Department of Veterans Affairs

Journal of Rehabilitation Research and Development Vol. 30 No. 1, 1993
Pages 39–48

Toward a theory of optimal hearing aid processing

James M. Kates, PhD
Center for Research in Speech and Hearing Sciences, City University of New York, New York, NY 10036

Abstract—An ideal hearing aid for a peripheral hearing loss would process the incoming signal in order to give a perfect match between the cochlear outputs of the impaired ear and a reference normal ear. As a first step toward this objective, a model of the normal and impaired peripheral auditory system was used to derive the optimal hearing-aid processing filter based on a minimum mean-squared error criterion. The auditory model includes the compression and suppression effects of the cochlear mechanics and the sensitivity of the neural transduction process. Simplifying assumptions were then incorporated into the processing to yield a practical frequency-dependent adaptive gain system. Processing examples of several individual speech sounds are presented for a flat hearing loss, and the results indicate that a three-channel compression system with adjustable gains and band edges will be close to the optimal solution for this case.

Key words: *frequency dependent adaptive gain system, hearing aid, hearing aid processing filter, three channel compression system.*

INTRODUCTION

An ideal hearing aid for a peripheral hearing loss would process the incoming signal in order to give a perfect match between the cochlear outputs of the impaired ear and those of a reference normal ear. Implementing this ideal system would require access to the complete set of neural fibers in the impaired ear and to a corresponding set of outputs from an accurate simulated normal ear. The simulated normal outputs could then be substituted directly for the neural responses of the impaired ear. In designing a hearing aid, however, one can only indirectly influence the neural outputs by modifying the acoustic input to the ear. Hearing-aid processing thus represents a compromise in which the acoustic input is manipulated to produce an average improvement in the accuracy of the assumed neural responses in the impaired ear.

Even though the ideal solution is not feasible, one can still consider an optimal hearing aid that modifies the acoustic signal to produce the best possible match between the outputs of simulated normal and impaired ears. An advantage of deriving an optimal system is that it requires an explicit mathematical statement of the problem and of the criterion being used for the solution. Thus the *ad hoc* nature of conventional hearing-aid design is replaced by a more rigorous procedure. But since the objective is to match the outputs of the impaired and normal ears, this procedure will still have embedded within it a set of assumptions about auditory behavior and auditory impairment. Therefore, true optimality may not be achieved due to limitations in the understanding of the auditory system.

The most obvious feature of impaired hearing is the shift in auditory threshold. Since the hearing loss tends to vary with frequency, there have been attempts to compensate for the threshold shift with wide dynamic-range compression systems that have behavior that also varies with frequency. While there has been some limited benefit reported for

Address all correspondence and requests for reprints to: James M. Kates, PhD, Center for Research in Speech and Hearing Sciences, City University of New York, Graduate Center, Room 901, 33 West 42nd St., New York, NY 10036.

Journal of Rehabilitation Research and Development Vol. 30 No. 1 1993

two-channel compression systems (1,2), multichannel compression systems have not demonstrated any significant advantage when compared with broadband compression (3,4,5). It thus appears that building more complicated compression systems does not automatically lead to improvements in speech intelligibility.

Improvements in speech recognition or sound quality will require a processing system that better matches the output of the impaired ear to that of the normal ear. The derivation of an optimal processor will indicate the best possible signal processing for those aspects of the problem that can be improved. By concentrating on auditory function, rather than on the characteristics of speech, the resultant signal processing will be effective independent of the auditory stimulus; the structure of the optimal signal processing will place an upper bound on the complexity of compression algorithms that should be implemented in practical devices.

In the next section of the paper, the problem of designing a hearing aid is stated mathematically, and the general form of the solution is derived. Specific features of the auditory system are then described; compression, the sensitivity of the neural transduction process, and two-tone suppression are incorporated into the auditory model. These features lead to an approximate solution that has a simple digital implementation, and examples of processed speech sounds are presented for the new compression algorithm. The relationships between the optimal algorithm and other compression systems are then discussed.

METHOD

Optimal Filter

The hearing aid and impaired ear form the system for which the optimal filter is being designed, with the normal ear as the reference. A block diagram of this system is presented in **Figure 1**. The input signal is X(k), where k is the frequency index at the output of a fast Fourier transform (FFT). This signal is processed by the auditory analysis filters $H_m(k)$ in the normal ear to produce the normal outputs $Y_m(k)$, where m is the channel index for the analysis filters. The filter function $H_m(k)$ includes the gain and frequency analysis of the

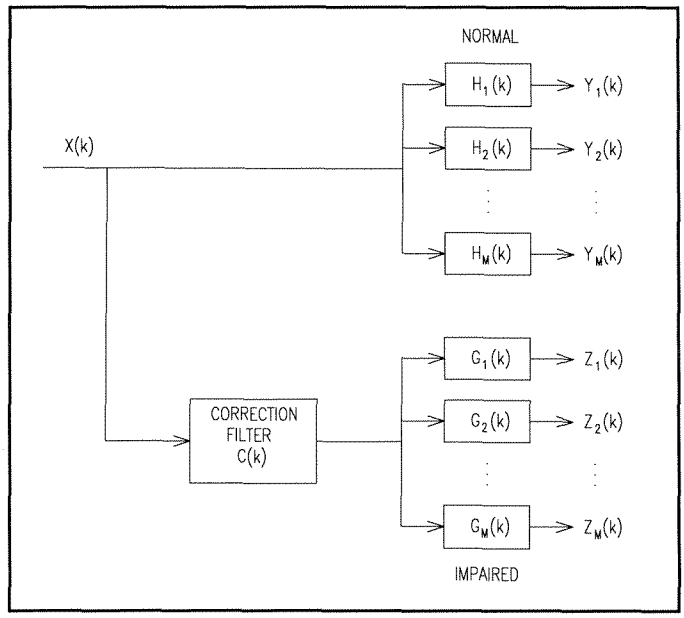

Figure 1.
Block diagram of the hearing-aid processing and the auditory system output.

