The basics of digital audio

By Steve Cawser

Many modern sound level meters have the ability to record digital audio as part of a measurement. This gives the ability to either carry out detailed post-processing on sections of recorded audio data, or enables users to carry out noise source identification for long-term unattended surveys. Both of these options have a multitude of benefits for acoustics consultants. However, even though advances in instrumentation have significantly simplified the operation of modern sound level meters, the audio recording options can introduce new settings with which some may be unfamiliar.

To explain the correct use of these settings, it is necessary to understand the basics of digital audio. Digital audio is essentially a method of representing an analogue signal using a series of numbers that can be understood by a computer. This requires converting the analogue signal into discrete steps in both time and level. The figure below gives a diagrammatic representation of this process, with the red dots representing the digital representation of the blue analogue signal.

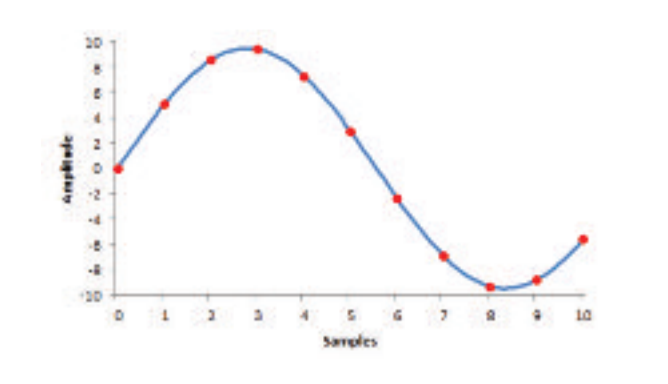

Analogue to digital conversion process

The process of capturing the signal at a discrete time interval is called sampling. This is the rate at which the amplitude of the analogue signal is acquired. The rate at which a signal is sampled is measured in terms of the number of samples per second, or frequency. The sample frequency is usually set to be high enough to ensure the samples accurately represent the signal to be acquired. This is traditionally a frequency at least twice the bandwidth of the signal being digitised. As an example, an audio compact disc has a sample frequency of 44,100 Hz.

The process of capturing the signal at discrete amplitude intervals is called quantisation. This process involves taking the analogue signal and assigning a number to represent the amplitude at the time the sample is taken. As an example, an audio compact disc uses 16-bit numbers for amplitude quantisation, which allows the signal to be represented by an integer number in the range -32768 to +32767, a total of 65536 (216) possible discrete amplitude levels.

So what does all this mean when setting up your sound level meter? The first thing to decide is for what purpose the audio will be used. If you are capturing short samples of audio for noise source identification in environmental monitoring, a sample rate of 8 kHz is often sufficient to decide whether you can hear distant road traffic noise or overflying aircraft.

If you are capturing audio for post-processing purposes, it is good practice to capture at a higher sample rate of either 24 kHz or even 48 kHz. If you are not sure, go for the highest sample rate to ensure you have the widest bandwidth in the recorded audio. However, it is worth noting that higher sample rates generate larger files, so if you double to the sample rate, you will also double the size of your files.

The other consideration is what quantisation level to use. Many sound level meters will give you the option to select either 16-bit or 24-bit audio. This is possible because many sound level meters use 24-bit processing to produce the parameters we ask it to measure, so can often natively handle 24-bit audio. However, there are two major considerations in the choice of bit depth.

The first relates to hardware that will be used to play back the audio. The sound level meter will store the recorded audio as a .wav file, which can be played back on your office computer. However, the audio hardware included in many office PCs may not be capable of playing back 24-bit audio. 24-bit audio is still considered to be a “professional” standard by many PC makers and on-board sound cards may not be able to play them natively. Therefore, if you want to play them back on a standard office PC, or are not sure if your sound card can handle 24-bit audio, it is safest to record your audio at 16-bit to ensure maximum compatibility.

However, this does have another consideration. The bit depth used to record the audio sets the dynamic range of the signal. For 16-bit audio, this dynamic range is approximately 90 dB ($20\text{x}\log_{10}(2^{16-1})$). For a 24-bit signal, this will increase to approximately 138 dB ($20\text{x}\log_{10}(2^{24-1})$). A 24-bit audio processor will have a dynamic range that far exceeds the analogue hardware which is placed in front of the digital hardware, which is likely to be less than 120 dB dynamic range. There are, of course, other considerations in the sound level meter specification that limit this potential dynamic range being fully realised.

It should also be considered that any .wav files recorded in 24-bit resolution will store the numbers as a 32-bit number in the .wav file, which means that a 24-bit audio file will be twice the size of a 16-bit file. Another confusing parameter is the type of WAV file. PCM format is “standardise”’ by Microsoft for 16-bit files for maximum compatibility with playback hardware. However, in 32-bit format, there is no similar standardisation. “IEEE extensible” is often used, but this is also not always supported by post-processing software, so it’s useful to have access to a WAV editor such as Goldwave or Audacity to enable transcoding to different WAV formats.

The other aspect of recording in 16-bit is that because the audio file will be stored using a lower dynamic range than the internal processing of the sound level meter, there will need to be some adjustment for the gain of the signal being recorded. This will need setting carefully because, as with all audio systems, if you set the gain too high, you will cause an overload and the signal will distort; set this too low and most of your signal will be lost inside the noise floor.

So in summary, if you are looking to simply identify noise with little or no post measurement analysis, then you may only need 16-bit data with a sample rate sufficient for identification purposes. However, if you need to record data with a view to carrying out detailed analysis, then higher sample rates and possibly 24-bit measurement may be required to give the recorded sound the appropriate level of accuracy.

Finally, it is worth noting that while your recordings may be good, the quality is likely to be significantly influenced by your playback system; don’t forget that the reproduction of data from a sound level meter costing thousands of pounds may be better than the £20 loudspeakers on your computer allow for.

Steve Cawser is Principal Acoustics and Vibration Consultant at Principal Acoustics and Vibration Consultant at URS Infrastructure & Environment UK.