

# Sound, signals and sampling

How does a sound become a list of numbers, and how do we get it back?

Video: [professorkettlewick.com/lectures/sound-signals-and-sampling/](https://professorkettlewick.com/lectures/sound-signals-and-sampling/) • An AI-generated course and presenter. Made by DigitalZoo.ai.

Checked by two AI reviewers against cited sources, 28 October 2026; lecturer review pending. Free to use, share and adapt for teaching and learning, with credit to Professor Kettlewick (DigitalZoo.ai); not for commercial use.

## Learning outcomes

### A Level

1. Read amplitude, period and frequency from a waveform and use  $f = 1/T$ .
2. Explain sampling, and use  $f_s$  and  $T_s = 1/f_s$ .
3. State and apply the sampling theorem,  $f_s > 2f_{\max}$ , and the Nyquist frequency  $f_s/2$ .
4. Explain aliasing, and why the anti-aliasing low-pass filter comes first in a PCM transmitter.

### First-year university

1. Show that  $f$  and  $f + kf_s$  give identical samples, and calculate any alias frequency by folding.
2. Explain why the anti-aliasing filter must be steep, and how oversampling eases it.
3. Explain reconstruction: why, ideally, a filtered DAC output is the original band-limited wave, not a staircase.

## Key equations

frequency and period

$$f = \frac{1}{T}$$

sampling period

$$T_s = \frac{1}{f_s}$$

the sampling theorem

$$f_s > 2f_{\max}$$

the Nyquist frequency

$$f_N = \frac{f_s}{2}$$

any alias,  $k$  the whole number that makes it smallest

$$f_{\text{alias}} = |f - kf_s|$$

# Notes: A Level

## What a signal is (video 0:21)


When a guitar string vibrates, it squashes and stretches the air beside it, and those squashes and stretches (compressions and rarefactions) travel outwards as a pressure wave at about 343 m/s in air at 20 °C. A microphone's diaphragm moves with the pressure, and the microphone turns that movement into a voltage: more pressure, higher voltage. That changing voltage is a **signal**, a quantity that changes over time and carries information.

The simplest signal is a **sine wave**, the sound of a tuning fork. Its **amplitude** is the height from the middle line to a peak, and a bigger amplitude sounds louder. Its **period**  $T$  is the time for one complete cycle, in seconds, and its **frequency**  $f$  is the number of cycles each second, in hertz (Hz). They are two views of the same thing:

$$f = \frac{1}{T}$$

The A that orchestras tune to is 440 Hz, so its period is  $1/440 \approx 2.27$  ms. Doubling a frequency raises a note by an octave; halving it lowers it by one.


**Phase** says where in its cycle a wave is at time zero, in degrees or radians. You cannot hear the phase of a single tone, but when signals are added together, phase decides whether they reinforce or cancel.



Period and frequency:  $f = 1/T$ .

## The range of hearing (video 3:11)

Young, healthy ears hear from about 20 Hz up to about 20 kHz. The top end falls with age, and many adults hear little above 15 or 16 kHz. That 20 kHz figure is what decides how fast audio has to be sampled.




Young, healthy hearing runs from about 20 Hz to about 20 kHz.

## Why go digital? (video 4:10)

The microphone's voltage varies smoothly and continuously. It is **analogue**: a likeness, an analogue, of the pressure wave. Every time an analogue signal is copied, stored or sent, a little noise and distortion creep in, and in general they cannot be taken out again; copy a tape, then copy the copy, and by the fifth generation you can hear the hiss.

A **digital** signal is a list of numbers, and a number can be copied a million times without changing. Numbers can also be processed by a computer: filtered, mixed, compressed and sent round the world.

Turning a voltage into numbers takes two steps. **Sampling** measures it at regular moments in time, and is this lecture's subject. **Quantisation** rounds each measurement to one of a fixed set of levels, and is Lecture 2's. Sampling chops up time; quantisation chops up the voltage.



**sampling: chops up time**

**quantisation: chops up the voltage (next lecture)**

Sampling chops up time; quantisation chops up the voltage.