cochlear mechanics and the sensitivity of the neural transduction process. The fine structure of the inner hair-cell neuron response is not included in this analysis, although it could be part of a more detailed solution that would attempt to match simulated neural firing patterns. The input signal X(k) also goes through the hearing-aid correction filter C(k), after which it is processed by the auditory analysis filters of the impaired ear, indicated by $G_m(k)$, to yield the impaired outputs $Z_m(k)$.

The design objective for the correction filter is the minimization of the differences between the outputs of the impaired ear and those of the normal ear. The criterion chosen is the mean-squared error between the sets of auditory outputs, summed across frequency for each analysis filter and averaged across all of the analysis filters that constitute the auditory model. The error is thus

$$E = \frac{1}{M} \sum_{m=1}^{M} \sum_{k=0}^{K-1} |Z_m(k) - Y_m(k)|^2 \qquad [1]$$

where M is the total number of auditory analysis filters in the model. Substituting the filtered inputs for the outputs yields

$$E = \frac{1}{M} \sum_{m=1}^{M} \sum_{k=0}^{K-1} |C(k)X(k)G_m(k) - X(k)H_m(k)|^2 \qquad [2]$$

Setting to zero the partial derivative of the error with respect to the correction filter coefficients yields the optimal filter, given by

$$C(k) = \sum_{m=1}^{M} H_m(k)G_m^*(k) \Big/ \sum_{m=1}^{M} |G_m(k)|^2 \qquad [3]$$

The filters in the ear change in response to the signal level, so adjusting C(k) may cause additional changes in the auditory filters $G_m(k)$ that represent the impaired ear. Thus several iterations may be required to converge to the solution.

The behavior of the correction filter is essentially to remove the effects of the impaired ear and to substitute the effects of the normal ear. The exact nature of the optimal solution, however, will depend on the specific model of normal and impaired cochlear function, since this determines the characteristics of the auditory filters $H_m(k)$ and $G_m(k)$.

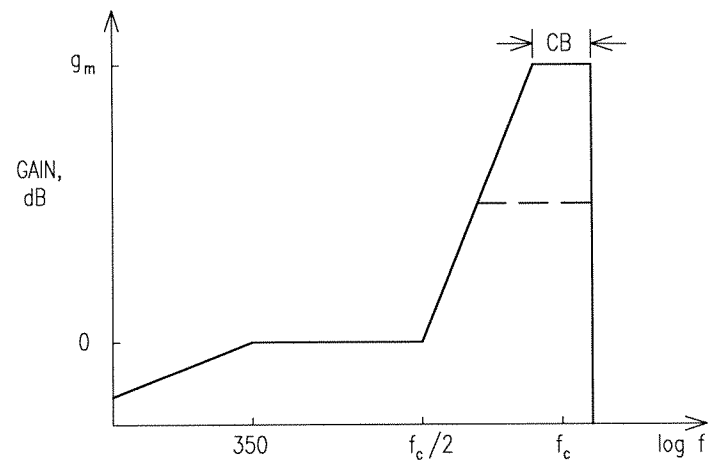

Figure 2.
Idealized auditory analysis filter shape for the condition of maximum filter gain g_m (solid line) and reduced gain due to increased signal level or outer haircell damage (dashed line). The filter bandwidth at auditory threshold is indicated by CB.

Auditory Model

The auditory model includes the effects of the mechanical behavior of the cochlear partition and the sensitivity of the neural transduction process. The purpose of the model is to represent the compression, gain, and frequency resolution in the normal and impaired ears. The model will be used to create a frequency-domain signal-processing system as shown in **Figure 1**, with a linear-phase correction filter C(k). The temporal features of the auditory processing, such as the variation of the auditory filter impulse-response duration with signal level (6), and the adaptation of the neural firing rate (7) are not included, even though they can be implemented in a more detailed model of auditory function (8). The features that are included in the simplified model presented here are those that will most strongly affect the average gain of the correction filter as a function of frequency.

Auditory Filters and Compression

The auditory analysis filters in the normal ear are narrow band-pass filters having high gain at low signal levels, and tend toward low-pass filters having low gain at high signal levels. The shape of the idealized filter used in the auditory model is approximated by the solid line in **Figure 2** for a stimulus at auditory threshold, and is based on the tuning curves measured physiologically in mammals (9,10), and psychophysically in humans (11). The peak of the filter response is at the characteristic frequency f_c, and the filter skirts decay to give a response of 0 dB at $f_c/2$ and decay even more rapidly above the characteristic frequency. Below 350 Hz, the response rolls off due to the transfer function of the middle ear (12).

