indigo™

AutoPro4™:

IMPROVED DIGITAL TECHNOLOGY DELIVERS SUPERIOR SOUND QUALITY

Executive Summary

Improved speech intelligibility and sound quality are the primary objectives of new hearing aid technology. Recent research reveals that the latest digital hearing instruments are having a positive impact in both of these areas. Notice in Figure 1 that six of the top seven performance factors reported by Kochkin (2005) relate to sound quality or clarity of speech. Developments in digital technology and hearing aid architecture have improved sound quality over the years. Newer digital processing platforms offer higher clock speeds, greater processing power and more programmable memory. In addition, sophisticated adaptive features, such as autoPro4™, provide intelligent switching between four independent sound destinations including quiet, group/party noise, traffic/intense noise and music. These adaptive features, as well as improvements to digital signal processing (DSP) platforms offer improved audibility and sound quality. But, what components of these new digital aids are now providing the benefits that have been promised for so long since the onset of the digital revolution in the industry?

unitron hearing™

Figure 1

Customer Satisfaction with Signal Processing

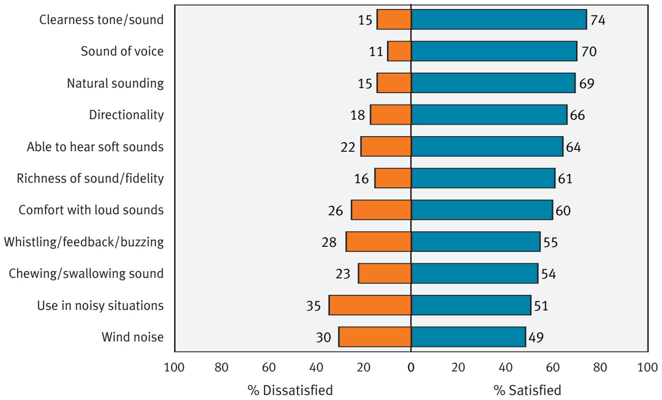

Results from MarkeTrak VII demonstrating customer satisfaction with signal processing for speech intelligibility and sound quality.

There is abundant literature to support the assertion that multichannel Wide Dynamic Range Compression (WDRC) improves speech intelligibility in quiet by providing greater audibility for soft sounds (Moore, Laurence et al. 1985; Benson, Clark et al. 1992; Jenstad, Seewald et al. 1999). Recent improvements in directional microphone technology can enhance speech perception in noise. Directional microphones improve the signal-to-noise ratio (SNR) for speech in the presence of diffuse or spatially separated backgrounds (Surr et al., 2002; Ricketts & Hornsby, 2003; Walden et. al., 2004). Furthermore, available literature supports the contention that adaptive noise reduction algorithms can improve comfort in background noise (Boymanns & Dreschler, 2000; Walden et.al., 2000). What about sound quality?

Improved speech intelligibility is described by listeners as providing clarity, but not necessarily quality. In fact, improvements to intelligibility, even in quiet, can often degrade sound quality by providing a sound that is described as sharp, tinny or just plain loud. Sound quality and comfort are more closely related to one another, but there is more to quality than mere comfort. Good sound quality is considered by most hearing aid wearers to be synonymous with full, rich, natural and clear sound. They want to be able to hear speech clearly. They also want the perception of the same quality of hearing while wearing hearing aids as they were used to when their hearing was normal.

Providing adequate amplification through ever smaller devices to a damaged auditory system, while enabling the listener to enjoy a natural, rich sound is a daunting task. Listeners still often describe their hearing aids as sounding unnatural. The technological advances during the first few years of the digital revolution did improve comfort and clarity; yet robust sound quality initially proved elusive. However, the recently released MarkeTrak VII survey shows that new hearing aid wearers report improved sound quality. The more recently they purchased their last set of hearing aids, the higher their satisfaction rating.

When digital technology first hit the hearing aid industry, most of the performance enhancements were achieved through the introduction of new digital features. It would be tempting to conclude that similar advances in adaptive features could explain this improvement as well. For example, multichannel amplification improved frequency/gain response shaping. Multimemory hearing aids allowed different sets of parameters to be optimized for specific listening situations (quiet listening, speech in noise, comfort in noise, music) rather than a one-size-fits-all approach. Digitally controlled microphones, noise reduction and feedback suppression algorithms all played a part. Each new feature improved speech intelligibility or overall listening comfort. These innovations were never directly linked with improved sound quality and they have been commonplace for a few years.

