A Shure Educational Publication

Selection

and

Operation

of

Audio

Signal

Processors

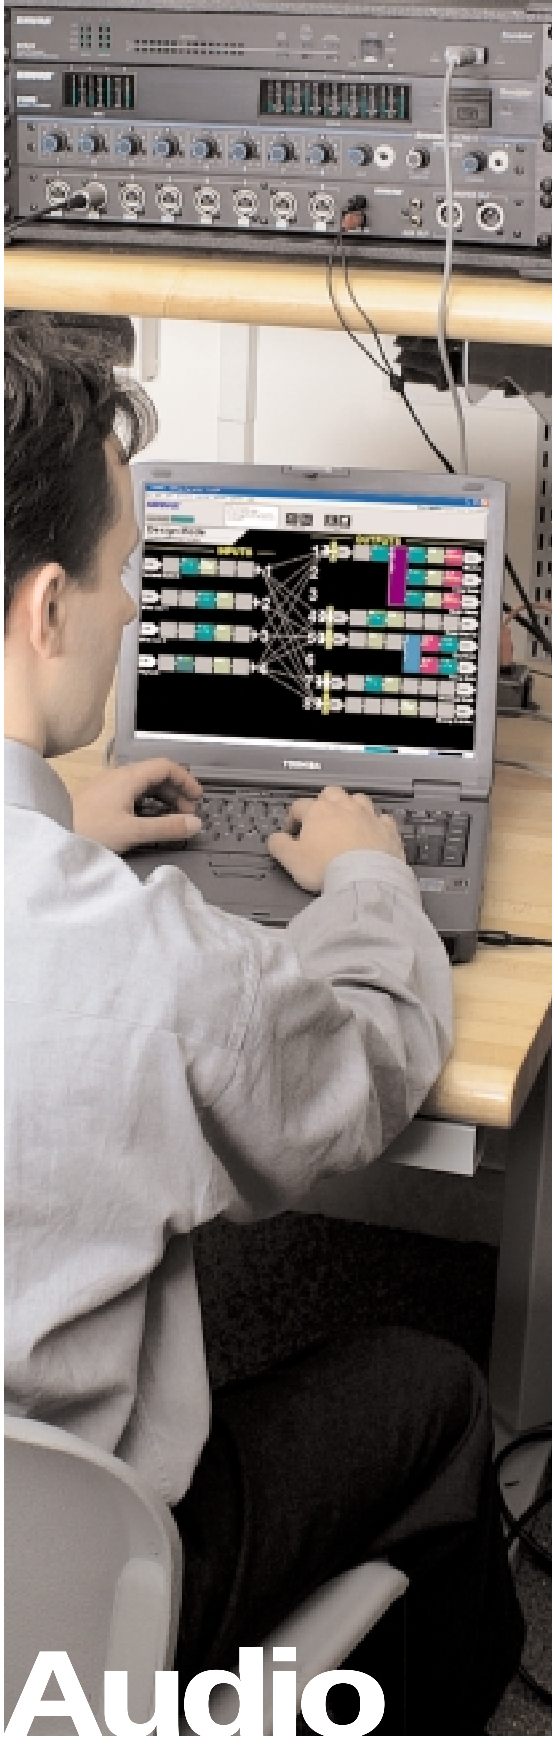

Audio Signal Processors

By Gino Sigismondi

SHURE®

Table of Contents

Introduction 4

What Are Audio Signal Processors? 4
What Types of Problems Can Benefit from Audio Processing? 5
Feedback 5

Chapter 1

Types of Audio Processors 6

1.1 Volume (Gain) Control 6
1.2 Filters and Equalization 6
1.3 Dynamics Processors 11
1.4 Delay 15
1.5 Adaptive Audio Processors 15

Chapter 2

Practical Applications For Audio Signal Processors 20

2.1 Maximizing Gain-Before-Feedback 20
2.2 Improving Speech Intelligibility 21
2.3 Sound System Gain Structure 23
2.4 Digital Signal Processing 25

Reference Information 26

Appendix 1: Sound Waves 27
Appendix 2: Potential Acoustic Gain (PAG) and
Needed Acoustic Gain (NAG) 29
Glossary 34
References 34
Acknowledgements 34
Biography 34
Additional Shure Educational Publications 34

INTRODUCTION

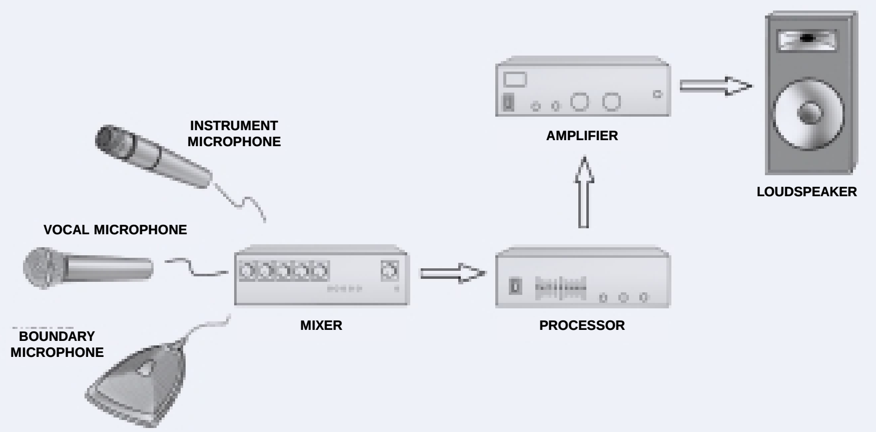

Figure 1-1: basic sound system

For any sound system, the primary goal is good sound. What, however, constitutes "good" sound? The three primary measures of good sound are audibility, intelligibility, and fidelity. Many factors contribute to the quality of the sound, including the quality of the sound sources, the sound system, and the room acoustics.

The audibility of speech or music at the furthest listener must be sufficient to achieve the desired effect: usually a comfortable listening level for speech, and more powerful levels for certain kinds of music. These levels should be attainable without distortion or feedback. Intelligibility is determined by the signal-to-noise ratio and direct-to-reverberant ratio at the listener's ear. The "signal" is the desired sound source (speech, musical instruments, etc.), while the "noise" is ambient sound in the room as well as electrical noise produced by the sound system. Maximum speech intelligibility requires a speech level of at least 20 dB above the noise floor at the listener's ear. The direct-to-reverberant ratio is determined by the directivity of the loudspeakers and the reverberation characteristics of the room. High levels of reverberation can severely degrade intelligibility by making it difficult to distinguish the end of one word and the start of the next. Finally, fidelity of sound is primarily defined by the overall frequency response of the sound arriving at the listener's ear. The frequency range must be sufficiently wide and relatively uniform in order to provide realistic and accurate reinforcement of speech and music. Every component in the signal chain contributes to this, and a limitation at any point will affect the fidelity of the entire system.

Other more subjective terms may be applied to good sound ("warmth", "punch", etc.), but these colloquialisms are not measurable in any meaningful way. Additionally, if the three primary measures are not satisfied, any subjective terms take on even less importance. Speech that is "warm" but unintelligible does the listener little good.

Audio signal processors offer a variety of tools to assist in optimizing a sound system for audibility, intelligibility, and fidelity. While not usually essential for a sound system to operate (i.e., provide high-level sound reinforcement of low-level sources), audio signal processors can be invaluable tools in sound system design. A basic sound system consists of four components:

- Input devices (microphones, CD players, etc)
- Mixers (to combine inputs, control levels, and provide preamplification, if necessary)
- Amplifiers
- Output devices (loudspeakers)

Audio signal processors are typically employed within or just after the mixer stage, but before amplification. (See Figure 1-1.) A processor can be used at the input stage, but since most processors are designed to operate with line level sources this is rare. Signal processors can be analog or digital, single- or multi-function, stand-alone devices or integrated with other components in the sound system. Most signal processors originated as stand alone devices designed for a specific purpose. Over time, integration of similar processors into one device became popular (e.g. compressor/limiters). The development of audio processors that operate in the digital domain allowed for further integration, leading to multi-function digital signal processors (DSP) that combine seemingly disparate functions into a single unit. Perhaps more importantly, DSP devices offer these functions at a cost that is a fraction of the purchase price of several individual processors.

