Introduction to

Sound Processing

Davide Rocchesso[1]

March 20, 2003

[1] Università di Verona
Dipartimento di Informatica
email: D.Rocchesso@computer.org
www: http://www.scienze.univr.it/˜rocchess

pinna
ear canal
ear drum
oval window
hammer
anvil
stirrup
cochlea
basilar membrane
scala vestibuli
scala timpani
base
apex
helicotrema
tectorial membrane
hair cells

Appendix C

Fundamentals of psychoacoustics

Psychoacoustics is a "discipline within psychology concerned with sound, its perception and the physiological foundations of hearing" [75]. A few concepts and facts of psychoacoustics are certainly useful to the sound designer and to any computer scientist interested in working with sound. Several books provide a wider treatment of this topic, at different degrees of depth [86, 105, 42, 111].

C.1 The ear

The human ear is usually described as composed of three parts. This system is schematically depicted in figure 1.

the outer ear: The pinna couples the external space to the ear canal. Its shape is exploited by the hearing system to extract directional information from incoming sounds. The ear canal is a tube (length $l \approx 2.6$cm, diameter $d \approx 0.6$cm) closed on the inner side by a membrane called the ear drum. The tube acts as a quarter-of-wavelength resonator, exciting frequencies in the neighborhood of $f_0 = \frac{c}{4l} \approx 3.3$kHz, where c is the speed of sound in air;

the middle ear: It transmits mechanical energy, received from the ear drum, to the inner ear through a membrane called the oval window. To do so, it uses a chain of small bones, called the hammer, the anvil, and the stirrup;

the inner ear: It is a cavity, called cochlea, shaped like a snail shell, which is shown rectified for clarity in figure 1. It contains a fluid and it is divided by the basilar membrane into two chambers: the scala vestibuli and the scala timpani. The length of the cochlea is about 3.5cm. Its diameter is about 2mm at the oval window (base) and it gets narrower at the other extreme (apex), where a narrow aperture (the helicotrema) allows the two chambers to communicate. On top of the basilar membrane, the tectorial membrane sustains about $16,000$ hair cells that pick up the transversal motion of the basilar membrane and transmit it to the brain.

The vibrations of the oval window excite the fluid of the scala vestibuli. By pressure differences between the scala vestibuli and scala timpani, the basilar

impedance of the tube
acoustic intensity

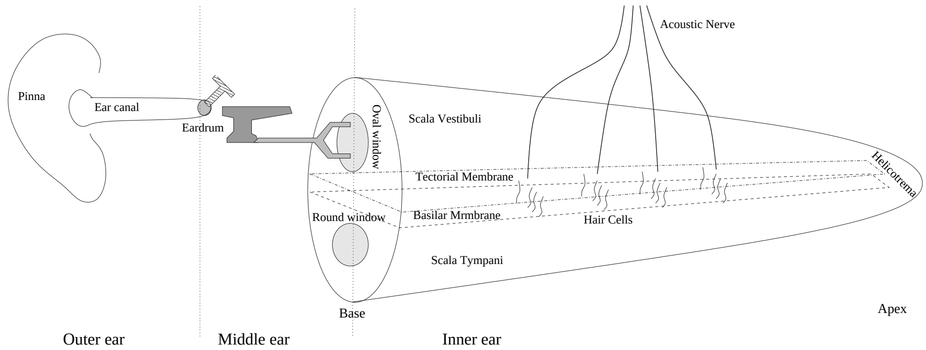

Figure 1: Cartoon physiology of the ear

membrane oscillates and transversal waves are propagated. The basilar membrane can be thought of as a string having a decreasing tension as we move from the base to the apex. This tension changes by about four orders of magnitude from base to apex. Along a string, the waves propagate at speed

$$c = \sqrt{\frac{T}{\rho_L}} = \sqrt{\frac{\text{Tension}}{\text{Linear density}}} \;, \tag{1}$$

and the wavelength associated with the component at frequency f is

$$\lambda = \frac{1}{f}\sqrt{\frac{T}{\rho_L}} = \frac{c}{f} \;. \tag{2}$$

The impedance of the tube is $z_0 = \sqrt{\rho_L T}$ and, if v_{max} is the peak value of transversal velocity, the wave power is

$$P = \frac{1}{2} z_0 v_{max}^2 = \frac{1}{2}\sqrt{\rho_L T} v_{max}^2 \;. \tag{3}$$

While a wave component at frequency f is propagating from the base to the apex, its wavelength decreases (because tension decreases) and, due to the physical requirement of power constancy, its amplitude increases. However, this propagation is not lossless, and dissipation increases with the amplitude, so that a frequency-dependent maximum region will emerge along the basilar membrane (see figure 2). Since the high frequencies are more affected by propagation losses, their characteristic resonance areas are cluttered close to the base, while low frequencies are more widely distributed toward the apex. About two thirds of the length of the cochlea is devoted to low frequencies (about one fourth of the audio bandwidth), thus giving more frequency resolution to the slowly-varying components.

C.2 Sound Intensity

Consider a sinusoidal point source in free space. It generates spherical pressure waves that carry energy. The acoustic intensity is the power by unit surface that

RMS
threshold of hearing
threshold of pain
intensity level
decibel
sound pressure level

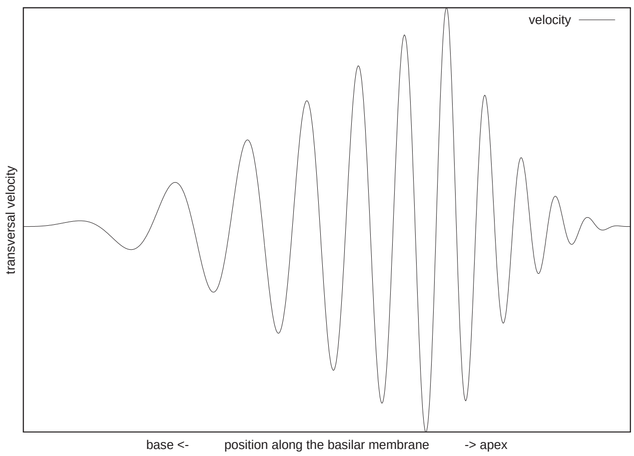

Figure 2: Cartoon of the transversal velocity pattern elicited by an incoming pure sine tone

is carried by a wave front. It is a vectorial quantity having magnitude

$$I = \frac{p_{max}^2}{2}\frac{1}{z_0} = \frac{p_{max}^2}{2\rho c} = \frac{p_{RMS}^2}{\rho c} , \tag{4}$$

where p_{max} and p_{RMS} are the peak and root-mean-square (RMS) values of pressure wave, respectively, and $z_0 = \rho c =$ density $\times$ speed is the impedance of air.