The bandwidth of the auditory filters varies with frequency, being approximately proportional to the filter center frequency (constant-Q) at high frequencies and approaching constant bandwidth at low frequencies. The psychophysical correlate is the critical band (13,14), indicated by CB in **Figure 2**. The equivalent rectangular bandwidth of the critical band corresponds to a distance of approximately 0.9 mm along the cochlea at low signal levels (15). Additional variation of the filter characteristics with frequency, such as the reduced tip-to-tail ratio of the auditory filters at low frequencies (9) and the reduced relative gain of the auditory filters at low frequencies and at very high frequencies (10), have been ignored in this approximate cochlear model.

The maximum gain of the auditory filter is g_m, which in a healthy cochlea can approach 60 dB for a sinusoid at auditory threshold. Increasing the signal level results in a decrease of gain and a broadening of the auditory filters (16), until at high levels the gain can be reduced to a level close to 0 dB. In a cochlea with extensive outer hair-cell damage, the filter shape and gain is similar at all levels to that of the healthy cochlea at high signal levels (17). The

approximate change in filter shape and gain due to outer hair-cell damage or to an increase in signal level is indicated in **Figure 2** by the dashed line truncating the filter response at a gain below g_m while preserving the rest of the filter shape. Since the filter bandwidth increases with increasing signal level, a distance of 1.2 mm along the cochlear partition was used in the model to approximate the bandwidth at signal levels typical of speech.

To estimate the compression ratio in the healthy cochlea, assume that an input at 0 dB sound pressure level (SPL) receives 60 dB of gain, while an input at 100 dB SPL receives 0 dB of gain. The resultant compression ratio is 2.5:1. A severely impaired cochlea, on the other hand, would have 0 dB of gain at all input signal levels due to the outer hair-cell damage, resulting in a linear system. Total outer hair-cell damage results in a threshold shift of no more than 60 dB, since that is the maximum amount of gain provided by the cochlear mechanics. Hearing losses greater than 60 dB must therefore be accompanied by damage to the neural transduction mechanism, and support for this is provided by Liberman and Dodds (18), who showed that inner hair-cell damage results in a threshold shift but no apparent change in the mechanical behavior of the cochlea. Thus outer hair-cell damage, in this model of hearing loss, causes a loss of sensitivity and a reduction in the compression ratio, while inner hair-cell damage causes a linear shift in sensitivity.

The gain in the correction filter depends on the gain in the normal and impaired ears for the incoming acoustic signal. Let the signal level in a given auditory analysis band be x dB. The gain in the normal ear will then be

$$G_{norm} = 60\,\frac{100-x}{100}\ \text{dB},\ 0 \le x \le 100\ \text{dB} \qquad [4]$$

given the auditory gain assumptions used above. The compressive gain in the impaired ear is assumed to be reduced proportionally by the amount of the outer hair-cell damage, with an additional linear reduction in gain due to the amount of inner hair-cell damage, giving a gain of

$$G_{imp} = 60\,\frac{100-x}{100}\,\frac{60-L_o}{60} - L_i,\ 0 \le x \le 100\ \text{dB} \qquad [5]$$

where L_o is the hearing loss in dB due to the outer hair-cell damage, $0 \le L_o \le 60$ dB, and L_i is the hearing loss in dB due to the inner hair-cell damage. The gain provided by the optimal correction filter, treating this band in isolation, is then given by $G_{norm} - G_{imp}$, and this leads to the family of processing input/output curves shown in **Figure 3**, where it has been assumed that a hearing loss of less than 60 dB is due exclusively to outer hair-cell damage and a hearing loss of greater than 60 dB is comprised of a 60 dB loss due to outer hair-cell damage with the remainder of the loss due to inner hair-cell damage. Hearing losses of less than 60 dB cause a change in the compression ratio, and losses greater than 60 dB add a linear shift in gain in addition to the maximum compression ratio of 2.5:1.

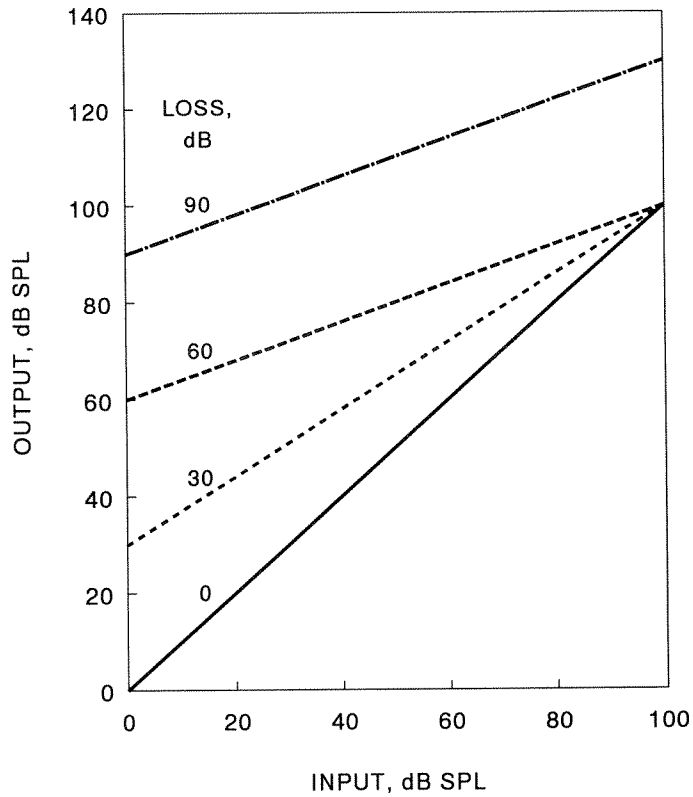

Figure 3.
Compression amplifier input/output curves within an auditory filter band for the model of hearing loss.

