Sampling of continuous-time signals

Alexandre Trilla, Xavier Sevillano

Departament de Tecnologies Mèdia
LA SALLE – UNIVERSITAT RAMON LLULL
Quatre Camins 2, 08022 Barcelona (Spain)
atrilla@salle.url.edu, xavis@salle.url.edu

2010

Abstract

This tutorial presents a thorough description of the sampling procedure of continuous-time signals in order to be processed by a discrete-time signal processing system.

1 Introduction

In the context of this tutorial, unidimensional digital signals are obtained from recording (and thus digitising) unidimensional analog signals that depend on time (i.e. evolve in time). As a physical example of this procedure the acoustic phenomenon is taken for convenience.

This tutorial is mostly based on [Oppenheim and Schafer, 2009], following its notation. Refer to it for further details.

2 Acoustic signal

An acoustic signal represents the temporal evolution of acoustic pressure variations in time, captured by the membrane of an electro-mechanical transducer (e.g., a microphone). These acoustic pressure variations are the local atmospheric pressure deviations caused by a sound wave, produced by some kind of acoustic source like the human voice or a musical instrument.

The electro-mechanical transducer represents sound in the form of a continuous-time (electrical) signal, and this signal needs to be sampled in order to be processed by a discrete-time signal processing system.

3 Analog-to-digital conversion

Once the acoustic signal is converted into an electric signal through the transducer, data can be acquired in order to switch to the discrete domain, where the discrete-time signal processing tools are of use. This step is accomplished by the analog-to-digital conversion (A/D conversion) process by sampling and quantising the signal. The A/D process provides a vector of quantised samples corresponding to the acoustic signal.

3.1 Periodic sampling

A sample $x_s(t)$ represents an instant value of a continuous-time signal $x_c(t)$. Generally, continuous-time signals are sampled periodically, i.e., the temporal distance between two consecutive samples is constant, which is called the *sampling period* T. The inverse of the sampling period is called the *sampling frequency* $f_s = \frac{1}{T}$ and it is measured in Hertz (Hz) or samples per second.

For convenience, no quantisation issues are treated in this section, therefore the samples $x_s(t)$ are just normalised in time, leaving $x[n]$. Note that $x[n]$ maintains the infinite precision of the signal's amplitude (as it is defined in a continuous domain), thus describing an ideal sampling process, aka continuous-to-discrete conversion process (C/D conversion).

C/D conversion is an idealised model of periodic sampling. C/D permits implementing continuous-time linear time-invariant (LTI) systems as discrete-time LTI systems if the input is bandlimited and its maximum frequency component is below the Nyquist rate (i.e. half the sampling rate).

3.1.1 Time-domain representation of sampling

The mathematical representation of sampling in C/D is shown in Eq. (1), and Figure 1 shows is graphical representation (noting that $s(t)$ corresponds to an impulse train with period T).

$$x[n] \leftarrow x_s(t) = \sum_{n=-\infty}^{\infty} x_c(t)\,\delta(t - t_n); \qquad t_n = nT = \frac{n}{f_s} \tag{1}$$

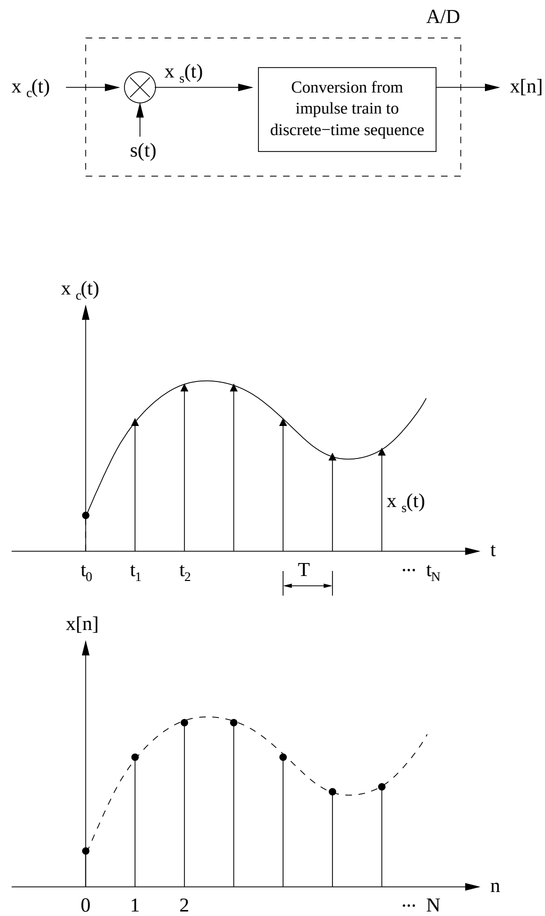

Figure 1: Time representation of the sampling process.