At 1000Hz the human ear can detect sound intensities ranging from $I_{min} = 10^{-12}\mathrm{W/m^2}$ (threshold of hearing) to $I_{max} = 1\mathrm{W/m^2}$ (threshold of pain).

Consider two spherical shells of areas a_1 and a_2, at distances r_1 and r_2 from the point source. The lossless propagation of a wavefront implies that the intensities registered at the two distances are related to the areas by

$$I_1 a_1 = I_2 a_2 . \tag{5}$$

Since the area is proportional to the square of distance from the source, we also have

$$\frac{I_1}{I_2} = \left(\frac{r_2}{r_1}\right)^2 . \tag{6}$$

The intensity level is defined as

$$IL = 10 \log_{10} \frac{I}{I_0} , \tag{7}$$

where $I_0 = 10^{-12}\mathrm{W/m^2}$ is the sound intensity at the threshold of hearing. The intensity level is measured in decibel (dB), so that multiplications by a factor are turned into additions by an offset, as represented in table C.2. Similarly, the sound pressure level is defined as

$$SPL = 20 \log_{10} \frac{p_{max}}{p_{0,max}} = 20 \log_{10} \frac{p_{RMS}}{p_{0,RMS}} \tag{8}$$

standing wave
Fletcher-Munson curves
equal-loudness curves
loudness level
phons
loudness
sones
standardized loudness scale

I	IL
×1.26	+1
×2	+3
×10	+10

Table C.1: Relation between factors in the linear intensity scale and shifts in the dB intensity-level scale

where $p_{0,max}$ and $p_{0,RMS}$ are the peak and RMS pressure values at the threshold of hearing. For a propagating wave, we have that $IL = SPL$. For a standing wave, since there is no power transfer and since IL is a power-based measure, the SPL is more appropriate.

Given a reference tone with a certain value of IL at 1kHz, we can ask a subject to adjust the intensity of a probe tone at a different frequency until it matches the reference loudness perceptually. What we would obtain are the Fletcher-Munson curves, or equal-loudness curves, sketched in figure 3. Each curve is parameterized on a value of loudness level (LL), measured in phons. The loudness level is coincident with the intensity level at 1kHz.

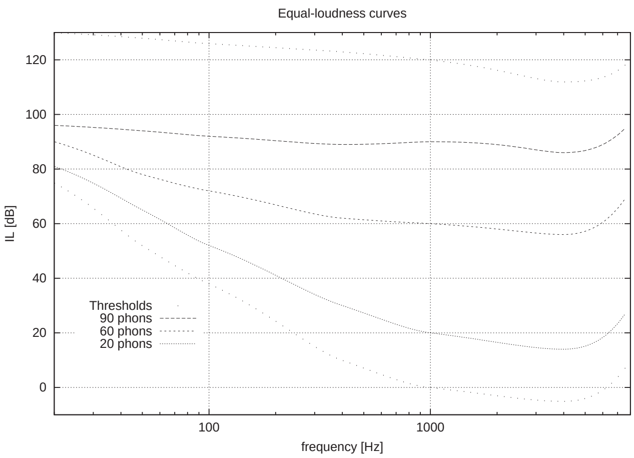

Figure 3: Equal-loudness curves. The parameters express values of loudness level in phons.

Even though the Fletcher-Munson curves are obtained by averaging the responses of human subjects, the LL is still a physical quantity, because it refers to the physical quantity IL and it does not represent the perceived loudness in absolute terms. In other words, doubling the loudness level does not mean doubling the perceived loudness. A genuine psychophysic measure is the loudness in sones, which can be obtained as a function of LL by asking listeners to compare sounds and decide when one sound is "twice as loud" as another. Somewhat arbitrarily, a LL of 40 phons is set equal to 1 sone. Figure 4 represents a possible average curve that may emerge from an experiment. The standardized loudness scale (ISO) uses the straight line approximation of figure 4, that corresponds

direct methods
pitch
frequency JND
subjective scale for pitch
mel

Experimental curves similar to that reported in figure 4 show in many cases significant deviations from (14). For instance, the relation between intensity and loudness is more similar to

$$L \propto \sqrt[3]{I}\ , \qquad (15)$$

as three doublings of intensity are needed for approximating one doubling in loudness.

Power laws such as the (15) are the natural outcome of the so called direct methods of psychophysical experimentation, where it is the sensation itself that is the unit for measuring other sensations. Such experimental paradigm was largely established by Stevens[3], and it is the one in use when the experimenter asks the subject to double or half the perceived loudness of a tone, or when a direct magnitude production or estimation is used.

C.3 Pitch

Periodic tones elicit a sensation of pitch, thus meaning that they can be ordered on a scale from low to high. Many aperiodic or even stochastic sounds can elicit pitch sensations, with different degrees of strength.

If we stick with pure tones for this section, pitch is the sensorial correlate of frequency, and it makes sense to measure the frequency JND using the tools of psychophysics. For instance, if a pure tone is slowly modulated in frequency we may seek for the threshold of modulation audibility. The resulting curve of average results would look similar to figure 5.

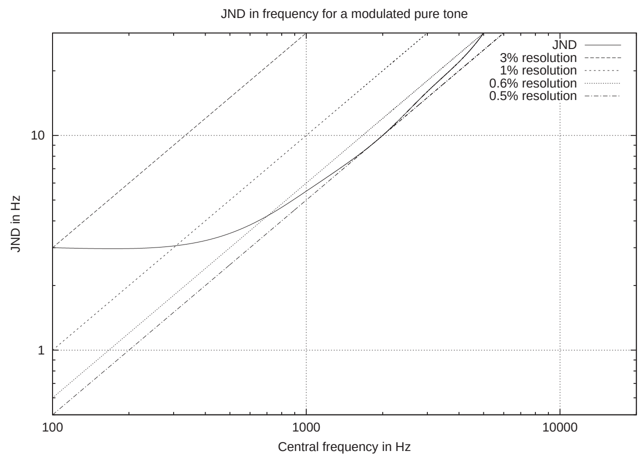

Figure 5: JND in frequency for a slowly modulated pure tone.