## Sampling (video 5:33)

Imagine photographing the voltage at regular intervals and writing down each value. Each measurement is a **sample**. The number of samples taken each second is the **sampling rate**, or sampling frequency,  $f_s$ , in hertz. The time between samples is the **sampling period**:

$$T_s = \frac{1}{f_s}$$

A CD samples at 44.1 kHz, 44 100 samples every second, so its samples are  $1/44\,100 \approx 22.7 \mu\text{s}$  apart. Once a signal has been sampled the smooth wave is gone and only the dots remain. The question for the whole lecture is whether the wave can be got back from them.

On average, each cycle of a wave gets  $f_s/f$  samples (any one cycle holds a whole number of them). At 44.1 kHz a 1 kHz tone gets 44.1 samples per cycle, and a 10 kHz tone gets 4.41: fewer, but still enough.



A sample at every tick of the sampling clock,  $T_s$  apart.

## The sampling theorem (video 7:33)

To capture a frequency you need more than two samples in each of its cycles. Put the other way round, the sampling rate must be more than twice the highest frequency in the signal:

$$f_s > 2f_{\max}$$


This holds for signals that start at 0 Hz, such as audio. (A signal confined to a narrow band far above 0 Hz can be sampled more slowly, by a technique called bandpass sampling.)

This is the **sampling theorem**, often called the Nyquist theorem or the Nyquist–Shannon theorem. Its discovery was shared: E. T. Whittaker found the rebuilding formula in 1915, Harry Nyquist's 1928 telegraph work gave the factor of two, Vladimir Kotelnikov stated and proved it for communications in 1933, and Claude Shannon's 1949 paper, which gave the full proof for engineers and called the result common knowledge, made it famous. Half the sampling rate,  $f_s/2$ , is the **Nyquist frequency**: the limit for a given sampling rate, since only frequencies below it can be captured.

Why more than two, and not exactly two? At exactly two samples per cycle the samples might land on the peaks, and the wave shows up, though its true size and phase still cannot be told. But they might equally land on the zero crossings, and then every sample is zero and the wave vanishes. Exactly two is not safe; just over two is enough.

Books differ in their words. Here, as in most audio work, the Nyquist frequency is  $f_s/2$ . Oppenheim and Schaffer instead call a signal's highest frequency its Nyquist frequency and twice that its Nyquist rate, and state the

theorem with  $\geq$  for a signal that has nothing at or above that frequency; Smith uses both names for  $f_s/2$ .



Below the Nyquist frequency,  $f_s/2$ , a tone can be captured.


### Why a CD samples at 44.1 kHz (video 8:51)

Hearing tops out at about 20 kHz, so audio needs a sampling rate above 40 kHz. A CD's 44.1 kHz gives a Nyquist frequency of 22.05 kHz, with a little room to spare, which the anti-aliasing filter needs (see below). The oddly precise number comes from history: the first digital recorders stored their samples on video tape, and 44 100 samples a second fitted the lines of a television picture exactly. In the European system that is 3 samples of each stereo channel on each of 294 usable lines, 50 fields (half-pictures) a second:  $3 \times 294 \times 50 = 44\,100$ . In the American system it is  $3 \times 245 \times 60$ , at the 60 fields a second of black-and-white television; colour American television runs at 59.94, which gives 44 056.

Different jobs use different rates:

| Use                   | Sampling rate | Nyquist frequency |
|-----------------------|---------------|-------------------|
| Traditional telephone | 8 kHz         | 4 kHz             |
| CD                    | 44.1 kHz      | 22.05 kHz         |
| Film and video        | 48 kHz        | 24 kHz            |
| Studio recording      | 96 kHz        | 48 kHz            |

A traditional telephone line carries about 300 Hz to 3.4 kHz, leaving room for the filter below its 4 kHz Nyquist frequency: enough to understand speech, but music loses its sparkle. Newer wideband ("HD voice") calls carry up to 7 kHz. Studios record at higher rates for more room while they work.




Sampling rates in use, each with its Nyquist frequency.

