

Digital Audio Coding

Lab Session 1 – WAVE

Sébastien Boisgérault,
Mines-ParisTech.

Dec. 10, 2012



This work is licensed under a [Creative Commons Attribution 3.0 Unported License \(CC BY 3.0\)](https://creativecommons.org/licenses/by-sa/3.0/). You are free to **share** – to copy, distribute and transmit the work – and to **remix** – to adapt the work – under the condition that the work is properly attributed to its author.

Synthesis of Pure Tones

We denote A_4 the analog audio signal defined by


$$A_4(t) = \cos(2\pi \times 440.0 \times t), \quad t \in \mathbb{R},$$

It's a pure tone with unit amplitude, frequency 440.0 Hz, and no phase.

Let Δf be the Compact Disc Digital Audio sample rate

$$\Delta f = 44.1 \text{ kHz}.$$

1. Define the variables `df` and `dt` that denote the sampling frequency in hertz and the sampling period in seconds that correspond to Δf .
2. Create the increasing array `t` of the multiples of Δt in the interval $[0, 3.0)$. Define the array `A4` of sampled values of A_4 at these instants.
3. Plot two periods of this sampled signal against the time.



4. Save the sampled signal **A4** as a WAVE file named "**A4.wav**".
5. Automate and generalize: the name “A₄” comes from the [scientific pitch notation](#), a convention that defines the frequency of symbols that are made of a letter followed by a number identifying the pitch octave. The following table displays such frequencies for the letter “A”:

Symbol	f (Hz)
A ₀	27.5
A ₁	55.0
A ₂	110.0
A ₃	220.0
A ₄	440.0
A ₅	880.0
...	...
A ₉	14080.0
A ₁₀	28160.0

Implement a function **make_tone** that given such a **symbol** argument:

- returns the values of a 3-sec. audio sample,
- creates a WAVE file whose name is **symbol** + ".wav".

Sound Pressure Level and Loudness

1. Create the pure tones **A0** to **A10** and the corresponding WAVE files.
2. Sound power:
 1. Implement a function **SPL** that given a 1-dim. array **x** argument returns the *sound pressure level* (SPL) of **x** in decibels:

$$L [\text{dB}] = 96.0 + 10 \log_{10} \langle \mathbf{x}^2 \rangle$$

where $\langle y \rangle$ is the mean value of y .

2. Check that the SPL of every tone is (approximately) the same.
 3. What is the analytical formula of this shared value ?
3. Listening tests:
1. Play the audio files "A0.wav" to "A10.wav" in sequence.
 2. Is the loudness of the sounds constant across the octaves ? What does it mean ?
 3. Does the frequency of "A10.wav" feel very different from the frequency of "A9.wav" ? Can you make sense of that ?

WAVE Format Header Analysis

The WAVEform audio file format – or WAVE for brevity – is a Microsoft and IBM digital audio file format. It is a subset of Microsoft's Resource Interchange File Format (RIFF) specification for multimedia file. The WAVE format supports several types of compressions, but we will deal only with the uncompressed format (also referred to as 16-bit linear PCM, for "Pulse Code Modulation").

1. WAVE file names often have a ".wav" extension, but this is not strictly mandatory: files content can be searched instead for a signature specific to this format.

Search the web for a description of the WAVE header, identify the WAVE file signature and implement an `is_wave` function that takes a filename argument and returns `True` if the file is a WAVE file and `False` otherwise.

2. Implement a function `wave_info` that given a WAVE filename returns informations about the audio data, namely:

- the number of channels
- the data sample rate.

This function should conform to the following example usage:

```
>>> wave_info("A4.wav")
{"num_channels": 1, "sample_rate": 44100}
```

Quantization and Signal-to-Noise Ratio

1. Use the `read` function of the `wave` Digital Audio Coding module to load the data of the audio file "`A5.wav`" as an array of floats in the $[-1.0, 1.0]$ range
2. Compute the *quantization error* – or *quantization noise* – `e`, defined as the difference between `A4_wav` and the original array `A4`.
3. Compute the quantization *signal-to-noise ratio* (SNR) in decibels:

$$\text{SNR [dB]} = 10 \log_{10} \text{SNR}^2 \quad \text{where} \quad \text{SNR}^2 = \frac{\langle \mathbf{A4}^2 \rangle}{\langle \mathbf{e}^2 \rangle}$$

Fading: Frames and Windows

The tones that we have generated so far do not start or end smoothly and the transition from one tone to another is not smooth either. Such fade in, fade out and cross-fading may be implemented with windows.

1. Extend the `make_tone` function so that it takes an optional argument `window`, a window factory : a function, that takes a length argument and returns a window of this length. A *window* is a one-dimensional array that is meant to be multiplied element-wise to the original signal.
Search the web for the definition of classic windows. Apply for example the “Hanning window” on the pure tone `A4` and listen to the result.
2. Consider the signals `A0` to `A10`. Create a single “blended” (crossfaded) signal that merges windowed versions of these signals in sequence, with a 1.5 sec (50%) overlap between successive signals.
3. What if the successive signals are actually *frames* extracted from a common signal ? When does the merge between the frames reconstruct exactly the original signal ? This condition is referred to as Constant-OverLap-Add (or COLA).
4. Show that the standard Hanning window satisfies only roughly this condition. Define the window (factory) `hanning2` that given a length `n` returns the first `n` values of the `hanning` window of length `n + 1`. Show that this window satisfies the COLA condition with a high precision.