Again, from the curve of figure 5 we notice a significant deviation from the Weber's law $\Delta f \propto f$. The physiological interpretation is that there is more internal noise in the frequency detection in the very-low range.

If we integrate $\frac{1}{\Delta f(f)}$ we obtain a curve such as that of figure 6 that can be interpreted as a subjective scale for pitch, whose unit is called mel. Conventionally 1000 Hz corresponds to 1000 mel. This curve shouldn't be confused with the

[3]Stanley Smith Stevens (1906-1973).

Musical scales
musical octave
chroma
place theory of hearing
virtual pitch
missing fundamental
temporal processing of sound
timbre

scales that organize musical height. Musical scales are based on the subdivision of the musical octave into a certain number of intervals. The musical octave is usually defined as the frequency range having the higher bound that has twice the value in Hertz of the first bound. On the other hand, the subjective scale for pitch measures the subjective pitch relationship between two sounds, and it is strictly connected with the spatial distribution of frequencies along the basilar membrane. In musical reasoning, pitch is referred to as chroma, which is a different thing from the tonal height that is captured by figure 6.

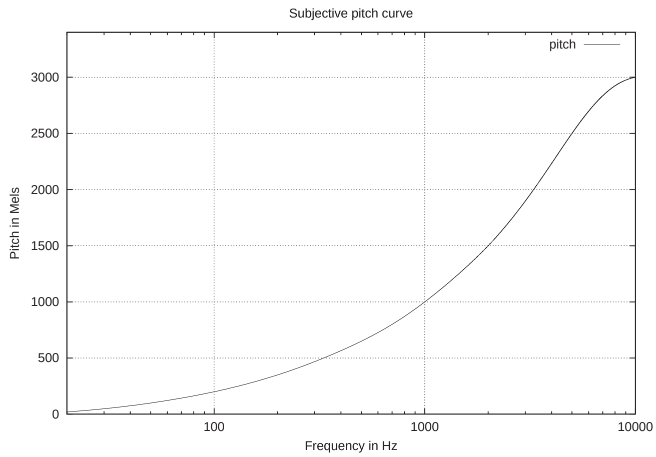

Figure 6: Subjective frequency curve, mel vs. Hz.

So far, we have described pitch phenomena referring to the position of hair cells that get excited along the basilar membrane. Indeed, the place theory of hearing is not sufficient to explain the accuracy of pitch perception and some intriguing effects such as the virtual pitch. In this effect, if a pure tone at frequency f_1 is superimposed to a pure tone at frequency $f_2 = \frac{3}{2} f_1$, the perceived pitch matches the missing fundamental at $f_0 = f_1/2$. If the reader, as an excercise, plots this superposition of waveforms, she may notice that the apparent periodicity of the resulting waveform is $1/f_0$. This indicates that a temporal processing of sound may occur at some stages of our perception. The hair cells convey signals to the fibers of the acoustic nerve. These neural contacts fire at a rate that depends on the transversal velocity of the basilar membrane and on its lateral displacement. The rate gets higher for displacements that go from the apex to the base, and this creates a periodicity in the firing rate that is multiple of the waveform periodicity. Therefore, the statistical distribution of neural spikes keep track of the temporal behavior of the acoustic signals, and this may be useful at higher levels to extract periodicity information, for instance by autocorrelation processes [86].

Even for pure tones, pitch perception is a complex business. For instance, it is dependent on loudness and on the nature and quality of interfering sounds [42]. The pitch of complex tones is an overly complex topic to be discussed in this appendix. It suffices to know that pitch perception of complex tones is linked to the third (after loudness and pitch) and most elusive attribute of sound, that is timbre.

kernel of the Fourier transform
spectrum
magnitude spectrum
phase spectrum
Z transform

A.8.2 The Fourier Transform

The Fourier transform of $y(t)$, $t \in \mathcal{R}$, can be obtained as a specialization of the Laplace transform in the case that the latter is defined in a region comprising the imaginary axis. In such case we define[14]

$$Y(\Omega) \stackrel{\triangle}{=} Y_L(j\Omega) , \tag{54}$$

or, in detail,

$$Y(\Omega) = \int_{-\infty}^{+\infty} y(t) e^{-j\Omega t} dt , \tag{55}$$

where $j\Omega$ indicates a generic point on the imaginary axis. Since the kernel of the Fourier transform is the complex sinusoid (i.e., the complex eponential) having radial frequency Ω, we can interpret each point of the transformed function as a component of the frequency spectrum of the function $y(t)$. In fact, given a value $\Omega = \Omega_0$ and considered a signal that is the complex sinusoid $y(t) = e^{j\Omega_1 t}$, the integral (55) is maximized when choosing $\Omega_0 = \Omega_1$, i.e., when $y(t)$ is the complex conjugate of the kernel [15]. The codomain of the transformed function $Y(\Omega)$ belongs to the complex field. Therefore, the spectrum can be decomposed in a magnitude spectrum and in a phase spectrum.

A.8.3 The Z Transform

The domains of functions can be classes of numbers of whatever kind and nature. If we stick with functions defined over rings, particularly important are the functions whose domain is the ring of integer numbers. These are called discrete-variable functions, to distinguish them from functions of variables defined over $\mathcal{R}$ or $\mathcal{C}$, which are called continuous-variable functions.

For discrete-variable functions the operators derivative and integral are replaced by the simplest operators difference and sum. This replacement brings a new definition of transform for a function $y(n)$, $n \in \mathcal{Z}$:

$$Y_Z(z) = \sum_{n=-\infty}^{+\infty} y(n) z^{-n}, z \in \Gamma \subset \mathcal{C} . \tag{56}$$

The transform (56) is called Z transform and the region of convergence is a ring[16] of the complex plane. Within this ring the transform can be inverted.

Example 3. The Z transform of the discrete-variable causal exponential is[17]

$$Y_Z(z) = \sum_{n=-\infty}^{+\infty} y(n) z^{-n} = \sum_{n=0}^{+\infty} e^{z_0 n} z^{-n} = \sum_{0}^{+\infty} (e^{z_0} z^{-1})^n = \frac{1}{1 - e^{z_0} z^{-1}} , \tag{57}$$

[14]Often the Fourier transform is defined as a function of f, where $2\pi f = \Omega$

[15]Exercise: find the Fourier transform of the causal complex exponential (48), with $s_0 = \alpha + j\Omega_0$, and show that it has maximum magnitude for $\Omega = \Omega_0$.