The compression model is consistent with fitting procedures, such as the half-gain rule, that have been developed for moderate hearing losses. For example, a narrow-band input at 50 dB SPL would get 30 dB of gain in the healthy cochlea given the compression action represented by Equation [4], while a 60 dB loss due exclusively to outer hair-cell

damage would result in 0 dB of gain for the same stimulus as indicated by Equation [5]. Thus, 30 dB of gain, or half the hearing loss, equalizes the levels that excite the inner hair cells within the auditory analysis band. For more severe losses, where inner hair-cell damage must also be assumed to exist, additional gain beyond the half-gain rule is needed according to the model given in Equation [5], and this has indeed been found to be the case in hearing-aid fittings (19).

Suppression

The second aspect of auditory function that strongly influences the gain in the cochlea is suppression. Physiological measurements of suppression in a normal ear (20,21,22) show that the neural response to a probe tone can be reduced by the simultaneous introduction of a second tone, termed the suppressor. The magnitude of the suppression is approximately linear, that is, increasing the suppressor amplitude by 10 dB will reduce the amplitude of the neural response to the probe to be approximately that of a probe having a 10 dB lower intensity. This effect, over a wide intensity range, depends only on the frequency and magnitude of the suppressor (21). At a stimulus level of 56 dB SPL, representative of speech one-third-octave band intensity levels, the effect of the suppressor has been shown in psychophysical experiments to extend over a frequency region from approximately one-half octave above to one octave below the probe frequency (23). Psychophysical effects for complex suppressors consisting of more than one sinusoid indicate that the suppression is dominated by the most intense sinusoidal component present in the complex (24,25).

The major features of the data cited in the above paragraph can be reproduced by a simple signal-processing model of suppression in the cochlea. When computing the compression gain for an auditory filter, the signal power within the filter is replaced by the maximum signal power observed over a frequency region extending approximately one-half octave above to one octave below the filter characteristic frequency. As a result of the signal-level substitution, the cochlear gain in the auditory model for any given frequency region will be determined by the most intense signal component within that region. Increasing the intensity of the strongest signal component will cause a reduction in the system gain due to the compression, and this gain change will affect all of the less intense signal components within the frequency region. Thus both the frequency- and gain dependence of the cochlear response to a probe in the presence of a suppressor are incorporated into the auditory model.

Outer hair-cell damage reduces the suppression effects in impaired ears (26,27), and this behavior is reproduced qualitatively in the suppression model. Total outer hair-cell loss, for example, results in a linear system given the compression model, and no suppression will be observed because the gain is a constant independent of the signal. Intermediate amounts of outer hair-cell damage will cause an intermediate reduction of the suppression in the model; the largest amounts of suppression will be observed at the higher compression ratios associated with mild hearing losses, and reduced amounts of suppression will be observed at the lower compression ratios associated with more severe hearing losses.

More detailed suppression behavior in impaired ears is not incorporated into the model. In particular, destruction of outer hair cells in the frequency region of the suppressor will reduce the magnitude of the suppression, even when the cochlear behavior in the region of the probe tone appears to be normal (26,28). This effect will be most pronounced for a region of normal hearing bordered above and below by regions of impaired hearing. The assumption made in the suppression model is that this type of hearing loss rarely occurs, and that the reduction in suppression in the impaired ear can normally be described by the hearing loss within the auditory filter. Thus the additional processing complexity of incorporating the details of suppression behavior in the impaired ear is not warranted given the small difference anticipated in the hearing-aid gain function.

Simplified Algorithm

The optimal hearing-aid correction filter given by Equation [3] requires that the output of every auditory filter in the model be computed at every frequency in the FFT, leading to a substantial amount of computation. The amount of computation can be greatly reduced, with only a small sacrifice in accuracy, by replacing the auditory filter of **Figure 2** with an equivalent rectangular bandpass filter; the auditory filter output is then computed

Journal of Rehabilitation Research and Development Vol. 30 No. 1 1993

over just the rectangular pass-band region, with the remaining frequencies set to zero. This simplification leads to a correction filter given by

$$C(k) = H_m(k)G_m^*(k) \;/\; |G_m(k)|^2 \qquad [6]$$

for frequencies k contained within auditory filter m. For a zero-phase hearing-aid system that ignores the auditory filter phase characteristics, the correction filter simplifies even further to yield

$$C(k) = H_m(k) \;/\; G_m(k) \qquad [7]$$

again for frequencies k contained within auditory filter m, and where the filters representing the normal and impaired ears now have a constant gain with zero phase shift within the pass-band and zero gain outside the pass-band. Since most of the auditory filter output power derives from signal components within the pass-band, the effect of the simplification on the optimal correction filter will be minimal, although there may be a slight increase in the low-frequency hearing-aid gain computed for the conditions of total outer hair-cell damage or very intense low frequency maskers since the simplified solution of Equation [7] ignores the effects of upward spread of excitation.

The correction filter given by Equation [7] is superficially similar to a multichannel system using wide dynamic-range compression. In most compression systems, however, the gain is computed independently in each channel (3). The suppression incorporated into the auditory model, on the other hand, links the computed gain values across channels to reduce the variation of system gain with frequency. Thus it is the auditory model that has the greatest influence on the simplified nearly optimal solution and provides its unique characteristics.

