Current Design Options and Criteria for Hearing Aids

Conceptions et critères actuels pour les prothèses auditives

William A. Cole
President, Etymonic Design Inc.

Key words: hearing aids, amplification, compression, AGC, programmable

Abstract

It is common to assume that, when hearing aids fail to meet the needs of their owners, either the hearing aid design criteria or the available technology was inadequate. This paper reviews several design options and presents, from a designer's perspective, the goals of these approaches, pointing out that performance inadequacies have often been the result of design compromises in response to market pressures rather than a lack of adequate design criteria or technology.

Résumé

Il n'est pas inhabituel de présumer que, lorsque des prothèses auditives ne répondent pas aux besoins de leurs utilisateurs, ni les technologies ni le modèle conceptuel des aides auditives étaient adéquats. L'auteur examine plusieurs modèles conceptuels possibles et présente, du point de vue du concepteur, les buts de ces approches. Il montre que les lacunes au niveau de la performance sont souvent le résultat de compromis effectués au niveau de la conception en réponse aux pressions du marché, plutôt que du manque de technologies et/ou de critères de conception appropriés.

Introduction

Hearing aid design is a process of striking compromises between size and performance. Most of the criteria for good hearing aid design have been known for decades. Over fifty years ago, Wengel (1940) stated that "It has been found that in the majority of cases of deafness the patient's ear is more deficient with respect to the higher frequencies than the lower ... It has also been found that the ears of certain deaf persons have rather narrow amplitude ranges." (p. 1)

Over 40 years ago, Poliakoff (1950) listed the attributes of a good hearing aid as:

1. Giving the patient his optimum volume in all reasonable conditions of use. In most cases this cannot be achieved without automatic volume control.
2. Avoidance of pronounced peaks.
3. Low case noise.
4. A nice looking response curve.

Both of these designers were describing the rationale for a K-Amp (R)* hearing aid, and the technology existed to build such a device even in 1950! However, in 1950, Poliakoff also noted that "No user wants to look like a military radio-set operator" (p. 274).

It is a given that every hearing aid designer would like to produce a hearing aid with no distortion, no internal noise, and a smooth frequency response. However, approaching these ideals invariably involves increased size. Noise and distortion are generally inversely related to power consumption and circuit complexity, and hence the size of the circuit and battery. Smoothness of response is related to the damping of the transducers and hence to their efficiency and thus, battery size. Technological advances over the past 50 years have repeatedly presented the hearing aid designer with two choices — make it smaller or make it better. We have only to compare a modern canal aid with a 1940 body aid to determine which choice the market rewarded. Given the available technology, each is as small as it can be — consistent with some acceptable level of performance. Fifty years of research have failed to define what this acceptable level of performance is. By default, it frequently has been defined by the hearing aid designer or by marketing departments interpreting the requirements of hearing impaired people as interpreted by hearing aid dispensers. Fifteen years ago, Pollock (1977) reported that 77% of audiologists and 98% of dealers reported that current hearing aids were meeting the needs of their clients. One cannot blame the hearing aid designer for concluding that the compromises being made resulted in an acceptable level of performance.

We are now at a unique period in hearing aid evolution. There is little market pressure to produce a hearing aid smaller than a canal aid, and technology is making it possible to achieve higher levels of performance in this small size. The need to compromise performance for size is finally vanishing, and hearing aid design now truly may be limited only by a lack of design criteria.

* K-Amp is a registered trademark of Etymotic Research Inc.

Design Options

Linear Amplification

Linear amplification is the most prevalent type of hearing aid fitting (Cranmer 1992). In its ideal form, it is simply the frequency dependent scaling of the input signal to any desired output level without corruption.

Rationale

The most fundamental characteristic of hearing impairment is a loss of hearing sensitivity, and the most obvious way to deal with this is to increase the level of the stimulus presented to the ear so that it becomes audible. Because hearing loss is a function of frequency, such linear amplification should reflect this frequency dependence.

Compromises

Linear hearing aids can depart from the ideal by failing to provide the desired gain frequency response, introducing noise and/or distortion, and/or by restricting the dynamic range. If size were not a factor, all of these deficiencies could be reduced to insignificance as they have been in modern audio equipment. However all of them are the subject of compromise in most linear hearing aid designs.