[16]A ring here is the area between two circles and not an algebraic structure.

[17]The latter equality in (57) is due to the identity $\sum_{n=0}^{+\infty} a^n = \frac{1}{1-a}$, $|a| < 1$, which can be verified by the reader with $a = 1/2$.

and it is convergent for values of z that are larger than $e^{\Re(z_0)}$ in magnitude[18].

Similarly to what we saw for continuous-variable functions, the Fourier transform for discrete-variable functions can be obtained as a specialization of the Z transform where the values of the complex variable are restricted to the unit circumference.

$$Y(\omega) = Y_Z(e^{j\omega}) \ , \tag{58}$$

or, in detail,

$$Y(\omega) = \sum_{n=-\infty}^{+\infty} y(n)e^{-j\omega n} \ . \tag{59}$$

In this book, we use the symbol ω for the radian frequency in the case of discrete-variable functions, leaving Ω for the continuous-variable functions.

###

binary digits
unsigned integer
bits
least significant bit
most significant bit
signed integers
two's complement representation

A.9 Computer Arithmetics

A.9.1 Integer Numbers

In order to fully understand the behavior of several hardware and software tools for sound processing, it is important to know something about the internal representation of numbers within computer systems. Numbers are represented as strings of binary digits (0 and 1), but the specific meaning of the string depends on the conventions used. The first convention is that of unsigned integer numbers, whose value is computed, in the case of 16 bits, by the following formula

$$x = \sum_{i=0}^{15} x_i \times 2^i \ , \tag{60}$$

where x_i is the i-th binary digit starting from the right. The binary digits are called bits, the rightmost digit is called least significant bit (LSB), and the leftmost digit is called the most significant bit (MSB). For instance, we have

$$0100001100100110_2 = 2^1 + 2^2 + 2^5 + 2^8 + 2^9 + 2^{14} = 17190 \ , \tag{61}$$

where the subscript 2 indicates the binary representation, being the usual decimal representation indicated with no subscript.

The leftmost bit is often interpreted as a sign bit: if it is set to one it means that the sign is minus and the absolute value is given by the bits that follow. However, this is not the representation that is used for the signed integers. For these numbers the two's complement representation is used, where the leftmost bit is still a sign bit, but the absolute value of a negative number is recovered by bitwise complementation of the following bits, interpretation of the result as a positive integer, and addition of one. For instance, with four bits we have

$$1010_2 = -(0101_2 + 1) = -(5 + 1) = -6 \ . \tag{62}$$

The two's complement representation has the following advantages:

[18] $\Re(x)$ is the real part of the complex number x

In order to obtain the signal coming out from a linear system it is sufficient to apply the convolution operator between the input signal and the impulse response.

sampling
sampling interval
spectrum
images
Sampling Theorem

1.2 The Sampling Theorem

In order to perform any form of processing by digital computers, the signals must be reduced to discrete samples of a discrete-time domain. The operation that transforms a signal from the continuous time to the discrete time is called sampling, and it is performed by picking up the values of the continuous-time signal at time instants that are multiple of a quantity T, called the sampling interval. The quantity $F_s = 1/T$ is called the sampling rate.

The presentation of a detailed theory of sampling would take too much space and it would become easily boring for the readership of this book. For a more extensive treatment there are many excellent books readily available, from the more rigorous [66, 65] to the more practical [67]. Luckily, the kernel of the theory can be summarized in a few rules that can be easily understood in terms of the frequency-domain interpretation of signals and systems.

The first rule is related to the frequency representation of discrete-time variables by means of the Fourier transform, defined in appendix A.8.3 as a specialization of the Z transform:

Rule 1.1 *The Fourier transform of a function of discrete variable is a function of the continuous variable ω, periodic*[3] *with period 2π.*

The second rule allows to treat the sampled signals as functions of discrete variable:

Rule 1.2 *Sampling a continuous-time signal $x(t)$ with sampling interval T produces a function $\hat{x}(n) \stackrel{\triangle}{=} x(nT)$ of the discrete variable n.*

If we call spectrum of a signal its Fourier-transformed counterpart, the fundamental rule of sampling is the following:

Rule 1.3 *Sampling a continuous-time signal with sampling rate F_s produces a discrete-time signal whose frequency spectrum is a periodic replication of the spectrum of the original signal, and the replication period is F_s. The Fourier variable ω for functions of discrete variable is converted into the frequency variable f (in Hz) by means of*

$$\omega = 2\pi f T = \frac{2\pi f}{F_s} \; . \tag{7}$$

Fig. 1 shows an example of frequency spectrum of a signal sampled with sampling rate F_s. In the example, the continuous-time signal had all and only the frequency components between $-F_b$ and F_b. The replicas of the original spectrum are sometimes called images.

Given the simple rules that we have just introduced, it is easy to understand the following Sampling Theorem, introduced by Nyquist in the twenties and popularized by Shannon in the forties:

[3] This periodicity is due to the periodicity of the complex exponential of the Fourier transform.

band-limited
reconstruction filter
Nyquist frequency
holder
sampler
sample and hold

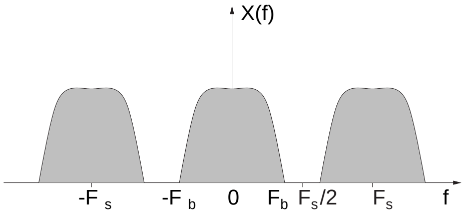

Figure 1: Frequency spectrum of a sampled signal

Theorem 1.1 *A continuous-time signal $x(t)$, whose spectral content is limited to frequencies smaller than F_b (i.e., it is band-limited to F_b) can be recovered from its sampled version $\hat{x}(n) = x(nT)$ if the sampling rate $F_s = 1/T$ is such that*

$$F_s > 2F_b \ . \tag{8}$$

It is also clear how such recovering might be obtained. Namely, by a linear reconstruction filter capable to eliminate the periodic images of the base band introduced by the sampling operation. Ideally, such filter doesn't apply any modification to the frequency components lower than the Nyquist frequency, defined as $F_N = F_s/2$, and eliminates the remaining frequency components completely.

The reconstruction filter can be defined in the continuous-time domain by its impulse response, which is given by the function

$$h(t) \stackrel{\triangle}{=} sinc(t) = \frac{\sin{(\pi t/T)}}{\pi t/T} \ , \tag{9}$$

which is depicted in fig. 2.