What Types of Problems Can Benefit from Audio Processing?

To understand the purpose of audio signal processing, it is necessary to examine the problems encountered in a typical sound system. Note that an audio processor cannot solve all the potential problems in a sound reinforcement system. The most common problems are listed on the next page:

Problems:	Remedies:
Feedback	Parametric Equalizer/ Automatic Mixer/ Feedback Reducer
Poor tone quality (subjective)	Graphic equalizer
Sound source too loud	Compressor/Limiter/AGC
Sound source too quiet	AGC
Varying signal levels from multiple sound sources	Compressor/Limiter/AGC
Unwanted noise	Noisegate/Downward expander
Unexpected transients	Compressor/Limiter/No overshot ("Look-ahead") Peak Limiter
Comb filtering due to open microphones	Automatic Microphone Mixer
Frequency response anomalies due to misaligned loudspeakers	Delay
Poor intelligibility	Parametric Equalizer/ Automatic Microphone Mixer
Acoustic echoes (in teleconferencing systems)	Acoustic Echo Canceller
Distortion (due to wide dynamic range)	Compressor/Limiter

Problems that cannot be solved by audio processing:

- Echoes because of poor room acoustics
- Poor sound due to excess room reverberation times
- Feedback caused by operating beyond the limits of PAG (see Appendix 2)
- Noise (amplifier hiss, ground buzz, etc.) due to improper system setup
- Distortion due to improper gain structure

The importance of good room acoustics cannot be underestimated. In any room where sound reinforcement will be used, excess reverberation times introduce a myriad of problems that cannot be solved by any audio processors. Reverberation time is the length of time that a sound persists in a room after the sound source has stopped. All attempts should be made to keep unwanted sounds from entering the microphone in the first place. The level of desired sound at the microphone should be at least 30 dB above any ambient sound picked up by the microphone. Proper microphone placement (a full discussion of which is beyond the scope of this publication) is also crucial. A good rule of thumb: always keep microphones as close as possible to the sound source.

Want to know more about proper microphone usage?

Shure offers the following educational guides free of charge:

- Microphone Techniques for Studio Recording
- Microphone Techniques for Live Sound Reinforcement

Visit shure.com or contact your local Shure office (see back cover) to request your complimentary copies.

Once sound energy is introduced into the acoustic space by the loudspeaker, processing no longer has any effect on the sound. Reverberation can be reduced only by absorptive acoustic treatment or structural modification; electronics cannot remove it. If additional acoustic treatment is not an option, directional loudspeakers allow the sound to be "aimed" toward the listener and away from reflective surfaces. Simply raising the level of the sound system will only aggravate the problem by raising the reverberation level as well. Long reverberation times severely reduce intelligibility. In audio teleconferencing systems, this results in a hollow, or "bottom-of-the-barrel" sound received by the remote site.

FEEDBACK

Feedback is characterized by a sustained, ringing tone, which can vary from a low rumble to a piercing screech. Echoes and reverberation caused by room acoustics, as well as ground buzz and other extraneous noises, are not the same thing as feedback, and cannot be cured in the same manner.

Feedback occurs whenever the sound entering a microphone is reproduced by a loudspeaker, picked up by the microphone, and re-amplified again and again. The familiar howl of feedback is an oscillation that is triggered by sound entering the microphone. The easiest way to (intentionally) create feedback is to point a microphone directly into a loudspeaker. Placing the microphone too close to the loudspeaker, too far from the sound source, or simply turning the microphone up too loud exacerbates feedback problems. Other contributing factors are too many open microphones, poor room acoustics, and uneven frequency response in either the microphones or loudspeakers.

The single easiest way to reduce feedback is to move the microphone closer to the desired sound source. Additionally, using a directional microphone (cardioid, supercardioid, etc.) will slightly increase the amount of gain-before-feedback. Reducing the number of open microphones with an automatic mixer will also improve the situation. Try to keep microphones and loudspeakers as far away from each other as possible. Lastly, acoustically treat the room to cover hard, reflective surfaces such as glass, marble, and wood. Realize, though, that in certain rooms long reverberation times may be desirable, such as a house of worship used for acoustic music performance.

If the system has been designed with careful consideration of these factors and feedback is still an issue, an automatic feedback reducer can be used to flatten the response at problem frequencies. These devices are discussed in Section 1-5.

CHAPTER 1

Types of Audio Processors

VOLUME (GAIN) CONTROL

Although often overlooked as an audio processor, a simple volume (or gain) control fits the definition. Volume adjustments can be made at several points within the sound system, from the microphone inputs on the mixer all the way to the inputs of the power amplifiers. Volume levels are typically manipulated in one of two ways: continuously variable adjustments, such as those made by rotary potentiometers or faders, or fixed attenuation such as that provided by a pad.

If adjusting a volume control adds amplification to the audio signal, it is said to be providing gain. The volume control that adjusts the amount of amplification added at a mixer's microphone input is sometimes referred to as a gain (or trim) control, since the volume potentiometer is controlling the gain of the microphone input's preamplifier. The function of this gain control is to match the input sensitivity of the device to the level from the source.

A second type of volume control acts as an attenuator, basically a continuously variable resistor that adjusts the amount of signal allowed to pass through it. No additional gain is provided by the volume control. The volume control on an electric guitar is an attenuator. These devices are often referred to as passive volume controls, since they do not require any power. Occasionally, a volume control will combine attenuation with gain. Faders on a mixing console typically provide attenuation below the "0" indication, and gain above that point.

Pads allow input stages to accommodate a variety of signal levels. Microphone inputs typically feature an input attenuation pad of some kind to reduce the sensitivity of the input beyond that of the preamplifier gain control, typically by 20 dB. A 50 dB pad is required for microphone inputs that are designed to accept either microphone or line level. The output stage of various devices can also employ pads, usually to prevent overloading of the input stage of the next device in the signal path. Care should be taken to use pads only when necessary. For example, using a 20 dB pad on a microphone input that does not need additional attenuation will require additional gain be added by the preamplifier, which adds more noise to the audio signal.

While volume controls are the simplest of all audio processors, they often the most misused. Correct calibration of the various volume controls in a sound system is known as proper gain structure. (See Section 2-3: Gain Structure.)

FILTERS AND EQUALIZATION

Filters are signal processors that affect frequency balance. At a basic level, filters are used to attenuate or boost the level of specific frequencies or frequency ranges. Designed originally to compensate for frequency-dependent loss in telephone lines, some form of frequency-dependent filtering (or equalization) is found in all but the most basic of sound systems. The simplest form of filter is the tone control, basically a filter that attenuates high frequencies above a predetermined frequency. Equalizers are typically characterized by combining several filter sets to offer more precise frequency response shaping. Historically, filters were passive devices capable of attenuation only. The frequency range and amount of attenuation were achieved with capacitors, inductors, or a combination of both. Favorably, passive filters do not require power and do not generate noise. The large size and expense of discrete components, however, precludes the ability to develop equalizers with larger numbers of filters and more precise control of frequency and level. Active filters allow for fast, easy tuning and the ability to add gain, using smaller components at lower cost. Tone controls employing active filters can be found on even the most inexpensive home stereo systems. In this scenario there are typically two controls, treble and bass, which correspond to filters that affect low frequency and high frequency response. Since they are active, these tone controls are capable of cut or boost.