Frequency Response. For decades hearing aids were designed with complete disregard for their real ear performance, and as a consequence most failed to meet the basic design criterion of restoring audibility. Even when this fact was known (Cole, 1975b) and solutions presented (Killion, 1978), hearing aid designers continued to compromise frequency response to save space, simplify manufacturing, and produce the largest gain and output figures for their data sheets.

The frequency response of most hearing aids is determined by the microphone and receiver (and associated "plumbing"). The microphones used in hearing aids for the past 25 years are nearly ideal devices having a smooth wideband frequency response and a noise level rivalling that of the normal ear (Killion, 1976). The drive toward small, highly efficient receivers in the 1960s resulted in devices of high acoustic impedance coupled to the ear with long tubing that produced peaks and valleys. Although Carlson (1974) demonstrated that a smooth insertion response was possible with these miniature receivers, few hearing aid manufacturers considered that the benefits outweighed the increased size and complexity that this implementation entailed. As a further deterrent, the existing hearing aid standards penalized hearing aids with smooth responses peaking in the 2 - 4 kHz region (desirable for restoration of normal ear canal resonance) and gave the biggest numbers to those with a peak at 1 kHz, which could be obtained without doing anything. Killion (1978) demonstrated that performance rivalling that of the best high fidelity speakers was possible with these miniature receivers in both ITE and BTE configurations, at the expense of power consumption, cerumen-collecting dampers, and stepped sound bores. No hearing aid designers considered that the improved performance was worth the price. In short, hearing aid designers have made and continue to make, the choice that minimizes size, cost, and manufacturing problems because these are the choices that get rewarded.

Noise. Noise is the result of the granular nature of matter. All amplifiers introduce noise, but some amplifier configurations (such as differential pairs) and some transistor types (such as MOSFETs) introduce more than others. Noise is inversely related to the current used in the preamplifier and is inversely related to the size of the coupling capacitors employed. In the interest of minimizing battery and amplifier size, designers may reduce preamplifier current or the size of coupling capacitors at the expense of higher noise levels. In the interest of reducing the number of components, inherently noisy circuit configurations may be employed or tone and volume controls may be placed where they add to the circuit noise.

Distortion. An ideal amplifier is able to scale any signal without other alteration. In the real world, amplifying devices are inherently non-linear and these non-linearities must be compensated effectively to reduce distortion. This compensation requires additional components and in some cases additional power consumption, so the hearing aid designer must decide what is an acceptable level of distortion. The meaningful characterization of circuit non-linearity is the subject of much current study (Kates, 1990), and its impact on the hearing aid user is at best a guess, so the hearing aid designer settles for a compromise that looks good on a data sheet (i.e., less than 10% total harmonic distortion).

Headroom. An amplifier non-ideality that has received much attention recently is signal clipping. This occurs when the amplified signal from the microphone exceeds the capability of some amplifier stage. If this stage is before the point of volume control, no amount of volume control reduction will reduce the resulting distortion. It is more common for this clipping to occur at the output stage of the amplifier, in which case reducing the volume control setting will reduce the distortion. This phenomenon has been called "the headroom problem" (Preves et al., 1990) and has been erroneously attributed to the low voltage of hearing aid batteries. A lack of headroom is a direct result of compromises made by the hearing aid designer and has nothing to do with battery voltage. Hearing aids exist that can deliver High Frequency Average (HFA per ANSI S3.22, 1987) outputs in excess of 133 dB SPL with a 1.3 volt battery. How much more headroom can the human ear stand?

The headroom problem results from "low SSPL90 combined with high gain" (Preves et al., 1990, p. 19). Clearly

Table 1. Available HFA-SSPL90 for 100 hour battery life with ideal amplifiers in an ITE configuration for various receiver and battery types.