Recent improvements in sound quality, by the newest generation of hearing aids, have not been achieved by introducing more clever signal processing strategies. They have been achieved by improvements to the underlying architecture of hearing aids. The central processor units (CPUs) run at faster clock speeds. They also have greater capacity to process more instructions during each clock tick. The hybrids, within which they are contained, pack much more random access memory (RAM) into a smaller space. This allows more complex sets of instructions to fit in the hearing aid. Just as computers have grown smaller, faster and smarter, so have hearing aids over the past two or three years. Developments in hearing aid architecture have improved sound quality and include:

General Increase in Capacity

The three main areas where the newest digital processing platforms have improved are:

1. Higher clock speeds – more computational cycles per second means more actions performed in the same amount of time.
2. Greater processing power – more instructions can be carried out during each computational cycle.
3. More available memory – more storage space is available to hold increasingly sophisticated instructions. This allows more intelligent decisions based on a larger number of acoustic parameters.

In each instance, the miniaturization of hardware that began with the home electronics industry has found its way into the hearing aid industry. The outcome in both cases can be described in terms of greater available capacity. However, capacity is not enough to improve sound quality. The improvements stem from how that increased capacity is put to use. Below are a few examples of how we have taken advantage of this capacity to improve the sound quality of all our hearing aids moving forward.

More Instructions per Second

Higher clock speeds and greater processing power both yield the same advantage: more instructions per second. When a digital hearing aid carries out more instructions every second, adaptive features can be updated much more frequently. Transitions to adaptive features usually sound to listeners as if they occur smoothly and continuously, but they actually happen at discreet intervals. For example, we assume an adaptive directional microphone will shift polar patterns seamlessly in a rapidly changing environment to keep placing a null at the azimuth from which the loudest noise source originates. However, the polar plot actually shifts in a series of discreet steps. Before each step, new information must be obtained about the current acoustic environment and the decision to update must be processed.

If information about the acoustic environment is obtained once each clock cycle, the rate at which the polar pattern can be updated is gated first by the aid's clock speed. A higher clock speed yields more updates per second. Since the polar pattern shifts in discreet intervals after each clock tick, the position of the null will generally lag slightly behind the loudest noise source in a rapidly changing environment, such as a busy city street. In other words, when the primary noise source is moving, there is an inherent inaccuracy in the placement of the null that suppresses it. The amount of inaccuracy is determined by how often data about the noise source is updated and acted upon. All things otherwise being equal, higher accuracy can be achieved with a faster clock speed because the processed information is updated more frequently, and therefore, the placement of the null is kept more current. But how does increasing processing power improve accuracy?

The ability to undertake more instructions for every clock tick does not inherently allow the polar pattern to update more frequently if each update requires a clock tick. However, a far more sophisticated polar pattern can be implemented without sacrificing speed and accuracy. For example, suppression of multiple noise sources using a multichannel directional system is now possible. If there is a noise source originating from one direction that is composed predominantly of low frequency energy, and a second noise source from a different direction that is predominantly high frequency in composition, an adaptive wideband system could only place a null on the louder of the two sources. However, a two channel beamformer split into high and low frequency channels could place nulls on both noise sources simultaneously. It may take at least twice the processing power (instructions per clock tick) to run both beamformers at the same update rate as the wideband system. Greater frequency specificity is achieved at the expense of double the processing. Our system provides sixteen channel resolution for all adaptive features including the directional microphones. Therefore, sixteen separate polar patterns must be updated with each clock tick requiring considerable processing power. The additional processing power allows the accuracy achieved using a faster clock speed to be maintained while running a vastly more sophisticated set of directional microphones. While greater processing power did not directly increase the speed of the directional microphones, it allowed far better frequency resolution in the spatial domain without compromising speed.

The effect on these two improvements can be seen in Figure 2. A recording was made of a running speech signal in a background of brief intense noise bursts. The speech signal was presented to the hearing aid from directly in front, 0° azimuth. Bursts of speech weighted noise (15 ms in duration) were presented from four separate speakers at 90°, 150°, 210°, 270° azimuth. The speech weighted noise bursts were designed to switch rapidly between each of the four speakers. There was a 15 ms burst from one speaker followed by 5 ms without noise, then the next 15 ms burst from another speaker and so on. The bursts were presented in order at 90°, 210°, 270°, 150° azimuth.