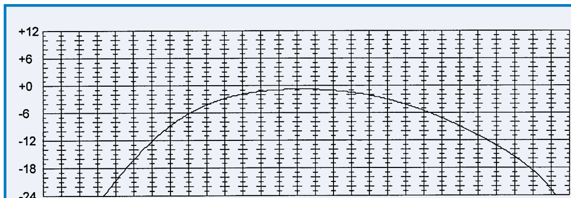

Low Cut: -6dB/octave below 125 Hz
High Cut: -6dB/octave above 2 kHz

Figure 1-2: low cut and high cut filters

Simple filters that affect a broad range of frequencies are divided into four types: high pass, low pass, band pass, and band reject. High pass filters, as the name implies, allow high frequencies to pass, and low pass filters do the same for low frequencies. It is often more convenient to think of these filters in terms of the frequencies that they cut instead. High pass filters are also known as low cut filters, and low pass filters are known as high cut filters, but their function is the same and these terms can be used interchangeably. (See Figure 1-2.) Low and high cut filters have an associated slope that defines how rapidly output declines below (or above) the filter frequency. Slope is typically defined in dB/octave. The span of an octave relates to a doubling (or halving) of frequency, for example, 50 to 100 Hz or 5 kHz to 2.5 kHz. A 6 dB/octave low cut beginning at 100 Hz, therefore, translates into 6 dB less output at 50 Hz, and 12 dB less at 25 Hz. Typical slopes are 6, 12, 18 and 24 dB/octave. The slope also has an associated rolloff characteristic, most commonly Bessel, Butterworth, or Linkwitz-Riley. See "Crossovers" for more information on filter slope types. The frequency that defines a filter is usually stated at its 3 dB down point (A low cut filter set to 100 Hz is actually 3 dB down at 100 Hz). A band pass filter allows only a certain range of frequencies to pass (called the passband). The same effect can be achieved by using a low cut and high cut filter together. The result is the similar to boosting the frequencies that comprise the pass band. A band-reject filter reduces a range of frequencies.

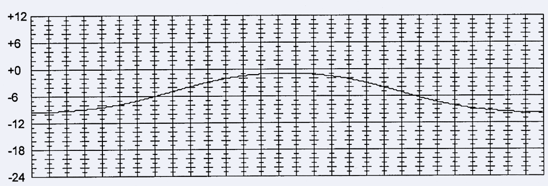

Figure 1-3: shelving equalizers

A further subdivision of high and low cut filters is the shelving equalizer. (See Figure 1-3.) Rather than continuing to decline at a certain dB/octave rate, attenuation flattens out at a certain fixed level, forming what appears as a "shelf" when observed on a frequency response chart. Unlike low or high pass filters, most shelving equalizers allow for boost as well as cut. Consumer bass and treble controls are typically shelving equalizers where increasing the amount of cut or boost often changes the frequency at which the EQ begins to take effect. More advanced shelving equalizers allow the user to select the frequency, the amount of cut, and occasionally the rate of cut (again in dB/octave).

Graphic Equalizers

The most common equalization tool for sound reinforcement is the graphic equalizer. A typical graphic equalizer consists of a bank of sliders (or faders), corresponding to specific frequencies, which can cut or boost the chosen frequency. (See Figure 1-4.) The center frequencies of these filters are identical for all graphic equalizers, regardless of manufacturer, because they are defined by ISO (International Standards Organization) documents. Since the position of the sliders roughly represents the amount of cut or boost, this type of equalizer offers an *approximate* visual representation of the frequency response alteration created by the equalizer, hence the term "graphic." The actual width of the filters, though, is wider than what is implied by the graphic equalizer, and the combined response of multiple filters will most likely be much more dramatic. Also, note that this is only the response imposed on the audio signal by the equalizer, not the actual frequency content of the audio signal. For example, if the audio signal is already 2 dB up at 2 kHz, using the EQ to add another 3 dB of boost at 2 kHz results in a total increase of 5 dB. However, the graphic equalizer only reflects the 3 dB boost. An analysis of the total frequency response of the sound system requires a measurement device, such as a Real Time Analyzer (RTA).

The number of filters available on a graphic equalizer can vary from as few as 5 (a 5-band graphic equalizer) to 31 (a 31-band graphic equalizer) or more. On a graphic equalizer, there is a direct correlation between the number of filters and the bandwidth of each filter. In general, more filters offer more precise control because the range of frequencies that each filter affects is smaller. The bandwidth of each filter is also a defining characteristic of the equalizer. Typical classifications are one octave, 2/3-octave, or 1/3-octave. Higher bandwidth filters, such as 1/6-octave, exist but are rarely encountered. 1/3-octave graphics are the most common, since they offer a fairly precise level of control with a manageable amount of sliders. The audible frequency range requires 30 or 31 1/3-octave filters; a 1/6-octave graphic requires at least 60 1/6-octave filters. Lower bandwidth devices, like one octave or 2/3-octave, are broadband in nature and typically used for overall tone shaping rather than precise control. A one-octave graphic equalizer usually has 7 or 8 filters, a 2/3-octave has 15. Note that some older equalizers use a rotary potentiometer rather than a vertical fader. This device is still termed a graphic equalizer, though the visual representation of frequency response created by pointers on knobs is far less intuitive.

Graphic equalizers can further be defined as combining or non-combining. (See Figure 1-5.) If a frequency that needs to be attenuated lies between two 1/3-octave band centers, those two filters can be cut to reach that frequency. In a non-combining equalizer, the area of overlap between two filters will be somewhat higher in level, requiring excessive cut to adequately reduce the level of the desired frequency. A combining equalizer, however, has a smoother transition between adjacent bands, requiring less overall cut to reach the same level

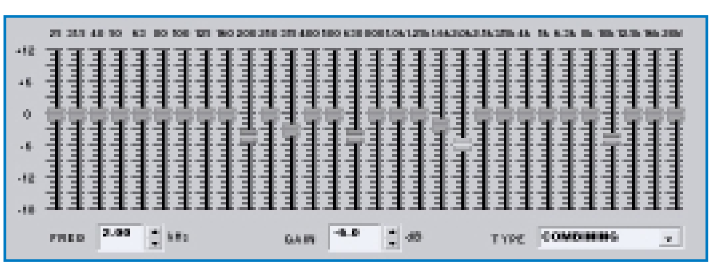

Figure 1-4: graphic equalizer

CHAPTER 1

Types of Audio Processors

of attenuation at said frequency. Additionally, there is less "ripple" in the overall frequency response, whether boosting or cutting. Due to this smoother response, graphic equalizers with combining filters are preferred for sound reinforcement applications. Generally, graphic equalizers use combining filters, unless otherwise specified.

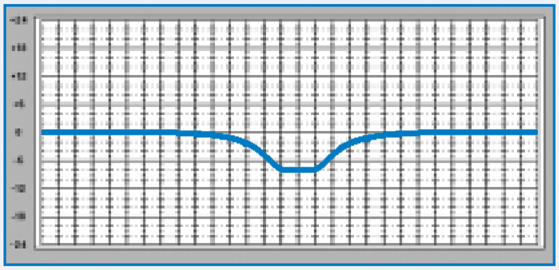

combining filters

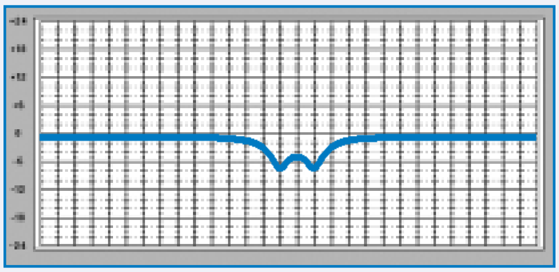

non-combining filters

Figure 1-5

Parametric Equalizer

The parametric equalizer offers a much greater degree of control than a graphic equalizer by giving the user more parameters to adjust. In addition to cut or boost of specific frequencies, a parametric equalizer also allows adjustment of the center frequency and bandwidth of the filter. (See Figure 1-6.)