The simplified solution of Equation [7] was implemented in a frequency-domain processing system. The incoming speech is sampled at a 20 kHz rate, and divided into segments of 512 samples (25.6 msec) having a 50 percent overlap and weighted with a triangular (Bartlett) window. The spectrum is computed using the FFT, and the signal magnitude is computed in fixed auditory analysis bands corresponding to a distance of 1.2 mm along the cochlea and having an overlap of 0.6 mm; the variation of filter bandwidth with signal level is ignored in this simplified system. The correction filter gain is then computed in decibels for each analysis band, using as a reference signal level the peak output of the auditory analysis bands over a region one octave below to one-half octave above the auditory band center frequency, as explained in the description of the suppression model. The gain calculations take into account the postulated outer and inner hair-cell damage for the impaired ear described in the compression model. The logarithmic gain values for each auditory analysis band are then interpolated across frequency, converted to linear gain values, and used to multiply the input spectrum to obtain the hearing-aid output using an overlap-add procedure.

Examples

The examples are individual phonemes excised from isolated consonant-vowel and vowel-consonant speech tokens produced by an adult male talker. The tokens are at a signal-to-noise ratio (SNR) of 25 dB, with the noise being multitalker babble from the SPIN test tape, and the root-mean-squared (rms) overall level of each speech token was adjusted to be 79 dB SPL. The impaired hearing was set to a flat 60 dB loss assumed to be due entirely to outer hair-cell damage, so the correction filter includes compression but not linear gain given the assumption of no inner hair-cell damage. Since the compression in the processing is wide dynamic range, a shift in the input speech level will cause a smaller compressed shift in the processed output level, but the shape of the hearing-aid frequency response will not change given the flat hearing loss.

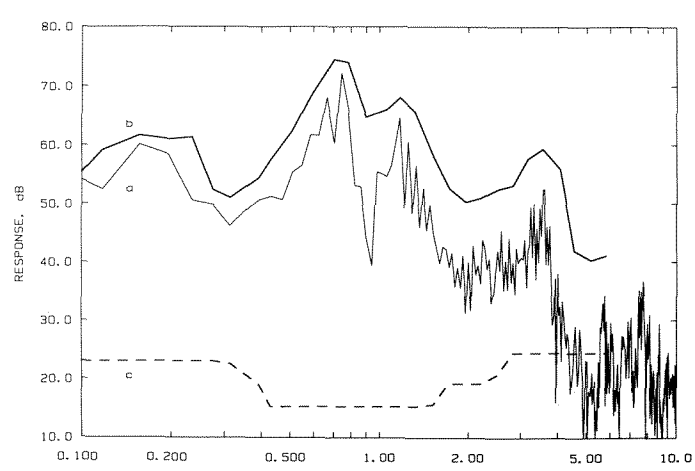

Figure 4.
/a/ in "ka": (*a*) speech spectrum, (*b*) spectrum in auditory analysis bands, and (*c*) computed hearing-aid gain for a flat 60 dB loss.

The first example is a steady-state portion of the vowel /a/ in "ka." The magnitude spectrum of the speech segment is shown in **Figure 4** as curve (*a*), along with the spectrum grouped in auditory analysis bands shown by curve (*b*) and the hearing-aid gain shown by curve (*c*). The vowel second formant at about 1,200 Hz is quite close to the first formant at about 750 Hz, so the first formant peak dominates the gain value computed in the region that contains both formants. The third formant is at about 2,800 Hz, which is more than an octave away from the first or second formants, and so the third formant controls its own distinct gain region. Since the level of the third formant is lower, it receives more gain than the first or second formants. Similarly, the low frequencies also form a region of constant gain determined by the peak at around 200 Hz.

The suppression model thus forms regions of constant gain around the major peaks of the spectrum. The hearing-aid gain function is much smoother than the speech spectrum, and can be represented by a set of plateaus separated by narrow transition regions. Formants that are close together will receive the same gain, thus preserving the details of the local spectral shape, while formants that are farther apart will receive separate gains that will amplify the weaker formant relative to the stronger, thus improving its expected detectability.

The remaining three examples are consonants that differ in their regions of primary spectral power density. The onset of the consonant /p/ in "pa" is shown in **Figure 5**. The spectrum, shown by curve (*a*), has most of its power at low frequencies, and the auditory band powers shown by curve (*b*) show the same pattern. The computed hearing-aid gain, shown by curve (*c*), is separated into a low-frequency region below 1,500 Hz and a high-frequency region above 2,500 Hz. The low-frequency gain is constant below about 1,200 Hz since the two low-frequency peaks in the auditory band spectrum have the same level. The high-frequency gain above 2,500 Hz reaches an asymptote determined by the background speech babble. As was noted for the vowel example of **Figure 4**, the hearing aid gain for the consonant also forms regions of constant gain due to the suppression model, with the gain within each region determined by the outer and inner hair-cell loss.

The second consonant example is just after the onset of /k/ in "ka." As shown in curve (*a*) of **Figure 6**, this consonant has a concentration of power in the midfrequencies around 1,600 Hz. The higher mid-frequency signal level in the auditory analysis bands, shown by curve (*b*), results in the reduced hearing-aid gain shown by curve (*c*) for this frequency region. The peak level of the consonant determines the gain over an octave-and-a-half range, so the shape of the spectral peak is preserved. The gain to either side of the peak is essentially flat, being set by the local peaks at low and high frequencies, so the hearing-aid response is neatly divided into three gain regions.