Since the bursts from the speakers were uncorrelated by a few milliseconds, the energy of the speech weighted noise rapidly changed azimuth between the four speakers. This simulates the ebb and flow of signal strength increases in a large group situation. In such an environment, the energy of conversation alternates rapidly between small groups of individuals conversing from different directions in the room. Energy levels rise and fall from different directions as talkers facing each other in several small groups take turns speaking.

Figure 2

Adaptive Directional Microphones

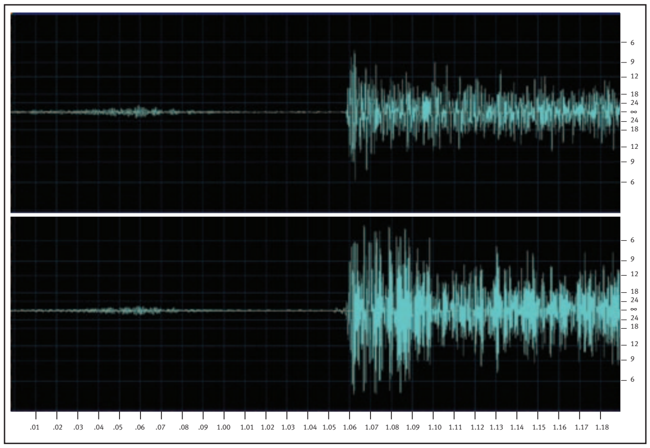

Adaptive directional microphones in speech and noise sampling the environment at 675 Hz (top) and at 125 Hz (bottom).

The bottom panel of Figure 2 is a recording of this speech in noise signal through an adaptive directional system commercially available for the past two years. The top panel shows the same recording as processed through the same directional microphones, but at a faster clock speed and with more processing power. The first (left) third of each panel shows a recording of speech processed though the aids during a pause in the noise bursts. The amplitudes (on the PCM scale) were approximately -39 dB [bottom] and -38 dB [top] respectively. In other words, the gain applied to the frontal speech signal by each aid during the recording was equivalent to within 1 dB. The onset of the noise burst occurs at the center of the figure. During the first 5 ms of the noise burst, the measured amplitudes in each panel were approximately -12 dB [bottom] and -14 dB [top].

Over the next 25 ms the slower adaptive microphones remained at approximately -12 dB [bottom] while the faster microphones reduced the noise level by an additional 5 dB to -19 dB [top]. During the next 100 ms of noise bursts, the amplitude of the microphones in the bottom panel averaged out to -17 dB whereas the top panel microphones averaged -21 dB.

The slower microphones update at a rate of 125 Hz or once every 8 ms. The faster microphones update at a rate of 625 Hz or once every 1.6 ms. A minimum of 3 update cycles is required for the microphones to converge on a noise target and place a null there to suppress it. Therefore, a minimum of 24 ms for the slower microphone and 4.8 ms for the faster microphones is required. This explains why the noise amplitudes through both microphones were roughly equivalent for the first 5 ms. This also indicates why the average noise level over the next 25 ms dropped by 5 dB for the faster microphones, but not for the slower ones. Interestingly, the faster microphone maintains an average 4 dB improved noise reduction compared to the slower microphone over the next 100 ms. In a steady noise, the two microphones would have settled out to be roughly equivalent after the first 30 ms or so. However, because the peak noise is changing azimuth every 20 ms from the back to the side and back again, the slower microphone never completely converges or places a null on the correct point in space. The slower update speed keeps the microphones in the bottom panel constantly out of sync with the moving noise source.

The reduction of short noise bursts has the effect of stabilizing the fluctuations in the background noise, keeping the noise floor constant and as low as possible. The listener can focus on the speech signal from in front without being so distracted by background noises that are rapidly fluctuating. The result is improved comfort, clarity and sound quality.

More Programmable Memory

Another advantage of this new architecture is the addition of considerably more programmable memory to the hearing aid. Increased memory space means that more sophisticated signal processing algorithms can be stored in and run on the device. Furthermore, when improved instruction sets can be processed at a higher speed by a more powerful central processor, the hearing aids become capable of more advanced detection and feature applications. These new capabilities have provided an upgrade path from our existing ClearPath detection system by adding two new sets of detectors to create autoPro4.

