Signal Processing Algorithms for a new, Software-based, Digital Hearing Device

By Brent W. Edwards, Zezhang Hou, Christopher J. Struck, Priya Dharan
ReSound Corporation
220 Saginaw Drive
Redwood City, CA 94063

INTRODUCTION

The recent development of commercial hearing devices with digital signal processing (DSP) capabilities opens the door for an explosive growth in device sophistication. This article reviews the state-of-the-art in software-based DSP technology for hearing instruments.

In addition to the multi-band dynamic range compressor, which was discussed in an article by Edwards et.al. in last month's Hearing Journal[1], the new ReSound digital hearing system incorporates dual-microphone directional processing, multi-band noise-reduction, and automatic feedback cancellation algorithms.

The basic function of a hearing device is to make sounds audible, yet not uncomfortably loud for the user. Furthermore, improvement of the signal to noise ratio (SNR) in difficult listening situations and effective feedback control schemes are important to both the hearing impaired patient and the hearing health care professional fitting the hearing device. The algorithms developed for the ReSound digital processor, taken together, make up a system that not only compensates for hearing loss and loudness recruitment, but also attempts to address the problems of feedback and speech understanding in noise. This article provides a functional description of these processing schemes and the basic building block for the compression and noise reduction algorithms, the Fast Fourier Transform (FFT).

FFT – THE BASIC BUILDING BLOCK

The Fast Fourier Transform is ubiquitous in the field of digital signal processing. It is an efficient algorithm for calculating the frequency spectrum of a signal and is useful for signal processing in the frequency domain. The FFT facilitates analysis of stationary periodic signals and is particularly well suited for the analysis of audio and speech signals.

An example of the use of FFT in high end audio is Dolby Digital coding, where, spectral processing is used to determine what frequency components of the audio signal are psychoacoustically masked. The FFT is also the basis for the design of most modern test and measurement equipment. Many noise-reduction techniques, such as spectral subtraction[2], require FFT processing since the spectrum of the unwanted background noise must be calculated and subtracted from the subsequent incoming signal.

The FFT is a very useful processing tool in a hearing device because the human auditory system itself performs frequency-based processing, and because much of the speech information is encoded in the spectral shape of the speech signal. Specific to speech processing for the hearing-impaired, the FFT is necessary for spectral enhancement techniques[3] that attempt to compensate for the reduced frequency resolution experienced by persons with sensorineural hearing loss.

Once the spectrum of a time signal is obtained, the values in the digital signal processor that correspond to the particular spectral samples can be altered to meet any processing requirement. For instance, formants of speech can be identified in a spectrum and more gain applied to the formant frequencies relative to the other frequencies in order to enhance their perception[4]. Also, the FFT is needed for spectral translation techniques, which have been shown to improve speech intelligibility for listeners with steeply sloping high-frequency hearing loss[5]. These and other advanced processing schemes to achieve higher fidelity can be very easily implemented in a digital processor that supports a fast and efficient FFT structure.

The first step in performing an FFT operation is to select a short segment or "window" of the acoustic signal. This window is typically only a few milliseconds in duration. The FFT is performed on that short segment producing a frequency spectrum. Thereafter, any frequency-dependent processing of the spectrum (e.g., multi-band syllabic compression) can be accomplished.

After the spectrum has been modified, an inverse FFT is applied to transform the modified spectrum back into a time domain signal, which is then routed to the receiver. Meanwhile, the next time signal segment is obtained and the process is repeated. Signal segments are generally processed in parallel using "overlapping" windows (i.e., the next record is started before the end of the first) in order to obtain a smooth output signal with minimal artifacts. The FFT processor can be thought of as a spectrum analyzer that constantly calculates the spectrum of short segments of the incoming audio signal.

ReSound's digital hearing system incorporates the FFT as the core of its signal processing. The FFT makes it easy to create a number of bands by simply combining spectral

From The Hearing Journal 1998; 51(9): 44-52.

values into groups. The level in each band can then be calculated and adjusted to result in compression. Also, the FFT can be used for compression with a variety of bandpass designs ranging from one broad-band filter to filters with several dozen bands. It is this flexible structure of the FFT that makes possible the implementation of Advanced Cochlea Dynamics™ processing as a complex, heavily overlapping, 14-band compression scheme in the new ReSound digital system[1].

The noise-reduction algorithm in the first-generation devices also takes advantage of the FFT structure. The DSP code has been optimized for this type of processing. The open nature of the DSP design will support future algorithms that use this FFT structure, many of which would be unrealizable otherwise.

