American Journal of
AUDIOLOGY
Vol. 2 No. 2 July 1993
American Speech-Language-Hearing Association
A Journal of Clinical Practice

The K-Amp Hearing Aid: An Attempt to Present High Fidelity for Persons With Impaired Hearing

Mead C. Killion
Etymotic Research, Elk Grove Village, IL

This paper is based on a presentation at the Mayo Audiology Video Conference, October 10, 1992.

Reprinted from the *American Journal of Audiology*, published by the American Speech-Language-Hearing Association, 10801 Rockville Pike, Rockville, Maryland 20852-3279

The K-Amp Hearing Aid: An Attempt to Present High Fidelity for Persons With Impaired Hearing

Mead C. Killion
Etymotic Research, Elk Grove Village, IL

The reason no one wants to be seen wearing hearing aids is probably that they don't work well, or at least they didn't. When the problems with such hearing aids are solved, a new/old problem arises: Background noises are often blamed on the hearing aid. The problem is that the user has lost ABONSO (automatic brain-operated noise suppressor option), and the problem persists until the user relearns how to recognize and localize background noises (at which time the brain automatically performs as a highly effective noise suppressor option). Ongoing attempts to replace the brain with a tiny circuit that will somehow reject noises we don't want to hear are unlikely to result in useful devices.

Why does no one want to be seen wearing a hearing aid? Why is a stigma associated with wearing a hearing aid? Edgar Villchur has speculated that it is because hearing aids didn't work. They left their wearer with an obvious handicap: inability to hear in noise. This is not the case with eyeglasses. If you see someone wearing eyeglasses, you presume that he or she can see perfectly well. You don't rush up with large print editions. But if you see someone wearing a hearing aid, your past experience tells you that you will have to raise your voice in order for you to communicate with that person if there is any noise present.

FIGURE 1. Signal-to-noise ratio for conversational speech in homes, hospitals, department stores, trains, and aircraft (after Pearsons, Bennett, & Fidell, 1976).

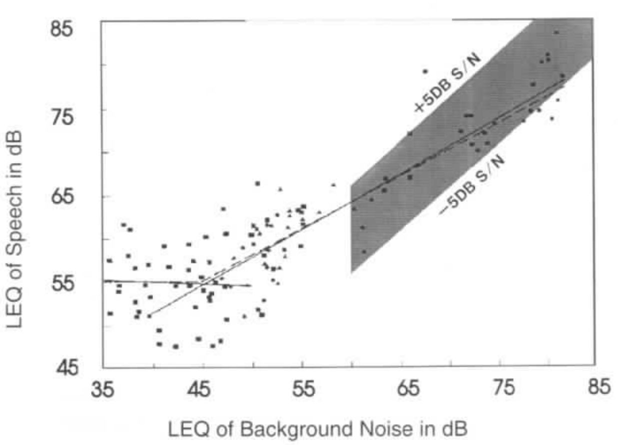

Why is this? It is increasingly clear that it is because hearing aids have made only a fraction of the speech cues audible. And why is that? There are several reasons, but let's start with some facts.

The Facts

People with normal hearing who are tested in noise at a -5 dB signal-to-noise ratio can get 50% correct on word lists. Catching 50% of the words, listeners can understand about 95% of conversational sentences, which means that at a party they can—with an occasional "what?"—carry on a conversation perfectly well.

We are all vocally lazy. When we speak in noise, we don't use enough effort to bring our voices well above the noise level. Pearsons, Bennett, and Fidell (1976) studied this phenomenon by putting a microphone and tape recorder on people and sending them out into homes, offices, hospitals, trains, cars, airplanes, and department stores.

They found that once the background noise level exceeds about 60 dB SPL, almost all normal conversations take place at a signal-to-noise ratio between -5 dB and +5 dB. A few conversations even dropped below -5 dB signal-to-noise ratio, as shown in Figure 1.

So in daily life, once the noise goes up, we're routinely presented with the problem of trying to understand someone whose vocal output is actually 5 dB below the sum of all the other people talking (or the other noise sources).

The surprising thing is that those with normal hearing can cope with the task. Once you have a hearing loss, however, you're no longer able to carry on as usual. A

1059-0889/93/0202-0052$01.00/0

series of studies by Tillman, Carhart, and Olsen, summarized in a 1970 paper, showed that whereas people with normal hearing can get by at a -5 dB signal-to-noise ratio, meaning they can get 50% of words and 95% of sentences correct, the people with typical presbyacusis required a +9 dB signal-to-noise ratio. This places their abilities outside the range of normal conversation in noise. They simply cannot participate unless others raise their voices well above the background noise.

Worse yet, Tillman et al. found that when they added the best hearing aid they could find back in the 1960's, it made things worse—so much worse that their subjects with impaired hearing required a +27 dB signal-to-noise ratio before they could get 50% of word lists correct. Figure 2 summarizes the situation.

Was the hearing aid really to blame? Yes. In fact, when Tillman et al. tested that same hearing aid on people with normal hearing, the normal-hearing subjects now required about a +7 dB signal-to-noise ratio for 50% correct. In other words, even subjects with normal hearing couldn't hear in noise with that hearing aid.

The Reasons

Why did hearing aids perform that way? Because they made only a fraction of the speech cues audible and clear.

Narrowband hearing aids were commonplace until recently. Figure 3 shows the estimated insertion response, calculated from the 2 cc coupler curve that appeared in their paper, for the hearing aid that Tillman, Carhart and Olsen used in their studies. This hearing aid passed only a narrow band of speech cues on to the listener.

Peaky response hearing aids were commonplace until recently, further limiting the usable bandwidth.

The other speakers today talked extensively about how hearing aids worked badly when they were distorting. Until recently, most hearing aids used so-called linear circuits, meaning nothing but peak clipping was used to reduce gain to prevent either discomfort or damaged eardrums. In many cases back then, especially if you went to a university to get a hearing aid recommendation, the saturation output of the hearing aid exceeded your discomfort level. The only choice in that case was to turn the gain down, reducing the audibility of quiet speech cues.

FIGURE 2. Effect of 1960s hearing aid on S/N required for normal conversation.

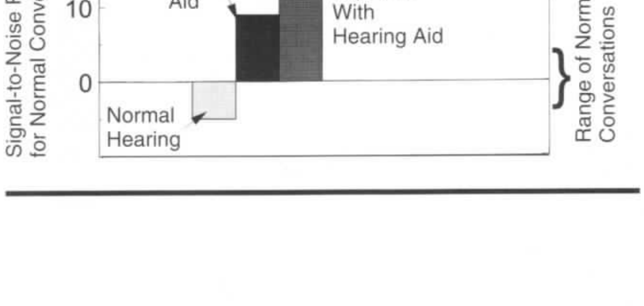

Why were hearing aids made that way? The apparent reason is that there was a lot of good market evidence that said users did not want them any other way. Manufacturers tried extending the high-frequency response of hearing aids and users rejected them as being harsh and unpleasant. Manufacturers tried extending the low-frequency response of hearing aids, and users rejected them as being muddy and unintelligible. Manufacturers tried automatic gain control circuits and compression circuits and again users rejected them, saying they preferred linear hearing aids. (You can still hear that occasionally on the lecture circuit.)

