A Digital Signal Processor for Musicians and Audiophiles

Published on Monday, 09 February 2009 09:54

-

The main focus of hearing aid research and development has been on the use of hearing aids to improve speech perception and intelligibility. Hearing aid designs have, naturally, evolved with this primary goal in mind.

Fortunately for hearing aid designers, measuring hearing aid performance solely using speech perception measures has been somewhat forgiving of design trade-offs. This is due to the fact that such measures of performance are relatively insensitive to some rather severe distortion.

A hearing aid that performs well with speech signals, however, may not perform well with music. Music signals are much more variable than

Jim Ryan, PhD, PEng, is the manager of algorithm development, and **Shailja Tewari , MBA**, is the director of marketing at Sound Design Technologies, Burlington, Ontario. Correspondence can be addressed to Jim Ryan at .

speech, and our perception of music is more sensitive to distortion.

Several authors have discussed the impact of hearing aid performance on music perception. Chasin[1,2] points out that a major influence is the peak input limiting level of the amplifier. Others have discussed the importance of audio bandwidth[3] and frequency response flatness.[4,5] These considerations are in addition to the usual restrictions on time delay.[6]

This article describes the performance aspects of hearing aid amplifiers with particular reference to digital amplifier design. It also introduces a new digital hearing aid amplifier, the Wolverine DSP by Sound Design Technologies, which is designed with the hard-of-hearing musician in mind, and the article uses this system as an example in design features that may provide improved audio performance for musicians.

Digital Amplifiers and Music

At the heart of every digital hearing aid lies an integrated circuit known as a digital signal processor (DSP). A DSP is a high-speed computer that manipulates the audio signals in a hearing aid numerically. It is this computational nature that gives DSPs their enormous flexibility in manipulating signals. There are some fundamental differences between digital and analog signal processing that have a direct impact on hearing aid operation. Signal bandwidth, input dynamic range, and time delay are all impacted by the choices made in designing a DSP as well as associated algorithms, and these are discussed below.

Audio Bandwidth

A DSP relies on the conversion of signals from their real-world analog format to digital format. This conversion process requires that continuous analog signals be sampled at discrete time intervals. To ensure an accurate signal representation, the sampling frequency must be at least twice as high as the bandwidth of the audio signal. This is a fundamental limitation of digital signal processing, and it leads to a sharper restriction on audio signal bandwidth than was required by the analog amplifiers of the past.

Of course, the audio bandwidth of a DSP device can be extended by increasing the sampling frequency. Unfortunately, a higher sampling frequency requires a faster DSP to handle the increased rate of audio samples and to allow for advanced features. A faster DSP, in turn, consumes more battery power, which is undesirable in a hearing aid.

The result is a trade-off between signal bandwidth and battery current. Often, designers of digital hearing aids will reduce the audio signal bandwidth to the minimum required for processing speech signals in order to minimize battery consumption. This can lead to poor performance for music, since the bandwidth of music signals can easily exceed that of speech.

The Wolverine DSP allows hearing aid designers to overcome audio bandwidth restrictions by offering an extremely low-power DSP. Wolverine's low power consumption is made possible through a patented, reconfigurable DSP architecture and an implementation in low-power, 90 nm CMOS technology.

This means a smaller battery current penalty for increasing the audio bandwidth for the user. For example, some studies have found that a bandwidth of 16 kHz is preferred when listening to music.[4] Extended bandwidth in hearing aids not only maintains "naturalness" of sound but also provides qualitative and quantitative improvements in alternate listening environments, such as music, speech, and spatial cues.[7] The Wolverine DSP is designed to operate at a nominal sample rate of 32 kHz providing an audio bandwidth of 16 kHz. At this sample rate, the circuit enables a digital hearing aid with an advanced, adaptive algorithm feature set with sub 1 mA current consumption.[8]

Input Limiting Level and Dynamic Range

In addition to the time sampling described above, analog-to-digital conversion requires amplitude sampling. The continuous-time analog waveform, sampled at discrete time intervals, is converted into a series of numbers by the analog-to-digital converter (ADC). The accuracy of the amplitude sampling is governed by the precision of some sensitive analog circuitry in the front end of the DSP.