KNOWLES RCVR TYPE	RCV SIZE (CC)	HFA-SSPL90 FOR 100 HR (ANSI) BATTERY LIFE #10	#312	#13	#675
CI	0.29	129	132	135	139
EF	0.19	125	129	132	136
ED	0.08	121	124	127	131
EH	0.056	114	118	121	125

then, it can be resolved by increasing the SSPL90 and/or reducing the gain. The SSPL90 is controlled by receiver type, its electrical impedance, and the acoustic coupling system — all chosen by the hearing aid designer. The gain is controlled by the number of amplifier stages — also a designer's choice. Sometimes hearing aid designers have tried to respond to the all too common request for a hearing aid with 70 dB of gain and a maximum output of 95 dB SPL (without automatic gain control). But more often, it is the demand for high gain in the smallest possible space that leads to the headroom problem. Table 1 shows the estimated available HFA-SSPL90 for 100 hours battery life with ideal amplifiers in an ITE configuration using various receiver types and zinc-air battery sizes. Battery life is calculated as the published battery capacities divided by the ANSI current drain.

It can be seen that choosing a #10 battery and an ED receiver limits HFA-SSPL90 to 114 dB in the ideal case. Assuming enough headroom to amplify 75 dB speech peaks without clipping gives a maximum HFA gain of 39 dB, while choosing a CI receiver and a #675 battery gives a maximum gain without clipping of 64 dB. For a given gain, receiver, acoustic coupling system, and battery life, the SSPL90 (and thus the headroom) can be increased only by improving the efficiency of the amplifier or increasing the energy capacity (not the voltage) of the battery. For a given physical volume, the energy capacity of a battery is determined primarily by the technology used (zinc-air, mercuric oxide, etc.). The zinc-air technology in current use produces an energy capacity double that of any other technology (including the high voltage lithium technologies). Developing high voltage batteries with lower energy capacities will make the headroom problem worse, not better.

Amplifier type plays a major role in the headroom problem because it determines the average power consumption for a given maximum acoustic output. This is directly related to battery life and hence to required battery size. The class B and D amplifiers have a considerable advantage over the class A amplifier (Carlson, 1989). This is because the class A circuit always draws the same power from the battery regardless of the power it is supplying to the receiver. The class B and class D amplifiers, on the other hand, only draw power from the battery when it is needed and the rest of the time operate at quiescent power, which may be 10 to 50 times lower than the peak requirement. In spite of this, the class A output stage was used extensively for ITE aids in the 1970s and early 1980s because it was cheaper and much smaller and simpler than the class B, which was reserved for power aids for which class A was totally unacceptable. A notable exception was the Dahlberg SHARP (R) circuit that was a class AB circuit made possible by hybrid packaging technology not generally available to other manufacturers of the time. In order to avoid the need for large batteries, other hearing aid designers simply "starved" the class A circuits. That is to say, they reduced the power consumption to the point at which the hearing aids clipped at very low input levels (i.e., created the headroom problem). The development of smaller class B amplifiers and of the Knowles integrated class D amplified receiver has eliminated the need for the compromises that led to the headroom problem in all but the smallest of hearing aids. However, the presence of peak clipping controls on hearing aids allows the clinician to reintroduce the headroom problem in the event that low distortion is not really a valid design criterion after all.

Amplification With Automatic Gain Control

Automatic gain control (AGC) adaptively changes the gain of a hearing aid in response to the level of the signal being processed. This has taken various forms and the terminology and classification has been inconsistent (Preves, 1991). In spite of this, it is possible to distinguish three types of AGC with three different objectives (Braida et al., 1979). These are High Level Limiting, Automatic Volume Control (AVC), and (Syllabic) Compression.

Figure 1 is a block diagram of a typical single channel AGC hearing aid. It consists of an amplifier with electronically controlled gain and a detector/controller that derives a controlling signal from some point in the circuit. The three input/output plots at the bottom show the result of attenuation at points A, B, and C. These have been stylized to show the operation of a High Level Limiter, but other input/output functions are possible. The three frequency response plots along the right hand side show the effect on both gain and AGC threshold of high pass filtering at points A, B, and C. The hearing aid designer may chose to place user or fitter controls or frequency shaping circuits (either fixed or adjustable) at any of these points or, through circuit simplification, may eliminate some or all of them.

Rationale

The rationale for High Level Limiting is simply to prevent amplified sound from becoming uncomfortable or hazardous

Figure 1. Block diagram of an AGC hearing aid. Static input/output functions for attenuation at points A, B, and C are shown at the bottom. The frequency dependence of gain and AGC threshold for high pass filtering at points A, B, and C is shown to the right.