## Aliasing (video 10:50)

Break the rule and something strange happens. A film camera samples the world, usually 24 times a second. When a wheel's spoke moves almost a full spoke-spacing between one frame and the next, the camera cannot tell that it went nearly all the way forwards; the nearest match is a small step backwards. So a fast wheel seems to turn slowly backwards.

Sound does exactly the same. Sample a 7 kHz tone at 8 kHz, far fewer than two samples per cycle, and a gentle 1 kHz wave passes through every one of the samples. From the samples alone there is no way to tell which wave was there, and playing them back gives the 1 kHz tone. This impostor frequency is an **alias**, and the effect is **aliasing**.

A tone rising steadily past the Nyquist frequency is heard to turn round and fall, even though the real tone keeps rising. Everything above the Nyquist frequency folds back down, like a reflection in a mirror at  $f_s/2$ ; above  $f_s$  the folding repeats, as the first-year part shows.




A 7 kHz tone sampled at 8 kHz: a 1 kHz wave passes through every sample, so we hear 1 kHz.

## The anti-aliasing filter (video 12:54)

Once aliasing has happened, it cannot be undone from the samples alone. The alias sits among the real sound, and nothing done afterwards can tell them apart. So the only cure is prevention: before a signal is sampled, it passes through a low-pass filter that removes everything above the Nyquist frequency (a real filter cuts it down until any alias is too quiet to matter). This is the **anti-aliasing filter**, the first block of an analogue-to-digital system.

A real filter cannot cut off like a cliff edge. It needs a range of frequencies, its **transition band**, to go from letting sound through to blocking it. On a CD that range is the gap between 20 kHz, the top of hearing, and 22.05 kHz, the Nyquist frequency. That is what the spare room above 20 kHz is for.

That is the safe, simple rule. Strictly, the filter can keep falling until 24.1 kHz, because anything between 22.05 and 24.1 kHz folds back to somewhere above 20 kHz, out of the band we hear; the first-year part shows why. Real converters use that extra room.




The anti-aliasing filter cuts down what is above  $f_s/2$ , before the sampler.

## The PCM transmitter (video 13:29)

In Eduqas (WJEC) A Level Electronics this is the transmitter of a **pulse code modulation (PCM)** system. Its blocks, in order:

1. **Low-pass filter:** removes frequencies above  $f_s/2$  so they cannot alias.
2. **Sampling gate**, driven by the sampling clock: takes a sample at each clock pulse.
3. **Analogue-to-digital converter (ADC):** turns each sample into a binary number.
4. **Parallel-in serial-out (PISO) shift register**, with its own clock: sends the number's bits out one at a time.

The low-pass filter comes first for exactly the reason above. The receiver mirrors the transmitter: a Schmitt trigger, a serial-in parallel-out (SIPO) shift register with its clock, a digital-to-analogue converter (DAC), and a second low-pass filter that smooths the output.



The PCM transmitter: filter, sampling gate, ADC, shift register.

# Notes: First-year university


## Frequencies that give the same samples (video 16:23)

Write a continuous tone as  $x(t) = A \cos(2\pi f t + \varphi)$  and keep it only at the sample times  $t = nT_s$ , where the integer  $n$  ( $\dots, -1, 0, 1, 2, \dots$ ) counts the samples. Since  $T_s = 1/f_s$ , the samples are

$$x[n] = A \cos \left( 2\pi n \frac{f}{f_s} + \varphi \right)$$

The samples depend only on the ratio  $f/f_s$ , the frequency measured in cycles per sample. Now replace  $f$  with  $f + f_s$ . The ratio gains 1, so the angle gains  $2\pi n$ : a whole number of complete turns, which changes nothing. So  $f$ ,  $f + f_s$ ,  $f - f_s$ , and  $f + kf_s$  for any integer  $k$  (positive, negative or zero), all give identical samples.

Cosine is an even function,  $\cos(-\theta) = \cos \theta$ , so a tone at  $-f$  gives the same samples as a tone at  $f$  with its phase reversed ( $\varphi$  becomes  $-\varphi$ ). Put the two facts together and  $f_s - f$  matches  $f$  as well, again with its phase reversed: the mirror at the Nyquist frequency. For a cosine with zero phase the samples are identical; for a sine they come out upside down.