The last consonant example is from a steady-state portion of /sh/ in "ish." This fricative has most of its power at the high frequencies, as shown

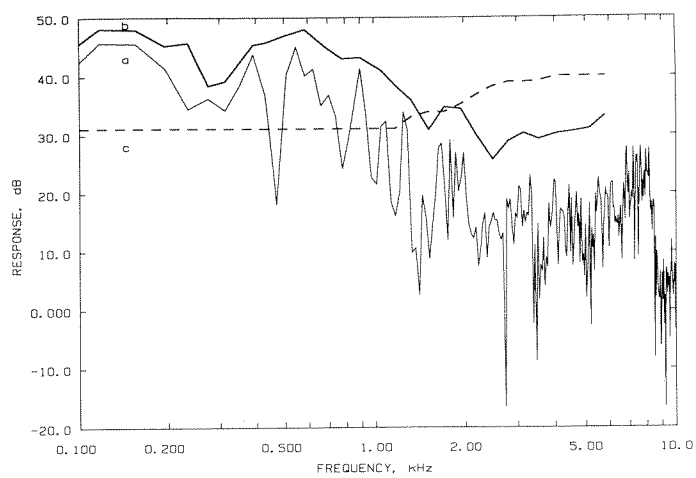

Figure 5.
/p/ in "pa": (*a*) speech spectrum, (*b*) spectrum in auditory analysis bands, and (*c*) computed hearing-aid gain for a flat 60 dB loss.

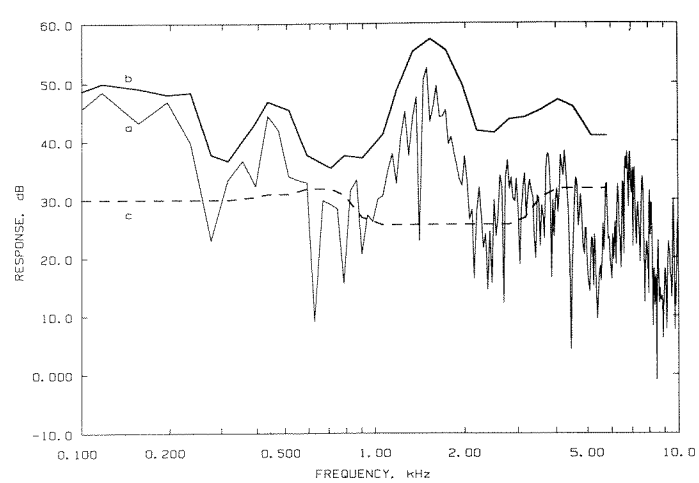

Figure 6.
/k/ in "ka": (*a*) speech spectrum, (*b*) spectrum in auditory analysis bands, and (*c*) computed hearing-aid gain for a flat 60 dB loss.

Journal of Rehabilitation Research and Development Vol. 30 No. 1 1993

in **Figure 7** by the spectrum of curve (*a*) and the auditory analysis bands of curve (*b*). The hearing-aid gain, shown by curve (*c*), has a region of constant gain at high frequencies extending half an octave below and one octave above the high frequency peak at about 2,600 Hz. The low-frequency gain is essentially flat, and there is a broad transition band between the low-frequency and high-frequency regions. The complete outer hair-cell damage assumed in the examples has resulted in the maximum compression ratio of 2.5:1 being used in the system, so the overall spectral contrast has been reduced by the processing while the local variations and the shape and sidelobes of the spectral peak have been preserved.

DISCUSSION

The examples for a flat hearing loss show that the optimal hearing-aid processing is characterized by a limited number of regions of constant gain. This is a direct result of the auditory behavior assumed in the processing, since the suppression model causes each significant peak in the speech spectrum to control the gain over a region of one to one-and-one-half octaves. Thus the hearing-aid gain curves show two or at most three regions of constant gain, separated by gradual transition regions, and this basic pattern holds for all of the speech sounds studied. This processing behavior is illustrated schematically in **Figure 8**. The optimal hearing aid is approximated by a three-band system, where the gain in each band and the location of the band edges describe the frequency response of the system for a flat loss. More than three bands are not needed due to the interaction of the speech spectrum peaks and the modeled effects of auditory suppression.

For a hearing loss that varies with frequency, the hearing aid response becomes more complicated. The curve of **Figure 8** now represents the regions of gain control related to the input speech signal, since the signal value used to compute the processing gain is determined by the speech peaks interacting with the suppression model. This control signal value is then used as the input to the gain computation based on the degree of outer and inner hair-cell loss at each frequency in the compression model. For a hearing loss that varies slowly with frequency, a three-band system with constant gains based on the central frequency in each band will still be a reasonable approximation. For a steeply sloping hearing loss, a level dependent frequency equalization curve will be required within each of the control bands indicated in **Figure 8**, and in this case the general FFT-based algorithm implemented in this paper or an equivalent multichannel system having cross-linked gains would be preferable.