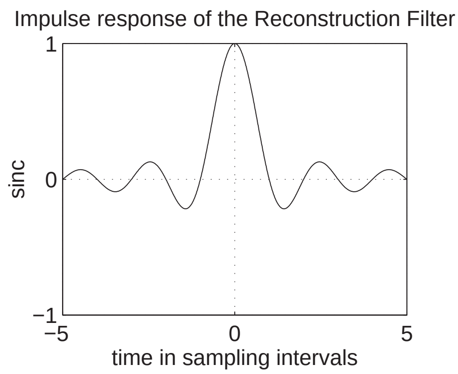

Figure 2: *sinc* function, impulse response of the ideal reconstruction filter

Ideally, the reconstruction of the continuous-time signal from the sampled signal should be performed in two steps:

- Conversion from discrete to continuous time by holding the signal constant in time intervals between two adjacent sampling instants. This is achieved by a device called a holder. The cascade of a sampler and a holder constitutes a sample and hold device.

aliasing
foldover
frequency modulation
digital frequencies
Discrete-Time Fourier Transform
DTFT

- Convolution with an ideal *sinc* function.

The *sinc* function is ideal because its temporal extension is infinite on both sides, thus implying that the reconstruction process can not be implemented exactly. However, it is possible to give a practical realization of the reconstruction filter by an impulse response that approximates the *sinc* function.

Whenever the condition (8) is violated, the periodic replicas of the spectrum have components that overlap with the base band. This phenomenon is called aliasing or foldover and is avoided by forcing the continuous-time original signal to be bandlimited to the Nyquist frequency. In other words, a filter in the continuous-time domain cuts off the frequency components exceeding the Nyquist frequency. If aliasing is allowed, the reconstruction filter can not give a perfect copy of the original signal.

Usually, the word aliasing has a negative connotation because the aliasing phenomenon can make audible some spectral components which are normally out of the frequency range of hearing. However, some sound synthesis techniques, such as frequency modulation, exploit aliasing to produce additional spectral lines by folding onto the base band spectral components that are outside the Nyquist bandwidth. In this case where the connotation is positive, the term foldover is preferred.

1.3 Discrete-Time Spectral Representations

We have seen how the sampling operation essentially changes the nature of the signal domain, which switches from a continuous to a discrete set of points. We have also seen how this operation is transposed in the frequency domain as a periodic replication. It is now time to clarify the meaning of the variables which are commonly associated to the word "frequency" for signals defined in both the continuous and the discrete-time domain. The various symbols are collected in table 1.1, where the limits imposed by the Nyquist frequency are also indicated. With the term "digital frequencies" we indicate the frequencies of discrete-time signals.

Nyquist Domain	Symbol	Unit	
$[-F_s/2 \;\ldots\; 0 \;\ldots\; F_s/2]$	f	[Hz] = [cycles/s]	
$[-1/2 \;\ldots\; 0 \;\ldots\; 1/2]$	f/F_s	[cycles/sample]	digital frequencies
$[-\pi \;\ldots\; 0 \;\ldots\; \pi]$	$\omega = 2\pi f/F_s$	[radians/sample]	
$[-\pi F_s \;\ldots\; 0 \;\ldots\; \pi F_s]$	$\Omega = 2\pi f$	[radians/s]	

Table 1.1: Frequency variables

Appendix A.8.3 shows how it is possible to define a Fourier transform for functions of a discrete variable. Here we can re-express such definition, as a function of frequency, for discrete-variable functions obtained by sampling continuous-time signals with sampling interval T. This transform is called the Discrete-Time Fourier Transform (DTFT) and is expressed by

$$Y(f) = \sum_{n=-\infty}^{+\infty} y(nT) e^{-j2\pi \frac{f}{F_s} n} . \tag{10}$$

digital filter
non-recursive filters
Finite Impulse Response
FIR
Infinite Impulse Response
IIR

Chapter 2

Digital Filters

For the purpose of this book we call digital filter any linear, time-invariant system operating on discrete-time signals. As we saw in chapter 1, such a system is completely described by its impulse response or by its (rational) transfer function. Even though the adjective digital refers to the fact that parameters and signals are quantized, we will not be too concerned about the effects of quantization, that have been briefly introduced in sec. 1.6. In this chapter, we will face the problem of designing impulse responses or transfer functions that satisfy some specifications in the time or frequency domain.

Traditionally, digital filters have been classified into two large families: those whose transfer function doesn't have the denominator, and those whose transfer function have the denominator. Since the filters of the first family admit a realization where the output is a linear combination of a finite number of input samples, they are sometimes called non-recursive filters[1]. For these systems, it is more customary and correct to refer to the impulse response, which has a finite number of non-null samples, thus calling them Finite Impulse Response (FIR) filters. On the other hand, the filters of the second family admit only recursive realizations, thus meaning that the output signal is always computed by using previous samples of itself. The impulse response of these filters is infinitely long, thus justifying their name as Infinite Impulse Response (IIR) filters.

2.1 FIR Filters

An FIR filter is nothing more than a linear combination of a finite number of samples of the input signal. In our examples we will treat causal filters, therefore we will not process input samples coming later than the time instant of the output sample that we are producing.

The mathematical expression of an FIR filter is

$$y(n) = \sum_{m=0}^{N} h(m)x(n-m) \ . \qquad (1)$$

In eq. 1 the reader can easily recognize the convolution (26), here specialized to

[1]Strictly speaking, this definition is not correct because the same transfer functions can be realized in recursive form

averaging filter
magnitude response
phase response

finite-length impulse responses. Since the time extension of the impulse response is $N+1$ samples, we say that the FIR filter has length $N+1$.

The transfer function is obtained as the Z transform of the impulse response and it is a polynomial in the powers of z^{-1}:

$$H(z) = \sum_{m=0}^{N} h(m) z^{-m} = h(0) + h(1) z^{-1} + \ldots + h(N) z^{-N} \ . \qquad (2)$$

Since such polynomial has order N, we also say that the FIR filter has order N.