Whole turns change nothing, so  $f$  and  $f + f_s$  give the same samples;  $f_s - f$  gives them too, with its phase reversed.


## Finding any alias (video 18:37)

To find the alias of any frequency, shift it by whole multiples of  $f_s$  until it lands between  $-f_s/2$  and  $+f_s/2$ , then take its size:

$$f_{\text{alias}} = |f - kf_s| \quad \text{with } k \text{ the whole number that makes it small}$$

At  $f_s = 48$  kHz, 30 kHz gives  $|30 - 48| = 18$  kHz, 50 kHz gives 2 kHz, and 47 kHz gives 1 kHz. A tone at exactly 24 kHz, the Nyquist frequency, maps to itself, but it is the edge case where the samples can vanish.

Drawn as the frequency that comes out against the frequency that goes in, the rule is a straight line up to  $f_s/2$ , then a zigzag: down to zero at  $f_s$ , up to  $f_s/2$  again at  $1.5f_s$ , and so on. Every input frequency, however high, lands somewhere between 0 and  $f_s/2$ .




At 44.1 kHz, 25 kHz is heard at  $|25 - 44.1| = 19.1$  kHz.

### How steep the filter must be, and oversampling (video 20:36)

At 44.1 kHz anything above 22.05 kHz folds back, and the fold reaches down into the band we hear: 24.1 kHz lands exactly on 20 kHz. So the anti-aliasing filter must pass everything up to 20 kHz untouched and all but silence everything above 24.1 kHz. That is a very steep filter.

Modern converters avoid building it in analogue by **oversampling**. They sample far faster, often 64 times faster, so that a gentle analogue filter is enough; then they filter digitally and keep one sample in every 64. Lecture 4 shows how.

oversampling:  
lecture 4



Oversampling 64 times: a gentle analogue filter is enough.

## Getting the wave back (video 21:32)

A digital-to-analogue converter can first produce a staircase, holding each sample's value until the next one arrives. The sharp edges of those steps are made of frequencies above the Nyquist frequency: copies of the signal, called **images**, around  $f_s$ ,  $2f_s$  and so on. A low-pass filter after the converter, the **reconstruction** or anti-imaging filter, removes them. Holding each value also tilts the top of the band down a little, by about 3 dB at 20 kHz for a CD, so real converters correct for that in the filter or oversample until the tilt is negligible. Provided the original held nothing at or above the Nyquist frequency, what is left is then not an approximation of the original wave: in the ideal case it *is* the original wave, slightly delayed and apart from the rounding of each sample, which is Lecture 2's subject.

To see why, give every sample a bump shaped like the **sinc function**,

$$\text{sinc}(x) = \frac{\sin \pi x}{\pi x} \quad (x \neq 0), \quad \text{sinc}(0) = 1$$


which is 1 at its own sample time and exactly 0 at every other sample time. Scale each bump by its sample's value and add them all up:

$$x(t) = \sum_{n=-\infty}^{\infty} x[n] \text{sinc}\left(\frac{t - nT_s}{T_s}\right)$$

The sum passes through every sample and wiggles no faster than the Nyquist frequency allows. This is the **Whittaker-Shannon interpolation formula**, and in theory it rebuilds a band-limited signal exactly from its samples, using every sample, past and future. Real filters approximate it.

So digital audio is not steppy: below the Nyquist frequency, ideal sampling loses nothing. It also explains why more samples are not automatically better. Once the sampling rate is safely above twice the highest frequency

you care about, extra samples make the filters easier to build, rather than the sound itself better.



The reconstruction filter passes the signal below  $f_s/2$  and strongly cuts everything above, where the images are.

## Questions from the video

**L1-Q1** A Level · video 2:20

What is the period of a 440 Hz note (concert A)?

**L1-Q2** A Level · video 6:37

How many samples per cycle does a 1 kHz tone get when sampled at 44.1 kHz?

**L1-Q3** A Level · video 7:08

How many samples per cycle does a 20 kHz tone get at 44.1 kHz?

**L1-Q4** A Level · video 10:29

What is the lowest sampling rate that can capture a bat's call reaching 100 kHz?