In C/D, the essential difference between $x_s(t)$ and $x[n]$ is that the former is, in a sense, a continuous-time signal (specifically a modulated impulse train) that is zero except at integer multiples of T, while the latter is a discrete-time signal in all senses. The sequence $x[n]$ is indexed on the integer variable n, which, in effect, introduces a time normalisation; i.e., the sequence of numbers $x[n]$ contains no explicit information about the sampling period T.

Furthermore, the samples of $x_c(t)$ are represented by finite numbers in $x[n]$ rather than as the areas of impulses as with $x_s(t)$.

3.1.2 Frequency-domain representation of sampling

Considering from Eq. (1) that

$$x_s(t) = x_c(t)\, s(t) \tag{2}$$

the transform of $x_s(t)$ is given by the Modulation or Windowing Theorem

$$X_s(\Omega) \rightarrow X_s(j\Omega) = \frac{1}{2\pi} X_c(j\Omega) * S(j\Omega) \tag{3}$$

where

$$S(j\Omega) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta(\Omega - k\Omega_s); \qquad \Omega_s = \frac{2\pi}{T} \tag{4}$$

it follows that

$$X_s(j\Omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} X_c(j(\Omega - k\Omega_s)) \tag{5}$$

A graphical representation of Eq. (5) is shown in Figure 2. Sampling representations rely only on the assumption of a bandlimited Fourier transform.

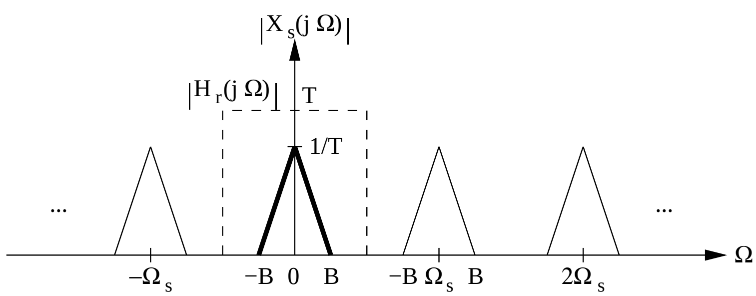

Figure 2: Frequency representation of the sampling process.

As long as the highest frequency component of the input signal remains $B \leq \Omega_s - B$, no aliasing distortion is produced and the original signal $x_c(t)$ may be recovered from $x_s(t)$ with an ideal lowpass reconstruction filter $H_r(j\Omega)$ with gain T (to compensate Eq. (5)) and cutoff frequency $B \leq \Omega_c \leq \Omega_s - B$. Otherwise, the copies of $X_c(j\Omega)$ overlap and the original signal is no longer recoverable by lowpass filtering. Hence, it is of utmost importance that no frequency component higher than half the sampling rate “poisons” the sampling procedure. By restricting the frequency content of the input signals that go into the sampler (prefiltering using continuous-time/analog lowpass filters), none of the copies of $X_c(j\Omega)$ can overlap. Therefore, it is ensured that no higher frequency continuous-time signal copy can produce the same output sequence of samples, and the aliasing distortion is avoided.

Eventually, it can be stated that

$$X_s(j\Omega) = X(e^{j\omega})|_{\omega=\Omega T} = X(e^{j\Omega T}) \tag{6}$$

where the frequency scaling or normalisation in the transformation from $X_s(j\Omega)$ to $X(e^{j\omega})$ is directly a result of the time normalisation in the transformation from $x_s(t)$ to $x[n]$.

3.1.3 Sample-and-hold

In a practical setting (i.e., A/D), the operation of sampling is implemented by a sample-and-hold system (S&H), which can be viewed as an approximation to the ideal C/D converter.

The S&H system holds the value of the continuous-time input signal during a certain period of time, namely T_Q. This single value stabilisation is needed by the quantiser that follows (in the A/D device) to output the corresponding binary code. This fact sets a limitation on the sampling frequency as it cannot be higher than $\frac{1}{T_Q}$.

The output of an ideal S&H system[1] is

$$x_{SaH}(t) = \sum_{n=-\infty}^{\infty} x_c(nT)\, h_0(t-nT) = h_0(t) * \sum_{n=-\infty}^{\infty} x_c(t)\, \delta(t-nT) \tag{7}$$

$$x_{SaH}(t) = x_s(t) * h_0(t) \tag{8}$$

where $h_0(t)$ is the impulse response of a zero-order-hold (ZOH) system, i.e.,

$$h_0(t) = \begin{cases} 1, & 0 < t < T \\ 0, & otherwise \end{cases} \tag{9}$$

$x_{SaH}(t)$ is a staircase waveform where the sample values are held constant during a sampling period, see Figure 3.

