

Bits, Bytes & Chips: Understanding Digital Instruments


By Brian Csermak & Steve Armstrong

In today's marketplace, there is one word which is guaranteed to attract attention to a product and instill in the consumer a sense of cutting edge technology. That word is *digital*. To maintain their share of the "high end" market, hearing instrument manufacturers are increasingly under pressure to have a digital product offering. In fact, it is certainly arguable that the "high end" is now defined by the word "digital."

All kinds of new technical specifications are appearing within digital hearing instrument advertisements and technical articles, including information on bits, sampling rates and MIPs. What do these terms mean and, more importantly, how do they affect the hearing-impaired patient? It can be difficult for the average hearing care professional to understand the terminology—let alone his/her customers—and ultimately to discern between what might be marketing hype and what is technically significant.

By obtaining a basic understanding of the components of a digital hearing instrument, the reader will be better equipped to sort out some of the new jargon. This article will focus on the first element of a DSP system: the analog-to-digital converter or A/D.

With the popularity of digital instruments comes a large amount of new terminology, like bits, sampling rates and MIPs. What do these terms mean and how do they affect customer satisfaction and benefit for the hearing-impaired patient? This article explores the first component, the A/D converter, in relation to its function and limitation of digital performance.



Components of a Digital Hearing Instrument

Fig. 1 presents a typical digital hearing instrument. The microphone takes the incoming acoustic signal and converts it into an electrical analog. This output is then fed into the *A/D converter*, which converts the analog electrical signal (through a process called *quantization*) into a digital electrical signal represented numerically by a series of 1's and 0's. These 1's and 0's, which now approximate the original acoustic waveform, are then fed into the DSP core where they are manipulated by any number of mathematical algorithms to apply gain, compression, filtering or other processing functions to the signal. This modified digital signal is then converted back into an analog electrical signal by the *digital-to-analog converter (D/A)*. Finally, the receiver converts this electrical signal into an acoustic output.

The All Important Front End

The popular old adage "garbage in equals garbage out" has become so well known in recent years that people in the computer industry refer to it as "GIGO." In the case of digital hearing instruments, GIGO has at least two implications:

1 The *microphone* must do a good job of accurately converting the audio signal into an analog electrical signal.

2 The *A/D* must do a good job of accurately converting this electrical analog signal into a *digital numerical signal*.

Any limitations in either of these components will restrict the amount of information extracted from the original audio signal and, therefore, limit the potential of the DSP core—and consequently the digital hearing instrument. This is analogous to the speech processing of humans: lost or damaged hair cells in the inner ear result in less information transmit-



Brian Csermak is applications engineer and Steve Armstrong is senior marketing manager at Gennum Corp., Burlington, Ontario, Canada.

ted to the brain which can cause errors in understanding speech. Similarly, limitations in the microphone and the A/D will result in less information (or more errors) in the signal transmitted to the DSP core and potentially inhibit an algorithm's effectiveness.

The performance limitations of the microphone are well established. The first limit is determined by the device's noise floor. The *noise floor* is the point at which the incoming signal is small enough that it can no longer be distinguished from the internally generated noise of the device. In many microphones, this occurs at approximately 23 dB SPL. The second limit is determined by the signal handling capability of the device. At some point, the incoming signal can be large enough that it



Fig. 1. Block schematic of a DSP hearing instrument with (l to r) microphone, A/D, DSP core, D/A and receiver.

causes the device to saturate. Signals above this magnitude can no longer be handled by the device without "clipping" the signal, thereby introducing distortion. With careful circuit design, achieving clean performance with inputs up to approximately 110 dB SPL are relatively easy. The difference between these two extremes is referred to as the *dynamic range*. For the microphone example in Fig. 1, we can see that we have approximately 87 dB dynamic range (110 dB SPL minus 23 dB SPL).

that has been converted by the microphone into an analog electrical signal. To convert this into a digital signal, the A/D captures snapshots or samples of this sine wave at evenly spaced, discrete points in time (Fig. 2b). The rest of

the sine wave is discarded. What remains is the original sine wave, still in analog form, waiting to be quantized into digital numbers (Fig. 2c). How accurately this conversion process can be done is dependent upon the number of bits of resolution we have to store the amplitude information of the samples.

For now, let's focus on the number of bits of resolution in the A/D. First, we must understand what a bit is. One bit is a binary digit that is similar to a decimal digit. However, while a decimal digit can be represented by




Fig. 2a-c. (Left to right) An arbitrary sine wave produced by the microphone (2a); a sine wave is sampled or captured at discrete points in time (2b); and analog samples within the A/D (2c).




Fig. 3a-c. (Left to right) Sampled sine wave by a 3-bit A/D converter with eight levels (3a); samples are rounded to the nearest level (3b); distorted approximated signal (3c).




Fig. 4. Noise floor of a hypothetical A/D converter versus frequency for various sampling rates (x = the audio band of interest).

Now that we know the microphone has already imposed a limitation on our digital hearing instrument, what about the A/D? It turns out that the A/D has very similar limitations to the microphone but due to slightly different mechanisms. The important question is, "Does it limit the potential of the digital hearing instrument further?" For that answer, the operation of an A/D converter must be examined.

Sampling and Bits of Resolution

Fig. 2a shows a sine wave representing some arbitrary pure tone

10 states (0 through 9) a binary digit (as the name implies) can only be represented by two states: 1 or 0. To take this comparison one step further, a 2 digit decimal number has 10^2 or 100 states while a 2 digit binary (or 2 bit) number has 2^2 or 4 states. Therefore, an A/D with n bits of resolution would have 2^n levels available to capture the amplitude information of a signal.