**L1-Q5** First-year university · video 19:05

At a 48 kHz sampling rate, what do 30 kHz, 50 kHz and 47 kHz tones alias to?



**L1-Q6** First-year university · video 19:32

At 44.1 kHz, what does a 25 kHz tone alias to?



**L1-Q7** A Level · video 15:33

A signal contains frequencies up to 12 kHz. What is the minimum sampling rate?

**L1-Q8** A Level · video 15:41

A system samples at 32 kHz. What is its Nyquist frequency, and could it capture an 18 kHz tone?

**L1-Q9** A Level · video 15:52

In a pulse code modulation transmitter, why must the low-pass filter come before the sampling gate, and not after it?

## Practice set

**L1-A1** A Level

A tone has a period of 0.5 ms. What is its frequency?

**L1-A2** A Level

A system samples at 48 kHz. What is its sampling period?

**L1-A3** A Level

A signal contains frequencies up to 15 kHz. What condition must the sampling rate satisfy?

**L1-A4** A Level

On average, how many samples does each cycle of a 5 kHz tone get at 44.1 kHz?

**L1-A5** A Level

A system samples at 22.05 kHz. What is its Nyquist frequency, and can it capture a 12 kHz tone?

**L1-A6** A Level

List the blocks of a PCM transmitter in order and say what each does.

**L1-A7** A Level

Explain why a filmed wheel can appear to turn slowly backwards.

**L1-U1** First-year university

At  $f_s = 48$  kHz, find the alias of 40 kHz, 70 kHz and 100 kHz.

**L1-U2** First-year university

Show that  $\cos(2\pi(f_s - f)t)$  gives the same samples as  $\cos(2\pi ft)$ . What changes if the tone has a phase  $\varphi$ ?

**L1-U3** First-year university

Write the first four samples of  $\cos(2\pi \cdot 3000 t)$  at  $f_s = 8$  kHz.  
Which other frequency below 8 kHz gives the same samples?

**L1-U4** First-year university

At  $f_s = 44.1$  kHz, which input frequencies alias into the audio band, 0 to 20 kHz, with  $k = 1$  (the copies around  $f_s$ )?

**L1-U5** First-year university

A sigma-delta converter oversamples 64 times a 44.1 kHz rate. What are its sampling rate and Nyquist frequency, and why does that help?

**L1-U6** First-year university

Why does a zero-order-hold (staircase) DAC output contain frequencies above  $f_s/2$ , and what removes them?

## Answers

**L1-Q1.**  $T = 1/f = 1/440 = 0.00227 \text{ s} = 2.27 \text{ ms}$ . Check the size: 440 cycles a second is a little under 500, so each cycle lasts a little over  $1/500 \text{ s} = 2 \text{ ms}$ .

**L1-Q2.** Samples per cycle, on average  $= f_s/f = 44\,100/1000 = 44.1$ .

**L1-Q3.**  $44\,100/20\,000 = 2.205$  samples per cycle on average: only just above two, the limit the sampling theorem sets.

**L1-Q4.**  $f_s > 2 \times 100 \text{ kHz}$ , so the sampling rate must be more than 200 kHz. 200 kHz itself, twice the highest frequency, is the boundary: sampled at exactly 200 kHz, a 100 kHz tone can vanish.

**L1-Q5.** The Nyquist frequency is 24 kHz, so all three fold. With  $k = 1$ :  $|30 - 48| = 18 \text{ kHz}$ ,  $|50 - 48| = 2 \text{ kHz}$  and  $|47 - 48| = 1 \text{ kHz}$ .

**L1-Q6.** 25 kHz is above the Nyquist frequency of 22.05 kHz. With  $k = 1$ :  $|25 - 44.1| = 19.1 \text{ kHz}$ , inside the audio band, 0 to 20 kHz.

**L1-Q7.**  $f_s > 2 \times 12 \text{ kHz}$ , so the sampling rate must be more than 24 kHz. 24 kHz itself, twice the highest frequency, is the boundary: sampled at exactly 24 kHz, a 12 kHz tone can vanish.