ClearPath technology was based on a three detector system. Those detectors were used to assess:

- Modulation Depth
- Modulation Frequency
- Transient Impulses/Short Signal Duration

Every hearing aid had one set of three detectors in each of the instrument’s sixteen channels. For a detailed description of the ClearPath processing system, refer to “ClearPath: A Three Part Approach to Improving Speech in

Noise" (2003). Within ClearPath technology, each detector provides a moment-to-moment update indicating the possible presence of speech, noise or speech mixed with noise on the basis of its single detection parameter. The outputs of the three detectors combine to generate a signal index number. An example of a signal index table is shown in Figure 3. A very low signal index indicates that noise is the dominant signal at that point in time within the given channel. Any channel where the dominant signal is exclusively composed of noise will have a signal index of approximately 10 or less. In this case, the orange index line is resting at 0 meaning that the signal within the channel is virtually all noise. A very high signal index indicates that speech is the predominant signal within the given channel. Channels that include only speech will have a signal index of 25 or above. Once an index of 25 is exceeded, the orange line asymptotes at a 100% probability that the signal within the channel is all speech. Mid-level signal indexes indicate a mixture of speech plus noise. As the signal within the channel becomes more noise-like, the signal index for that channel drops from 25 down toward 10. In other words, as the SNR gets progressively worse within a given channel, the signal index number for that channel will also drop.

Figure 3
Speech Enhancement

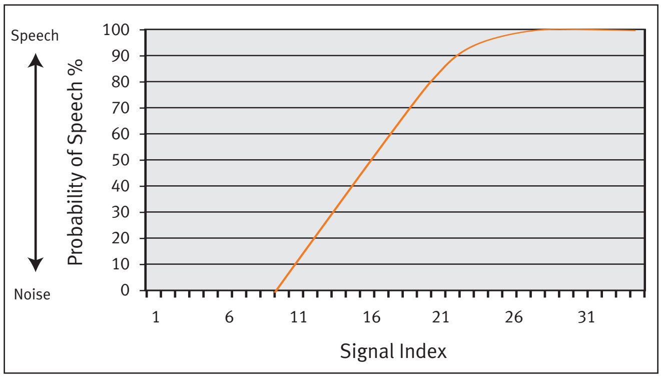

A signal index number above 20 disables noise reduction and engages speech enhancement.

The existing detectors have historically been used exclusively to control noise reduction and speech enhancement algorithms. The decision to enable noise reduction is made once the signal index drops below 20. If the signal index continues to fall within the channel, the noise reduction algorithm adapts by cutting gain more aggressively. Lower index numbers result in greater gain reduction within each channel. Conversely, a signal index number above 20 disables noise reduction and engages speech enhancement. Above 20, speech is the predominant signal and a gain boost of 2 dB – 4 dB is applied.

The three detector system has been found to very reliably categorize acoustic energy into speech, stationary noise, pseudo-stationary noise and transients. However, limited processing power and lack of internal memory has always limited what could be done with that information to the fairly basic operations described above. The development of a new hardware platform has opened the door to new uses for that information and two new detectors that were not previously available.

AutoPro4™: Intelligent Sound Classification Across Multiple Channels Simultaneously

Sophisticated Detectors Across 16 Channels

The precursor to autoPro4 was the classification scheme composed of three detectors which generated a speech and noise signal index number. Decisions about the reaction of some adaptive features such as noise reduction and speech enhancement were based on that index number. This concept has now been taken to the next level. Not only have two new detectors been added, but the output of the intelligent sound classification is used to continuously adapt the hearing aids to as many as four different listening environments.

AutoPro4 compares results from five detectors:

- Modulation Depth
- Modulation Frequency
- Transient Impulses/Short Signal Duration
- Initiation
- Spectral Variance

As with earlier models, the first three detectors make decisions about speech and noise on the basis of the acoustic signal within each of the sixteen channels. The two new detectors make decisions about the acoustic signal as it is detected across the sixteen channels. The Initiation and Spectral Variance detectors examine changes in the acoustic signal across multiple channels simultaneously. Speech cues do not limit themselves to individual 1/3 octave bands. Therefore, detection of the characteristic features typical of speech can be improved by observing changes to the acoustic environment beyond the borders of individual channels.