The center frequency is defined as the point at which the maximum amount of boost or cut occurs. The bandwidth, as stated above, indicates the actual range of frequencies affected by the filter. A semi-parametric (or sweepable) equalizer allows selection of center frequency and boost or cut, but the bandwidth is fixed. (See Figure 1-7.) This is a common feature on more affordable mixing consoles. Most modern mixing consoles with an EQ section have at least one sweepable filter (usually for midrange). More advanced mixers include two or more bands of fully-parametric EQ. Concentric potentiometers are often used to help save real estate on the console, typically to control frequency and bandwidth. Stand-alone, analog parametrics also are limited by space requirements, since each band requires three controls. They are typically available with 5 to 7 filters. Digital parametric equalizers, on the other hand, are limited only by the processing capabilities of the device. More processing power means more filters. Many digital signal processing (DSP) equalizers allow the user to deploy the filters using a display that shows a graphical representation of the filters. This type of EQ combines the visual advantage of a graphic equalizer with the more precise control of a parametric.

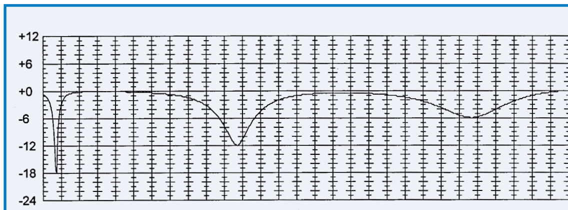

Low: 1/40 octave, -18dB @ 30Hz
Middle: 1/3 octave, -18dB @ 300Hz
High: 1 octave, -18dB @ kHz

Figure 1-6: parametric filters

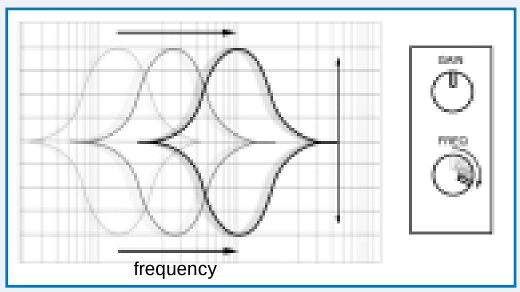

Figure 1-7: sweepable filter

The main advantage to adjustable bandwidth is less effect on adjacent frequencies when applying corrective equalization. (See Figure 1-8.) For example, in a sound system where 500 Hz needs to be attenuated by 6 dB, using a 1/3-octave graphic equalizer results in approximately 3 dB of attenuation at 400 Hz and 630 Hz. By using a parametric equalizer and reducing the bandwidth to 1/10 of octave, the same frequencies are barely affected. Conversely, employing a wider bandwidth allows several adjacent frequencies to be intentionally cut (or boosted) with only a single filter.

Different devices express bandwidth using one of three measures: fractions of an octave, Q, or number of Hertz. (See Figure 1-9.) At 1 kHz, a filter with 3 dB of cut and a 1-octave bandwidth corresponds to a Q of 1.41 and covers approximately 709 Hz. For the purpose of defining Q, bandwidth is measured from the 3 dB up or down points (depending on whether there is boost or cut). Dividing the center frequency by this bandwidth in

Hz gives the Q, which stands for "Quality Factor." Q gives an indication of how tightly the filter is focused near the center frequency. In this example, the –3 dB points for a one-octave filter are approximately 645 Hz and 1355 Hz, a difference of 710 Hz, therefore:

Q = 1000 Hz/710 Hz
Q = 1.41

Note that when determining Q, the 3 dB points are defined relative to the peak or trough, not the audio pass band. This sometimes leads to confusion, because the effective bandwidth of a filter is sometimes also defined as the difference in frequencies at 3 dB points relative to unity gain, rather than the center frequency. Unfortunately, the meaning of the term bandwidth can change with context.

While significantly more powerful than graphic equalizers, parametrics do require a greater level of understanding on the part of the user, particularly when adjusting bandwidth. A graphic equalizer provides simple operation for general tone shaping and on-the-fly tweaks. With proper application, the parametric equalizer is a powerful tool for surgical adjustment of frequency response anomalies and problematic feedback frequencies. Also, note that a parametric filter can be adjusted to duplicate the function of an individual graphic EQ filter.

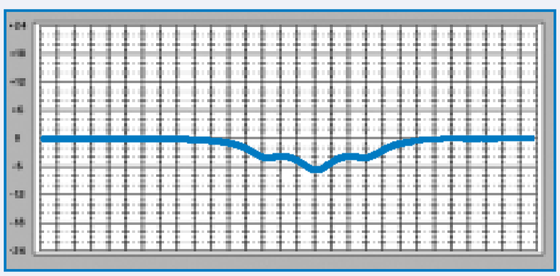

1/3 octave graphic EQ:
-3 dB @ 500 Hz, -6 dB @ 1000 Hz, -3 dB @ 2 kHz

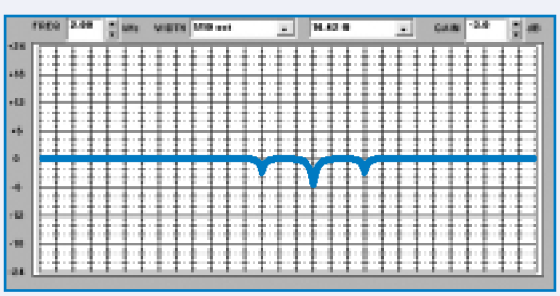

1/10 octave parametric EQ:
-3 dB @ 500 Hz, -6 dB @ 1000 Hz, -3 dB @ 2 kHz

Figure 1-8

Applications Tip: ***Graphic EQ vs. Parametric EQ***

Many audio professionals differentiate the two main types of equalizers in this way:

Parametric EQ: *The "problem solver." Use the parametric equalizer to correct response peaks in the sound system. Microphones and loudspeakers, in particular, introduce many irregularities into the overall frequency response. With the appropriate audio measurement device, these irregularities are easily identified and corrected by a parametric equalizer.*

Graphic EQ: *The "tone control." Use the graphic equalizer to make broad changes to the sound system's frequency response. Once the parametric equalizer has flattened the frequency response of the system, the graphic equalizer serves as a tool for subjective shaping to achieve "pleasing" sound quality.*

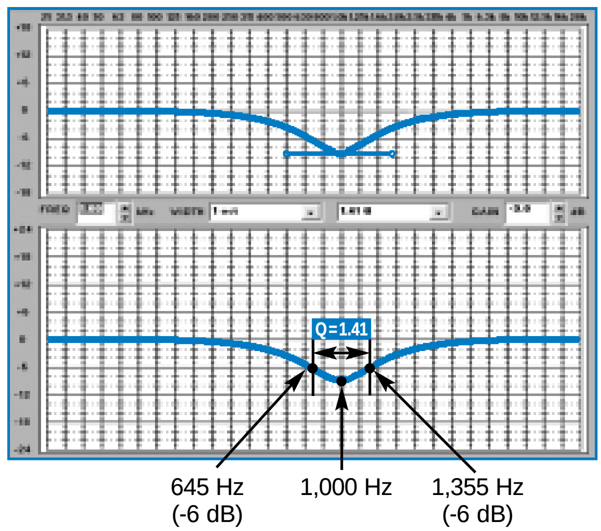

Figure 1-9: 1 octave filter, -9 dB @ 1kHz

Crossovers

To understand the purpose of a crossover, it is helpful to understand the frequency response characteristics of a typical loudspeaker. When measured with a pink noise source, it becomes apparent that any given loudspeaker can, at best, only reproduce a decade of frequency response without compromise. Whereas an octave represents a doubling of frequency, a decade (from the Greek *deca*) represents a factor of ten. The range from 100 Hz to 1 kHz is a decade. Therefore, to accurately reproduce the entire audible range for humans requires at least three loudspeakers, each theoretically optimized for the following frequency ranges:

20 – 200 Hz
200 – 2 kHz
2 kHz – 20 kHz

Chapter 1

Types of Audio Processors

In reality, most loudspeakers will not have exactly these specifications, due to compromises that must be made in the design of loudspeaker systems. Very few sound systems actually produce much output below 40 Hz, especially for speech applications. Therefore, two-way loudspeakers with acceptable fidelity are possible, and quite popular. The frequency response of this type of loudspeaker typically extends only as low as 80 Hz. In this case a subwoofer could be used to provide extended low frequency response, if necessary.

