Fa-Long Luo
and Arye Nehorai

Recent Developments in Signal Processing for Digital Hearing Aids

Approaches in signal processing research on digital hearing aids fall into four areas, which cover signal acquisition, amplification, transmission, measurement, filtering, parameter estimation, separation, detection, enhancement, modeling, and classification. The first area uses advanced signal processing techniques to characterize and compensate for various hearing impairments such as loudness and frequency selectivity loss. The second area consists of effective target signal enhancement and noise reduction, which includes adaptive microphone array technologies, spectral subtraction algorithms, and blind source separation methods. The third area focuses on the real-world use of hearing aids and addresses issues such as flexibility, convenience, feedback cancellation, and artifact reduction. The fourth area is devoted to expanding hearing aid technology into devices that are also able to perform other functions, such as mobile phones and music players. In this area, issues such as echo cancellation, bone-conductive microphones, and wireless voice link are of interest [1], [2]. Because of the limitations imposed by the hardware requirements, computational speed, power consumption, and other practical factors, the development and implementation of signal processing techniques for digital hearing aids has been a challenging and active research area over the past decade.

This article provides an overview on recent developments in these areas from an engineering and applications point of view and outlines three topics in signal processing techniques for hearing aids, namely, wind-noise reduction, feedback cancellation, and bone-conductive microphones. In turn, their solutions are potentially the most beneficial for meeting the increased demands placed on the performance of hearing aids.

WIND-NOISE REDUCTION

Wind noise is a major problem when hearing aids are used outdoors as well as in many applications that use acoustic microphones, such as mobile phones and outdoor recordings. A regular hearing-aid microphone can pick up more than 100 dB SPL, even in a relatively light breeze. This level can be even higher in hearing aids that use directional microphones. Frequency analysis and measurements have shown that wind noise has two properties: 1) The energy of wind noise is mainly in the frequency band lower than 1 kHz, and 2) cross-correlation coefficients have values that are much lower when two nearby microphones pick up wind noise than when they pick up acoustic signals. On the basis of these two properties, several wind-noise reduction methods have been proposed and used in some hearing-aid products [3], [4].

The key issue in these wind-noise reduction schemes is the design of an accurate and effective wind-noise detector. The wind-noise detection methods that have been proposed are mainly based on the cross-correlation difference between wind noise and acoustical sound. There is a high degree of cross-correlation when two nearby microphones pick up acoustical sounds. In contrast, cross-correlation is very low when wind noise passes through these two nearby microphones, mainly because the propagation speed of connective air-flow is much less than that of the acoustical signals [3]. With this property, the cross-correlation of two microphone signals is first calculated (or estimated) and then a comparison of the estimated cross-correlation with a predetermined threshold is made. If the estimated cross-correlation is less than the threshold, the presence of the wind noise is assumed. Otherwise, wind noise is considered to be absent. For better performance, comprehensive sufficient statistics can be employed to replace the cross-correlation. For instance, both auto-correlation and cross-correlation of two microphones can be combined to get a more effective sufficient statistic. Estimating the cross-correlation of two microphone signals is computationally intensive and very time consuming. To avoid the time-consuming cross-correlation estimation, the power difference (or ratio) of the summed signals and subtracted signals between two microphones can be used as the sufficient statistic for the wind-noise detection. This is motivated by the fact that the subtracted signal for wind noise between two nearby microphones is of similar power to the summed signal, but signals propagating at acoustic speed will result in a relatively large difference in the summed and the subtracted signal powers. Additionally, if one of these two microphones is a directional microphone, the power difference (ratio) of these two microphones can be directly used as the sufficient statistic, because the directional microphone signal performs like a subtracted signal of two omni-microphones. These wind-noise detection ideas can be extended to the case with more than two microphones.

FEEDBACK CANCELLATION

Acoustic feedback signals in hearing aids come mainly from the leakage of the receiver to the microphone. Feedback signals are not only annoying, but they also prevent the hearing-aid compression amplification from reaching the desirable gains for compensating for the loudness loss. Hearing aids with traditional methods to reduce feedback mainly consist of a microphone, the physical feedback path, the core hearing-aid processing unit, a receiver, and an adaptive finite impulse response (FIR) filter.