The Initiation detector samples the acoustic signal looking for patterns typical of speech; primarily these include sudden onset, broadband phonemes. A good example of this type of signal can be seen using the word “cod” in Figure 4. There is a brief cessation of energy up to time frame 1.022 as the lips are closed in preparation for the burst beginning the “c”. This is followed by a broadband (200 Hz – 2000 Hz) burst of energy (1.022-1.027 s time frame) as turbulent air is forced through the narrow constriction created between the lips. There are a few milliseconds of turbulent air flow before the onset of voicing at about time frame 1.06. At that time the voicing of “o” creates a substantial energy increase up to 1200 Hz and a smaller increase for higher level formants from 2600 – 4400 Hz. After the cessation of voicing, there is one more broadband initiation that takes place just after the 1.12 time frame. This is the burst of energy signifying the onset of the ‘d’ at the end of “cod”. Small parts of each energy change can be observed by the detectors within individual channels of the hearing instruments. But a detector looking for the initiation of this type of energy burst across multiple channels makes the recognition that this is a speech signal far more accurate. Collecting more data over a broader frequency range in conjunction with information from the original three sets of detectors makes the entire system more robust. Over time a clean speech input will have far more multichannel initiations than either noise or music. Speech is heavily laden with temporal gaps that do not occur in a steady background noise. Rapidly occurring broadband initiations of the type in the “cod” example are very unlikely in music that does not contain singing. Therefore, the rapid fluctuations of a typical speech signal can be very reliably detected using this approach.

Figure 4
Spectrograph

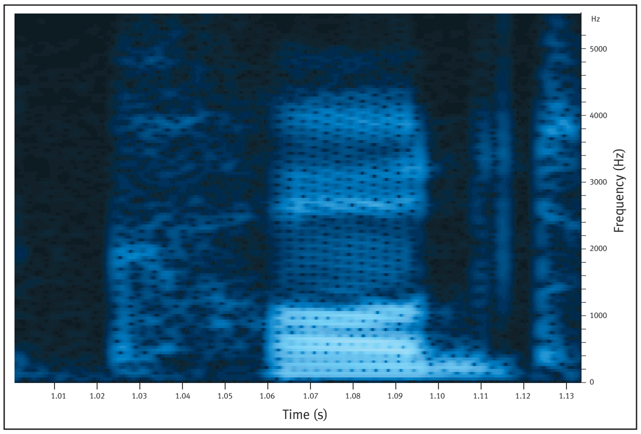

Spectrographic representation of the word /cod/.

The second new multichannel detector looks for variations in the frequency domain that are also consistent with a speech signal. The Initiation detector is sensitive to sharp peaks and dips in the temporal envelope of speech. The Spectral Variance detector is sensitive to rapid fluctuations in the spectral domain. The top section of Figure 5 below shows a spectrographic image of 150 ms of jet noise. The axes are time (s) and frequency (Hz). Intensity is indicated by color. Darker portions indicate little or no signal amplitude and brighter portions show areas of greater intensity (amplitude). The lower panel shows the word /spain/ spectrographically in a background of jet noise over the same 150 ms time frame. Notice that the jet noise is very consistent over time relative to the speech signal. The highest energy occurs below about 700 Hz and there are no large scale fluctuations over either time or frequency.