Furthermore, since loudspeaker drivers are generally optimized to reproduce a particular band of frequencies, a given loudspeaker may be subject to damage if driven with a high-level signal that contains a frequency range it was not designed to handle. This situation is particularly true for high frequency transducers, such as compression drivers, that have very little response below 1 kHz.

A crossover divides the audio signal into two or more frequency bands. (See Figure 1-10.) The frequency at which the division occurs is the crossover frequency. A crossover can be either active or passive, and is described using parameters similar to those found in low pass and high pass filters, namely: frequency, slope, and filter type. The most common filter types found in crossovers are Bessel, Butterworth, and Linkwitz-Riley, with slopes of 6, 12, 18, or 24 dB per octave. While providing a minimal amount of phase shift, a 6 dB per octave slope results in significant overlap between the frequency ranges fed to the loudspeaker components, and may not provide enough protection for high frequency drivers. Historically, the 18-dB per octave Butterworth filter has been a sound system standard, though the 24 dB/octave Linkwitz-Riley crossover has eclipsed its popularity. Besides the advantage of minimal overlap at the crossover point, the Linkwitz-Riley filter provides an in-phase relationship at the crossover outputs.

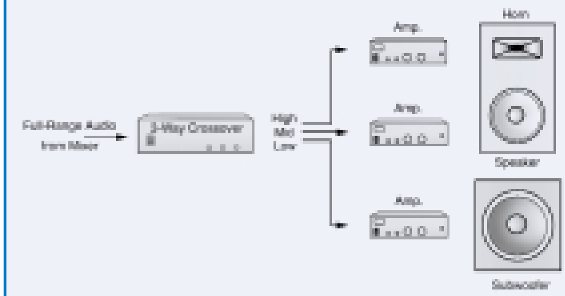

Figure 1-11: 3-way active crossover

A passive crossover, basically a combination low-pass and high-pass filter, is typically the last processor encountered before the loudspeaker; often integral to the design of the loudspeaker itself. Passive crossovers do not require power to operate and are normally invisible to the user. The crossover frequency is fixed, optimized by the designer for that particular loudspeaker. Passive crossovers are often referred to as high-level, since they operate with speaker-level signals. Unfortunately, the full output of the amplifier may not be delivered directly to the loudspeaker since some power if absorbed by the crossover. Also, the electrical components required for passive crossovers may dictate physically large devices, and production tolerances in these devices can vary, affecting the accuracy of the crossover.

The active crossover, also known as a low-level or electronic crossover, provides several significant advantages over the passive design. These advantages include increased amplifier headroom and more efficient use of available amplifier power. Low frequencies tend to place the greatest demands on amplifier power. If a single amplifier is used to drive a multi-way loudspeaker, any distortion due to overload at he amplifier input is reproduced by every transducer in the system. This situation can result in audible clipping, especially of high frequency material. By dividing the audio signal with an active crossover, a separate amplifier is used for each frequency band, (see Figure 1-11) thereby reducing the likelihood of audible distortion. If low frequency energy causes the woofer amplifier to clip, the other amplifier, and the loudspeaker connected to it, will not be affected. This is known as a bi-amplified sound system. Similarly, a three-way

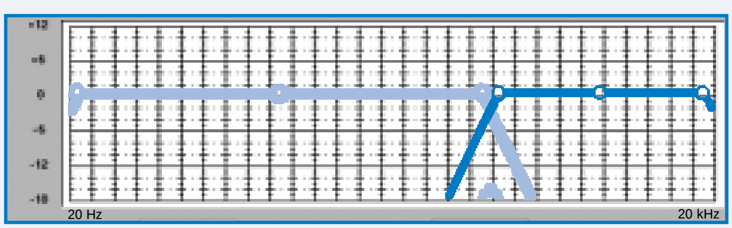

2-way crossover

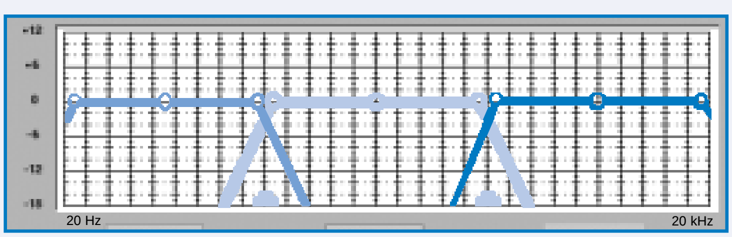

3-way crossover

Figure 1-10

crossover feeding three power amplifiers is called a tri-amplified system. If clipping occurs in the low frequency amplifier, the higher frequency harmonics created by the clipping are reproduced only by a woofer that has very low output at high frequencies, thus reducing the audibility of the distortion. The use of active components also offers smaller size and more repeatable production due to better tolerances.

Quite often, a sound system combines elements of both passive and active crossover networks. These types of systems typically use an active crossover to provide a separate subwoofer output for low frequencies, while a passive crossover in a two-way loudspeaker divides mid- and high frequencies. This could be described as a three-way, bi-amplified sound system.

Most active crossovers allow for control of crossover frequency and level at each output. DSP-based crossovers typically offer greater adjustment, providing the user with selectable filter slope, filter type, and polarity reversal.

DYNAMICS PROCESSORS

The term *dynamics* refers to the wide variations in signal levels commonly encountered in sound systems. Every sound has a dynamic range, defined as the difference between the loudest and quietest levels. A signal that varies greatly in level, such as speech, is described as having a wide dynamic range. A noise source (such as pink noise) that is held to a consistent level has a narrow dynamic range. (See Figure 1-12.) Music sources typically fall somewhere in between speech and noise, although some music sources can have a dynamic range much greater than speech. Used properly, a dynamics processor can manipulate the level variations in a signal to increase audibility and reduce undesired noise in a sound system. Common dynamics processors include compressors, limiters, expanders, noise gates, and speech levelers (a.k.a. automatic gain controls.)

