

Laboratory for week 6:

Digital Signals & Systems

Introduction

In order to process an acoustic signal on a computer the **continuous analogue waveform** has to be converted into a **discrete digital waveform**. ‘Analogue’ means that at any point in time, a value of the wave is available; ‘discrete’ means that only a limited number of values are available over time. In other words: a digital signal consists of a limited number of samples per second taken from the analogue signal. This number of samples is called the **sampling rate**. The standard sampling rate for CD quality is 44100 samples/second. Sampling rates that are too low have a negative effect on the sound quality.

For accurate reproduction, the sampling rate has to be at least twice as high as the highest frequency component present in the wave. Otherwise, *aliasing* will occur (i.e. frequency components in the original wave become different frequencies in the sampled wave). A sampling rate of 44100 samples/second will therefore only be able to represent frequencies less than 22050 Hz. All higher frequencies will be aliased. The highest accurately recordable frequency of a particular sampling rate is called the Nyquist frequency, and is simply one-half the sampling rate.

But a digital signal has to be discrete not only on the time axis, but also on the amplitude axis. Only a limited number of amplitude values is possible in a digital system. This discreteness is referred to as the **quantisation** level. The current quantisation standard is 16 bits (*binary digits*), which means that the whole range of the amplitude is divided into 65536 possible values ($2^{16} = 65536$). Lower quantisation levels cause again changes to the signal shape and thus introduce distortion to the signal.

Scientific Objectives

To get acquainted with the effects of changes in sampling rate and quantisation level to the signal.

Learning Objectives

- To learn to use the software package Praat
- To learn how signals can be varied in sampling rate and quantisation level.
- To experience the effect of the Nyquist frequency and aliasing on a signal.
- To gain familiarity with processing signals on a computer.
- To gain familiarity with creating systems on a computer.

Apparatus

You will use the software Praat on the computers in the lab. It is downloadable for free from the internet under www.praat.org.

(You will work with two options in Praat that have especially been created for this lab (‘Create sine wave’ and ‘Re quantise’) and are therefore not included in the Praat version you can download from the internet. If you are interested in having these two extra options for your own computer, see the web site.

1.) Sampling

Use the option ‘Create sine wave...’ which you can find under ‘New’ in the Praat menu (Praat is in the ‘speech analysis tools’ folder on the desktop of your PC):

- Specify your name in the first entry field.
- Set the sampling rate to 44100 (CD quality)
- Click on o.k.

- Make a printout of the picture you are receiving in the picture window (under 'File' → 'Print...' in the 'Praat picture' window – not 'Praat objects'!). This printout is a souvenir from this lab which will help you to remember what you did. You can also take further notes on it. You do not measure anything in the printout. You are welcome to make further printouts of other measurements you take if that helps you.

You have created two files with your operation that appear in the list of objects on the left hand side of the 'Praat objects' window. Now do the following:

- Click on the file 'SpectrumYOURNAME' and then on 'Edit' in the list of options on the right hand side. You can now see the spectrum of the window.
- Zoom in on a frequency band from about 0 to about 600 Hz so that you can view the area of your sine wave better. You do that by marking the area with the mouse (click left button and move mouse over area to be selected – like in a word processor) and then click on 'sel' (bottom left of the window).
- Measure the frequency of the spectral peak by clicking on it (zoom in more in case you can't get it properly). The frequency will appear in red above the cursor. Your frequency should be 200 Hz, hardly astonishing since you specified a 200 Hz waveform. But don't worry: things will change soon.
- Fill in the following table in the same way you did your first measurement:

frequency	amplitude	phase	sampling Rate	measured frequency (rounded to integer)	level of measured frequency
200	1	0	44100		
200	1	0	5000		
200	1	0	1000		
200	1	0	500		
200	1	0	410		
200	1	0	401		
200	1	0	400		
200	1	0	399		
200	1	0	350		
200	1	0	310		
200	1	0	290		
200	1	0	250		
200	1	0	210		
200	1	0	200		

2.) Quantisation

For the following table create a sinusoid of 200 Hz, and 0° phase with a sampling rate always of 22050 samples/s but varying the number of bits according to the values in the table. Please do this for the amplitude values of 1mV and ½ mV. Then:

- Count the number of quantisation levels (For 8 bits and higher it will be difficult, actually impossible, to count the quantisation levels but maybe you can calculate them!)
- Measure the level of the fundamental frequency component in the spectrum in the 'Praat objects' window (not in the 'Praat picture' window!). By clicking on the respective peak values in the spectrum you will be able to obtain their level in dB (in red) to the left of the spectrum.
- When the wave shape of a sinusoid is changed this can only be done by introducing new frequencies (remember: simple vs. complex waveform). Requantisation of a signal leads to waveform changes in the signal. The newly introduced harmonics are referred to as **harmonic distortion**. The strength of harmonic distortion can be measured by calculating the difference in dB between the peak of the fundamental and the peak of the largest introduced harmonic.

amplitude = 1 mV		
number of bits	quantisation levels	harmonic distortion (in dB)
1		
2		
4		
8		
12		
16		

- 1.) How does harmonic distortion develop as a function of the number of quantisation levels? Why?
-
-

3.) Requantising real speech

So far so good: We have seen that the level of quantisation can actually change the shape of a sine wave significantly and thus introduce strong harmonic distortion, highest in case of the 1 bit example. But, as you may already know, sine waves are a sort of artificial concept and they almost never occur in the real world.

With complex waveforms, low quantisation levels introduce another source of error: since many samples lying on the border between quantisation levels randomly fall into the higher or lower quantisation level, a low amplitude noise is produced that is called the **quantisation noise**.

In the next section, you will re-quantise a real speech example by changing the number of bits in its representation. You do that by the use of the option 'Re quantise...' under 'New' in the 'Praat objects' window. Apart from the number of bits you do not need to change any of the settings in this window. Create a version of the test file for each quantisation level possible and answer the following questions after you have listened to your examples:

- 1.) At what quantisation level does speech become non-intelligible? Would you have expected the result you came up with?
-
-

- 2.) At what levels of quantisation are you able to perceive quantisation noise in the signal?
-
-