Figure 5
Spectrographic Representation of Speech and Noise

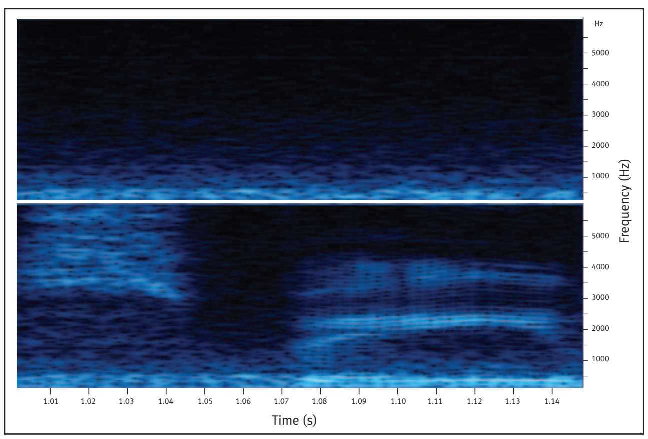

Top=jet noise, Bottom=jet noise plus the word /Spain/

The bottom panel (Figure 6) has three sections; each section shows a Fast Fourier Transform (FFT) of the jet noise (blue) and one phoneme of the word “spain” (orange). In these figures the frequency scale (Hz) runs along the bottom and intensity (dB) is scaled along the side. Each FFT shows the average amplitude over a section in time from the spectrograph that corresponds to the given phoneme. The left hand FFT covers the time frame 1.005 – 1.04 s corresponding to the /s/ phoneme. As expected, virtually all of the jet noise energy is concentrated in the low frequencies while the bulk of the /s/ energy is in the high frequencies above 3000 Hz. The middle frame is the FFT that corresponds in time to the /p/ phoneme (1.05 – 1.07 s). The window for this FFT began only 10 ms after that for the /s/ and yet the spectral distribution has completely changed. This time the weak energy of the /p/ is obscured by the jet noise and so the frequency profile for the jet noise is unchanged while the speech has shifted downward substantially to match it. The final FFT frame (1.09 – 1.13 s) shows the frequency distribution has changed again for the speech during the /ai/ vowel. The jet noise has a constant low frequency energy distribution. Meanwhile, the speech signal is composed of powerful formant peaks at approximately 400 Hz, 2200 Hz and 3800 Hz each where the 400 Hz formant is actually larger than the jet noise.

Figure 6
Fast Fourier Transform (FFT)

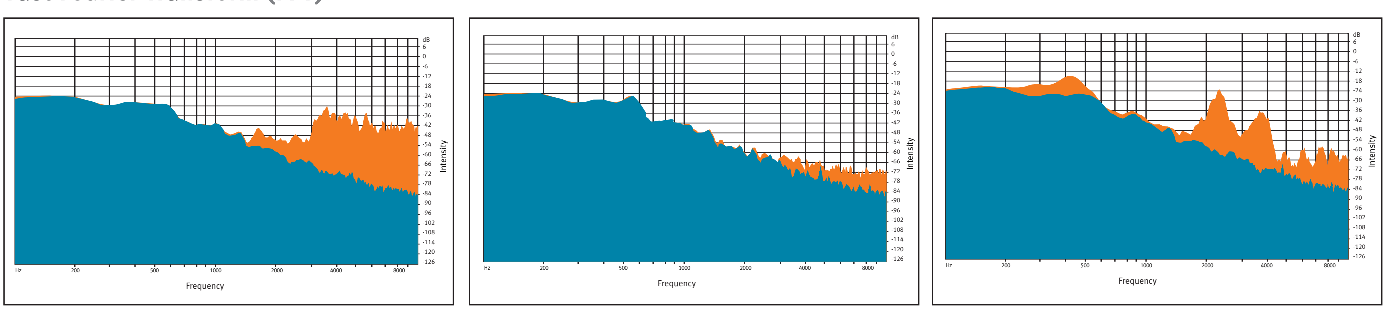

FFT–Blue=jet noise, orange=/s/ in Spain. *FFT–Blue=jet noise, orange=/p/ in Spain.* *FFT–Blue=jet noise, orange=/ai/ in Spain.*

This example shows how the input spectrum to the hearing aid varies over time for a typical speech signal compared to a common noise. The spectral variance detector will monitor changes to the spectral input across multiple channels to tell the intelligent sound classifier if the dominant signal is speech or noise. For example, during this 150 ms recording, the detector will calculate the mean frequency of maximum amplitude and the amount of variance to that mean over time. The three FFT panels in the example illustrate that the mean energy peak of the jet noise remains very constant. The jet noise energy is confined to a consistent narrow band of frequencies below 600 Hz, and therefore, has only a small amount of variance. However, the mean frequency of maximum energy for the speech signal varies constantly over just a few hundred milliseconds. In the first panel, the mean energy peak is well above 3000 Hz and the range of the /s/ is fairly large (300 Hz – 10000 Hz). In the second panel, the energy of the /p/ is obscured by the jet noise so the spectral energy distribution for that panel is the same as for the jet noise, the narrow range below 600 Hz. In the final panel, the vowel energy is distributed across three formant frequencies. The mean of this energy distribution will fall in the mid frequencies. Thus, across three phonemes in one word, the mean energy peak moves considerably relative to that of the stationary noise. Furthermore, the distribution of the energy changes substantially from one frame to the next. First it is a high frequency signal with a broad range, then a narrow band of low frequency energy, and finally a mid frequency signal composed of three separate energy peaks rising out of the noise. Thus, the signal variance of the speech signal is also quite large compared to the stationary noise. Consequently, the spectral variance detector can use the spectral mean and spectral variance characteristics over time to recognize the presence of speech and noise.