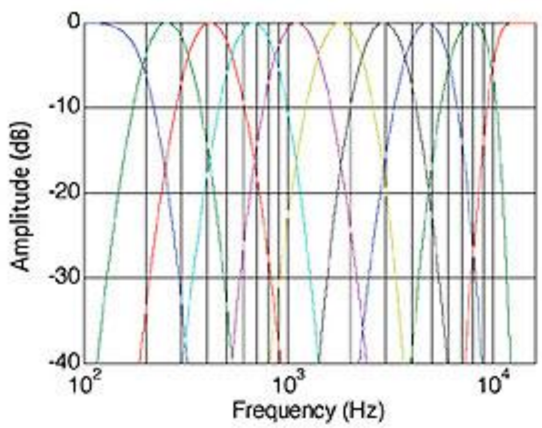

FIGURE 1. Individual band shapes of a logarithmically spaced, 10-band filterbank as implemented using the Wolverine DSP.

Increasing the dynamic range of the conversion process requires higher precision analog circuitry. Typically, however, this requires an increase in the power consumption, which is disproportionate to the increase in dynamic range. As a result, the dynamic range of a hearing aid ADC is usually limited to roughly 80 dB. Since this is less than adequate to cover all situations, other means are usually provided to increase the signal-handling ability of the hearing aid. Typically, this input range is increased by providing a programmable-gain amplifier in front of the ADC. This allows the hearing aid performance to be tuned for specific situations by adjusting the fixed gain of the preamplifier. Unfortunately, this method does not increase the dynamic range, since the preamplifier gain does not change with time.

The Wolverine architecture overcomes the input dynamic range limitation with its patented HRX system, which dynamically adjusts the input range of the ADC. With HRX enabled, the front end is able to handle signal levels covering the full acoustic range of a typical hearing aid microphone, providing a 96 dB input dynamic range. This has been shown to provide demonstrable benefits when listening to music.[1,2]
Of course, the numerical precision of the subsequent DSP affects the system dynamic range. To understand why, consider that a DSP

manipulates audio signals through digital computations using the binary number format. In the binary system, numerical precision is measured in binary digits, or bits. A well-known rule of thumb is that each bit of numerical precision represents approximately 6 dB of dynamic range. Thus, a 16-bit digital word, as used in the CD audio format, results in a dynamic range of approximately 96 dB.

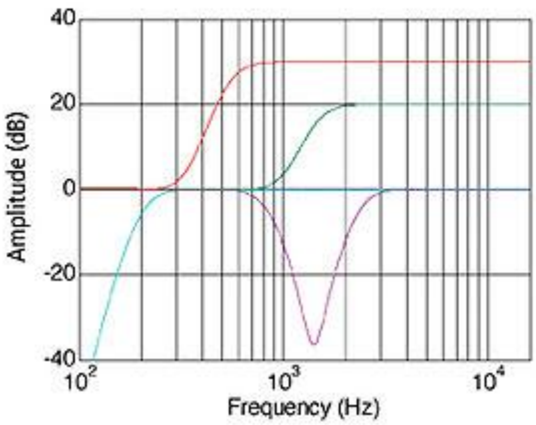

FIGURE 2. Various frequency responses achievable by varying the in-band gains for a 10-band filterbank.

When two numbers are multiplied within a DSP, the product contains twice the number of bits compared to the multiplicands. Because the number of bits cannot grow beyond the DSP's native word length, rounding must be used—reducing precision. The rounding process introduces a small error in the signal representation that is manifested as a noise added to the audio signal. For repeated operations on the same signal, rounding errors accumulate and increase the noise by the same amount *each time*. For each doubling of the number of rounding operations, the noise increases by 3 dB, reducing dynamic range by the same amount.

It does not make sense, therefore, to apply a 16-bit DSP to the output of an ADC with a 96 dB dynamic range. After only four rounding operations, the dynamic range of the system would be reduced by 6 dB, preventing complicated algorithms from being applied.