The inclusion of suppression in the auditory model provides a system that reduces spectral contrast between broad, widely separated regions of the spectrum, but preserves the spectral shape within each region. The regions are determined by the input signal spectrum, while the gains within them are determined by the hearing loss. Thus the processing continuously adapts to both the signal spectrum and

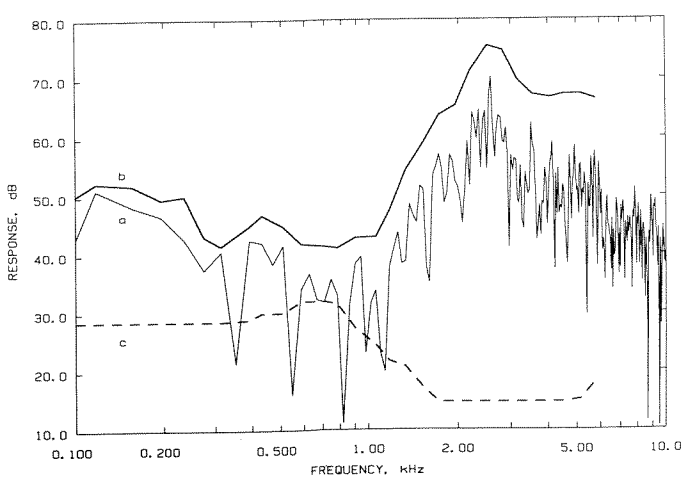

Figure 7.
/sh/ in "ish": (*a*) speech spectrum, (*b*) spectrum in auditory analysis bands, and (*c*) computed hearing-aid gain for a flat 60 dB loss.

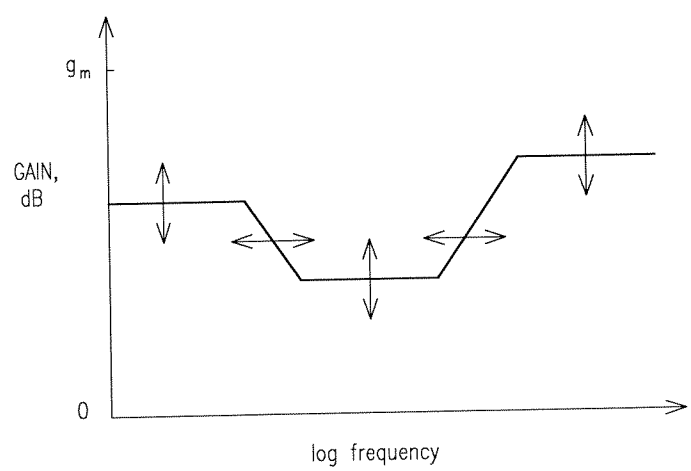

Figure 8.
Frequency response for a three-channel compression system that approximates the optimal hearing aid.

the auditory impairment. This system may offer benefits in comparison with two-channel compression systems having fixed filter crossover frequencies (2); a system having fixed filter band edges may provide inadequate amplification to a weaker speech sound located within the same frequency channel as a more intense sound, since the gain within the channel is determined by the total signal power. The optimal system, however, will provide additional gain for the weaker sound if its spectral power peak is sufficiently distant from that of the more intense sound.

The optimal system may also work better than many recent multichannel systems. Systems having independent gains in each channel (3) can cause wide differences in gain between nearby channels, and the concomitant flattening of the spectrum may reduce intelligibility. The broad regions of uniform gain in the optimum algorithm avoid this problem. The algorithm presented in this paper may also be more effective than compression systems that globally reduce the spectral contrast (4,5), since in the optimal processing the less intense regions of the spectrum receive additional gain, but there is no change in the spectral shape of a significant peak or in the relative amplitude of its sidelobes. The gains determined by the optimal processing algorithm indicate that only a modest amount of spectral modification is needed for the hearing-impaired; large spectral changes may increase rather than reduce the error between the auditory outputs in the impaired and normal ears that is used as the design criterion.

The relevance of the mathematically optimal solution is dependent on the accuracy of the optimization criteria and the assumptions built into the problem formulation and solution. Other optimization criteria, such as minimizing a mean-squared error formulation more representative of average neural firing rates, may yield better processing since the use of the linear outputs favors the higher signal levels. This change, however, would not cause any substantial differences in the simplified algorithm. Weighting the error would also influence the processing; the correction filter can match one auditory-filter output exactly or many outputs approximately, so a weighting function would apportion the accuracy of the match across auditory filter channels. It may also prove advantageous to provide a compression threshold in the processing algorithm in order to limit the gain for low-level sounds, thereby reducing the annoyance of some background noises.

The optimal processing system may also have applications beyond hearing aids, especially since no assumptions about speech or any other specific signal characteristics have been designed into the processing. Broadband compression for audio signals, for example, can be accomplished by setting a fixed compression ratio (e.g., 2:1) in the algorithm, which would be equivalent to selecting a mild-to-moderate flat loss due exclusively to outer hair-cell damage. Stereophonic signals would require selecting the same processing gain in both signal channels; one approach would be to control the gain with a composite spectrum made up of the larger of the left and right channel outputs in each auditory analysis band.

CONCLUSIONS

This paper has presented a new hearing-aid processing approach based on minimizing the mean-squared error between the outputs of simulated normal and impaired cochleas. An optimal solution was derived, and the salient features of the solution were preserved in a simplified frequency-domain algorithm having an efficient digital implementation. The resulting hearing-aid processing system is dependent solely on the characteristics of the modeled cochlea, and not on any assumptions about the nature of the speech signal. Thus the processing should work equally well on speech, music, and environmental sounds. Since the structure of the algorithm is clearly dependent on the optimization criterion and on the characteristics of the model used to represent normal and impaired hearing, inadequacies in the auditory model will limit the effectiveness of the signal processing. Even though the system is promising, it will still be imperfect due to the inability to completely model the auditory periphery and the effects of hearing loss.

