is the so called "sweet spot". The sweet spot is the focal point of the loudspeakers (or headphones). The sweet point and the two sound sources (loudspeakers or acoustic sound sources) must ideally describe an equilateral triangle (see Fig. 14.2.).




Figure 14.2. ideal stereo listening – sweet point with two sound sources

Stereo sound systems can be divided into two forms: true or natural stereo and artificial or pan-pot stereo. In the case of true stereo, a live sound is captured by an array of microphones, so the sound image contains the direct and all the reflected sounds dispersed in the space. Artificial or pan-pot stereo is when a single-channel (mono) sound is moved between the left and right loudspeakers. Pan-pot stereo is based on interaural intensity differences, and the relative amplitude of the two channels is varied by a device called a pan-pot (panoramic potentiometer).

If the intensity is equal between the two loudspeakers we hear it as one sound source directly in front of us and not as two separate sounds. By changing the intensity ratio between speakers we perceive the sound moving on one side. This technique exploits the interaural intensity difference phenomenon described earlier to create the illusion of moving sound.

14.2. Practical Exercises

SLApp14 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp14 Windows](#), [SLApp14 Mac OS X](#).

14.2.1. How SApp14 works




Figure 14.3. Layout of Ex01

- 1 – Sound On-Off button.
- 2 – Selection of excercises: four exercises, four tests
- 3 – Selection of soundfiles
- 4 – Selection of soundfile playback mode
In the pop-down menu you can select whether you want to hear one soundfile repeted or a random selection.
- 5 – Preset buttons to trigger selected sound positions
- 6 – Sound position indicator
- Test interface
- 7 – Test
- 8 – Show test
- 9 – Hide test
- 10 – Test interface

Here there are two ways of testing your recognition of sound localization. The first you are familiar with. Pressing the !New Sound! button, listening the sound, pressing repeat if you want to hear again and than clicking on a location, Left100, etc. You can make multiple guesses. In the new Test you only have one chance to answer each time. Begin the Test by clicking "Start", then click "New Sound". Listen to it (you can click Repeat if you need to hear it again), that click on the location you think is correct. If it is correct the first led in the row of ten will be green if not, red (when the test is very hard the led will become yellow when your guess is close to the correct answer.) Then click new sound to move on to the next. After answering ten times illuminating all the leds you will see the percentage of correct answers.

14.2.2. Using SLApp14

SLApp14 is designed to explore localization of sound. We start with learning to recognize fixed monophonic sounds in fixed positions from left to right (Ex01). In Ex02 we pan the same monophonic sounds back and forth across different widths of the stereo field slowly. In Ex03 we use a new noisy the sound and pan it back and forth much quicker to imitate the experience of a true stereo field. In Ex04 we take a number of stereo sounds (some environmental field recordings, some electronic music compositions) and change the width of the stereophonic field.

We suggest listening to the sounds both on headphones and different quality loudspeakers.

Ex01




Figure 14.4. Ex01

Listen to the fixed spatial positions of your selected or randomly played soundfiles by clicking on the grey numbered buttons.

Test01.1-01.3




Figure 14.5. Test01.1




Figure 14.6. Test01.2




Figure 14.7. Test01.3

The three tests are essentially the same but with an increasing resolution of positions within the stereo field. Test01.1 has five positions and Test01.3 has seventeen. Specific details of the test can be found above at 14.2.1.

Ex02




Figure 14.8. Ex02

Ex02

Here we explore slow panning monophonic sounds. This time to trigger sounds you need to click on the brown bars. These indicate the width of the panning.

Test02 - use as test 01




Figure 14.9. Test02

Ex03




Figure 14.10. Ex03

Here we explore the rapid panning of a single monophonic sounds. Again you trigger the sounds by clicking on the brown bars. These indicate the width of the panning.

Test03 - use as Test01





Figure 14.11. Test03

Ex04

*Figure 14.12. Ex04*

This time we work with static stereo soundfiles. Again you trigger the sounds by clicking on the brown bars which indicate the widths and positions of the sounds.