In the frequency domain, a transform of a convolution in time domain is a product of transforms

$$X_{SaH}(j\Omega) = X_s(j\Omega)\, H_0(j\Omega) \tag{10}$$

[1] It is emphasised that the S&H system of this discussion is ideal in the sense that the electronic equipment of the A/D device that implements it introduces little deviations due to physical restrictions.

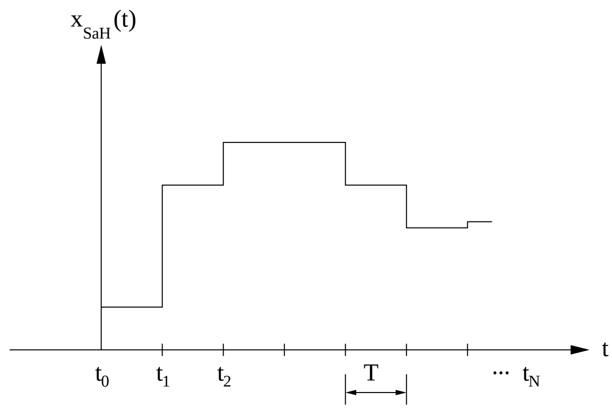

Figure 3: Time representation of the sample-and-hold system.

Note that the ideally sampled signal $X_s(j\Omega)$ (as it would be obtained with C/D) is now distorted by $H_0(j\Omega)$ due to the ZOH system of the S&H, where

$$H_0(j\Omega) = \frac{2\sin\left(\frac{\Omega T}{2}\right)}{\Omega} e^{-j\Omega\frac{T}{2}} \tag{11}$$

Figure 4 shows the frequency representation of $H_0(j\Omega)$. The gain of the ZOH only drops to $\frac{2}{\pi}$ (or -4dB) at $\Omega = \frac{\pi}{T}$. Hence, a reconstruction filter $H_r(j\Omega)$ is needed to remove the high-frequency components introduced by the discontinuities of the staircase waveform. Note that wrt the reconstruction filter used in the ideal sampling scenario, this one does not need to be amplified by T because this sampling method does not introduce the former $\frac{1}{T}$ attenuation. Nevertheless, $H_r(j\Omega)$ might have a special shape to compensate the gain drop introduced by the ZOH. Similar to $x_s(t)$ in C/D, $x_{SaH}(t)$ is discretised in time into $x[n]$.

3.2 Quantisation

Sample quantisation is a nonlinear process that consists in representing / encoding the infinite precision value of a sample $x[n]$ into one of a finite set of prescribed values (finite precision) defined by a binary codebook, thus obtaining $x_q[n]$. Given that $x[n] \in \mathbb{R}$ and $Q+1$ bits define a codebook range of $2^{Q+1} \in \mathbb{N}$ symbols (amplitude levels), quantising a sample implies approximating its real value to one of these possible integer values, see Figure 5.

Quantisers may be defined with uniform or nonuniform quantisation levels Δ, yielding linear or nonlinear quantisers. The Δ parameter determines

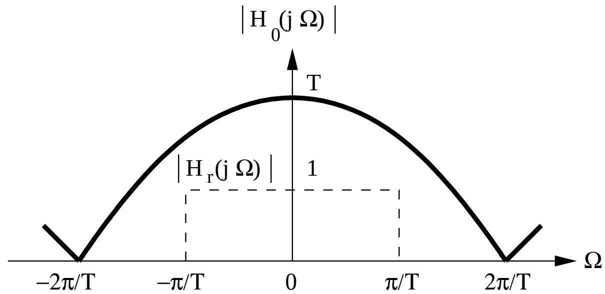

Figure 4: Frequency representation of the zero-order-hold system and the reconstruction filter.

the full-scale level of the A/D converter. In general,

$$\Delta = \frac{2X_m}{2^{Q+1}} = \frac{X_m}{2^Q} \tag{12}$$

where $Q+1$ represents the number of bits of the words of the code, including the sign of the value, and X_m represents the maximum signal peak allowed. Note that the range of the codebook is adjusted to the dynamic range of the A/D, i.e. $2X_m$. Regarding the actual code that represents the values, it is generally preferred a binary code that permits doing arithmetic directly with the code words as scaled representations of the quantised samples, e.g., the two's complement code.