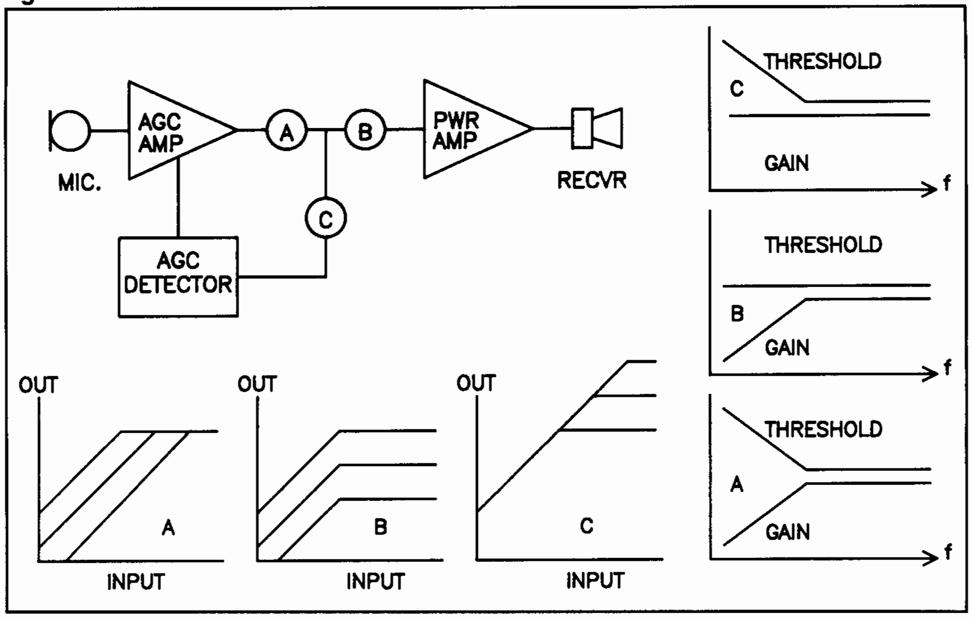

(without the distortion associated with clipping), thereby encouraging the use of higher gain settings. Automatic Volume Control is intended to maintain the output of the hearing aid at a comfortable level regardless of input level variations, thereby reducing the need for frequent volume control adjustments. Early hearing aid designers discovered that many hearing impaired individuals had restricted dynamic ranges (Wengel, 1940) and that they had well defined optimum listening levels for speech (Poliakoff, 1950). AVC circuits were a natural response to this discovery and they had the additional benefit of preventing clipping (thereby solving the headroom problem before it was discovered). The rationale for syllabic compression lies in the observation that low energy consonants are often inaudible to the hearing impaired listener. Syllabic compressors are intended to improve the consonant-to-vowel ratio (CVR), thereby enhancing intelligibility. It is not often appreciated that in improving the CVR, such systems will also degrade signal to noise ratio and signal to reverberation ratio.

Although these rationale seem fairly straightforward, many researchers and hearing aid designers have failed to grasp the distinctions among them or the fact that they may all need to exist in the same hearing aid. This has resulted in possibly every combination of dynamic and static characteristics and user and fitter adjustments imaginable (Nabelek, 1973, 1975) and widely varying claims of success (Braida et al., 1979). Given that the operating characteristics of any of these combinations also can be greatly altered by design compromises, it is little wonder that the promise of AGC often has not been fulfilled.

Compromises

AGC amplifiers are used daily in the broadcast and recording industry to reduce dynamic range. The dynamic range of an FM radio station is 40 dB and that of an AM station is only 20 dB, yet this reduction is accomplished without perceptible degradation of speech or music (Blesser, 1969). The fact that this is not the case in hearing aids is due to both intentional signal processing and design compromises. Many of the techniques used in broadcast applications require high voltage supplies, multiple AGC stages, and multiple trimpots to compensate for component tolerances, making them inappropriate for use in hearing aids. The complexity of AGC circuitry provides both the need and the opportunity for compromise in order to meet space and power constraints.