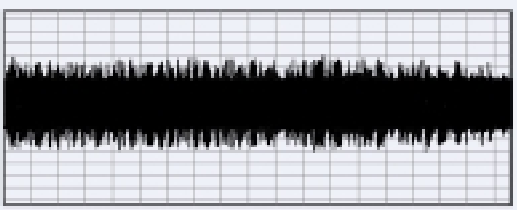

pink noise - narrow dynamic range

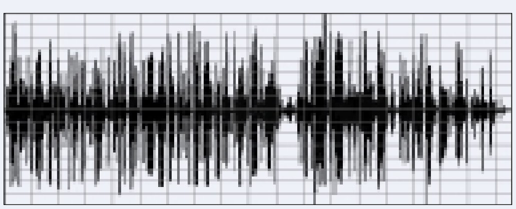

uncompressed speech - wide dynamic range

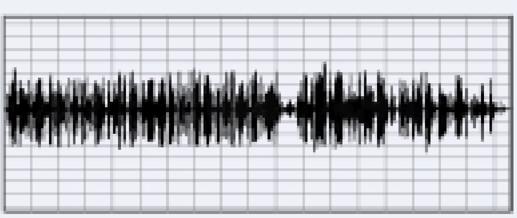

compressed speech

Figure 1-12

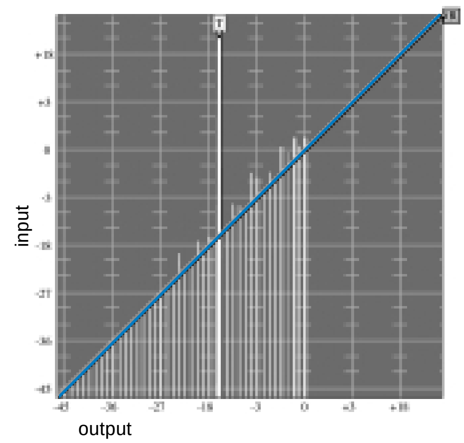

uncompressed
vs.
compressed

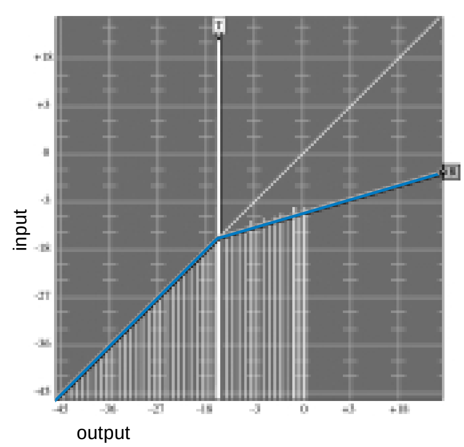

Figure 1-13

CHAPTER 1

Types of Audio Processors

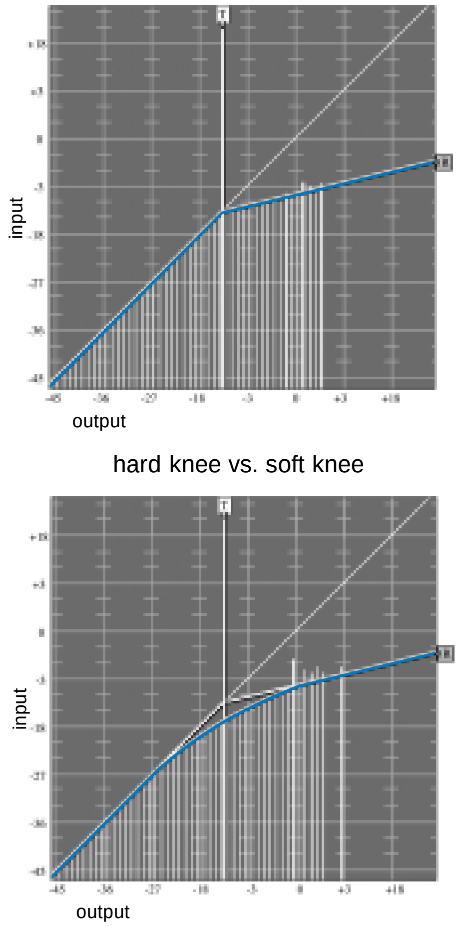

Figure 1-14

Compressors

Perhaps the most commonly encountered dynamics processor, a *compressor* reduces (or "compresses") the dynamic range of an audio signal. A compressor functions by reducing the level of all signals above a user-defined point (the threshold), by a specified amount. (See Figure 1-13.) A ratio defines the amount of reduction that occurs above the threshold. A ratio of 2:1, for example, will allow an audio signal to exceed the threshold by only half as much as what it would have without compression. Assuming a threshold setting of 0 dB, a +10 dB signal is output at +5 dB. Similarly, a 4:1 setting will reduce the output by one-quarter of the original signal level. This reduction limits variation between the lowest and highest signal levels, resulting in a smaller dynamic range. A common myth concerning compressors is that they make quiet signals louder. While this may be the perceived effect, reducing the dynamic range of a signal allows the user to boost the overall level of the signal, yet keeps loud signals from getting "too loud" and causing distortion further down the audio chain - or simply annoying listeners. The compressor itself does not boost lower signal levels, but simply allows them to be perceived closer in level to louder signals.

Other compressor settings include attack, release, and decays. A compressor's attack time relates to how quickly the compression takes effect once the signal exceeds the threshold. Shorter attack times offer greater transient control. Longer attack times generally sound more natural, and are often employed in musical applications. Too long an attack time can cause the compressor to miss signals that otherwise should be compressed. Release refers to the time it takes for the compressor to return the signal level to its original value after the level drops below the threshold. Too short a release time can result in "pumping" and "breathing" with signals that have rapid level changes. Too long a release time can render quieter passages inaudible since gain reduction is still being applied to the audio signal.

A compressor's knee is the point where the signal crosses the threshold. Standard compression schemes reduce the signal by the same amount once the signal has passed the threshold. This is known as hard knee compression. Some compressors allow the user to select soft knee compression instead, where the onset of compression near the threshold occurs more gradually than the more aggressive hard knee compression. (See Figure 1-14.) The compression ratio near the threshold is actually less than specified. Audibly, soft knee compression creates a more gradual transition from uncompressed to compressed signals, making the compression less noticeable.

Applications Tip: Compressor vs. Loudspeaker

Here is a common complaint made by the owner of a damaged loudspeaker: "How could I have blown a loudspeaker? I have a compressor!" Unfortunately, while compressors and limiters help prevent audio transients from causing clipping or possibly damaging a loudspeaker, high-level transients are not the only cause of damaged loudspeakers. In fact, over-compression of the audio signal can contribute to premature loudspeaker failure.

It is standard practice to use an amplifier with a power rating at least twice the loudspeaker's continuous power rating (e.g. use a 200 watt amplifier for a 100 watt loudspeaker). The extra headroom afforded by the larger amplifier allows for peaks in the program material to be delivered to the loudspeaker without clipping. The majority of the amplifier power goes largely unused since the average level of an uncompressed audio signal is considerably lower than the peaks. Highly compressed signals have an average level much closer to the peak level. If the level of the compressed signal is raised to take advantage of the additional amplifier power (thereby making it louder for the audience), the average power delivered to the loudspeaker may be more than the continuous power rating of the loudspeaker, overheating the loudspeaker's voice coil and causing failure.

As with all audio processors, using a compressor does not eliminate the need for proper system operation. Though a compressor or limiter is essential for reducing transient peaks, excessive compression is the enemy of the loudspeaker.

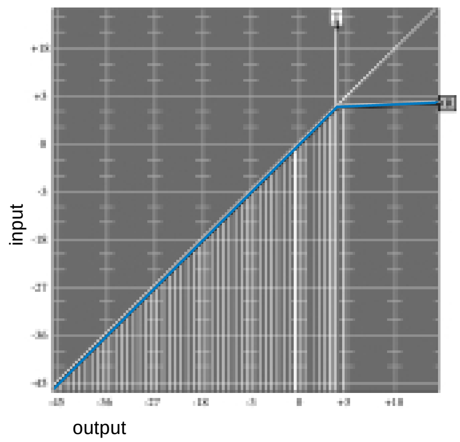

Figure 1-15: limiter

Limiters

A *limiter* functions in much the same way as a compressor, differentiated more by its application than its operation. Similar to a compressor, a limiter also reduces signals that pass a threshold by a certain ratio. The ratios used by limiters, though, tend to be much greater than those used by compressors. Typical limiter ratios can range anywhere from 10:1 to ∞:1 (infinity:1, where the threshold setting dictates the maximum signal level). (See Figure 1-15.) The goal of a limiter is usually system protection, by preventing transient audio peaks from causing distortion further up the audio chain or, worst case, damaging loudspeaker components. Typically, limiter threshold settings are also much higher than on compressors; low threshold settings on a limiter lead to excess compression. Limiters also share other parameters with compressors, including attack and release.

To further illustrate the difference between compressors and limiters, imagine someone jumping on a trampoline in a low-ceilinged room. The up and down motion of our trampoline artist represents the varying voltage of an audio signal; the ceiling represents the threshold of either the compressor or limiter. If the ceiling is made of thin rubber, it will give when the trampoline artist hits it, allowing the person to pan beyond the ceiling (or "threshold"). But not by as much as he would if there were no ceiling there at all. A hard plaster ceiling, however, is analogous to a limiter. When the artist hits the ceiling, no further travel beyond it is possible.