3.2.1 Analysis of the effects of quantisation

Since the quantisation procedure is basically an approximation $x_q[n] \sim x[n]$, it always produces an error signal $e_q[n]$:

$$e_q[n] = x_q[n] - x[n] \tag{13}$$

In general,

$$|e_q[n]| < \frac{\Delta}{2} \tag{14}$$

except when $x[n]$ is outside the allowed range of the quantiser ($x[n] > X_m$). Such samples are said to be clipped, and the quantiser is said to be overloaded. Note that Eq. (14) has this form due to the use of the rounding function in implementing the quantisation.

Given that $e_q[n]$ is generally unknown, a simplified model of the quantisation process is used considering this uncertainty: $e_q[n]$ is modelled with a random variable e assuming that

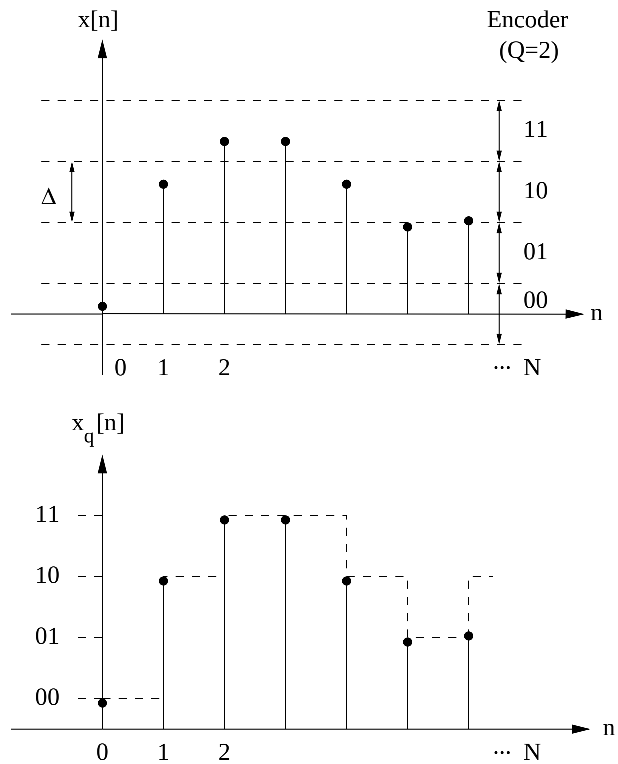

Figure 5: Quantisation process.

1. e is a sample sequence of a stationary random process, modelling the error process.

2. e is uncorrelated with $x[n]$.

3. The error process is a white-noise process.

4. The probability distribution of the error process is uniform over the range of the quantisation error: $P_e = U[-\frac{\Delta}{2}, \frac{\Delta}{2}]$.

The preceding assumptions lead to an effective analysis of quantisation that is useful to predict the system's performance, especially when the input signal is complicated and fluctuates in an unpredictable manner. The error signal e is referred to as *quantisation noise*.

A common measure of the amount of degradation of a signal by additive noise in general is the signal-to-noise ratio (SNR), defined as the ratio of signal variance (power) to noise variance

$$SNR = 10\log\left(\frac{\sigma_x^2}{\sigma_e^2}\right) \tag{15}$$

Regarding the average power of quantisation noise, i.e., the second order moment of an uniform distribution

$$E(e^2) = \int_{-\Delta/2}^{\Delta/2} e^2\,\frac{1}{\Delta}\,de = \frac{\Delta^2}{12} = \sigma_e^2 \tag{16}$$

note that by increasing one bit the number of bits used for linear quantisation, the average power of quantisation noise decreases 6dB approximately, thus increasing 6dB the SNR:

$$10\log\sigma_e^2 = Q\cdot 10\log 0.25 + 10\log X_m^2 - 10\log 12 \tag{17}$$

3.2.2 Quantisation with a compander

Companding consists in compressing and expanding the dynamic range of an analog signal. The name is a portmanteau of compressing and expanding. In digital systems, companding can reduce the quantisation error (hence increasing the signal to quantisation noise ratio).

The reason of their practical use is related to the properties of everyday harmonic signals, like the voice. These typical signals carry a lot more information in the amplitude range of low values than in the range of high values; their sample value histogram is thus centred at zero amplitude. Therefore, it is sensible to raise the presence of low-valued samples versus high-values samples before entering a linear quantiser.