**L1-Q8.** Nyquist frequency  $= 32/2 = 16 \text{ kHz}$ . 18 kHz is above it, so no. Without an anti-aliasing filter it would alias to  $32 - 18 = 14 \text{ kHz}$ ; with the filter, it is cut down before sampling until that alias is too quiet to matter.

**L1-Q9.** Sampling is what creates the aliases. Once a frequency above  $f_s/2$  has been sampled, its alias cannot be separated from the real signal, so a filter after the sampling gate would be too late. Frequencies above  $f_s/2$  must be removed before sampling.

**L1-A1.**  $f = 1/T = 1/(0.0005 \text{ s}) = 2000 \text{ Hz} = 2 \text{ kHz}$ .

**L1-A2.**  $T_s = 1/f_s = 1/48\,000 = 20.8 \mu\text{s}$ .

**L1-A3.**  $f_s > 2 \times 15 \text{ kHz}$ : more than 30 kHz. 30 kHz itself, the boundary, is not enough.

**L1-A4.**  $44\,100/5000 = 8.82$  samples per cycle, on average.

**L1-A5.** Nyquist frequency  $= 22.05/2 = 11.025 \text{ kHz}$ . 12 kHz is above it, so no. Without an anti-aliasing filter it would appear as an alias at  $22.05 - 12 = 10.05 \text{ kHz}$ .

**L1-A6. Low-pass filter:** removes frequencies above  $f_s/2$  so they cannot alias. **Sampling gate with sampling clock:** takes a sample at each clock

pulse. **ADC:** turns each sample into a binary number. **PISO shift register with its clock:** sends the number's bits out one at a time.

**L1-A7.** A film camera usually samples 24 times a second. If a spoke moves just under one spoke-spacing between frames, the nearest match from one frame to the next is a small step backwards, so the fast forward motion looks like a slow backward one: aliasing in time.

**L1-U1.** 40:  $|40 - 48| = 8$  kHz. 70:  $k = 1$  gives  $|70 - 48| = 22$  kHz ( $k = 2$  gives 26, larger). 100:  $k = 2$  gives  $|100 - 96| = 4$  kHz.

**L1-U2.**

$\cos(2\pi(f_s - f)n/f_s) = \cos(2\pi n - 2\pi fn/f_s) = \cos(-2\pi fn/f_s)$ , since adding  $2\pi n$  changes nothing for an integer  $n$ ; and  $\cos(-\theta) = \cos \theta$ , so it equals  $\cos(2\pi fn/f_s)$ . With a phase,  $\cos(2\pi(f_s - f)n/f_s + \varphi) = \cos(2\pi fn/f_s - \varphi)$ : the samples of a tone at  $f$  with its phase reversed. So a sine at  $f_s - f$  gives the samples of a sine at  $f$  turned upside down.

**L1-U3.** The angles  $0^\circ, 135^\circ, 270^\circ, 405^\circ$  give 1,  $-0.707$ , 0,  $0.707$ . The same samples come from 5 kHz ( $8 - 3$ ): its angles  $0^\circ, 225^\circ, 450^\circ, 675^\circ$  give 1,  $-0.707$ , 0,  $0.707$ .

**L1-U4.** Those within 20 kHz of  $f_s$ :  $|f - 44.1| \leq 20$  gives 24.1 kHz to 64.1 kHz. This is why the anti-aliasing filter must be deep by 24.1 kHz.

**L1-U5.**  $64 \times 44.1$  kHz = 2.8224 MHz; its Nyquist frequency is 1.4112 MHz. Aliases now come only from far above the audio band, so a gentle analogue filter is enough; the steep filtering is done digitally before the rate is reduced.

**L1-U6.** Each step's sudden jump needs high frequencies to form; in the spectrum they appear as images of the signal around  $f_s, 2f_s$  and so on. The reconstruction (anti-imaging) low-pass filter after the DAC removes them. Holding each value also droops the top of the band: the hold's response is  $\text{sinc}(f/f_s)$ , about  $-3.2$  dB at 20 kHz for a CD. So the filter also corrects that droop, or the converter oversamples until it is negligible; what is left is then the band-limited original.