Test04 - use as Test01




Figure 14.13. Test04

15. The multidimensional nature of timbre – interpolations between different states of band filtering

15.1. Theoretical background

In traditional music when thinking about melody, it's possible to identify two dimensions (the so called primary musical parameters), pitch structure and rhythm. For melody to exist change has to occur in both dimensions over time. We naturally perceive melody as a whole, but to have an understanding of how it works we need to consider the details of what changes along these two dimensions. Separating melody into the dimensions of pitch and rhythm requires skills which are developed thru conventional music education or training.

What we are concerned with understanding in this series of lessons is not melody but the nature of the sound itself, its quality, its texture, its timbre. In the natural environment, in electronic music as in melody, change is essential. To create timbre-based music we have to change parameters in many more dimensions than just the two of melody-based music. In order to do this we have to identify these dimensions. Normally we hear timbre as a whole just as we hear melody as whole. Analyzing melody in this way has a long history and tradition. Analyzing timbre in this way has very short history, which could be said to have begun with [Helmholtz](#). The development of computers in the second half of the 20th century has enabled much more research and experiment revealing the microscopic details that cause changes within sound. With this knowledge we can begin to understand and manipulate these previously hidden dimensions and create sound which is dynamically organic. The first steps of understanding are to hear, identify and name these dimensions.

In the previous chapters we have explored how sound can be built from sine waves (additive synthesis) how it can be shaped from white noise (subtractive synthesis). For the purposes of understanding the basic principles, the sounds we created were static, the parameters we used didn't change over time. But to understand and create "living" sound we have explore these parameters dynamically.

In the present chapter we return to using filtered white noise. Already you have noticed that we can produce a surprisingly varied range of sounds thru the simple process of filtering. What we are going to explore is essentially moving from one filtered state to another in a process known as interpolation.

This time we are using eight bandpass filters simultaneously. Each filter has three parameters:

- center frequency
- bandwidth

the bandwidth is controlled here with a new parameter known as Q (Quality factor). Without going into complex technical details here you need to know, that when the value of Q is high, the bandwidth will be narrower (so the pitch content of the timbre will increase) and vice versa, if Q is smaller, the bandwidth will be wider (so the spectrum will be noiser).

- amplitudes of the filtered bands.

Of these three parameters, two will be changed over time: the amplitudes and the Q of the bands. The center frequencies of the 8 bands remain constant. Their position can be changed to try out different source sounds. By interpolating between different values of amplitudes and bandwidths we can alter timbre dimensions like brightness, pitch-positions of the bands, noisiness, harmonicity, and tonality.

With timbre it is not always possible to identify which parameter is causing which change (unlike pitch, where we can say the sound we hear is higher or lower than the previous sound, because pitch is one-dimensional). Hopefully the graphic representations will help you to see and hear which parameters are causing the changes in the sound.

15.2. Practical Exercises

SLApp15 is downloadable for Windows and Mac OS X platforms using the following links: [SLApp15 Windows](#), [SLApp15 Mac OS X](#).

15.2.1. How SLApp15 works




Figure 15.1. Layout of SLApp15 (Ex01)

- 1 – Selection of exercises - in the pop-down menu one exercise and one test can be selected.
- 2 – Play button1 - this button triggers the sound without interpolation.
- 3 – Fundamental frequency
The fundamental frequency of the spectra can be selected from the keyboard and the number box.
- 4 – Duration
- 5 – Envelope
You can select here between a continuous sound (longnote) and a series of short repeated sounds (impulses)
- 6 – Partial ratios
The partial ratios of the eight filter bands will multiply the fundamental frequency to specify the center frequencies of the filter bands. There are two presets triggering a harmonic and an inharmonic source spectrum.
- 7 – Parameters of the eight band-pass filters.