You will get those answers every time if you start out with a low-fidelity hearing aid. Those are the right answers for low-fidelity hearing aids. In fact, they make low-fidelity hearing aids tolerable. The problem is you simply cannot hear in noise with that type of hearing aid.

Amplifier Distortion

Today seems to be the day for linear bashing and amplifier-distortion bashing, and I am certainly delighted to join in. In fact, of the four hearing aid engineers in the room, you've heard from two of them about the severity of the problem, you're about to hear from me, and during the break Bill Johnson told me that back in the 1970s he tried wearing some hearing aids and went immediately to his supervisor and said "We should stop selling these things; they don't have any headroom!"

Amplifier distortion reduces intelligibility two ways. One is directly, because the harmonic distortion products and intermodulation distortion products mask the quieter speech cues. (Both types of distortion generate energy at frequencies not present in the input signal, that is, they effectively generate noise. Severe distortion simply fills up the auditory space with this noise and buries all the quieter, typically high frequency, speech cues.)

Amplifier distortion reduces intelligibility indirectly, more subtly and insidiously, because it produced all the wrong answers about what hearing aids should be like! It caused us to believe things that were absolutely false unless you were going to make low-fidelity, high-distortion hearing aids.

FIGURE 3. Estimated insertion response of 1960's hearing aid used by Tillman, Carhart, & Olsen (1970).

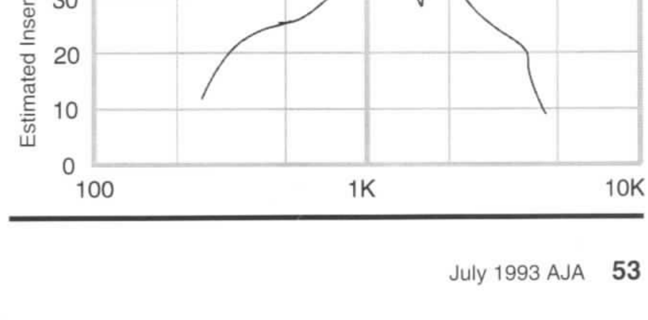

FIGURE 4. Maximum undistorted output of Knowles BP-series earphone with well-damped conventional earmold (after Killion, 1980).

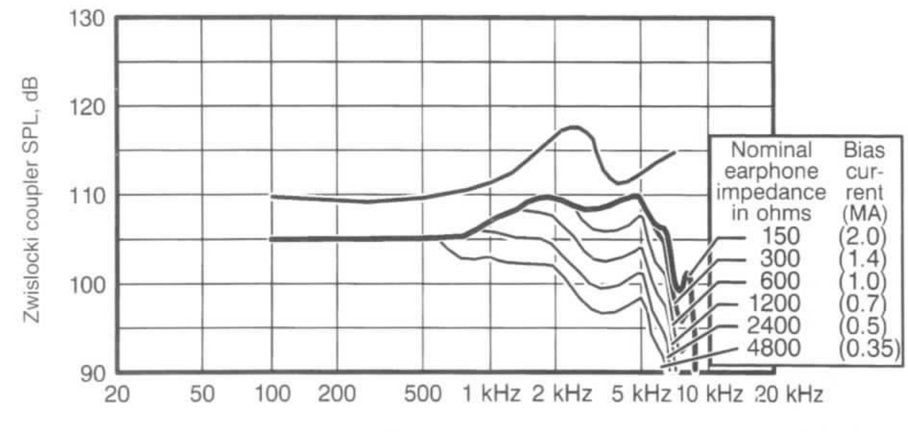

Limited by earphone overload (—), amplifier current clipping (▬), and amplifier voltage clipping (—)

Mas Harada was one engineer in the industry who was sort of the lone voice in the wilderness, crying out that the important problem with hearing aids was distortion. And he was strong enough in his company so that his hearing aids, even his in-the-ear hearing aids, used class B amplifiers even though it made them bigger. And the people in his marketing department used terms such as "large" and "more for your money" to describe them. But those Qualitone hearing aids *worked*. In most other companies the marketing department had more influence, listened intelligently to what people in the marketplace were saying, and the users were saying "I'd rather have it small." Unfortunately, the users rarely heard a comparison between small-and-distorted vs. slightly-bigger-and-clear.

So most companies went to class A amplifiers for their in-the-ear products because they were the only things that would fit in the smaller in-the-ear and in-the-canal hearing aids. Once you make that class A decision, you are faced with another problem, which is battery life. In a canal aid using a 10A cell, for example, you need to bias the Class A amplifier at only a few tenths milliampere in order to obtain acceptable battery life. (A week is the generally accepted minimum). In order to prolong the battery life, you are forced to starve the Class A output amplifier of bias current.

At first glance, you would think your maximum output SPL would drop too low to be practical with such a "starved Class A" output, but there is a trick. You can keep the low-frequency undistorted output up as high as you need by asking Knowles to wind a receiver coil with very

FIGURE 5. Setup for and results of Olson listening-test experiment. Plan and elevation views of the schematic arrangement of the apparatus for direct testing of frequency-range preference for speech and music. A sectional view, acoustical network, and response-frequency characteristic of the acoustical filter used in the test. Results for speech and music are on the right (after Olson, 1957).

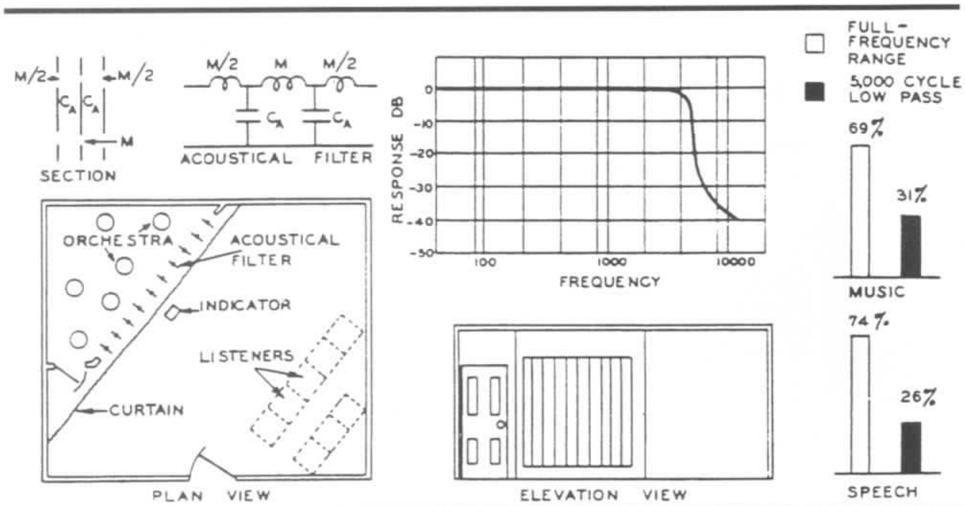

fine wire and a lot of turns, choosing one of the curves in Figure 4. You can even mostly maintain your SSPL 90 specifications with that tiny A battery drain. The hearing aid will sound plenty loud enough, because loud sounds will create overload distortion that makes things sound louder. The hearing aid may even sound louder because it distorts so badly. No one will complain that your hearing aid isn't loud enough, and you've achieved a low battery drain.

So starved Class A circuits became the rule. Unfortunately, these sacrifice nearly all of the undistorted high frequency output capabilities of the hearing aid. (Compare the lowest curve in Figure 4 to the higher ones). Thus, instead of having plenty of high-frequency output capability to handle the high-frequency emphasis in most hearing aids, typically you find that you have nothing but intermodulation distortion for high-intensity, high-frequency speech sounds and noise.