DUAL-MICROPHONE DIRECTIONAL PROCESSING

Hearing-impaired persons find it more difficult understanding speech in noise than normal-hearing persons and may require up to a 10 dB higher signal-to-noise (or in this case, speech-to-noise) ratio than normal listeners for the same recognition performance[6]. A directional microphone can be used with hearing devices in order to provide direction-dependent gain. Multi-microphone arrays in hearing instruments enable users to switch between an omnidirectional mode for low ambient level or relatively quiet environments and a directional setting for high noise level situations. A two-microphone directional processing system can provide a significant amount of directionality that results in improved speech understanding in noisy environments[7].

Directional systems typically assume that the target signal is in front of the hearing device user and therefore are designed to maximize the attenuation for diffuse or random-incidence background noise. This configuration can be optimized to obtain the highest Directivity Index (DI) which is defined as the ratio (in dB) of sensitivity along a reference axis (typically frontal) to the sensitivity to random-incidence signals. The DI can be thought of as an SNR. In free field, the measured polar pattern corresponds directly to a given DI. For a two-sound-inlet system, the highest DI is obtained for a hypercardioid polar pattern.

Directionality is obtained by introducing an acoustic delay between the microphones and subtracting the rear microphone signal from the front microphone signal. The ReSound digital system, which has two omnidirectional microphones, achieves the necessary additional delay to the rear-microphone signal by delaying the second analog-to-digital (A/D) converter output.

The distance between the two microphones determines the acoustic delay. The size of the housing limits the maximum separation between microphones to approximately 0.6 inches or 1.5cm in the case of behind-the-ear devices. An in-the-ear configuration limits the physical microphone separation to approximately 0.4 inches or 1cm. This small distance between microphones limits the directional performance in the low frequency region. Without a large separation, low frequency directionality is entirely dependent upon the response matching (magnitude and phase) of the front and rear microphone signal paths.

Directionality becomes less effective when there are phase and amplitude mismatches between the two microphones on a directional instrument, and the effect of the mismatch is more significant as the distance between microphones decreases (see Figure 1). In the new ReSound system, a digital filtering technique is used to achieve the critical matching between the front and rear microphones.

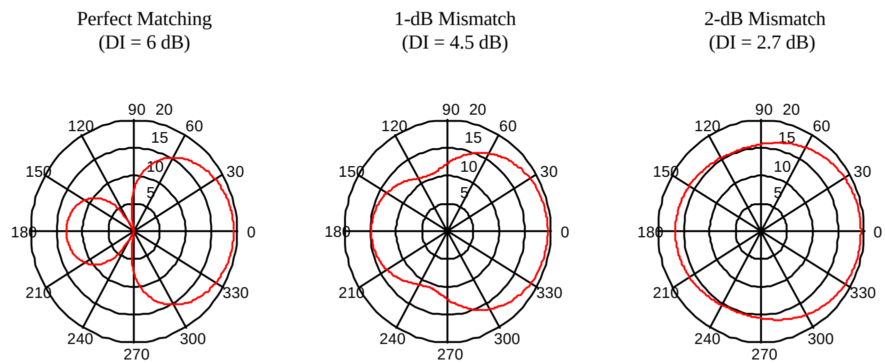

Figure 1 *Effect of mismatch in the characteristics of the front and rear microphones in a directional system. The Directivity Index or DI decreases as degree of mismatch increases and the polar pattern looks more like that of an omnidirectional system.*

The device has two A/D converters. The output of each microphone is sampled and digitized, the rear microphone signal is delayed, and then subtracted from the front microphone signal. By subtracting the response of the two microphones digitally instead of electrically before the A/D converters, digital filters can be used to match the two microphones so that their amplitude and phase responses are as close as possible.

With directional processing, the on-axis (frontal) frequency response typically has a high-pass characteristic, rolling off at 6-dB/octave. However, a moderate low-frequency boost can be provided to compensate for the roll-off (and otherwise “tinny” sound quality) and to keep the low-frequency signals audible.

At high frequencies, the acoustic effect of the head and pinna provides a significant amount of directionality. The microphone in canal and completely-in-the-canal devices is located so that it can take advantage of these head and pinna effects. In the case of a behind-the-ear instrument (with an omnidirectional microphone located on top of the ear) and to a lesser extent an in-the-ear device (with the microphone located in the plane of the pinna), the directionality effects provided by the head and pinna are somewhat diminished.