However, these two new detectors are not used to determine how to directly process speech within a given band of the instrument like the three original detectors. Instead, they work on a higher level of processing to help autoPro4 determine which listening environment the wearer is in, and subsequently decide, which of the destinations autoPro4 should select: quiet, group/party noise, traffic/intense noise or music.

Determining Destination Probability

Information gathered within each channel is analyzed by three separate detectors and used to generate a signal index. The signal index determines how aggressively the adaptive parameters, (e.g., noise reduction) enabled within that program destination should react. However, the adaptive parameters cannot be turned on or off at the channel level. Instead, this occurs at the listening environment level. Take for example an AutoMic program that switches back and forth from omnidirectional to directional performance in a basic hearing aid. In this case, the aid is automatically switching between two microphone states. If other adaptive parameters are turned on in the algorithm, they will operate the same way regardless of the microphone state. Each adaptive parameter functions independently. However, autoPro4 provides intelligent switching between four independent destinations: quiet, group/party noise, traffic/intense noise and music. Each destination comprises a fully configurable listening environment. Any or all adaptive features may be turned on or off within each listening environment.

What makes autoPro4 so powerful is the ability to correctly assess the listening environment and intelligently switch to the appropriate destination. By utilizing the information from the five detectors, autoPro4 calculates the probability that the current listening environment can be classified into one of the four destinations. Once the probability of a single "correct" destination rises above the probabilities of the other three, autoPro4 will switch the aid into the "correct" destination. The top portion of Figure 7 below shows how autoPro4 has calculated the probabilities of the four listening environments over the course of two minutes for a varying acoustic situation.

Figure 7

Destination Probability and AutoPro4 Sound Destinations

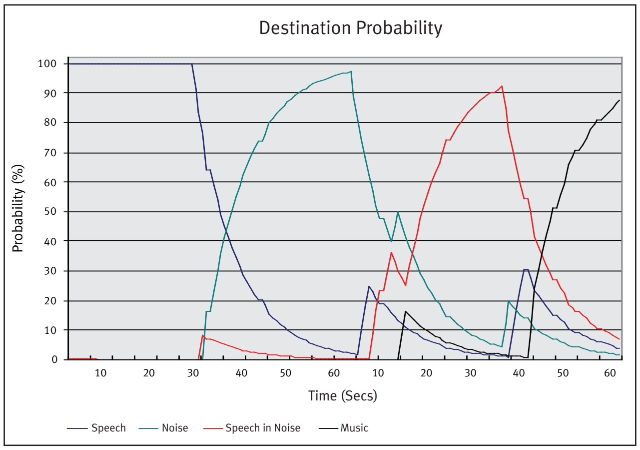

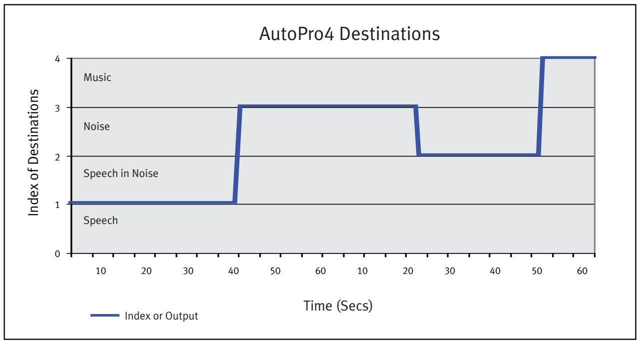

AutoPro4 calculates four probabilities to determine the best sound destination for the particular listening environment.

The four colored lines represent the moment-to-moment probabilities that the hearing aid is in one of the four destinations. At each second, the probability that the aid is in a clean speech environment is shown by the blue line. The green line shows the probability that the listening environment is only noise. The red line represents speech in noise and the black line represents the probability of music. These values were obtained by reading the probabilities directly from the hearing aid while presenting 30 seconds of speech, 35 seconds of noise, 30 seconds of speech in noise and 25 seconds of classical music. Notice that the probability that the environment is speech equals 100% for the first 30 seconds. Speech was already being presented to the aid before the beginning of these recordings. At 30 seconds, the input was abruptly switched from clean speech to the sound of intense rain. At this time, the probability of speech (blue) drops precipitously while that of noise (green) rises rapidly. Concurrently, the incorrect probability of speech and noise rises slightly for about two seconds and then rapidly declines. At 65 seconds, the intense rain noise is replaced by male speech in a background of cocktail party noise. While the probability of the noise environment drops sharply at the changeover, this complex signal triggers rapid fluctuations in the other three probabilities. For about five seconds, both the speech, and speech in noise probabilities rise. Then the clean speech probability drops off and the speech in noise probability continues to increase. After five more seconds, some background music in the cocktail party noise causes a short spike in the music probability that is quickly extinguished. At one minute and forty seconds, the speech in noise signal is replaced by a classical music piece. The speech in noise probability immediately begins to drop. However, the music probability does not rise immediately; the clean speech probability begins to rise first. This occurred because the music piece started slowly and softly. After the sudden release from the cocktail party noise, a brief quiet interval is more likely to contain speech next than to contain music. After a few seconds, the music begins to get louder. At that time, the speech probability drops quickly and the music probability begins its ascension.