The Wrong Answers Explained

Why did we have all of these low fidelity answers and low fidelity aids? Let me first just play a recorded reminder of what hearing aids used to sound like. (The listeners heard pre-recorded examples of a low-fidelity hearing aid at a symphony concert and at a night club.) No one would want to wear such an aid personally, of course, but an embarrassing number of them have been sold. You can still buy hearing aids that sound like that, and save about $100 an aid to boot.

How is it that we got the answers that no one wanted any more highs? I believe distortion, subtly, fathered this wrong answer. It was a replay of the 1940s, when high-fidelity console radios and phonographs first became available. Published studies showed that if you actually did

FIGURE 6. Explanation of results of Olson's predecessors: amplifier distortion. Experimental results of subjective tests of reproduced speech and music depicting objectionable, tolerable, and perceptible nonlinear distortion for various high-frequency cutoffs. The numbers in the distortion data points correspond to the numbers that label the response frequency characteristics (after Olson, 1957).

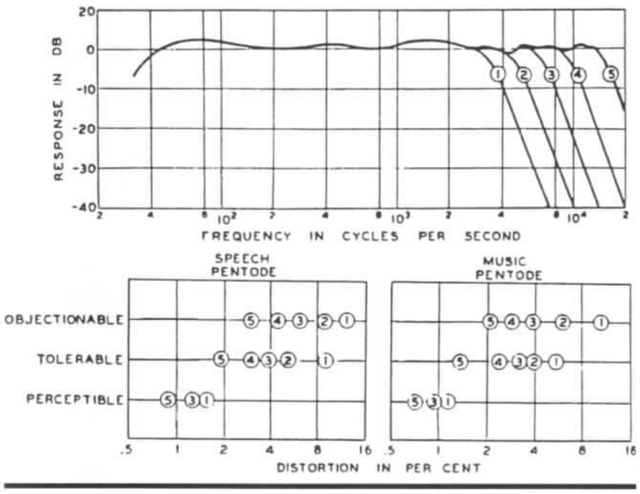

listening tests, the person on the street wanted you to roll off the high frequencies starting at about 5 kHz. The engineers might like to listen to a full 15 kHz response, but to sell those consoles to the public you had to roll them off at 5 kHz.

Harry Olson, chief engineer at RCA, became irritated enough with this answer that he set up a live listening test in which he had an orchestra behind a gauze screen so you couldn't see what was happening, and arranged organ louvers as shown in Figure 5 (Olson, 1957) to seal off the orchestra stage. With the louvers closed, the sound of the orchestra would have been blocked off, except Olson drilled holes through each louver to form an acoustic 5 kHz low-pass filter. (The mass of air in the properly sized holes with the compliance of air in the space between the louvers gave the acoustic filter whose equivalent analog and frequency response are shown in Figure 5). Olson set up live listening tests with both a live orchestra and live voice and with real listeners in the room, and he found two-thirds to three-fourths of the subjects preferred the full-frequency-response reproduction.

Olson went on with a series of experiments to demonstrate that the real problem was not bandwidth, but amplifier distortion. The results are shown in Figure 6.

The reason listeners preferred the 5 kHz roll-off in previous studies was that you needed to roll off the response above 5 kHz if you were going to allow 8%–10% amplifier distortion. With that kind of amplifier distortion, the sound was objectionable unless you rolled it off above 5 kHz. If you cleaned up the amplifier and got the distortion below 1.5%, on the other hand, then you could go the full 15 kHz. In fact, people actually preferred the full frequency response as long as distortion was low.

I think exactly the same thing happened with hearing aids. As long as clipping is the means of limiting the output, as has been true for most hearing aids purchased for several decades, it means that every time things got reasonably loud, you heard nothing but distortion products at high frequencies. Indeed, the way most hearing aids were designed, the input circuit itself would overload once you got above 85 to 90 dB SPL, so you got distortion even if you turned down the volume control.

Once you have a lot of distortion, the higher-frequency distortion products themselves become more intense than the high-frequency portions of speech and music (which are typically weaker than the low-frequency portions). If you extend the bandwidth, all you hear in that extended bandwidth is the noise of the higher-frequency distortion products. The lower-level speech and music signals in the extended bandwidth are not audible because they are covered up by the more intense distortion products. So you can extend the bandwidth of the hearing aid, but the user reasonably says, "That's no improvement for me. I'd rather listen to it the other way."

Figure 7 shows the insertion response of two hearing aids. The wideband hearing aid actually has a better amplifier, a cleaner amplifier, than the narrowband hearing aid. All those who hear the comparison, however, agree that they would rather listen to the narrowband hearing aid. The listener would much rather have those distortion

FIGURE 7. Frequency response of two hearing aids. The narrowband response survives clipping distortion much better (from Killion, 1984).

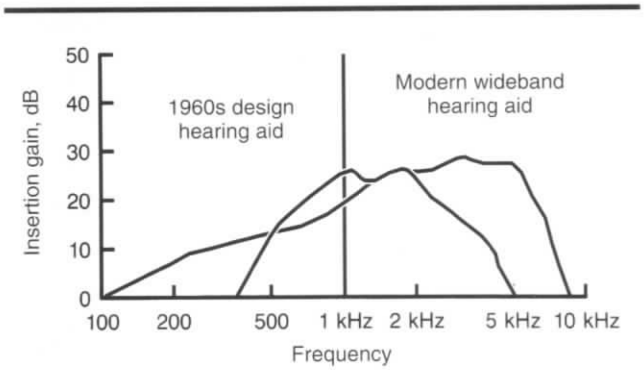

products filtered out. (Recorded comparison at a symphony concert.) Neither hearing aid sounds great in overload, but the narrowband aid is easier to listen to. Obviously, under these circumstances the right thing to do is to roll off the highs! There are no useful high-frequency speech cues in a wideband hearing aid when distortion covers them up.

What about the belief that we should always roll off the lows? I think the same explanation holds as for rolling off the highs: distortion. Back in the '40s it was well known that if you were going to use clipping as a means of limiting the output, you needed to use high-frequency emphasis before you clipped if you want to retain intelligibility. Licklider and Pollack (1948) showed in their experiments that you could use infinite peak clipping (nothing but square waves coming out) as long as you first rolled off all the low-frequency sounds. Even the Harvard Report (Davis, Hudgins, Peterson, & Ross, 1946) said, "[With low-frequency emphasis] the scores fell rapidly toward zero when the input to the limiting stage was increased above the clipping threshold. Intelligibility was best maintained with 6 dB/octave upward tilt."

In starved Class A hearing aids, clipping distortion was and is the rule. The way they are designed, they are distorted whenever it gets loud. Thus you either roll off the lows or destroy intelligibility.

So we had very good evidence, accurate evidence, that

FIGURE 8. Villchur's electronic simulation of a listener with severely impaired hearing subjected to a typical 1960/1970s ITE hearing aid.