2.1.1 The Simplest FIR Filter

Let us now consider the simplest nontrivial FIR filter that one can imagine, the averaging filter

$$y(n) = \frac{1}{2} x(n) + \frac{1}{2} x(n-1) \ . \qquad (3)$$

In appendix B.1 it is illustrated how such filter can be implemented in Octave/Matlab in two different ways: block processing or sample-by-sample processing. The simplest way to analyze the behavior of the filter [97] is probably the injection of a complex sinusoid having amplitude A and initial phase ϕ, i.e. the signal $x(n) = A e^{j(\omega_0 n + \phi)}$. Since the system is linear we do not loose any generality by considering unit-amplitude signals ($A = 1$). Since the system is time invariant we do not loose any generality by considering signals with initial zero phase ($\phi = 0$). Since the complex sinusoid can be expressed as the sum of a cosinusoidal real part and a sinusoidal imaginary part, we can imagine that feeding the system with such a complex signal corresponds to feeding two copies of the filter, the one with a cosinusoidal real signal, the other with a sinusoidal real signal. The output of the filter fed with the complex sinusoid is obtained, thanks to linearity, as the sum of the outputs of the two copies.

If we replace the complex sinusoidal input in eq. (3) we readily get

$$y(n) = \frac{1}{2} e^{j\omega_0 n} + \frac{1}{2} e^{j\omega_0 (n-1)} = (\frac{1}{2} + \frac{1}{2} e^{-j\omega_0}) e^{j\omega_0 n} = (\frac{1}{2} + \frac{1}{2} e^{-j\omega_0}) x(n) \ . \qquad (4)$$

We see that the output is a copy of the input multiplied by the complex number $(\frac{1}{2} + \frac{1}{2} e^{-j\omega_0})$, wich is the value taken by the transfer function at the point $z = e^{j\omega_0}$. In fact, the transfer function (2) can be rewritten, for the case under analysis, as

$$H(z) = \frac{1}{2} + \frac{1}{2} z^{-1} \ , \qquad (5)$$

and its evaluation on the unit circle ($z = e^{j\omega}$) gives the frequency response

$$H(\omega) = \frac{1}{2} + \frac{1}{2} e^{-j\omega} \ . \qquad (6)$$

For an input complex sinusoid having frequency ω_0, the frequency response takes value

$$H(\omega_0) = \frac{1}{2} + \frac{1}{2} e^{-j\omega_0} = (\frac{1}{2} e^{j\omega_0/2} + \frac{1}{2} e^{-j\omega_0/2}) e^{-j\omega_0/2} = \cos(\omega_0/2) e^{-j\omega_0/2} \ , \qquad (7)$$

and we see that the magnitude response and the phase response are, respectively

$$|H(\omega_0)| = \cos(\omega_0/2) \qquad (8)$$

lowpass filter

and

$$\angle H(\omega_0) = -\omega_0/2\ . \qquad (9)$$

These are respectively the magnitude and argument of the complex number that is multiplied by the input function in (4). Therefore, we have verified a general property of linear and time-invariant systems, i.e., sinusoidal inputs give sinusoidal outputs, possibly with an amplitude rescaling and a phase shift[2].

If the frequency of the input sine is thought of as a real variable ω in the interval $[0, \pi)$, the magnitude and phase responses become a function of such variable and can be plotted as in fig. 1. At this point, the interpretation of such curves as amplification and phase shift of sinusoidal inputs should be obvious.

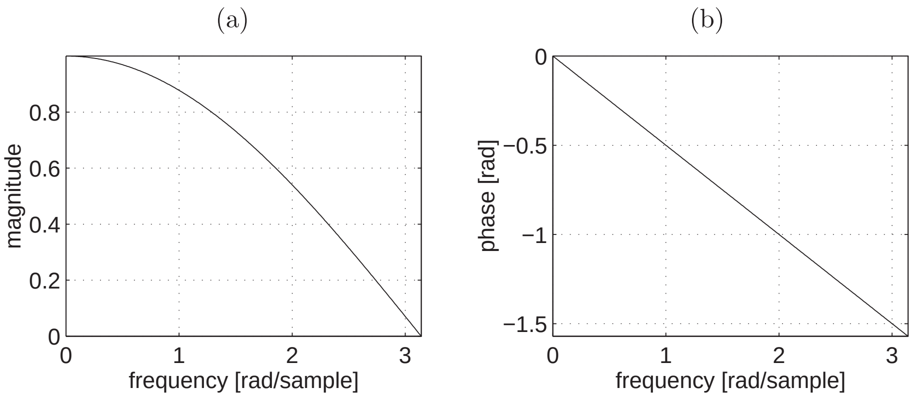

Figure 1: Frequency response (magnitude and phase) of an averaging filter

In order to plot curves such as those of fig. 1 it is not necessary to calculate closed forms of the functions representing the magnitude (8) and the phase response (9). Since with Octave/Matlab we can directly operate on arrays of complex numbers, the following simple script will do the job:

```
global_decl; platform('octave');
w = [0:0.01:pi];                      % frequency points
H = 0.5 + 0.5*exp(- i * w );          % complex frequency response
subplot(2,2,1); plot(w, abs(H));      % plot the magnitude
xlabel('frequency [rad/sample]');
ylabel('magnitude');
eval(myreplot);
subplot(2,2,2); plot(w, angle(H)); % plot the phase
xlabel('frequency [rad/sample]');
ylabel('phase [rad]');
eval(myreplot);
```

The averaging filter is the simplest form of lowpass filter. In a lowpass filter the high frequencies are more attenuated than the low frequencies. Another way to approach the analysis of a filter is to reason directly in the plane of the complex variable z. In this plane (fig. 2) two families of points are marked: the

[2]The reader can easily verify that this is true not only for complex sinusoids, but also for real sinusoids. The real sinusoid can be expressed as a combination of complex sinusoids and linearity can be applied.

circular buffer
signal flowgraphs
taps
tapped delay line
Auto-Regressive Moving Average
filter order

The three memory words are put in an area organized as a circular buffer. The input is written to the word pointed by the index and the three preceding values of the input are read with the three preceding values of the index. At every sample instant, the four indexes are incremented by one, with the trick of beginning from location 0 whenever we exceed the length M of the buffer (this ensures the circularity of the buffer). The counterclockwise arrow indicates the direction taken by the indexes, while the clockwise arrow indicates the movement that should be done by the data if the indexes would stay in a fixed position. In fig. 13 we use small triangles to indicate the multiplications by the filter coefficients. This is a notation commonly used for multiplications within the signal flowgraphs that represent digital filters. As a matter of fact, an FIR filter contains a delay line since it stores N consecutive samples of the input sequence and uses each of them with a delay of N samples at most. The points where the circular buffer is read are called taps and the whole structure is called a tapped delay line.