In practice, the operation of a limiter is not quite this absolute. A standard limiter cannot have an attack time of zero. An unexpected, loud transient could pass through the output before the limiter circuitry has a chance to act on it. To provide maximum loudspeaker protection, a limiter needs the ability to anticipate peaks. DSP-based limiters can accomplish this by inserting a small amount of delay (normally 1ms) into the signal path. By delaying the signal, the limiter is able to act on transient peaks that would otherwise escape before the limiting occurs. The attack time of this limiter is effectively zero. These types of limiters are commonly known as *look-ahead* or *peak stop* limiters. They are often the last device in the signal path before the power amplifier and are typically assigned a very high threshold. Since the nature of the look-ahead limiter is last-resort system protection, such a limiter may have less than pleasing audio quality. A high threshold assures the limiter will not affect the audio signal unless absolutely necessary.

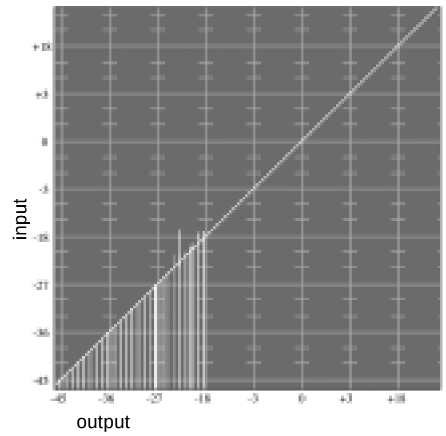

low level unprocessed signal

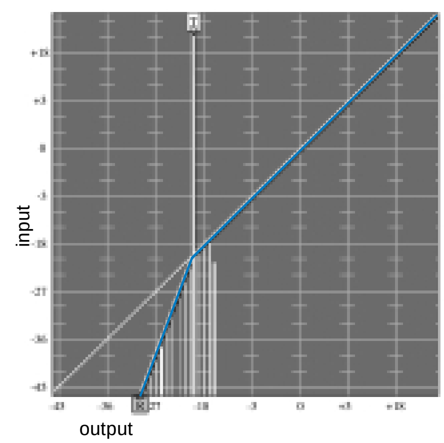

expander processing – smoother transition

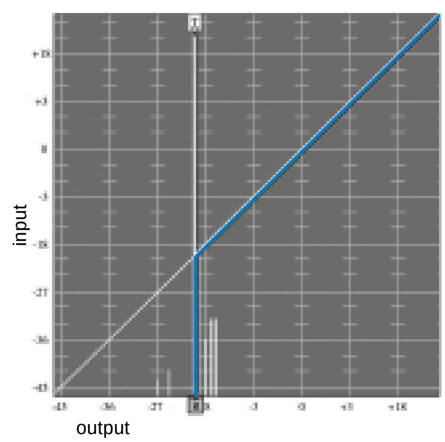

noise gate processing – more abrupt transition

Figure 1-16: limiter

CHAPTER 1
Types of Audio Processors

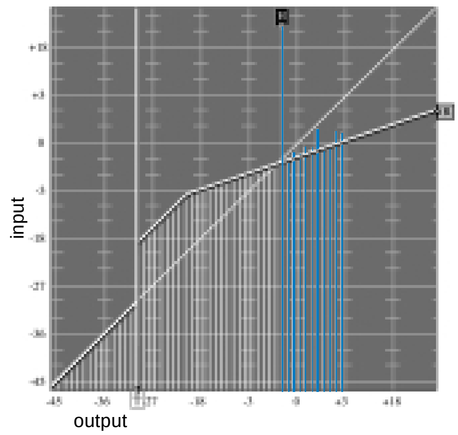

AGC compresses above hinge point (H) vs.
AGC boosts below hinge point

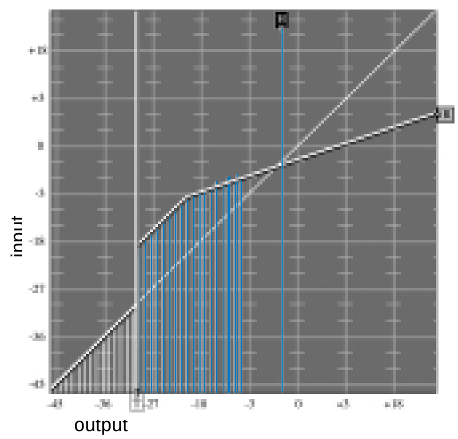

Figure 1-17

Expanders and Noise Gates

An *expander*, as the name implies, functions as the reverse of a compressor by expanding the dynamic range of an audio signal. An expander works by raising signals that pass above the threshold and, in some cases, by also attenuating signals that are below the threshold. As in a compressor, the ratio dictates how much gain is added to the signal. A downward expander, conversely, only reduces signal levels below the threshold, again using a ratio. The same set of adjustments (attack, decay) also apply to expanders. The applications for true expanders in sound systems are limited. They are often used in conjunction with a compressor to form a compander, a circuit commonly used in noise reduction systems and wireless microphone systems. Compression is normally employed in the transmitter of a wireless system to prevent the radio frequency signal from deviating beyond (usually government imposed) bandwidth limitations. An upward expander in the receiver serves to restore the original dynamic range of the audio signal. In a noise reduction system, frequency-dependent companding is used to reduce unwanted hiss and tape noise. For sound system applications, the downward expander can be used to reduce unwanted background noise when there is no program material present. A system with multiple open microphones benefits greatly from downward expansion.

A downward expander with a ratio setting of ∞: 1 becomes a noise gate. (See Figure 1-16.) When signal level drops below the threshold, the output is essentially turned off (or "gated"), also preventing build-up of undesired noises. The audible effect of a noise gate can be somewhat more disturbing than a downward expander, since the transition to the "off" state is more abrupt, audibly similar to manually muting or un-muting an audio channel. The downward expander sounds more like a rapidly raised (or lowered) fader – a much less jarring transition. The terms *noise gate* and *expander* are often used interchangeably, since many noise gates have an adjustable ratio rather than solely infinite attenuation. The gate circuit found in some automatic mixers allows the user to select an "off-attenuation" setting that uses a fixed amount of gain reduction, such as –15 dB or ∞ (off), rather than a ratio.

Automatic Gain Control (Speech Leveler)

A unique case of dynamics processor, the *automatic gain control* (AGC) either adds or reduces gain, depending on the strength of the incoming signal. The term *speech leveler* more accurately describes the function of the AGC. A properly adjusted AGC should compensate for differences in level between loud and soft talkers, again fulfilling a dynamics processor's purpose of increasing audibility. A typical AGC has a hinge point. Gain is added to signals below the hinge, while signals above the hinge are reduced in level. Another way to think of the hinge is as the unity gain point, where no addition or subtraction of gain occurs. The hinge point should be set at the desired output, or target level, for the given sound source. The threshold sets the level where the AGC begins to take effect. Signals below the threshold are not processed. (See Figure 1-17.) Similar to the compressor, the attack setting adjusts the speed at which the AGC takes effect, and decay sets how long the AGC takes to release. AGCs typically use longer attack and decay times than other dynamics processors, in part to emulate the reaction time it would take for a human to make similar gain adjustments. The AGC is one of the only processors that can raise the volume of the sound system to compensate for soft talkers. To use an AGC, the sound system must have high enough gain-before-feedback to accommodate the maximum gain setting of the AGC.

Applications Tip: *Use AGC to compensate for different talkers.*

Automatic gain controllers tend to work best when used to compensate for differences in level from various talkers, rather than from a single talker. Attempting to level a single talker requires relatively short attack, hold, and release times to create a noticeable effect. These shorter times can lead to undesirable pumping and breathing in the audio as the AGC continuously raises and lowers the gain to keep up with the rapidly changing levels characteristic of speech signals. When used to make gain adjustments for different talkers, an AGC with longer attack, hold, and release times results in smoother transitions and less false triggering.