Broadcast AGC systems have complex input/output functions with regions of expansion, moderate compression, and limiting. This may be contrasted with hearing aids which frequently have only a linear region and a region of limiting. While the rational for this simplification lies largely in the reduced circuitry required to achieve it, it can also be justified for AVC on the grounds that the reduced auditory area of the impaired listener makes it important to maintain the average signal level within a narrow range. It should also be noted that, because complex signals are not steady state, the dynamic properties of an AGC system have a much greater

Figure 2. (a) Frequency response characteristics of BILL and TILL type level dependent filters. (b) Block diagram of feedforward level dependent filter. (c) Block diagram of feedback level dependent filter.

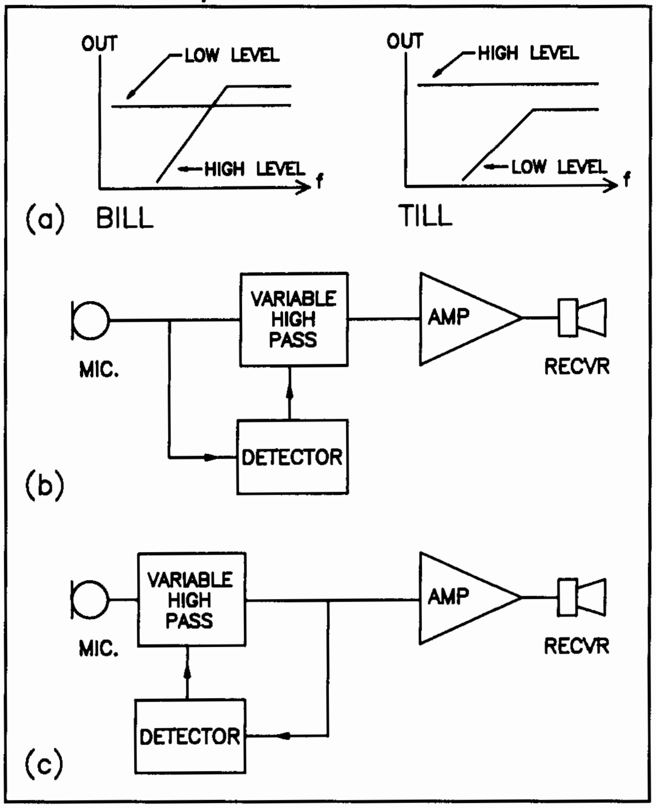

effect on the nature of the compression than does the input/output function (Blesser, 1969; Dillon, 1988).

The dynamic properties of broadcast AGC systems are carefully optimized for the dynamic range of the channel and the program material. The attack time frequently is chosen to be in the 5-10 msec range to avoid overemphasizing loud transients and to prevent peaks from controlling the gain (Blesser, 1969). The widespread use of standard integrated circuits for AGC hearing aids has tended to produce more consistent stable attack transients than those reported by Nabelek in 1973, and they tend to fall in the 5-10 msec range. Recovery characteristics, on the other hand, are very different for broadcast AGC systems than for most hearing aids. Broadcast systems have multiple release times that depend on the peak to average ratio of the program material, the interval between peaks, and the number of peaks occurring per second. Although some hearing aids with adaptive release time have been marketed (Newby, 1979; Teder, 1991), the majority have had single release time constants for reasons of circuit simplicity. This has given rise to the common complaints, familiar to hearing aid dispensers, of "dropouts," "breathing," and "pumping," particularly in the presence of background noise.

One of the most common compromises in AGC hearing aid design is to derive the controlling signal from the output of the power amplifier rather than the output of the AGC amplifier. A significant reduction in the number of parts required is realized by combining the AGC amplifier with the power amplifier and connecting the detector to the output where signal levels are highest. However, this eliminates controlling point B (Figure 1) and its related input/output and frequency response functions. This is commonly called Output AGC and it is characterized by a volume control that adjusts gain independent of maximum output. The availability of controlling point B provides the possibility of a volume control that controls both gain and maximum output together (Input AGC) allowing the user to set and maintain a comfortable listening level.