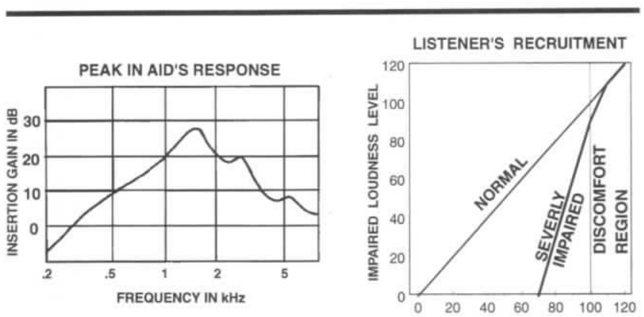

said you had to roll off the highs and you had to roll off the lows in a practical hearing aid. And it was largely because we had these starved class A circuits.

(A technical aside: It is important to note that there is nothing wrong with Class A amplifiers per se. If you use enough bias current in a class A amplifier, you can make it as clean as you want, just as clean as a Class B or Class D amplifier. With the right receiver coil, you can also have plenty of undistorted high-frequency output. It is not Class A that is the problem, but the fact that a good Class A amplifier takes so much current that the battery dies quickly. Thus most designers are forced to starve the Class A amplifiers to obtain reasonable battery life, resulting in the exaggerated high-frequency distortion problems we've been discussing. A well-designed Class B amplifier takes only about twice the battery drain of a Class D amplifier, so a much more reasonable coil compromise can be made. The main problem with Class B amplifiers historically was that they were usually too big to fit into the smaller aids)

There is one other problem that is also a little subtle, and that is the loss of audible speech cues caused by a peak in the frequency response. A large peak will cause the hearing aid to fail what I call the DHTTDGTLWI (dit-ta-dug-attle-we) test, which stands for "Doesn't Have To Turn Down Gain To Live With It." I want to play a simulation of what it sounds like, a recording made using the electronic hearing-loss-simulation equipment described by Villchur.

The hearing loss simulated is of someone with a severe impairment. That individual's loudness growth function hasn't reached normal until it rises into the discomfort region, as seen in Figure 8.

The loss may be severe, but the frequency response shown in Figure 8 is typical of almost every ITE hearing aid produced for about two decades. They used the BK series receivers that I helped design back in 1964. We didn't know how bad the real-ear high-frequency roll-off was back then, because we didn't have probe microphone measures. We knew you had to roll off the lows, however, so we rolled off the lows. We rarely used damping back then, so the standard ITE insertion response showed a large peak at the 1700 Hz resonance of the BK-series receiver, with a roll-off above and below that frequency. You can measure hearing aid after hearing aid of this type, and this is roughly the real-ear frequency response you'll see, with perhaps some variation in low-frequency slope.

Here's what it sounds like to have a severe impairment and have to listen through that frequency response. [Demonstration plays.] What you heard was that every time a vowel format frequency coincides with that 1700 Hz peak frequency, you get whacked on the ear. It hurt, and your automatic reaction was almost certainly to turn the gain down. And that's what users did. They turned down the gain, and then all that remained audible was the small region of speech information right around 1700 Hz. So once again you had a hearing aid that you could not hear with in noise.

Fortunately, we now have solid evidence as to just what happens to someone with a hearing loss using the sort of narrowband hearing aid we have been describing. Skinner and Miller (1983) did a series of experiments in which they looked at the effect of bandwidth on word recognition scores in quiet and noise. Their listeners had moderate to severe loss. (There was plenty of evidence for listeners with normal hearing that you need wide bandwidth, but for many years we thought that perhaps it would be more merciful to the people with moderate to severe hearing loss to limit the bandwidth and not cause confusion. I can remember Hugh Knowles making basically that argument back in the 1960s.) What Skinner and Miller found was that all of their listeners did best with the widest audible bandwidth. (Villchur [1973] reported an example of a subject who had no useful hearing above 3 kHz, and who did better with the high-frequency response of her experimental hearing aid rolled off. During audiometric testing, if asked, such subjects may report that high-frequency tones sound like a buzz or hum, but we often forget to ask. Obviously these subjects need to be excluded from the generalizations above. Fortunately, they are rare.)

Figure 9 illustrates Skinner and Miller's results for listener #6, who could get 48% correct in noise with the proper wideband frequency response. Getting 48% of words typically means 95% of sentences; listener #6 could carry on a more or less normal conversation in that noise with the full bandwidth response. When tested with a bandwidth comparable to the bandwidth of hearing aids such as the one in Figure 8, listener #6 got only 4% correct in noise.

Skinner and Miller's listener #1 shown in Figure 10 did somewhat better, dropping from 66% correct with a full bandwidth response to 26% correct with a narrow bandwidth. Even listener #1 would be unable to carry on a conversation in noise with the typical hearing aid of the last two decades. Taking away some of the redundant speech cues leaves these listeners with insufficient cues to understand speech in noise.

With all of this hindsight, it seems obvious why hearing aids didn't work in noise. That, in turn, is almost certainly the reason your family doctor—whose mother had a hearing aid that didn't work in noise—told everybody with "nerve deafness" not to get a hearing aid because it wouldn't help in noise. He was right; it didn't.

But you can get totally different answers to all of these questions if you start out with high-fidelity hearing aid designs and use appropriate automatic adjustments to the gain and frequency response in order to keep all speech cues audible.

High Fidelity Is Possible in Hearing Aids

For years, no one thought high fidelity was possible. Ironically enough, high fidelity has been possible in hearing aids for about 15 years. In the late 1960s, broadcast and

FIGURE 9. Effect of bandwidth on the ability of Listener 6 to understand speech in noise. Word recognition scores given in percent for speech in quiet (left numbers) and noise (right numbers) (after Skinner & Miller, 1983).

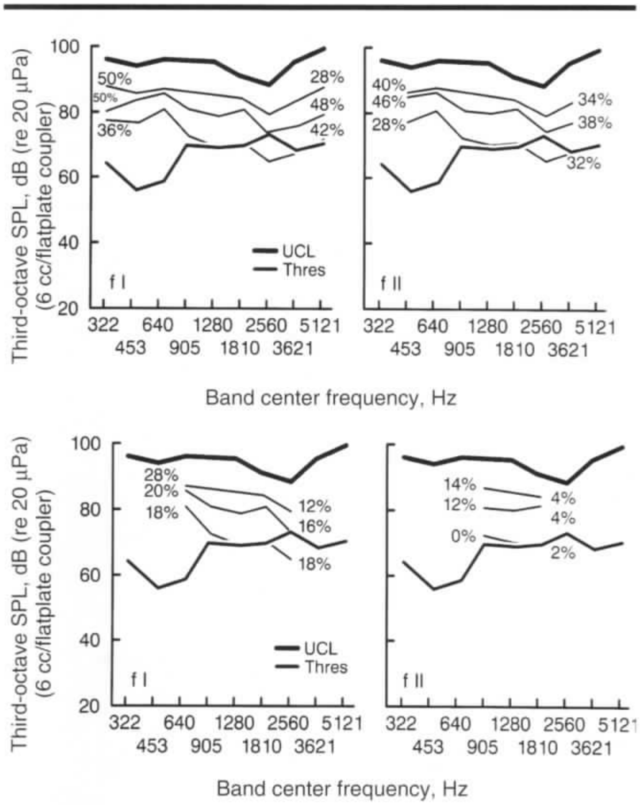

FIGURE 10. Effect of bandwidth on the ability of Listener 1 to understand speech in noise. Labels same as Figure 9 (after Skinner & Miller, 1983).