- The center frequency, the Q value and the amplitude of each filter can be set here.
- 8 – Q and Amp presets
Q presets will load different saved Q values for all eight band-pass filters
Amp presets will load different saved amplitude values for all eight band-pass filters
 - 9 – Play button2 - this button triggers the sound with interpolation.
 - 10 – Amplitude interpolation display - shows the interpolation between the different amplitude settings.
 - 11 – Q interpolation display - shows the interpolation between the different Q values.
 - 12 – Amplitude interpolation presets
Each preset button shows you the initial and the final state. Click to select the chosen preset. Pressing "Play button2" will trigger the interpolation.
 - 13 – Q interpolation presets
Each preset button shows you the interpolation curve. Click to select the chosen preset. Pressing "Play button2" will trigger the interpolation.
 - 14 – Recording
Pressing the "open" button, naming the file and clicking on the record button you can record the created sounds. Afterwards stop the recording by pressing the recording button again.
 - 15 – Volume slider and indicator
 - 16 – Sound On-Off button.

15.2.2. Using SLApp15

In SLApp15 you can explore much more complex sounds using an eight band-pass filter where you can change the fundamental frequency, the partial ratios, the bandwidth of each filter and then apply interpolations to these settings.

Ex01 (Fig.15.1.)

First play with the top part of the interface. To hear the results click on "Play button1" (2). Try the different presets to hear the range of sounds that can be created thru simultaneously filtering whitenoise with eight band-pass filters. Try entering your own values for the eight filters to create as many sounds as possible. Note which parameters cause different timbres.

When you have explored, and got a sense of what the "instrument", the subtractive synthesizer can do, move on to the lower part of SLApp15 and experiment with interpolation. First select one of the amplitude interpolation presets. The buttons show you the initial and end states. Similarly select a Q interpolation preset. It is curve showing how the bandwidth of the filtering will change. Note that if the values of the curve are high at the beginning and low at the end, this will mean that a narrow, pitched sound will change into a wideband noisy sound.

Test01




Figure 15.2. Test01

Here you can test your ability to recognize the different interpolation settings and curves.

Press the button 'NEW SOUND!' to hear one of the many combinations of the eighteen presets. To answer, click on two preset buttons, one for amplitude and one for Q. Then click on the button "Play my sound" to listen to your guess. There are three led buttons to indicate if your answer was correct: If you have chosen the correct amplitude interpolation, the lefthand small led will be green. If you have chosen the correct Q curve, the righthand small led will be green. If both your answers are correct the middle led will become green. If you need to hear the sound again, press the 'REPEAT!' button.

15.2.3. Practicing strategies

1. We recommend you start with the first Q curve. The fact that it is a straight line indicates that in this case the value of Q will not change. Its value will produce pitched sounds in the lower frequency bands and noise in the higher ones. Using this fixed curve, listen to all the Amplitude interpolations. Try to identify which filterbands are responsible for the noisiness.

2. Then move on to the second straight line. Its values are lower so the bandwidth is wider and therefore noisier. Notice how the different amplitude presets will influence the darkness or brightness of the resulting sound.
3. Now choose the first Amplitude preset and listen to what the different Q curves produce. Try learning their characteristics. Systematically work thru all the amplitude presets.
4. Then try combining the Amplitude and Q presets predicting the resulting interpolations. You can create your Q interpolation curves. By clicking on the curve displayed you can change its shape, adding nodes to create really complex new curves (though these can not be saved, but of course you can record the sound).
5. It is important to try all of this with both envelope settings: longnote and impulse. You will notice how this will influence the resulting sound.
6. Try changing fundamental frequency and the Partial ratios of the filters. Put them close together, far apart, etc.
7. Try shortening the duration to a 100 msec. You will notice that the resulting interpolation sounds very different. One of the most interesting and surprising phenomenon is that interpolation also works for extremely short durations, like one tenth of a second. Experiment with a whole range of different durations from very short (a few milliseconds) to very long.