The signal to the core hearing-aid processing part is the difference signal between the output of the microphone and the output of the adaptive filter. It is obvious that the impact of the feedback problem can be minimized if we make the output of the adaptive filter be as close as possible to the signal that the feedback path provides to the microphone. Ideally, if the output of the adaptive FIR filter were equal to the feedback signal in the microphone, the feedback could be canceled completely.

For the sake of simplicity, one may assume that the receiver is an amplifier with unity gain. This setup suggests that the frequency response of the adaptive filter should always be close to the frequency response of the acoustical feedback path so as to continuously and adaptively cancel the feedback signal. Now the problem is how to adaptively update the coefficients of the adaptive FIR filter, which includes the following two key issues [1], [2], [5].

SELECTION OF FILTER LENGTH

How many taps are required for the adaptive FIR filter to approximate the acoustical feedback path, which is, in effect, equivalent to a time-varying infinite impulse response (IIR) filter? If the FIR filter is too long, the computational complexity could be very high, which would increase the power burden on the hearing aids. In addition, the convergence will be very slow if gradient-decent based adaptive algorithms such as the least mean squared (LMS) algorithm are used in the long-tap FIR filter, which will in effect degrade the feedback cancellation performance. On the other hand, if the FIR filter is too short, there would be a large difference in frequency response between the adaptive FIR filter (all-zeros filter) and the real acoustical feedback path (filter with both zeros and poles), even if the optimized coefficients of the adaptive FIR filter were obtained under some optimization criteria. Thus, the performance of the feedback cancellation is limited, and the additional amplification gains for hearing aids resulting from feedback cancellation are limited as well.

To achieve a good compromise among performance, convergence speed, and computational complexity, two-stage filtering has been employed in current hearing-aid devices [6], [7]. Two-stage filters consist of an adaptive FIR filter and a fixed IIR filter. The fixed IIR filter has fixed coefficients, and it is used to roughly approximate the acoustical feedback path. These fixed coefficients can be designed in the offline fitting stage of the hearing aids, and the values of these coefficients can be different for different hearing-aid wearers. The adaptive FIR filter is used to model the difference between the fixed IIR filter and the time-varying acoustical feedback path. Because the variation of this difference is much smaller than that of the acoustical feedback path itself, the length of adaptive FIR can be much shorter than it would be if the fixed IIR filter were included. Short-tap adaptive FIR filtering can result in simple computations in real-time implementation and also a fast convergence.

However, there will be a significant variation of the acoustical feedback path if telephone headsets are placed to the aided ear. Measurements have shown that the amplitude of the transfer function of the feedback path could change by as much as 10 dB with the placement and removal of telephone headsets to the aided ear [6], [7]. The above two-stage filtering techniques will fail in such a situation. To overcome this problem, third-stage processing using an adaptive gain (one-tap FIR filter) can be included after the second-stage adaptive FIR filter. With this adaptive gain processing, significant shift of the amplitude of the transfer function of the physical feedback path can be tracked and compensated for. Another idea to attack the large variation of the feedback path is that different predetermined fixed IIR filters can be chosen to correspond with different situations in which large variations of the acoustical feedback path could occur. For example, one IIR filter can be predetermined at the normal wearing position of the hearing aid. Another IIR filter can be predetermined by placing telephone headsets to the aided ear. Hearing-aid users could switch these two different predetermined IIR filters as needed. Alternatively, a detector that can detect the placement and removal of a telephone headset can be used to switch automatically between these two different IIR filters, which would offer more convenience for users.