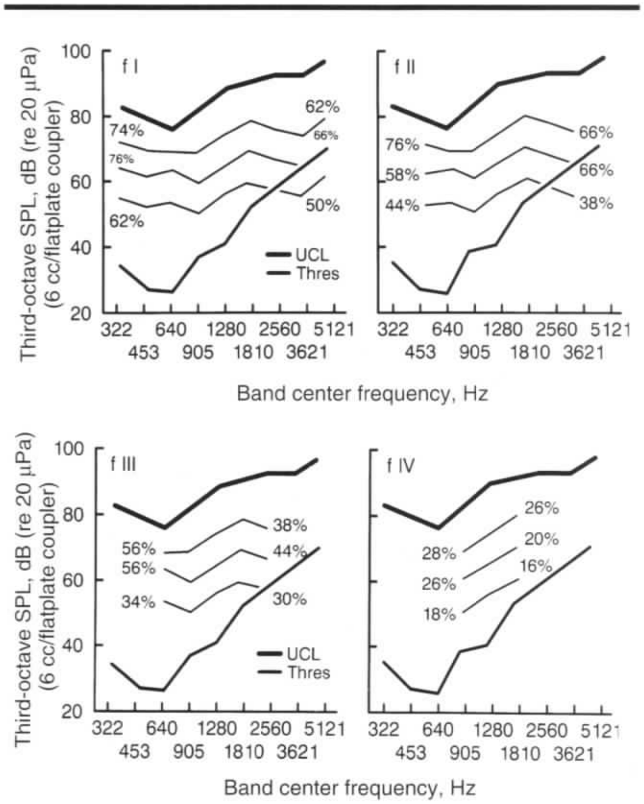

FIGURE 11. Experimental high-fidelity hearing aid with 16 kHz bandwidth (from Killion, 1979).

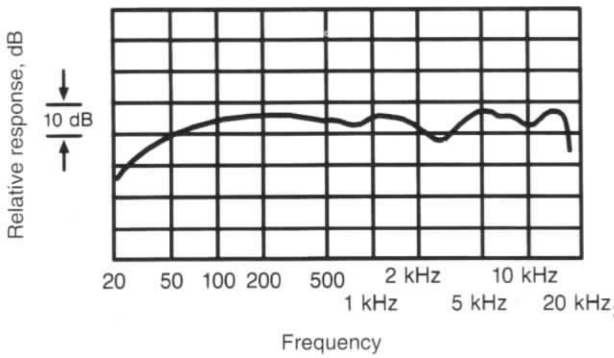

recording studios were already using a flat-frequency-response version of one of the Knowles hearing aid microphones. Hearing aid receivers capable of a 16 kHz bandwidth have been available for almost as long: the Knowles BP-series receivers were introduced in the 1970s. Etymotic Research has just introduced to the hi-fi market a pair of stereo earphones we call the ER-4, ear-phones that I believe are the highest fidelity earphones available, with the most accurate real-ear response around. These earphones use the same basic Knowles ED-series receiver mechanism used in the K-Amp hearing aid designs.

My PhD thesis back in 1979 was basically a demonstration that you could make a high-fidelity hearing aid, at least as judged by someone with normal hearing (Killion, 1979). Figure 11 shows the 16 kHz frequency response of an experimental in-the-ear aid used in that study. I prepared KEMAR recordings of loudspeakers, ear phones, and experimental hearing aids. These were judged by one of three groups of listeners: a "man on the street" panel of 12 men and 12 women (average people, no engineers), a panel of 5 "golden ears" (people like Julian Hirsch, who does stereo equipment rating for *Stereo Review*), and 6 trained listeners (people who have spent hours and hours listening

FIGURE 12. Fidelity ratings from three listening panels: "Man-on-the-street" (12 male, 12 female), "golden ears," and trained listeners (from Killion, 1979).

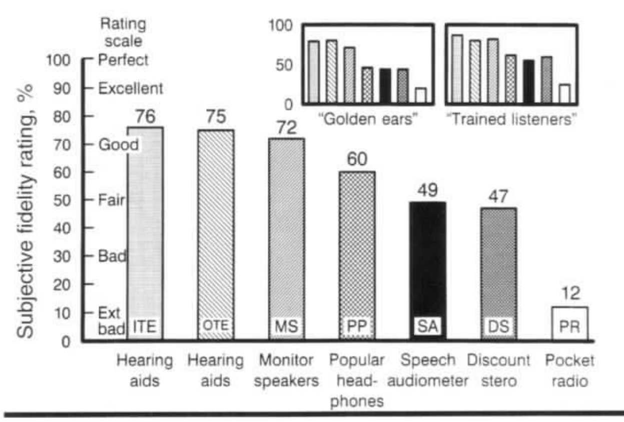

FIGURE 13. Undistorted SPL output requirements for high fidelity hearing aid.

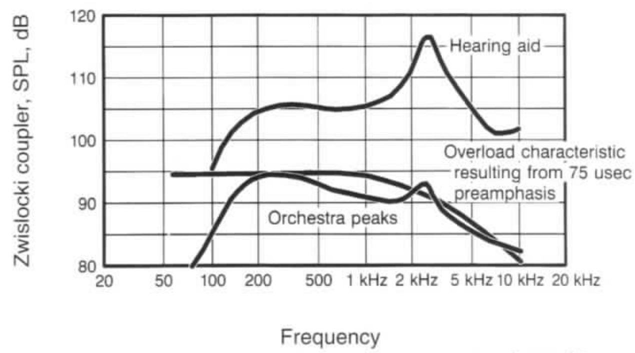

Orchestra peaks correspond to instantaneous peaks of 105 dB SPL or peak readings of 95 dB SPL on sound level meter

to some type of psychoacoustic experiments).

All three panels gave basically the same answers, shown in Figure 12. Indeed, so did an unreported panel consisting of nine of my relatives who were forced to "take the listening test" as they visited throughout one summer. So the answer was pretty solid: Any reasonable group of people with normal hearing would judge the experimental high-fidelity hearing aids to be slightly better (but not statistically significantly so) than the expensive pair of high-fidelity stereo loudspeakers that had been chosen because they were the most popular recording-studio monitor speaker used in Chicago at the time.

There was one practical problem in rushing into production with this hearing aid design. I had used 5 mA of battery drain in my class A output stage in order to provide the peak undistorted output capabilities needed to handle a live symphony orchestra concert. What is needed from a hearing aid if it is to handle an orchestra concert is shown in Figure 13. I can verify from my own measurements that the Chicago Symphony can produce 104 to 106 dB SPL readings on a sound-level meter located at a seventh-row, first-balcony seat in Orchestra Hall. This means that the instantaneous peaks run something like 115-118 dB. Now if you want the hearing aid to never, never distort, the hearing aid has to be capable of those outputs, which

FIGURE 14. Enlarged view of Knowles Class D amplifier chip.