Wolverine offers a native word size of 20 bits, resulting in a dynamic range of 120 dB—24 dB higher than a 16-bit DSP. What it means is that over 200 additional rounding operations can be applied to the audio signal while maintaining the same quantization noise. This is sufficient to support the most complex signal processing algorithms in hearing aids today.

Time Domain Processing

DSP algorithms can be categorized into two approaches: time domain (or sample-based) and frequency domain (or block-based). In a time domain implementation, each signal sample is processed immediately as it arrives in the DSP and a corresponding new output sample is produced. In a frequency domain implementation, audio samples are collected into blocks and transformed into the frequency domain using a Fast Fourier Transform (FFT). Here, the block of samples is processed together until a new block of output samples is computed.

Frequency domain implementations typically result in fewer computations per audio sample than time domain implementations. This can often lead to lower power consumption, especially for a general-purpose, instruction-based DSP.

Unfortunately, frequency domain algorithms also incur longer time delays and suffer from higher noise, due to the larger number of rounding operations required. Both of these facts are clear disadvantages for hearing aids, particularly when worn by musicians.

While the Wolverine DSP efficiently supports frequency domain algorithms, it is equally effective in the time domain. The circuit is able to offer the advantages of time domain processing, low delay, and quantization noise, with extremely low power consumption. This is due to its blend of programmable DSP cores and hard-wired coprocessors. Numerically intensive time domain operations that might be power hungry on a programmable DSP are, instead, provided as dedicated hardware blocks to minimize battery drain.

In addition, performance of these dedicated coprocessors is scalable: simpler algorithms require only a portion of the coprocessor's capacity. In turn, they consume less power, incur less time delay, and generate less quantization noise. In contrast, a frequency domain algorithm must always incur the delay and noise overhead associated with the FFT calculation.

Time Delay

Time delay in a digital hearing aid is due to both the analog-to-digital (and digital-to-analog) conversion process and the signal processing algorithms.

ADC delay arises due to the aggressive low-pass filtering that must be applied to the analog signal in order to restrict its bandwidth for digital sampling. In a typical high-quality audio converter, time delays of several milliseconds are common, owing to the nature of the filtering used. Such delays would be unacceptable for a digital hearing aid since there would be very little time left to implement advanced features. Consequently, different filtering strategies are used to minimize converter delay.

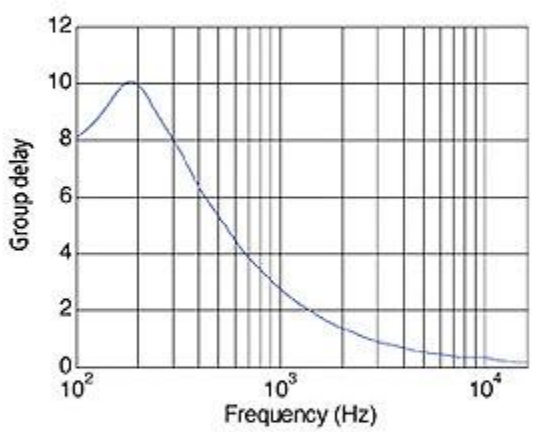

FIGURE 3. Group delay response for the 10-band filterbank. Group delay does not vary with in-band gain setting.

Signal processing delays are typically dominated by a filterbank that forms the core of many advanced audio features, such as noise reduction or dynamic range compression. As described above, filterbanks can be implemented using either a time domain or frequency domain approach. In either case, there is a fundamental relationship between frequency resolution and time delay. A filterbank with fine frequency resolution (narrowband filters) incurs a longer time delay.

For digital hearing aids, it is generally accepted that the total system delay should be less than 10 ms. The tolerable limit for musicians may even be lower due to the need for accurate rhythm control. To stay within this budget, any delays incurred by the ADC process reduce the amount of time allowed for signal processing. At the same time, many advanced features, such as adaptive noise reduction, require a narrowband filterbank for maximum effectiveness. In a DSP design, therefore, it is advantageous to minimize the conversion delays to allow maximum time for signal processing.