SELECTION OF FILTERING AND THE REFERENCE SIGNAL

In real-time applications, microphones pick up the summation of the feedback path signal and the target signal, but the acoustical feedback path signal alone is unknown. As a result, we cannot immediately use traditional adaptive algorithms to obtain the optimized coefficients of the adaptive FIR filter that should minimize the difference of the adaptive filter output signal (the reference signal) from the signal provided by the acoustical feedback path. However, if the input of the adaptive filter is uncorrelated to the system input (the target speech), then minimizing the difference of the adaptive filter output signal from the feedback path signal equates to minimizing the difference of the adaptive filter output signal from the total microphone signal (the summation of the feedback path signal and the target speech signal). As a matter of fact, this assumption can be made true if we add some extra delay processing in any stage between the output of the hearing-aid core-processing unit and the input of the adaptive filter.

However, the performance of the system with delay processing will degrade if the input signals are narrowband. For narrowband signals, the delay processing will not effectively decorrelate the input speech signals from the signal in the reference channel, which might either cancel the input target signal (known as the signal cancellation problem) or cause a large mismatch between the adaptive feedback path model and the real acoustical feedback path. To overcome this problem, several methods have been proposed, such as multidelay processing, notch filtering, and constrained adaptive filtering. Of all these proposed methods, constrained adaptive filtering seems to be the most effective. Constrained adaptive filtering methods can be divided into the band-limited filtering [5] and coefficient-constrained filtering methods [2], [6], [7].

The band-limited filtering method is based on the fact that, if the adaptive filter is constrained to operate only in the regions that contain all possible oscillation frequencies, then the filtering would be more efficient in suppressing oscillation incurred by feedback. The performance of the band-limited method depends mainly on the identification of the oscillation frequencies, the selection of the band-limit filters, and related adaptive algorithms.

The coefficient-constrained filtering method first selects a reference set of adaptive filter coefficients under controlled conditions and then uses these reference coefficients as a constraint to limit the mismatch between the adaptive filter and the real acoustical feedback path while the adaptation is taking place. Although this constraint-based algorithm is effective in overcoming artifacts and signal cancellation problems, the ability to model large deviations from the initial feedback path, such as the placement and removal of a telephone handset, is limited. This drawback is similar to that encountered in the above two-stage filtering method. One simple idea to overcome this problem is to introduce several (at least two) sets of reference coefficients and use these different reference coefficients according to their corresponding situations. For example, we can select two sets of reference coefficients: One is for the normal situation of the aided ear, and the other is for placement of a telephone headset to the aided ear. This solution will cause some inconvenience for practical use because hearing-aid wearers would need to switch two different parameters. An automatic switch scheme has been proposed that first detects whether the large feedback path variation (such as placement or removal of telephone headsets) occurs and then indicates the adaptive algorithm to use the proper reference set coefficients.

BONE-CONDUCTIVE MICROPHONES

Microphones discussed in the previous section are, in effect, regular air-conductive microphones. Recently, bone-conductive microphones have received increasing attention and have found uses in many aspects of voice signal processing [8], [9]. Unlike regular air-conductive microphones, bone-conductive microphones receive vibrations only from the bones of the human talker and from nearby air. In other words, the output of the bone-conductive microphone is only the voice of the human talker. Taking this property of bone-conductive microphones, a novel hearing-aid system has been proposed and is shown in Figure 1. In comparison with traditional air-conductive microphone-based hearing-aid systems, the proposed scheme offers significant advantages in feedback cancellation and the reduction of the hearing-aid wearer's own voice, which is unnecessarily amplified by compression amplification units. Also with this proposed scheme, processing delay and its impact on speech quality can be reduced. It should be noted that the air-conductive microphone in Figure 1 is either one single microphone or a microphone array.

The configuration in Figure 1 provides two physical feedback paths, one for the air-conductive microphone (which could be a directional microphone, an omni-microphone, or a combination of both), the other for the bone-conductive microphone. The output of the air-conductive microphone is the summation of the target signal and the signal through one feedback path. The output of the bone-conductive microphone is the signal from only the other feedback path, and this feedback signal is highly correlated with the feedback signal in the air-conductive microphone. Following the bone-conductive microphone, an adaptive FIR filter is added. The input of the core hearing-aid processing unit is the difference signal between the output of the adaptive FIR filter (the reference signal) and the output of the air-conductive microphone (the primary signal). If the coefficients of the adaptive FIR filter are updated by minimizing the power of the difference between the primary signal and the reference signal, the output of the adaptive filter can track the part of the signal provided by the feedback path through the air-conductive microphone, and the