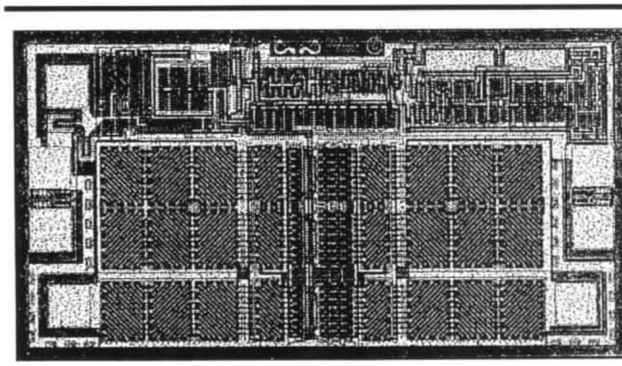

means you can't even *think* of using a starved class A output amplifier. You have to feed a class A amplifier very well if you want it to handle those outputs. "Very well" in this case means 5 mA of battery drain. But this means that a 10A cell in a canal aid would last about 10 hours: Install a fresh battery in the morning and it would need to be replaced at dinner time. You will not sell many hearing aids that eat batteries like that.

The breakthrough solution to that problem was new amplifier technology. Elmer Carlson got me interested in Class D amplifiers back in 1964, and I built a fist-sized 7-watt amplifier that year. That one would obviously not fit in a hearing aid, and the hearing aid version became a project that was pursued and canceled several times. Finally, after I left Knowles, Hugh Knowles gave me one more chance to see if we could make a practical low-battery-drain, high-fidelity Class D hearing aid amplifier. With a lot of help from our friends, we finally succeeded. The resulting CMOS integrated-circuit chip is now in all of the Knowles class D receivers. Figure 14 is a photomicrograph of that chip (whose actual size would drop nicely in the "u" in the word actual.) Please notice the mustache at the top of the chip.

This chip cuts the Gordian knot. It is now possible to have a hearing aid output amplifier that will idle at .17 mA bias, and yet provide the undistorted output of a class A amplifier biased at several mA's. Oddly enough, the bonus is that if you use the Class D receiver with a low-current preamplifier such as the K-Amp circuit, the entire hearing aid only takes about .35 mA. Thus, with less battery drain than the starved Class A output amplifiers consumed by themselves, you can have 3–4 weeks of 13A battery life in a complete hearing aid that never overloads—even if you go to a Chicago Symphony concert. You can "sit in" on piano with a Dixieland band—with the drummer only a few feet away—and the hearing aid simply does not overload. (We assume here that the volume control has been set to produce unity gain for loud sounds.)

The performance advantage of the Class D receiver was obvious. The big question was cost. Adding the Class D amplifier roughly doubled the cost of the receiver, which meant adding about $50 to the wholesale price of the hearing aid. The obvious question everyone asked was: "Is anyone willing to pay for that?" One way to answer that question was to do a research study. Catherine Palmer (Palmer, Killion, Wilber, & Ballad, 1993) did just such a study under Knowles sponsorship. She compared the sound of two hearing aids, each set to 15 dB gain, one using a starved class A amplifier (biased at 0.4 mA), and the other using a class D receiver. She made KEMAR-based recordings using three different source materials (orchestra, piano trio, and male voice), and using a range of input SPL's. She played them back in a randomized form to the listener-judges, giving them roughly the following instructions: Assume you need a hearing aid; assume you are ready to buy a hearing aid and have enough money to buy a hearing aid, and that hearing aids cost at most $700. What would you pay for a hearing aid that sounded like ——. The listeners circled, on a ruler going from $0 to $700, what they would pay. She also used a quality scale, asking the listeners to rate, on a 0 to 100% scale, the overall quality of the sound they had heard.

Figure 15 shows Palmer's results. The laboratory answer to the question of whether or not anybody was willing to pay substantially more for the Class D hearing aid was clearly yes. Listeners indicated they would pay $100 to $200 more over a wide range of input SPL's. Just as interesting, and as far as I know for the first time in the history of the world, Palmer was able to quantify what quality is worth. She found that if she multiplied each quality point by $6.75, she obtained transformed quality-judgment curves nearly identical to the dollar curves. So we now know what quality is worth; it's worth $6.75 a percentage point. That is a pretty high and exciting number. It justifies some improvements in hearing aids that we used to think were not justifiable.

High Fidelity for People With Hearing Loss

Where are we? We have wide bandwidth, we have a smooth frequency response, we have experimental hearing aids that people with normal hearing judge as high fidelity, we have a low battery drain, and we have no audible distortion. We're up to the crucial question that Hugh Knowles asked me some 15 years ago: What does high fidelity mean for someone with a hearing impairment? We

FIGURE 15. Quality judgments for hearing aids using Class D vs. "starved Class A" power amplifiers.

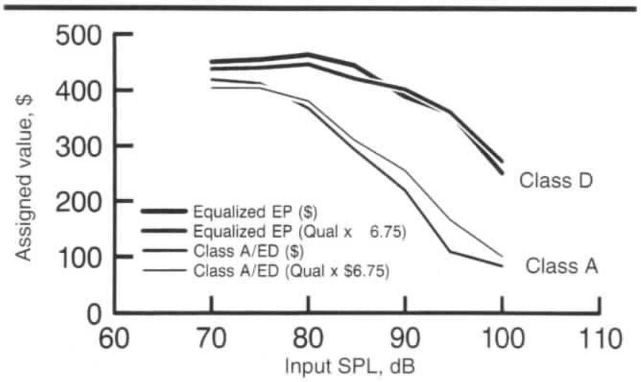

FIGURE 16. Loudness function for 1000 Hz tone (after Steinberg & Gardner, 1937).

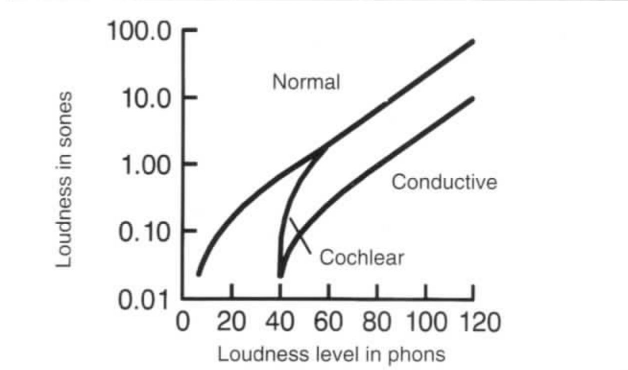

still haven't solved the real problem: providing a high-fidelity hearing aid for someone with a hearing impairment. What do you do now?

Recruitment

There is some reason to think we can answer that question. If we consider the loudness growth curves shown in Figure 16, we see that someone with a mild-to-moderate cochlear hearing loss shows normal loudness perception for high-intensity sounds. These cochlear losses are usually due to a loss of outer hair cell function. This is nice because at higher levels we see almost complete recruitment; at higher levels the outer hair cells probably don't affect the motion of the basilar membrane anyway. Complete recruitment means normal loudness perception, and typically covers a large region of moderate- and high-level input sounds. Not only is normal loudness obtained, but in many cases normal psychoacoustic tuning curves, difference limens for frequency, etc. You can do all sorts of tests, and as long as you have only a mild loss, perhaps a mild-moderate loss, you can get normal function over quite a region of high-level sounds.

I personally believe that one bad definition in psychoacoustics has set back hearing aid design enormously. We learned in school that recruitment was an abnormal growth of loudness. This may be true in a very narrow sense, but it provides exactly the wrong perspective on the problem to be solved.