The auditory model includes cochlear suppression effects, and this provides an important difference between the new algorithm and previous multichannel compression systems. Because of the modeled suppression, the hearing-aid gain is typically divided into two or three frequency regions, with the gain in each region governed by the most

intense spectral peak within the region. This results in a gain function that varies smoothly with frequency, increasing the relative gain of the less intense spectral regions while preserving the details of the spectral shape within each region. The speech examples for a flat hearing loss indicate that a three-channel system with variable channel gains and crossover frequencies can be designed to have performance close to that of the optimal system. Additional independent channels do not appear to be necessary in a hearing aid and may even be counterproductive, since additional spectral modifications may increase rather than reduce the error between the auditory outputs in the impaired and normal ears that is the hearing-aid design criterion.

ACKNOWLEDGMENT

This work was supported by the National Institute on Deafness and Other Communication Disorders (NIDCD) under Grant 2POlDC00178.

REFERENCES

1. Villchur E. Signal processing to improve speech intelligibility in perceptive deafness. J Acoust Soc Am 1973;53(6):1646-57.
2. Moore BCJ. Design and evaluation of a two-channel compression hearing aid. J Rehabil Res Dev 1987;24(4):181-92.
3. Walker G, Byrne D, Dillon H. The effects of multichannel compression/expansion amplification on the intelligibility of nonsense syllables in noise. J Acoust Soc Am 1984;76(3):746-57.
4. Bustamante DK, Braida LD. Principal-component amplitude compression for the hearing-impaired. J Acoust Soc Am 1987;82(4):1227-42.
5. Levitt H, Neuman A. Evaluation of orthogonal polynomial compression. J Acoust Soc Am 1991;90(1):241-52.
6. Møller AR. Frequency selectivity of single auditorynerve fibers in response to broad-band noise stimuli. J Acoust Soc Am 62(1):135-42.
7. Harris DM, Dallos P. Forward masking of auditory nerve fiber responses. J Neurophys 1979;42(4):1083-1107.
8. Kates JM. A time-domain digital cochlear model. IEEE Trans Acoust Speech Sig Proc 1991;39(12):2573-92.
9. Kiang NYS, Moxon EC. Tails of tuning curves of auditory nerve fibers. J Acoust Soc Am 1974;55(3):620-30.
10. Kiang NYS. Processing of speech by the auditory nervous system. J Acoust Soc Am 1980;68(3):830-5.
11. Nelson DA, Freyman RL. Broadened forward-masked tuning curves from intense masking tones: delay-time and probe-level manipulations. J Acoust Soc Am 1984;75(5):1570-7.
12. Lynch TJ, III, Nedzelnitsky V, Peake WT. Input impedance of the cochlea in the cat. J Acoust Soc Am 1982;72(1):108-30.
13. Zwicker E, Terhardt E. Analytical expressions for critical-band rate and critical bandwidth as a function of frequency. J Acoust Soc Am 1980;68(5):1523-5.
14. Moore BCJ, Glasberg BR. Suggested formulae for calculating auditory-filter bandwidths and excitation patterns. J Acoust Soc Am 1983;74(3):750-3.
15. Greenwood DD. A cochlear frequency-position function for several species—29 years later. J Acoust Soc Am 1990;87(6):2592-2605.
16. Johnstone BM, Patuzzi R, Yates GK. Basilar membrane measurements and the travelling wave. Hear Res 1986;22:147-53.
17. Harrison RV, Aran J-M, Erre J-P. AP tuning curves from normal and pathological human and guinea pig cochleas. J Acoust Soc Am 1984;69(5):1374-85.
18. Liberman MC, Dodds LW. Single neuron labelling and chronic cochlear pathology III: stereocilia damage and alterations of threshold tuning curves. Hear Res 1984;16:55-74.
19. Byrne D, Parkinson A, Newall P. Hearing aid gain and frequency response requirements for the severely/profoundly impaired. Ear Hear 1990;11(1):40-9.
20. Sachs MB, Kiang NYS. Two-tone inhibition in auditory-nerve fibers. J Acoust Soc Am 1968;43(5):1120-8.
21. Javel E. Suppression of auditory nerve responses I: temporal analysis, intensity effects and suppression contours. J Acoust Soc Am 1981;69(6):1735-45.
22. Delgutte B. Two-tone suppression in auditory-nerve fibers: dependence of suppressor frequency and level. Hear Res 1990;49:225-46.
23. Duifuis H. Level effects in psychophysical two-tone suppression. J Acoust Soc Am 1980;67(3):914-27.
24. Stelmachowicz PG, Small AM, Jr., Abbas PJ. Suppression effects for complex stimuli. J Acoust Soc Am 1982;71(2):410-20.
25. Weber DL. Combined effects of two suppressors. J Acoust Soc Am 1984;75(5):1563-9.
26. Dallos P, Harris DM, Relkin E, Cheatham MA. Two-tone suppression and intermodulation distortions in the cochlea: effect of outer hair cell lesion. In: Van den Brink G, Bilsen F, editors. Psychophysical, physiological, and behavioral studies in hearing. The Netherlands: Delft Univ. Press, 1980: 242-9.
27. Schmiedt RA. Acoustic injury and the physiology of hearing. J Acoust Soc Am 1984;76(5):1293-1317.
28. Robertson D, Johnstone BM. Primary auditory neurons: nonlinear responses altered without changes in sharp tuning. J Acoust Soc Am 1981;69(4):1096-8.