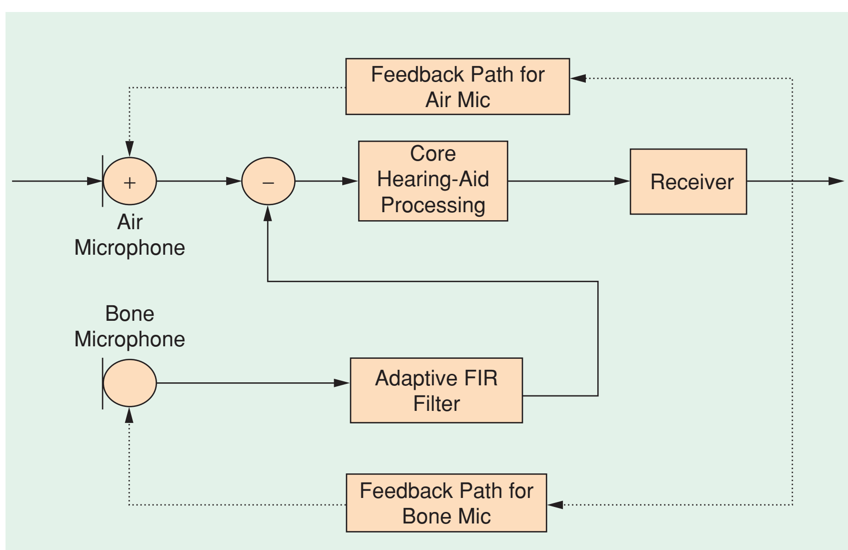

[FIG1] Block diagram of the use of bone-conductive microphone in hearing aids.

effect incurred by the feedback can be minimized. In the ideal case, the output of the adaptive filter approximates the feedback signal through the air-conductive microphone so that the difference signal can approximate the real target signal without any feedback signal in the receiver.

If we analyze this scheme in the frequency domain by denoting the frequency responses of the physical paths in the air-conductive microphone, the bone-conductive microphone, and the adaptive filter as $A(w)$, $B(w)$, and $W(w)$, respectively, we see that the frequency response of the adaptive filter after the convergence is the ratio of $A(w)$ to $B(w)$. In contrast, the optimized frequency response $F(w)$ of the adaptive FIR filter of the traditional feedback cancellation scheme is equal to $A(w)$. It is easy to see that $W(w)$ will have a simpler response than $F(w)$. Hence it would be easier to implement $W(w)$ than to implement $F(w)$. In other words, the adaptive FIR filter following the bone-conductive microphone is used to model the decibel difference between the time-varying acoustical feedback paths of the air-conductive microphone and the bone-conductive microphone. Because the variation of this difference is much smaller than that of the acoustical feedback path itself in the air-conductive microphone, the length of the adaptive FIR filter [$W(w)$] used in the proposed scheme of Figure 1 can be much shorter than that of the adaptive filter [$F(w)$] used in the traditional scheme. In the ideal case, $W(w)$ can simply be a gain factor. With these, this proposed scheme can provide us with a simple alternative tool in feedback cancellation.

As a matter of fact, the physical feedback path of the bone-conductive microphone in Figure 1 serves as the IIR filtering in the two-stage filtering discussed previously. The difference is that the physical feedback path of the bone-conductive microphone is no longer fixed but varying as the physical feedback path of the air-conductive microphone. In other words, the processing in Figure 1 can be considered to include two-stage adaptive processing: adaptive FIR filtering and the adaptive IIR filtering provided by the bone-conductive microphone. As a result, the problems listed previously will no longer exist in this bone-conductive microphone-based scheme. For further illustration, let us discuss the placement of a telephone headset to the aided ear. In this situation, the feedback signal picked up in the air-conductive microphone will increase about 10 dB, and the scheme with the adaptive FIR filtering plus the fixed IIR filtering would fail, as pointed out in a previous section, since the variation is too large. However, the feedback signal picked up in the bone-conductive microphone could also increase about 10 dB, and this would make the decibel difference variation of the two feedback path signals very small. The adaptive FIR filter following the bone-conductive microphone can still easily track this small variation.