2.2 IIR Filters

In general, a causal IIR filter is represented by a difference equation where the output signal at a given instant is obtained as a linear combination of samples of the input and output signals at previous time instants. Moreover, an instantaneous dependency of the output on the input is also usually included in the IIR filter. The difference equation that represents an IIR filter is

$$y(n) = -\sum_{m=1}^{N} a_m y(n-m) + \sum_{m=0}^{M} b_m x(n-m) \; . \tag{24}$$

Eq. (24) is also called Auto-Regressive Moving Average (ARMA) representation. While the impulse response of FIR filters has a finite time extension, the impulse response of IIR filters has, in general, an infinite extension. The transfer function is obtained by application of the Z transform to the sequence (24). In virtue of the shift theorem, the Z transform is a mere operatorial substitution of each translation by m samples with a multiplication by z^{-m}. The result is the rational function $H(z)$ that relates the Z transform of the output to the Z transform of the input:

$$Y(z) = \frac{b_0 + b_1 z^{-1} + \ldots + b_M z^{-M}}{1 + a_1 z^{-1} + \ldots + a_N z^{-N}} X(z) = H(z)X(z) \; . \tag{25}$$

The filter order is defined as the degree of the polynomial in z^{-1} that is the denominator of (25).

2.2.1 The Simplest IIR Filter

In this section we analyze the properties of the simplest nontrivial IIR filter that can be conceived: the one-pole filter having coefficients $a_1 = -\frac{1}{2}$ and $b_0 = \frac{1}{2}$:

$$y(n) = \frac{1}{2} y(n-1) + \frac{1}{2} x(n) \; . \tag{26}$$

second-order filter

Instead of starting from the transfer function or from the difference equation, in this case we begin by positioning the two poles in the complex plane at the point

$$p_0 = re^{j\omega_0} \tag{38}$$

and at its conjugate point $p_0{}^* = re^{-j\omega_0}$. In fact, if p_0 is not real, the two poles must be complex conjugate if we want to have a real-coefficient transfer function. In order to make sure that the filter is stable, we impose $|r| < 1$. The transfer function of the second-order filter can be written as

$$\begin{aligned} H(z) &= \frac{G}{(1 - re^{j\omega_0}z^{-1})(1 - re^{-j\omega_0}z^{-1})} \\ &= \frac{G}{1 - r(e^{j\omega_0} + e^{-j\omega_0})z^{-1} + r^2 z^{-2}} = \frac{G}{1 - 2\, r\, cos\omega_0 z^{-1} + r^2 z^{-2}} \\ &= \frac{G}{1 + a_1 z^{-1} + a_2 z^{-2}} \end{aligned} \tag{39}$$

where G is a parameter that allows us to control the total gain of the filter.

As usual, we obtain the frequency response by substitution of z with $e^{j\omega}$ in (31):

$$H(\omega) = \frac{G}{1 - 2\, r\, cos\omega_0 e^{-j\omega} + r^2 e^{-2j\omega}} \; . \tag{40}$$

If the input is a complex sinusoid at the (resonance) frequency ω_0, the output is, from the first of (39):

$$H(\omega_0) = \frac{G}{(1 - r)(1 - re^{-2j\omega_0})} = \frac{G}{(1 - r)(1 - r\cos 2\omega_0 + j\, r \sin 2\omega_0)} \; . \tag{41}$$

In order to have a unit-magnitude response at the frequency ω_0 we have to impose

$$|H(\omega_0)| = 1 \tag{42}$$

and, therefore,

$$G = (1 - r)\sqrt{1 - 2r\cos 2\omega_0 + r^2} \; . \tag{43}$$

The frequency response of this normalized filter is reported in fig. 16 for $r = 0.95$ and $\omega_0 = \pi/6$. It is interesting to notice the large step experienced by the phase response around the resonance frequency. This step approaches π as the poles get closer to the unit circumference.

It is useful to draw the pole-zero diagram in order to gain intuition about the frequency response. The magnitude of the frequency response is found by taking the ratio of the product of the magnitudes of the vectors that go from the zeros to the unit circumference with the product of the magnitudes of the vectors that go from the poles to the unit circumference. The phase response is found by taking the difference of the sum of the angles of the vectors starting from the zeros with the sum of the angles of the vectors starting from the poles. If we move along the unit circumference from dc to the Nyquist rate, we see that, as we approach the pole, the magnitude of the frequency response increases, and it decreases as we move away from the pole. Reasoning on the complex plane it is also easier to figure out why there is a step in the phase response and why the width of this step converges to π as we move the pole toward the unit circumference. In the computation of the frequency response it bandwidth

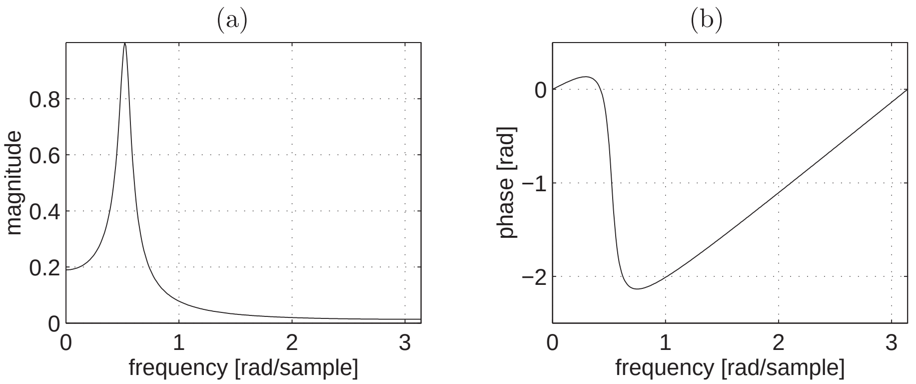

Figure 16: Frequency response (magnitude (a) and phase (b)) of a two-pole IIR filter