Many compromises can and have been made in both the AGC amplifier and the detector. These may result in limited AGC range, undesirable input/output functions, and poor dynamic performance. However, the advent of the integrated circuit AGC amplifier (Cole, 1975a) has reduced the number of such compromises by making it feasible to incorporate more sophisticated circuitry into designs that work from a single battery. The advent of the K-Amp (R) (Killion, 1988) with dual time constant release has, for the first time, provided hearing aid designers with an AVC type amplifier without compromises.

Level Dependent Frequency Response (LDFR)

Hearing aids with level dependent frequency response (LDFR) adaptively change their frequency gain characteristic in response to the level and/or spectrum of the signal being processed. The two most common types have been identified as BILL (Bass Increases at Low Levels) and TILL (Treble Increases at Low Levels) (Killion et al., 1990) and their characteristics are diagrammed in Figure 2(a).

LDFR is achieved by varying the characteristics of an electronically controlled filter in response to the level of the signal at some point in the circuit (Figure 2). (It may also be achieved using multi-channel AGC described in a later section.) The signal that controls the filter characteristics may be derived before the controlled filter as in Figure 2(b) (called feedforward control) or from a point after the filter as in Figure 2(c) (called feedback control). (The term feedback as used here should not be confused with its use to mean self-oscillation.) The configuration of Figure 2(b) tends to be more complex because the detector must be more sensitive than in 2(c), which is sensing the signal after amplification, and it must also function with a wider range of signal levels.

BILL type systems are designed to attenuate low frequencies when they are present but otherwise to provide a wideband response. The controlled filter normally has a wideband response but this becomes a high-pass as the detector is activated. The configuration of Figure 2(b) allows a wide range of variation in the low frequency response to occur for a small change in the low frequency content of the signal. The simpler configuration of Figure 2(c) is rather unresponsive because the low frequencies needed by the detector to adapt the controlled filter are excluded from it. A problem that arises with BILL type systems is the need to differentiate between the fundamental of the male voice and background noise so that the response is not altered by a speech signal in quiet. This may be done by filtering the signal to the detector so that it responds only to non-speech frequencies and/or choosing detector time constants so that it does not respond to rapidly fluctuating signals (Gebert, 1988).

TILL type systems are designed to amplify high frequencies selectively and to become wideband at high levels. The controlled filter normally provides high frequency emphasis but becomes a wideband low gain amplifier as the detector is activated. The configuration of Figure 2(c) is used, which eliminates the problem of differentiating between low frequency noise and low frequency speech because noise rarely has significant high frequency energy and so the response and gain are only altered by the high frequency components of speech.

Rationale

BILL systems resulted from the observation that hearing aid wearers often performed better in a noisy environment when the hearing aid response was altered to reduce low frequency gain. However, this frequency response produced complaints of a "tinny" sound and was judged unacceptable in a quiet environment. The solution is to provide a hearing aid with a wideband gain for normal inputs and low frequency reduction for noisy inputs.

TILL systems resulted from the observation that the most common form of hearing loss is characterized by a loss of high frequency hearing sensitivity but little change in the uncomfortable listening level. The solution to this problem is a hearing aid with high frequency gain for low level inputs and acoustic transparency for high level inputs (Killion, 1979). A second rationale is that high frequency gain for low level inputs excludes much of the low frequency background noise in the absence of speech, while a broadband low gain response triggered by high frequency speech components produces a high quality sound with improved signal/noise ratio due to reduced forward and backward masking (Cole, 1986).

Compromises

BILL type systems have been the subject of several compromises to reduce size. These include the use of the configuration of Figure 2(b) at the expense of reduced low frequency attenuation and the use of simplified noise/speech differentiation methods at the expense of some unintentional response change for male voices in quiet. The only commercially available TILL system is the K-Amp (R), which only has a 6 dB/octave high frequency emphasis. This compromise limits its range of application to those with moderate sloping losses.

Multi-Channel Hearing Aids

The multi-channel hearing aid divides the audio band into two or more channels with gain control in each. One or more of the channels may have AGC with adjustable static and dynamic characteristics. The gain in the AGC channels may be controlled by the signal level in the channel or by some combination of levels in other channels.