To illustrate, let's look at a 3-bit A/D converter. In this case, the A/D is capable of characterizing the amplitude of each sample in 8 states (or 2^3 levels). Fig. 3a shows that most of the samples taken by the A/D lie

somewhere in between the eight available levels and are therefore rounded to the nearest level (Fig. 3b). In other words, this rounding has introduced an error (which is equal to the original signal minus the rounded approximation) and results in a distorted signal (Fig. 3c). This distortion can be reduced by increasing the number of bits in the A/D. The greater the number of bits, the greater the number of available levels which results in smaller errors.

In addition to determining the amount of distortion introduced to the signal, the number of bits of resolution also determines the *noise floor of the A/D*. As stated earlier, the noise floor is the level at which signals can no longer be accurately distinguished from the internal noise of the device. In our previous example, all samples which lie at a point halfway between the two lowest levels (i.e., 0 and 1) or less would be indistinguishable from one another and would be rounded down to 0. Since adding more bits of resolution increases the number of states available to use and, therefore, reduces the amount of rounding, the addition of more bits pushes the effective noise floor to lower levels and expands the dynamic range of the A/D. It turns out that each bit of resolution adds 6 dB to the A/D's dynamic range. Therefore, our example of a 3-bit A/D would have only 18 dB of dynamic range, whereas a 14-bit A/D would have 84 dB of dynamic range.

Oversampling and the 1-bit Sigma-Delta converter

The previous example suggests that one of the properties of a good A/D converter is having many bits of resolution. However, each bit of resolution consumes valuable space and power that is at a premium in hearing instrument applications. The straight-forward techniques used for the conversion of analog quantities to equivalent numerical values becomes much more challenging as we try to obtain precision beyond about 14 bits. This has enticed hearing instrument designers to explore an alternative method using what is known as a *1-bit Sigma-Delta converter*. How can a 1-bit A/D give acceptable per-

formance? The answer is that it uses a technique called *oversampling*.

Long ago, a mathematician named Nyquist determined that (amazingly) only 2 samples needed to be taken per sine wave cycle to accurately represent a particular sine wave. What this implies is that the sampling rate of an A/D converter must be *at least twice the highest frequency component in the signal that you are interested in capturing*.

In other words, the sampling rate must be twice the bandwidth in which you are interested. This leads to the understanding that, if a 10 kHz bandwidth in a digital hearing instrument is desired, the A/D must sample at a rate of at least 20 kHz.

But what happens if we sample

at a rate higher than twice the desired bandwidth? Fig. 4 plots a hypothetical A/D's noise floor versus frequency for various sampling rates. As stated earlier, in order to obtain a bandwidth of x we must sample at a rate of at least $2x$. At this sampling rate the noise floor of a 1-bit A/D is very high. However, as we increase the sampling rate, the graph indicates that the noise floor decreases. In fact, this occurs at a rate of approximately 3 dB per doubling of the sampling rate. Why does this happen? Regardless of the rate at which we sample, the total noise energy remains the same. However, as the sampling rate increases, we *increase* the bandwidth over which that same noise energy is spread. If we sample at sufficiently high rates, much of this noise energy is pushed into frequencies greater than 20 kHz—where it is inaudible to the human ear. (Note: A further technique called *noise shaping* can also be employed to place even more of the noise energy at inaudible frequencies, but this is beyond the scope of this article.)

Does the fact that we are sampling at very high rates mean we are achieving excellent results? By necessity, an oversampling 1-bit Sigma-Delta converter needs very high sampling rates in order to get acceptable noise performance. Sampling at 1 MHz and beyond is not uncommon and doesn't necessarily imply better performance than a more traditional A/D sampling at 20 kHz. Conversely, a 1-bit Sigma-Delta converter may be capable of achiev-

ing the same dynamic range as a 14-bit A/D implemented some other way. It is the combination of bandwidth and dynamic range (sometimes known as the number of "effective bits") that should be used to determine the merit of one A/D over another.

Combining the Microphone and the A/D

Typically, when a hearing instrument designer hooks his/her A/D converter to the microphone, disappointment results. The reason is usually quite straight-forward. The input referred noise of the hearing instrument is too high, owing to the noise levels of the A/D. The solution at first seems simple: add more gain before the A/D to make sure the microphone noise dominates the results. However, this extra gain can create a headroom problem for loud sounds which will overload the A/D. Some designers have chosen to use a high-level AGC limiter as a means to avoid the distortion that would otherwise occur. Ideally, we would like to increase the dynamic range of the A/D. The general rule of thumb is to use an A/D with at least 6 dB more dynamic range than that of the microphone (i.e., 87 dB + 6 dB = 93 dB or approximately 16-bit effective resolution) in order to prevent significant degradation of the front-end performance.

Summary

When discussing digital signal processing techniques, there are two main parameters that define the strength of the A/D link in the processing chain.

1 The sampling rate that determines the bandwidth of the hearing instrument; and

2 The number of effective bits tells how accurately the analog samples are approximated in the digital realm, directly impacting the input dynamic range.

Earlier, we asked the question, "Do the typical A/Ds presently available further limit the potential of digital hearing instruments beyond the limits already imposed by the microphone?" The better A/Ds used in today's digital hearing instruments have the equivalent of 13 bits of resolution. This translates into 78 dB of dynamic range which is about 9 dB less than that of the microphone. Clearly, there are challenges remaining for circuit designers to tackle. ♦

Correspondence can be addressed to HR or Brian Csermak, Gennum Corp., P.O. Box 489 Sta. A, Burlington, Ontario, Canada L7R 3Y3; e-mail: brian_c@gennum.com.

The limitations of presently available A/Ds do impose a reduce dynamic range on current digital hearing instruments. Clearly, challenges remain for circuit designers.