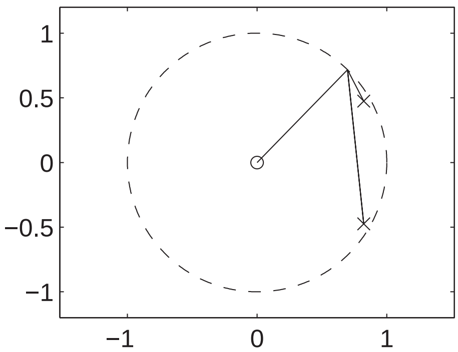

Figure 17: Couple of poles on the complex plane

is clear that, in the neighborhood of a pole close to the unit circumference, the vector that comes from that pole is dominant over the others. This means that, accepting some approximation, we can neglect the longer vectors and consider only the shortest vector while evaluating the frequency response in that region. This approximation is useful to calculate the bandwidth $\Delta\omega$ of the resonant filter, which is defined as the difference between the two frequencies corresponding to a magnitude attenuation by $3dB$, i.e., a ratio $1/\sqrt{2}$. Under the simplifying assumption that only the local pole is exerting some influence in the neighboring area, we can use the geometric construction of fig. 18 in order to find an expression for the bandwidth [67]. The segment $\overline{P_0A}$ is $\sqrt{2}$ times larger than the segment $\overline{P_0P}$. Therefore, the triangle formed by the points P_0AP has two, orthogonal, equal edges and $AB = 2P_0P = 2(1-r)$. If AB is small enough, its length can be approximated with that of the arc subtended by it, which is the bandwidth that we are looking for. Summarizing, for poles that are close to the unit circumference, the bandwidth is given by

$$\Delta\omega = 2(1-r)\ . \tag{44}$$

The formula (44) can be used during a filter design stage in order to guide the pole placement on the complex plane.

subtractive synthesis
excitation signal
allpole filter
pitch shifting
time stretching
data reduction
digital oscillator

5.1.3 LPC Modelling

As explained in section 4.2, the Linear Predictive Coding can be used to model piecewise stationary spectra. The LPC synthesis proceeds according to the feed-forward scheme of figure 5. Essentially, it is a subtractive synthesis algorithm where a spectrally-rich excitation signal is filtered by an allpole filter. The excitation signal can be the residual e that comes directly from the analysis, or it is selected from a code book. Alternatively, we can make use of voiced/unvoiced information to generate an excitation signal that can either be a random noise or a pulse train. In the latter case, the pulse repetition period is derived from pitch information, available as a parameter.

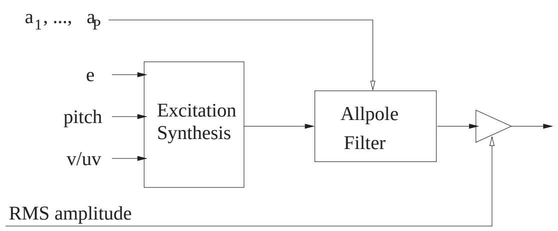

Figure 5: LPC Synthesis

Between the analysis and synthesis stages, several modifications are possible:

- pitch shifting, obtained by modification of the pitch parameter;
- time stretching, obtained by stretching the window where the signal is assumed to be stationary;
- data reduction, by model order reduction or residual coding.

5.2 Time-domain models

While the description of sound is more meaningful if done in the spectral domain, in many applications it is convenient to approach the sound synthesis directly in the time domain.

5.2.1 The Digital Oscillator

We have seen in section 5.1.1 how a complex sound made of several sinusoidal partials is conveniently synthesized by the FFT^{-1} method. If the sinusoidal components are not too many, it may be convenient to synthesize each partial by means of a digital oscillator.

From the obvious identity

$$e^{j\omega_0(n+1)} = e^{j\omega_0} e^{j\omega_0 n} \ , \tag{6}$$

said $e^{j\omega_0 n} = x_R(n) + jx_I(n)$, it is evident that the oscillator can be implemented by one complex multiplication, i.e., 4 real multiplications, at each time step:

$$x_R(n+1) = \cos\omega_0 x_R(n) - \sin\omega_0 x_I(n) \tag{7}$$

$$x_I(n+1) = \sin\omega_0 x_R(n) + \cos\omega_0 x_I(n) \ . \tag{8}$$

The initial amplitude and phase can be imposed by scaling the initial phasor $e^{j\omega_0 0}$ and adding a phase shift to its exponent. It is easy to show[2] that the calculation of $x_R(n+1)$ can also be performed as

wavetable
wavetable oscillator
increment

$$x_R(n+1) = 2\cos\omega_0 x_R(n) - x_R(n-1) \;, \tag{9}$$

or, in other words, as the free response of the filter

$$H_R(z) = \frac{1}{1 - 2\cos\omega_0 z^{-1} + z^{-2}} = \frac{1}{\left(1 - e^{-j\omega_0 z^{-1}}\right)\left(1 - e^{j\omega_0 z^{-1}}\right)} \;. \tag{10}$$

The poles of the filter (10) lay exactly on the unit circumference, at the limit of the stability region. Therefore, after the filter has received an initial excitation, it keeps ringing forever.

If we call x_{R1} and x_{R2} the two state variables containing the previous samples of the output variable x_R, an initial phase ϕ_0 can be imposed by setting[3]

$$x_{R1} = \sin(\phi_0 - \omega_0) \tag{11}$$
$$x_{R2} = \sin(\phi_0 - 2\omega_0) \;. \tag{12}$$

The digital oscillator is particularly convenient to perform sound synthesis on general-purpose processors, where floating-point arithmetics is available at no additional cost. However, this method for generating sinusoids has two main drawbacks:

- Updating the parameter (i.e., the oscillation frequency) requires computing a cosine function. This is a problem for audio rate modulations, where to compute a modulated sine we need to compute a cosine at each time sample.

- Changing the oscillation frequency changes the sinusoid amplitude as well. Therefore, some amplitude control logic is needed.

5.2.2 The Wavetable Oscillator

The most classic and versatile approach to the synthesis of periodic waveforms (sinusoids included) is the cyclic reading of a table where a waveform period is pre-stored. If the waveform to be synthesized is a sinusoid, symmetry considerations allow to store only one fourth of the period and play with the index arithmetic to reconstruct the whole period.

Call `buf[]` the buffer that contains the waveform period, or wavetable. The wavetable oscillator works by circularly accessing the wavetable at multiples of an increment I and reading the wavetable content at that position.

If B is the buffer length, and f_0 is the frequency that we want to generate at the sample rate F_s, the increment has to be set to

$$I = \frac{B f_0}{F_s} \;. \tag{13}$$

[2]The reader is invited to derive the difference equation 9

[3]The reader can verify, using formulas (29–32) of appendix A, that $x_R(0) = \sin\phi_0$, given $x_R(-1) = x_{R1}$ and $x_R(-2) = x_{R2}$.