In this sense, the μ-law algorithm is taken as an example of companding algorithm. The compressing function is shown in Eq. (18) and its inverse function, the expanding function, is shown in Eq. (19).

$$f(x(t)) = \text{sgn}(x(t))\,\frac{\ln(1+\mu|x(t)|)}{\ln(1+\mu)} \qquad |x(t)| \leq 1 \tag{18}$$

$$f^{-1}(y(t)) = \text{sgn}(y(t))\,(1/\mu)\,((1+\mu)^{|y(t)|}-1) \qquad |y(t)| \leq 1 \tag{19}$$

Note how the μ-law algorithm actually amplifies low-valued signals. For example, observe that the derivative of the compressing function

$$\frac{d}{dx}f(x(t)) = \frac{\mu}{\ln(1+\mu)\,(1+\mu x(t))} \qquad x(t) > 0 \tag{20}$$

amplifies $x(t)$ when it takes low amplitude values, i.e. Eq. (21), otherwise it flattens $x(t)$ (the gain becomes lower than the unit).

$$x(t) < \frac{1}{\mu}\left(\frac{\mu}{\ln(1+\mu)} - 1\right) \tag{21}$$

Also note that the average power of quantisation noise in Eq. (17) assumes that the input signal $x(t)$ covers the whole dynamic range of the quantiser $2X_m$, where all the quantising discontinuities affect the signal. In the compander scenario, this assumption is accurate and favourable only when

$$x(t) < f(x(t)) \leq X_m \tag{22}$$

Alternatively, in the sole linear quantiser scenario, less quantising discontinuities cover the range of the same low-valued input signal $x(t)$. This could be interpreted as a decrease of bits needed to quantise, and according to Eq. (17), it increases the average power of quantisation noise.

4 Digital-to-analog conversion

The digital-to-analog conversion or D/A conversion is the inverse process of A/D, i.e., the reconstruction of a bandlimited original signal $x(t)$ from its discrete-time samples $x[n]$. Note that 1) a high sampling frequency and 2) many bits per sample are desirable to obtain a good signal reconstruction. With (1) more frequency components are captured, and with (2) the quantisation noise is minimised.

4.1 Bandlimited interpolation

The original continuous-time signal $x_c(t)$ may be recovered from its samples $x[n]$ with knowledge of the sampling period T. To this end, the impulse train modulation provides a convenient means for understanding this reconstruction process.

If the condition $B \leq \frac{\Omega_s}{2}$ imposed by the sampling theorem (in order to avoid aliasing) is met, and the modulated impulse train is filtered by an appropriate ideal lowpass filter $H_r(j\Omega)$ with cutoff frequency $\Omega_c \leq \frac{\Omega_s}{2}$ and gain T, then the Fourier transform of the output of the filter $X_r(j\Omega)$ will be identical to the Fourier transform of the original continuous-time signal $X_c(j\Omega)$, and thus, the output of the filter will be $x_c(t)$, see Figure 6. This is an idealised D/A model of bandlimited interpolation.

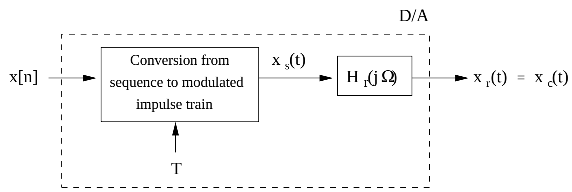

Figure 6: Idealised D/A model of bandlimited interpolation.

4.2 Linear interpolation

Given that ideal lowpass filters for signal reconstruction cannot be implemented, in some cases simple interpolation procedures are just adequate. In linear interpolation, the reconstructed signal $x_r(t)$ between two original consecutive samples $x_1 = x[n]$ and $x_2 = x[n+1]$ is described by a linear equation

$$x_r(t) = \frac{x_2 - x_1}{T}\, t + x_1 \qquad 0 \leq t < T \tag{23}$$

The piecewise linear approximation of the original signal described in Eq. (23) is attained with a first-order-hold (FOH) system

$$x_{FOH}(t) = x_s(t) * h_{lin}(t) \tag{24}$$

where $h_{lin}(t)$ is a triangular function

$$h_{lin}(t) = \begin{cases} 1 - \frac{|t-T|}{T}, & |t - T| < T \\ 0, & otherwise \end{cases} \tag{25}$$

$H_{lin}(j\Omega)$ is determined using the convolution property of the Fourier transform and the transform of the rectangular function with duration T and amplitude $\frac{1}{T}$

$$H_{lin}(j\Omega) = \mathrm{sinc}^2\left(\frac{T}{2}\Omega\right)\, e^{-j2\pi\Omega T} \tag{26}$$

Note that the FOH is more enhanced than the ZOH for signal reconstruction as its transitions in time are a little smoother, and therefore less energy is dispersed to higher frequencies.

References

[Oppenheim and Schafer, 2009] Oppenheim, A. V. and Schafer, R. W. (2009). *Digital Signal Processing*. Prentice–Hall, third edition.