Recruitment is not an abnormal growth of anything, but an abnormal loss of sensitivity for quiet sounds. (Indeed, we now have solid physiological evidence that such is exactly the case: The loss of outer hair cells in the cochlea produces a loss of sensitivity for quiet sounds). If you sit mired in the region of sensitivity loss and look at the problem, the obvious solution is to provide gain. And then having done that you go home. But you've left the hearing aid wearer with that gain for all input levels, because the only thing you see to solve is the loss of sensitivity.

If you view the problem from the high-level region where recruitment is complete, you see that you don't want any gain—you don't even want a hearing aid—for much of this range of input. This is a region in which you have normal hearing, and you see that nothing is growing abnormally, as long as you don't wear a hearing aid that gives you gain you don't need.

Viewed from the input sound levels where the individual mostly lives, the problem is an abnormal loss of sensitivity for quiet sound. It's the equiv-alent of night blindness: You don't have any trouble when it's bright, you've just lost the ability to see in the dark.

If you look at the problem from the standpoint of high-level inputs, for those inputs you want the hearing aid to do nothing. The hearing aid should not stand in the way of your normal hearing, and it certainly shouldn't provide amplification for things that are already loud. Down in the low-level-input region, on the other hand, the person needs some gain or otherwise he or she wouldn't have come into your office. This person came because he or she can't hear quiet sounds.

K-Amp Processing

The presumed ideal answer to the overall problem is the hearing aid characteristic shown in Figure 17. (For identical reasons, Barfod [1978] recommended an essentially similar characteristic.) For high-level sounds it gives unity gain. It's as close as we can come to having the hearing aid sit on your shoulders, out of your way for loud sounds, and then when a quiet sound comes along having it jump into your ears to give you the gain you need. For high-level inputs (80 dB HL and above), the hearing aid should be as transparent as we can make it. For very-low-level inputs (0-20 dB HL), it should produce the gain you need for quiet sounds, to solve the problem that brought you in to get help.

The only remaining question is, how do you get from the high-level region to the very-low-level region? You certainly don't want to suddenly switch gain at some level.

FIGURE 17. Presumed ideal input-output function for a hearing aid intended for an individual who has relatively normal high-level hearing.

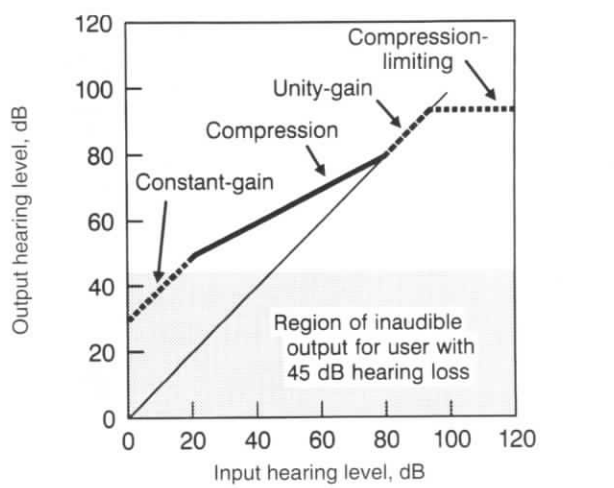

FIGURE 18. Input-output curves illustrating two approaches to providing a 25 dB threshold improvement.

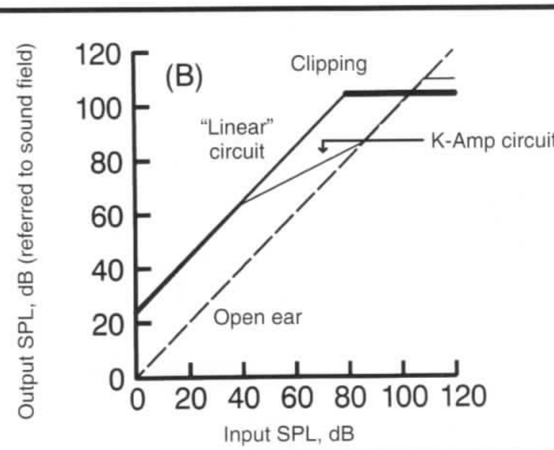

(That was one of the problems with the Intellitec circuit. If you've ever listened to it, it would suddenly switch 15 dB of gain in and out in a given channel.) No, for the least obtrusive operation, the highest subjective fidelity, you want to change gain as gradually as you can, which means wide dynamic range compression (sometimes called logarithmic compression). This makes it possible to gradually change from 20–30 dB of gain for quiet sounds to 0 dB of gain for loud sounds, spreading the gain change out over a wide region (over 50 dB in the case of the K-Amp circuit).

Figure 18 shows a comparison between the input-output characteristics of the K-Amp circuit and a linear circuit. The linear circuit operates perfectly well for very quiet sounds, but it provides progressively too much gain for louder sounds, until mercifully it starts clipping at about 80 dB SPL input. But clipping means it's distorting. So a linear circuit works fine in quiet, but the rest of the time it's pretty bad: too much gain for moderately loud sounds and you can't hear anything in noise above 80 dB.

If the Dave Hawkins "80% of hearing aids sold are linear" statistics are correct, that seems to me to mean that four of five people listening to this broadcast should turn in their hearing aid dispensing licenses. But then I've been thinking about that statement, and probably the people who come out on Saturday to listen through an all-day hearing aid fitting broadcast are not the problem; it's the four of five people who didn't come out to listen who should turn in their licenses.

Everyone knows that several research studies confirm that linear is just as good as compression ... if the user can adjust the volume and tone controls fast enough! In a given situation, it is nearly always possible to adjust the volume and tone control so you score as well with a linear aid as with a compression aid on a word list, especially if the word list is a standard constant-presentation-level list (which Villchur calls a "precompressed" word list).

In the real world, however, you would have to devote your entire energy to that task in order to keep the volume and tone controls at optimum with normally varying input conditions. No one does that, so you have a choice if you're dispensing linear hearing aids: You can either use clipping to keep the thing from hurting someone's eardrums or becoming uncomfortably loud, or you can use so much power that the hearing aid never clips. The latter hearing aid is the type that Darrell Rose loves to talk about, one that will reduce your IQ 10 points in 3 minutes. Those are your choices with linear aids. You're either going to clip and destroy intelligibility or you're going to have so much power that you're going to be in trouble with discomfort.

FIGURE 19. Gain-vs.-input curves illustrating two approaches to providing a 25 dB threshold improvement.

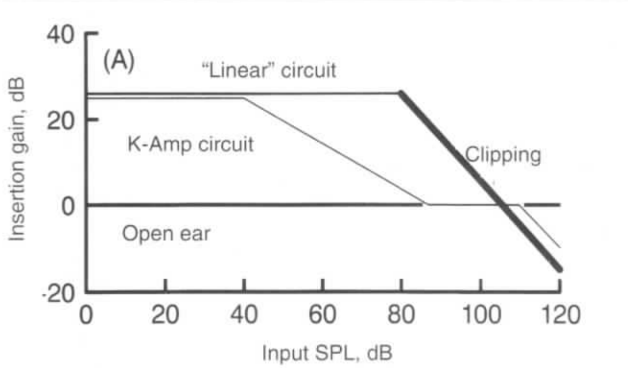

FIGURE 20. 2 cc coupler curves of "FFR" (Fixed Frequency Response) version of K-Amp hearing aid. Volume control adjusted to produce 0 dB REIR for loud sounds.