Rationale

The multi-channel AGC hearing aid was originally proposed as a syllabic compressor (Villchur, 1973) to compress the dynamic range of speech for those hearing impaired with severe cochlear losses. Twenty years of research has failed to provide solid support for the utility of syllabic compression but has reinforced the importance of audibility. This has motivated some hearing aid designers to employ multi-channel amplification to provide more accurate frequency shaping and high level limiting. Other designers have proposed multi-channel schemes to reduce background noise (Kates, 1989). Multi-channel AGC hearing aids represent an alternative (more complex) method to those previously described for achieving level dependent frequency response and the same rationale apply.

Compromises

Most multi-channel schemes have been implemented either on paper or on laboratory scale equipment so compromise has not been necessary. However some wearable units have entered the market (Craigwell, Resound, Triton) and their designers have been faced with making the classic size/performance trade-offs. These may include a reduction in the number of channels, the elimination of AGC in some or all channels, high internal noise levels, reduced dynamic range, non-optimum control functions, or limited adjustability. As complexity is increased, the adjustment possibilities become large and some form of electronic programmability becomes essential.

Digital Control of Analog Hearing Aids

This option was first proposed by Mangold et al. (1979) to control a multi-channel hearing aid and first implemented in a headworn hearing aid in 1985 (Cole, 1985). It replaces screwdriver adjustments with digitally controlled electronic switches and potentiometers controlled by electronic memory. The electronic memory is programmed by an external computer (via a cable, RF, infrared, or ultrasonic link), which may be controlled by the hearing aid fitter, the hearing aid wearer, or both. A variety of hearing aid characteristics can be altered by changing the settings stored in the electronic memory, and several sets of settings may be stored and recalled via a user operated switch or hand-held remote programmer. The long-term memory need not be in the hearing aid but may be in the hand-held programmer with parameters loaded into a short-term memory in the hearing aid via a wireless link each time a change is made.

Rationale

The hearing aid industry has suffered from a "low tech" image because much of the high technology in a hearing aid is deliberately invisible to the general public. This author became involved in the development of digitally programmable analog hearing aids because a hearing aid company needed a "high tech" image. While other companies may have developed programmable hearing aids for other reasons, the need for a computer connection has undoubtedly been a motivating force. Mangold et al. (1979) proposed digital control of analog circuitry as a way to allow for more adjustability without increasing the size of the hearing aid by adding more screwdriver adjustments. An additional benefit is the possibility of making adjustments at several points in a circuit at once to achieve a desired effect. This is simply not possible with trimpots.

Digital programmability opens the door to computer assisted fitting, which not only can save time, but also becomes a necessity as increased adjustability creates a multidimensional optimization problem. As hearing aids become smaller and the controls more difficult for human fingers to operate, remote programmability provides a convenient way to operate volume and tone controls and makes it possible to provide different operating characteristics for different listening situations.

Compromises

Any of the compromises possible in trimpot controlled analog circuitry can also be made with digital control. The compromises that are truly unique to this design option are related to the memory, its programming, and its interface to the external computer. Early programmable instruments employed small 3 volt lithium batteries to maintain the memory, and these were prone to failure because they had not been designed for this application. Newer lithium batteries and memory technologies that require no batteries have eliminated memory loss problems. Interfacing to the programmer is commonly accomplished via a cable, and the connectors are a compromise between size and reliability. When wireless links are used, the compromise is made between simplicity (ultrasonic or infrared) and the user's ability to operate the control out of sight. RF links do not require close proximity to the hearing aid being controlled but are more complex.

Conclusions

The most basic function of a hearing aid is to compensate for a loss in hearing sensitivity. Hearing aid designers have been aware of this goal for decades and also have been aware of the need to accommodate a reduced dynamic range. The technology has been available to meet these needs but always at the expense of size and cost. Hearing aids from each era have been designed to be as small as possible consistent with some set of performance compromises. These compromises, necessary to promote the use of the product, often have led to unsatisfactory performance. Now technology is making it possible to build hearing aids without these compromises in the smallest useful package. As a consequence we are now able to address the fundamental issue of making speech both audible and comfortable without facing market rejection for cosmetic reasons.

Address all correspondence to: William A. Cole BASc, PEng., President, Etymonic Design Inc., 41 Byron Ave., Dorchester, Ontario, N0L 1G0

References

Blesser, B.A. (1969). Audio dynamic range compression for minimum perceived distortion. *IEEE Transactions on Audio and Electroacoustics, AU-17* (1), 22-32.