The Wolverine system minimizes conversion delay using minimum-phase anti-aliasing filters. This eliminates a large part of the conversion delays encountered in many other audio converters, and it provides a conversion delay just under 0.5 ms at 1 kHz. This allows maximum time for implementing signal processing features.

Phase Distortion

Phase distortion refers to a nonlinear variation of phase with frequency. In contrast to a pure time delay (where phase varies linearly with frequency), phase distortion can alter the character of a sound and is audible for certain types of input signals, particularly music.[9]

Phase distortion can be avoided in DSP systems by the use of linear phase filters. However, linear phase filters suffer from two drawbacks. First, the total phase shift introduced by a linear phase filter exceeds that for a comparable nonlinear phase filter. This can lead to restrictions on achievable frequency response targets if a 10 ms delay limit is to be maintained. Second, for a given frequency response, a linear phase filter typically requires more computations than a nonlinear phase filter. This, in turn, leads to higher power consumption.

Nonlinear phase filters can overcome these limitations, but care must be exercised to ensure that phase distortion is not a problem. This is particularly true for adaptive algorithms since the frequency response, and phase distortion, can vary with time. Due to the possibility of introducing audible artifacts and the difficulty of predicting and controlling phase distortion in real time, audio system designers try to avoid such time-varying phase distortion.

The main source of potential phase distortion in a digital hearing aid is the filterbank used to split the audio signal into different frequency bands. Filterbanks can introduce phase distortion even when the in-band gain settings call for a flat overall response. Generally speaking, narrower frequency bands lead to more severe phase distortion; however, narrower bands are often desirable for implementing advanced features.

The Wolverine DSP is designed to help system engineers overcome phase distortion problems by providing dedicated hardware for a time-domain, low-delay filterbank. This filterbank provides up to 10 independent frequency bands with adjustable cross-over frequencies. The filterbank itself is based on a unique, nonlinear phase design that maintains a constant phase response for all band gain settings. This eliminates the problem of time-varying phase distortion for adaptive algorithms.

As an example, consider the filterbank settings depicted in Figure 1. This configuration implements a 10-band, logarithmically spaced filterbank. The individual band filter shapes are depicted by the various colors. Figure 2 shows several different frequency responses that can be achieved by adjusting the various in-band gains using this configuration. Despite the large variations in frequency response, the group delay remains constant as shown in Figure 3.

In summary, advancements in silicon processing have enabled the design of highly versatile DSP platforms that meet the demanding needs of musicians.

References

1. Chasin M. Music and hearing. *Hear Jour.* 2003;56(7):36-41.
2. Chasin M, Russo F. Hearing aids and music. *Trends Amplif.*2004;8(2):35-47.
3. Beck DL, Olsen J. Extended bandwidths in hearing aids. *Hearing Review.* 2008;15(11):22-26.
4. Moore BCJ, Tan C-T. Perceived naturalness of spectrally distorted speech and music. *J Acoust Soc Am.* 2003;114(1):408-419.
5. Killion M, French J, Viranyi S, Preves D. A novel method of improving sound quality and reducing acoustic feedback in hearing aids. *Jour Acoust Soc Am.* 2002;111:2381.
6. Stone M, Moore BCJ. Tolerable hearing aid delays. III. Effects on speech production and perception of across-frequency variation in delay. *Ear Hear.* April 2003:175-183.
7. Byrne D, Noble W. Optimizing sound localization with hearing aids.*Trends Amplif.* 1998;3(2):51-73.
8. Sound Design Technologies. Inspiria Ultimate (GA3285) data sheet. Burlington, Ontario: Sound Design Technologies; 2008.
9. Deer JA, Bloom PJ, Preis D. Perception of phase distortion in all-pass filters. *J Aud Eng Soc.* 1985;33(10):782-786.

Citation for this article:

Ryan J, Tewari S. A digital signal processor for musicians and audiophiles.*Hearing Review.* 2009;16(2):38-41.