When evaluating the benefit of a two-microphone directional system, it is interesting to compare the amount of directionality for the directional system to that of an omnidirectional microphone system, when both are in-situ. Figure 2 shows the simulated in-situ polar patterns in the horizontal plane for a two-microphone directional system and a single omnidirectional microphone. Both systems are in BTE devices measured on a manikin.

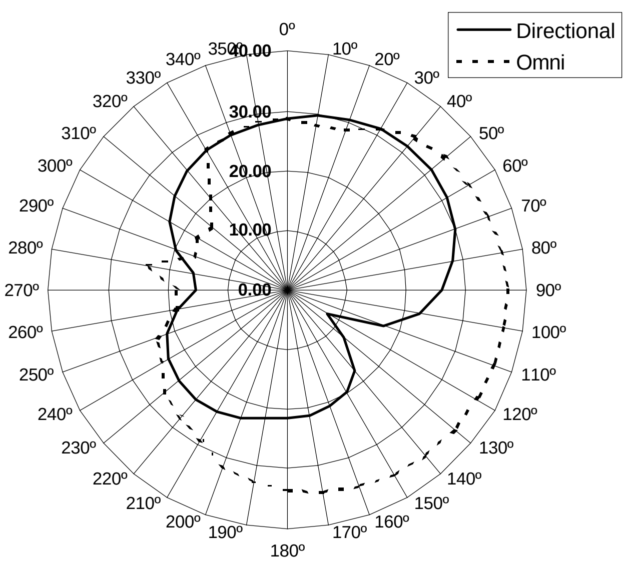

Figure 2 *Polar plots for a behind-the-ear device with a two-microphone directional array and one with a single omnidirectional microphone, measured with the devices on a manikin.*

Notice that the directional pattern shows an attenuation of up to nearly 30 dB (at 120°) compared to the omnidirectional pattern. The amount of attenuation varies with direction and the overall benefit is calculated as the integrated average over a sphere. The approximate improvement in speech intelligibility for a monaural fitting is 5 dB when the speech signal originates in front of the listener and the background noise is diffuse. This means that the SNR is 5 dB less than that needed for the same intelligibility with a single omnidirectional microphone system. To take into account the importance of different frequency regions to speech intelligibility, an Articulation Index-weighted Directivity Index can be calculated to obtain an overall estimate of the benefit of the directionality of a device with respect to understanding speech in noise[7].

MULTI-BAND NOISE REDUCTION

The noise-reduction algorithm in the new ReSound digital system is designed to provide the subjective benefit of better sound quality or listening comfort. It detects the presence of speech by measuring the modulation or level fluctuations of the signal envelope in each of the 14 compressor bands. If the measured statistics of the signal envelope are similar to what is expected for speech in quiet, then no attenuation is applied in that band. But, as the band statistics approach that of noise, whether it is stationary environmental noise or speech babble, the gain in that band is reduced in a systematic manner.

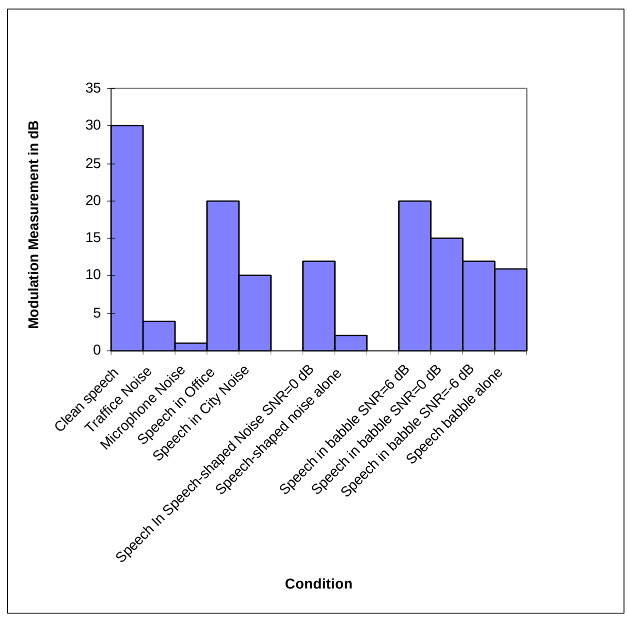

Figure 3 *Modulation in dB for different signals including clean speech and speech in different background noise environments.*

Speech in quiet has the largest measured modulation. The modulation or dynamic range of level fluctuations of noise alone is very low. The measured modulation in dB for clean speech as well as for some common background noise environments is depicted in Figure 3. Clearly, the modulation of noise-contaminated speech is less than that of clean speech. For the same type of background noise, the modulation decreases with decreasing SNR (e.g., speech babble in Fig. 3). This fact has been used to determine the dependence of the gain reduction or attenuation function on modulation (see Fig.4).