This example shows how the probabilities calculated by autoPro4 correctly followed each of four listening environments. However, what is of more concern is how well the hearing aid converged upon the correct destination based on the probabilities that were calculated. Using data obtained from the hearing aid at the same time as the probabilities (lower section of Figure 7), we can show which destination the aid was pointed to over the course of the two minute recording. Notice that the aid always converged on the correct destination. However, there was always a few seconds delay between the sudden change in the listening environment and the selection of a new destination. This is because the destination cannot be allowed to fluctuate rapidly during those intervals of competing probabilities just after the environment changes. When one probability clearly exceeds the other three, that new destination is chosen. This is done to maintain stability in the autoPro4 program.

Summary

New hearing aid technology focuses on the primary benefits of improved speech intelligibility and sound quality. Recent improvements in the quality of sound provided by the newest generation of hearing aids have been achieved by improvements to the underlying architecture of hearing aids. Faster and more powerful processors execute more instructions per second. More instructions per second allow the adaptive parameters to be updated more frequently. Consequently, more frequent updates in a digital system yield a more precise representation of the listening environment and more timely responses to changes in that environment. Greater memory capabilities provide room for more complete instruction sets. Better instructions to hearing aids, that can be processed closer to real time, allow for faster and more accurate responses to the listening environment. All of this extra capacity has led to more sophisticated detectors that enable autoPro4 to correctly predict an appropriate destination for a given listening environment. Sophisticated adaptive features improve audibility by maximizing: the perception of a wide range of acoustic inputs; speech intelligibility by preserving speech while suppressing multiple noise sources; and comfort by permitting awareness of the total environment and allowing for great noise suppression. All of these new features result in the best possible true-to-life sound quality across diverse listening environments.

Bibliography

Arndt, Horst and Nancy Tellier. (2003). ClearPath™: A Three Part Approach to Improving Speech in Noise. Unitron Hearing.

Benson, D., T. Clark, et al. (1992). "Patient experiences with multiband full dynamic range compression." Ear and Hearing 13(5): 320-330.

Boymans, M. and W. Dreschler (2000). "Field trials using a digital hearing aid with active noise reduction and dual-microphone directionality." Audiology 39: 260-8.

Jenstad, L., R. Seewald, et al. (1999). "Comparison of linear gain and wide dynamic range compression hearing aid circuits: aided speech perception measures." Ear and Hearing. 20(Apr): 117-26.

Kochkin, Sergei. "MarkeTrak VII: Hearing Loss Population Tops 31 Million People." The Hearing Review. 12(7), July 2005, pgs 16-29.

Moore, B., R. Laurence, et al. (1985). "Improvements in speech intelligibility in quiet and in noise produced by two-channel compression hearing aids." British Journal of Audiology 19(Aug): 175-87.

Ricketts T.A., and B.W. Hornsby. (2003). Distance and reverberation effects on directional benefit. Ear and Hearing. 24:472-84.

Surr, R.K., B.E. Walden, M.T. Cord, L. Olson, (2002). "Influence of environmental factors on hearing aid microphone preference". Journal of the American Academy of Audiology 13: 308-322.

Walden B.E., R.K. Surr, M.T. Cord, O. Dyrlund. (2004). Predicting hearing aid microphone preference in everyday listening. Journal of American Academy of Audiology. 15:365-96.

Walden B.E., R.K. Surr, M.T. Cord, B. Edward, L. Olson. (2000). Comparison of benefits provided by different hearing aid technology. Journal of American Academy of Audiology. 11:540-560.

Author

Don Hayes, PhD, Director of Audiology

06-003 028-5176-02