Braida, L.D., Durlach, N.I, Lippmann, R.P., Hicks, B.L., Rabinowitz, W.M., & Reed, C.M. (1979). Hearing aids — a review of past research on linear amplification, amplitude compression, and frequency lowering, *ASHA Monographs Number 19*, Rockville. MD: American Speech-Language-Hearing Association.

Carlson, E.V. (1974). Smoothing the hearing aid frequency response. *Journal of the Audio Engineering Society, 22*, 426-429.

Carlson, E.V. (1989). The Knowles amplified receiver. *Knowles Electronics Report 10676-1.*

Cole, W.A. (1975a). Variable gain electronic amplifier. *Canadian Patent 963544*

Cole, W.A. (1975b). Hearing aid gain – a functional approach. *Hearing Instruments, Oct.*, 22-24.

Cole, W.A. (1985). Dahlberg programmable hearing aid project. *Etymonic Design Inc.* (Internal Report). Dorchester, ON: Etymonic Design Inc.

Cole, W.A. (1986). Noise reduction system. *US Patent 4,622,692.*

Cranmer, K.M. (1992). 1992 Hearing instruments dispenser survey results. *Hearing Instruments, 43* (6), 8-15.

Dillon, H. (1988). Compression in hearing aids. In R. Sandlin (Ed.), *Handbook of hearing aid amplification, volume I* (121-145). Boston: College-Hill Press.

Gebert, A.M. (1988). Hearing aid noise suppression system. *US Patent 4,750,207.*

Kates, J.M. (1989). Hearing aid signal-processing system. *US Patent 4,852,175.*

Kates, J.M. (1990). A test suite for hearing aid evaluation. *Journal of Rehabilitative Research and Development, 27*, 255-278.

Killion, M.C. (1976). Noise of ears and microphones. *Journal of the Acoustical Society of America, 59*, 424-433.

Killion, M.C. (1978). *Problems in the application of wideband hearing aid earphones.* Presented at the Conference on Acoustical Factors Affecting Hearing Aid Measurement and Performance, City University of New York, June 15.

Killion, M.C. (1979). *Design and evaluation of high-fidelity hearing aids.* Unpublished Ph.D. Thesis, Northwestern University, University Microfilms, Ann Arbor, MI.

Killion, M.C. (1988). An "acoustically invisible" hearing aid. *Hearing Instruments, 39* (10), 39-44.

Killion, M.C., Staab, W.J., & Preves, D.A. (1990). Classifying automatic signal processors. *Hearing Instruments, 41* (8), 24-26.

Mangold, S., & Leijon, A. (1979). Programmable hearing aid with multi-channel compression. *Scandinavian Audiology, 8* (2) 121-126.

Nabelek, I.V. (1973). On transient distortion in hearing aids with volume compression. *IEEE Transactions on Audio and Electroacoustics, AU21* (3), 279-285.

Nabelek, I.V. (1975). Amplitude compression in hearing aids. *Journal of the Audio Engineering Society, 23* (3), 213-216.

Newby, P.K. (1979). The importance of dynamic AGC action. *Unitron Industries* (Internal Report), June 28.

Poliakoff, A. (1950, August). Hearing aid design. *Wireless World*, pp. 274-276.

Pollòck M.C. (1977). Hearing aids: needs and developments. *Hearing Journal, 10*, 29-30.

Preves, D.A., & Newton, J.R. (1990). The headroom problem and hearing aid performance. *Hearing Journal, 42*(10), 19-26.

Preves, D.A. (1991). Output limiting and speech enhancement. In G. Studebaker, F. Bess, & L. Beck (Eds.), *The Vanderbilt hearing aid report II* (pp. 35-52). Parkton, MD: York Press.

Teder, H. (1991). Hearing instruments in noise and the syllabic speech-to-noise ratio. *Hearing Instruments 42* (2), 15-18.

Villchur, E. (1973). Signal processing to improve speech intelligibility in perceptive deafness. *Journal of the Acoustical Society of America, 53* (6), 1646-1657.

Wengel, A.M. (1940). Hearing Aid Device. *United States Patent 2,192,669.*