In this example, there is no attenuation applied if the modulation is greater than 24 dB. As the modulation decreases, the amount by which the gain is reduced increases linearly. Maximum attenuation (-12 dB) is applied when the modulation is 0 dB. In the ReSound digital hearing system, this relationship between attenuation and measured modulation has been optimized to maximize the benefits of the noise-reduction processing for a wide variety of background noise environments.

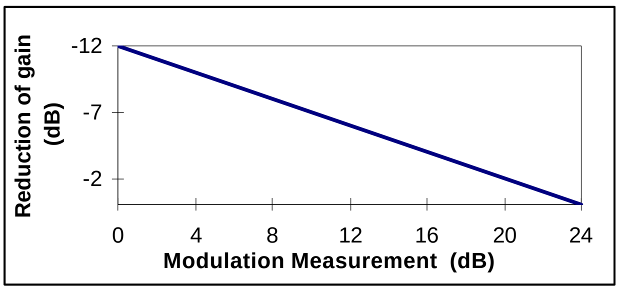

Figure 4 *Example of a gain reduction or attenuation function relative to measured modulation (both are in dB).*

The time constants for modulation detection are considerably longer than the time constants for the syllabic compressor. Using different time constants minimizes the interaction of the two processors so that one does not interfere with the other. The slower time constant for noise reduction processing also eliminates unwanted artifacts, which has been the major problem with many other fast-acting adaptive noise reduction schemes[2,8].

The ReSound noise reduction technique also addresses a common hearing device problem resulting from audible microphone noise in quiet environments when the compressor is programmed for maximum gain. The modulation measurement for microphone noise alone is very low (close to 0 dB), which will result in the noise reduction algorithm applying nearly maximum attenuation (-12 dB in Fig.4) when there is no sound input. The result is a much quieter system.

The multi-band noise-reduction algorithm takes advantage of spectral mismatches between the noise and speech. This is particularly significant when the competing background noise has a narrow-band response. In this situation, the algorithm reduces the gain only in the frequency region where noise is present. In contrast, broad-band systems or systems with only a small number of bands, reduce the gain not only in the frequency region where noise is present but also in other (potentially important) frequency regions where noise is not present.

AUTOMATIC FEEDBACK CANCELLATION

Feedback or "whistling" is a significant problem for many hearing device users. Acoustic leaks around an earmold can result in feedback. Venting is often used to reduce the "occlusion effect" or stuffed up feeling perceived by device users, but the larger the vent, the more the potential for acoustic feedback. An instrument in use can also feedback if an object such as a telephone receiver, a hat or even a hand is brought near that aided ear.

One solution that has often been used is simply to reduce the gain in a broad frequency band around the peak of the feedback signal thereby lowering the level of the feedback. However, such broad-band feedback-management schemes inadvertently limit the amount of gain delivered by the device to sub-optimal levels, i.e., to less than what is necessary for proper hearing loss compensation. Notch filters have also been used to solve this problem and are better than broad-band methods because they reduce the gain in a narrower frequency band. The drawback, however, with both these solutions is that lowering the gain to reduce feedback also reduces the level of all speech cues in that frequency band, which may make them inaudible to the hearing device user.

The ReSound device uses digital signal processing for adaptive feedback cancellation. During the fitting process the hearing professional can initiate an estimation of the acoustic feedback path between the receiver and the microphone. A brief, internally generated noise burst is applied directly to the receiver, and the level and frequency of the receiver signal appearing at the microphone are determined. Using correlation techniques, an initial estimate of the feedback path is obtained.

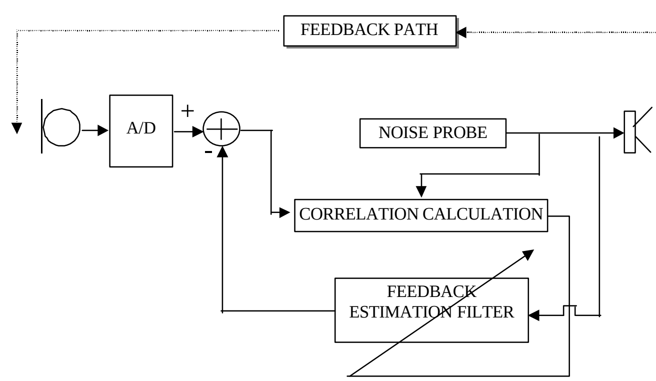

Figure 5 *Functional block diagram of feedback path estimation where a noise burst is applied to the receiver and correlation techniques are used*

to determine the parameters of the digital feedback estimation filter. The characteristics of this filter mimic those of the feedback path.