With traditional hearings aids, the wearer's own voice is picked up by the air-conductive microphone and is unnecessarily amplified. It is highly desirable to effectively reduce this amplified voice. The scheme in Figure 1 very powerfully attacks this problem. The wearer's voice is picked up by both the air-conductive microphone and the bone-conductive microphone; hence, the voice to be amplified by the compression amplification processing is the difference of two microphone channels instead of the part picked up in the air-conductive microphone of traditional hearing aids alone. More importantly, the optimized coefficients of the adaptive filter are obtained by minimizing the difference signal between the output of the air-conductive microphone and the output of the adaptive filter following the bone-conductive microphone. This means that the wearer's own voice inputted to the amplification processing unit can also be minimized, and its impact on the target signal can be correspondingly minimized as well. Ideally, the wearer's own voice through the microphone channel can be reduced to zero.

AUTHORS

Fa-Long Luo (f.luo@ieee.org) holds the title of chief scientist of Anyka, Inc., San Jose, California. He has been a Senior Member of the IEEE since 1995. He has authored two books and more than 100 technical articles, and he holds 14 patents (some pending) on multimedia signal processing, parallel computing, DSP implementation, and machine learning.

Arye Nehorai (nehorai@ese.wustl.edu) is the chair of the Department of Electrical and Systems Engineering at Washington University in St. Louis and the inaugural holder of the Eugene and Martha Lohman Professorship of Electrical Engineering. He was editor-in-chief of *IEEE Transactions on Signal Processing* (2000–2002) and vice president, Publications, of the IEEE Signal Processing Society (2003–2005). He is a Fellow of the IEEE.

REFERENCES

[1] B. Edwards, "Hearing aids and hearing impairment," in *Speech Processing and the Auditory System, Springer Handbook of Auditory Research*, vol. 18. New York: Springer, 2004.

[2] F.-L. Luo, B. Widrow, and C. Pavlovic, "Editorial: Special issue on speech processing for hearing aids," *Speech Commun.*, vol. 39, no. 1, pp. 1–3, 2003.

[3] G.W. Elko, "Reducing noise in audio systems," U.S. Patent application 2003/0 147 538, July 16, 2003.

[4] W.B. Coney, G.L. Duckworth, and J.C. Heine, "Systems and methods for adaptive wind noise rejection," U.S. Patent 6 859 420, Feb. 2005.

[5] H.-F. Chi, S.X. Gao, S.D. Soli, and A. Alwan, "Band-limited feedback cancellation with a modified filtered-X LMS algorithm for hearing aids," *Speech Commun.*, vol. 39, no. 1–2, pp. 147–161, Jan. 2003.

[6] F.-L. Luo, J. Yang, C. Pavlovic, and A. Nehorai, "Two-stage adaptive feedback cancellation scheme for hearing instruments," U.S. Patent 6 754 356, June 2004.

[7] J.M. Kates, "Constrained adaptation for feedback cancellation in hearing aids," *J. Acoust. Soc. Amer.*, vol. 106, no. 2, pp. 1010–1020, 1999.

[8] M. Graciarena, H. Franco, K. Sonmez, and H. Bratt, "Combining standard and throat microphones for robust speech recognition," *IEEE Signal Processing Lett.*, vol. 10, no. 3, pp. 72–74, Mar. 2003.

[9] Z. Zhang, Z. Liu, M. Sinclair, A. Acero, L. Deng, J. Droppo, X. Huang, and Y. Zheng, "Multi-sensor microphones for robust speech detection, enhancement and recognition," in *Proc. 2004 Int. Conf. Acoustics, Speech Signal Processing, (ICASSP2004)*, 17–21 May, 2004, vol. 3, pp. 783–786.

[SP]