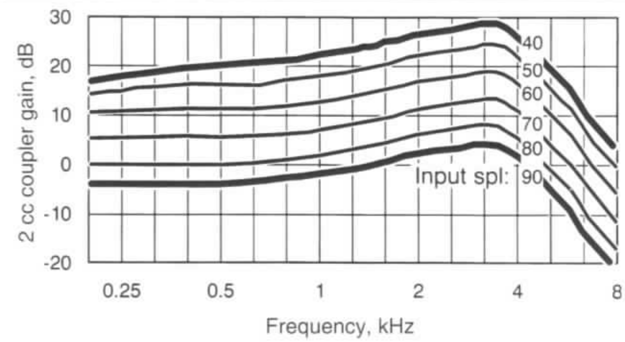

Figure 19 shows the comparison of linear and K-Amp circuits in terms of gain. As someone said, "I sell gain, not output." Both have the required 25 dB of gain for quiet sounds. The difference between the linear and K-Amp circuits is that the linear circuit only has the right amount of gain for quiet sounds, too much gain for louder sounds, and then it mercifully goes into clipping for loudest sounds. The K-Amp circuit has the same gain for quiet sounds but gradually turns down the gain to unity or zero dB gain for loud sounds, and doesn't clip until inputs of 110 or 115 dB SPL. Unless you work near jet aircraft, it never clips.

Figure 20 shows the 2 cc coupler curves at various input levels for the K-Amp circuit with the FFR (Fixed Frequency Response) option. Here you have about 25 dB of gain for 40 dB SPL and lower. As you increase the input, the gain keeps dropping (although the frequency response doesn't change) until for 90 dB SPL input you have a 2 cc coupler curve that corresponds to about 0 dB real-ear insertion response (about -3.5 dB gain at low frequencies on a 2 cc coupler). This FFR characteristic works very well for the flat losses shown in Figure 21 and, surprisingly enough, it also works very well for reverse slope losses such as indicated in Figure 21. Several patients with reverse slope losses have said that their K-Amp hearing aids were by far the best thing they've ever had. (It's possible to wire the K-Amp circuit for a "BILL" response to handle this a little better, but we haven't gotten around to trying it because everyone with this loss has been so happy with the FFR response.)

High-Frequency Loss

The FFR characteristic left us without a solution for the more common problem, which is someone with a high-frequency hearing loss. The loss shown in Figure 22 is frequency dependent. At low frequencies we see very little

FIGURE 21. Low-risk fittings for K-Amp hearing aids: flat and reverse-slope losses have a high probability of a successful fitting without special effort.

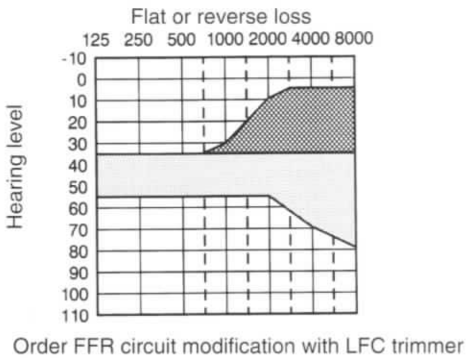

loss, little need for amplification, but we see a substantial loss at high frequencies.

If you believe Barfod's (1978) data, the individual with the audiogram of Figure 22 has normal or near-normal hearing at all frequencies once you make things loud enough. This means that a flat frequency response for very loud sounds makes sense, but some high-frequency emphasis is required for quiet sounds.

Fortunately, Margo Skinner (1980) studied exactly that question. She had subjects with normal or near-normal low-frequency hearing but with substantial high-frequency loss, and she tested these subjects with the five different frequency responses shown in Figure 23, using a wide range of sound pressure levels. She did an enormous amount of testing (she complained that she only had about 25 hours of word responses from each subject!) but obtained definitive answers. Figure 24 shows data for one subject and two of Skinner's frequency responses. (The rest of the curves pretty much fill in between.) Skinner found, not surprisingly, that for quiet sounds you got a much higher word recognition score if you gave a lot of high-frequency emphasis. But, by the time you got up to 50 dB input with that amount of high-frequency gain, the subjects said ouch, they could not take it any more, it was just too loud.

FIGURE 22. High-frequency loss with region of normal loudness calculated from Barfod (1978) data.

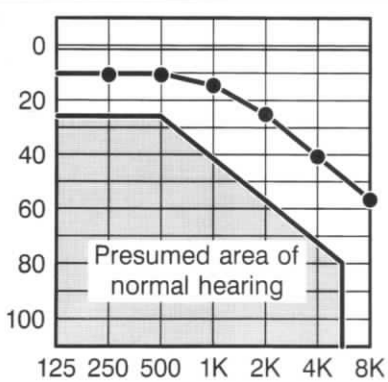

FIGURE 23. Frequency responses used in Skinner (1980) study.

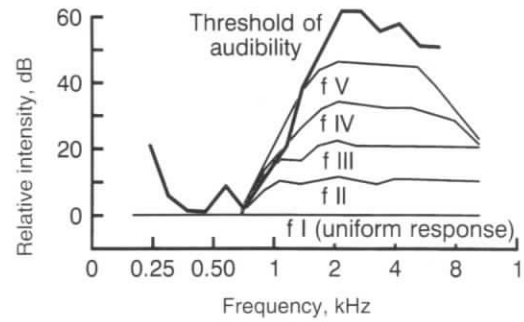

On the other hand, for very loud sounds, Skinner found that a flat frequency response gave about as good word recognition scores. And a flat frequency response for loud sounds makes some sense, because all the speech cues still fit in the auditory area. A flat response for loud sounds generally sounds much better, also, especially on music. (Some subjects obtained their best scores with some high-frequency boost. See Skinner, 1980, for more information.)

K-Amp Level-Dependent Frequency Response

So quiet sounds need high-frequency emphasis, and loud sounds need a flat frequency response. What is obviously needed is a level-dependent frequency response modification to that FFR curve that we had, one that will provide more gain and treble boost for quiet sounds than for loud sounds, and will in fact give a flat frequency response and no gain at all for loud sounds. Such a set of curves is shown in Figure 25, illustrating the level-dependent gain characteristic of the K-Amp circuit.

Figure 26 shows the real-ear output curves for a complete hearing aid using that K-Amp circuit. The volume control has been set to about 1/3, and an input of 90 dB SPL produces a 90 dB SPL output (give or take a little because the real-ear insertion response isn't perfectly flat). And 100 dB in gives 100 dB out; 110 dB in gives 110 dB out. (If speech and music had a flat spectrum, a 110 dB input would cause overload distortion at high frequencies. As it turns out, the spectrum of speech and music rolls off at high frequencies, so an overall input of 110 dB doesn't cause audible distortion. When the Chicago Symphony hits 105 dB, a lot of the energy is due to the kettle drum; the frequency band above 1500 Hz band typically has less than 95 dB SPL in it.)

In Figure 26, note that when you get down to quiet sounds, 40 dB in does not produce 40 dB out. Instead, 40 dB in receives gain and a treble boost, coming out something like 70 dB at high frequencies. So you have the high-frequency gain you need for quiet sounds with a flat, zero-gain operation for loud sounds.

Figure 27 shows the integrated circuit chip that goes into the K-Amp hearing aid. You'll note that it also has a mustache on it, along with the initials of the nine other technical people who made this chip possible. (Posters of this chip are available incidentally, and if you write or