This information is then used to calculate the parameters of a digital filter whose characteristics essentially mimic those of the feedback path (see Fig. 5). This internal filter block generates a compensation signal that is equal to the actual feedback signal. During normal operation of the hearing device, the signal to the receiver is passed through the digital feedback cancellation filter and the resulting compensation signal is subtracted from the microphone output to cancel the feedback (see Fig. 6).

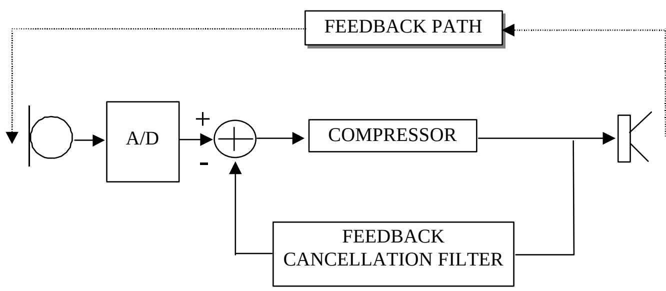

Figure 6 *Functional block diagram of the automatic feedback-cancellation system. The feedback-cancellation filter generates a compensation signal, which is subtracted from the microphone output in order to cancel the feedback.*

In addition to static feedback cancellation the filter also functions as an "adaptive" cancellation tool. Its response is continuously adjusted during the normal use of the device in order to track variations in the feedback path. These changes in the feedback path may be caused by the placement of a telephone near the aided ear or a change in position of the earmold. Unlike techniques in use today, the advantage of this feedback cancellation scheme is that it has minimal effect on gain and frequency response of the device.

CONCLUSIONS

The ReSound digital hearing system takes advantage of the power of digital signal processing and the FFT structure of the processor to support advanced software-based signal processing algorithms. Special schemes like multi-band noise reduction, automatic feedback cancellation and dual-microphone directional processing are designed so as to not counter the effect of the Digital Cochlea Dynamics multi-band compression scheme for hearing loss compensation. While improved speech intelligibility is one of the primary goals for design of hearing devices, algorithms that improve subjective perception such as sound quality are equally important. The digital product line also has diagnostic capabilities and supports an improved Loudness Growth in Octave Bands (LGOB) test for measurement of loudness growth functions.

REFERENCES

1. Edwards B.W., Struck C.J., Dharan P. and Hou Z. "New digital processor for hearing loss compensation based on the auditory system", *The Hearing Journal*, August 1998.
2. Boll S.F. "Suppression of acoustic noise in speech using spectral subtraction". *IEEE Transaction on Acoustics, Speech, and Signal Processing* 27, 113-120, 1979.
3. Stone M.A. and Moore B.C.J. "Spectral feature enhancement for people with sensorineural hearing impairment: Effects on speech intelligibility and quality". *Journal of Rehabilitation Research and Development* 29: 39-56, 1992.
4. Ribic Z., Yang J. and Latzel M. "Adaptive spectral contrast enhancement based on masking effect for the hearing impaired", ICASSP-96, Atlanta, 937-939, 1996.
5. Turner C.W., Kwon, B., Tanaka, C., Knapp, J. "Further studies in high-frequency amplification", Second Biennial Hearing Aid Research and Development Conference, Bethesda, MD, 1997.
6. Plomp R. "Auditory handicap of hearing impairment and the limited benefit of hearing aids". *Journal of the Acoustical Society of America* 63: 533-549, 1978.
7. Killion M., Schulein R., Christensen L., Fabry D., Revit L., Niquette P and Chung K. "Real-world performance of an ITE directional microphone", The Hearing Journal, 51: No. 4, 24-38, 1998.
8. Graupe D., Grosspietsch J.K., and Basseas S.P. "A single-microphone-based self-adaptive filter of noise from speech and its performance evaluation". *Journal of Rehabilitation Research and Development* 24: 119-126, 1987.

From The Hearing Journal 1998; 51(9): 44-52.