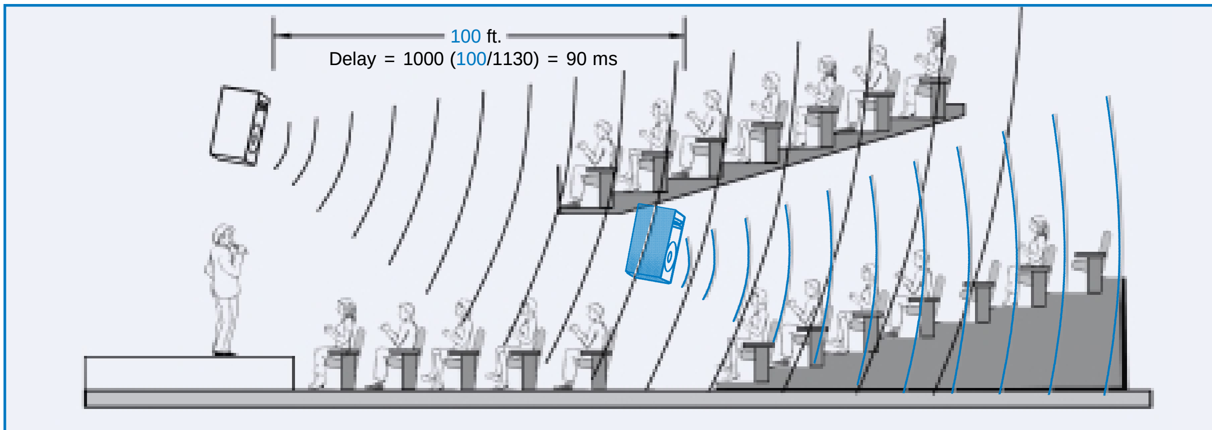

Figure 1-18: under balcony loudspeaker delayed to arrive with main loudspeaker

DELAY

A third type of audio signal processor works in the time domain, by delaying the incoming audio signal by some user-defined amount. The primary function of a *delay* unit in sound systems is loudspeaker alignment, either to align drivers within the main loudspeaker array or align remote speakers with the main PA. Within a given loudspeaker, the individual drivers are often physically offset, causing phase anomalies due to the differences in time arrival from the drivers. In a sound system where every driver is mounted in its own cabinet, this problem can be corrected by moving the boxes until the drivers are aligned. In most two- or three-way loudspeakers, the drivers cannot be moved. A few milliseconds of delay applied to the "forward-mounted" driver are usually sufficient to restore proper alignment. Note that this method of alignment requires a bi-amplified system with an active crossover, since the signal for each individual driver must be delayed independently.

In larger sound systems, delayed loudspeakers are used to provide additional coverage to remote areas. (See Figure 1-18.) Larger houses of worship and theaters will often have loudspeakers mounted above or under balconies. Outdoor concerts sometimes use delay towers. Since the distance between the main PA system (which is typically mounted on or near the stage) and the remote loudspeaker is significant, the signal sent to the remote loudspeaker must be delayed. Without delay, the audience will experience a degradation of sound quality that, depending on the distances involved, could range from comb filtering to an audible echo. Use the following formula to determine the proper amount of delay:

Delay (milliseconds) = 1000(D (feet)/1130)

The speed of sound varies with environmental conditions, but 1130 feet per second is commonly used in calculations. If D = 100 feet, the required delay is 90 ms. Delaying the signal by an additional 10 ms or so may help increase the perception that the sound is originating from the stage and not the remote loudspeaker. This approach takes advantage of the *precedence effect*, a psychoacoustic phenomenon in which listeners perceive sound as coming from the direction of the first sound arrival, even if it is somewhat lower in level than a sound that arrives a short time later. Keep in mind that air temperature, humidity, and elevation above sea level all have an effect on the speed of sound in air. Delay times may need to be adjusted by a few milliseconds to compensate. DSP-delays can usually calculate delay times if the required distance is known, and most algorithms are able to take air temperature into consideration. In general, the speed of sound increases as the temperature rises.

ADAPTIVE AUDIO PROCESSORS

Adaptive audio processors perform real-time, automated functions to optimize sound systems, ideally without the intervention of an operator. Three of the most commonly employed adaptive processors are automatic microphone mixers, feedback reducers, and acoustic echo cancellers.

Automatic Microphone Mixers

Automatic microphone mixers, also known as *voice-activated* or *sound-activated microphone mixers*, have one fundamental function: to attenuate (or reduce in level) any microphone that is not being addressed, and conversely, to rapidly activate any microphone that is addressed by a talker. The operation of a well-designed automatic mixer should be transparent to the audience of the sound system. In general, any speech sound reinforcement system that uses four or more microphones should employ an automatic mixer. To fully understand the advantages of an automatic mixer, it is necessary to examine in some detail the audio problems caused by multiple open microphones. These problems are:

1. Excessive pickup of background noise and reverberation
2. Reduced gain-before-feedback
3. Comb filtering

Chapter 1

Types of Audio Processors

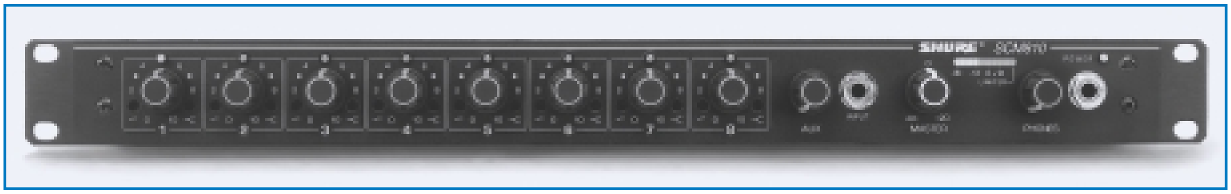

Shure SCM810 Automatic Microphone Mixer

The first problem of multiple open microphones is the excessive pickup of background noise, which adversely affects the audio quality of the sound system. Consider a city council with eight members and eight microphones. For this example, only one member is talking at a time. If all eight microphones are open when only one microphone is needed, the audio output will contain the background noise and reverberation of all eight microphones. This means the audio signal will contain substantially more background noise and reverberation than if only the talker's microphone was open. Speech clarity and intelligibility always suffer as background noise and reverberation increase. In the city council example, the audio output of eight open microphones would contain 9 dB more background noise and reverberation than a single open microphone. To the human ear, the noise would sound roughly twice as loud when all eight microphones were open.

In addition to only activating microphones that are being addressed, an automatic mixer uses a NOMA (Number of Open Microphones Attenuator) circuit, or equivalent, to help minimize the build-up of background noise and reverberation. This circuit proportionally reduces the overall output of the mixer whenever the number of open microphones increases. A well-designed automatic mixer maintains a consistent level of background noise and reverberation, regardless of how many or few microphones are active.

The NOMA circuit also plays a major role in controlling the second major problem with multiple open microphones, reduced gain-before-feedback. Acoustic feedback, characterized by an obnoxious howling or screeching sound, can be a problem with any sound system using microphones. Most sound systems are operated below the point where feedback occurs. The margin for stable (feedback-free) operation reduces every time another microphone is opened. Each doubling of the number of open microphones results in 3 dB less gain-before-feedback. Open one open microphone too many, and feedback occurs. By keeping the overall system gain constant no matter how many microphones are open, the NOMA circuit helps prevent feedback. Assuming all microphones are equidistant from loudspeakers, an automatic mixer ensures that if there is no feedback with one microphone open, then there will not be any feedback even if all the microphones are open.

Comb filtering is phase cancellation that occurs when a single sound source is picked up by more than one microphone at different distances from the source, and those signals are combined at the mixer. Since sound travels at a finite speed, the resultant frequency response of the combined microphone signals is considerably different from that of a single microphone. The frequency response chart of the combined signals resembles the teeth of a hair comb, thus the name. (See Figure 1-19.) The aural result sounds hollow, diffuse, and thin. An automatic mixer significantly reduces comb filtering by keeping the number of open microphones to an absolute minimum. Certain models of automatic mixers further reduce comb filtering by employing a circuit that will only allow one microphone to turn on for a given sound source.

Most popular automatic mixers belong to one of two categories, either some form of gated mixer or a gain-sharing automatic mixer.

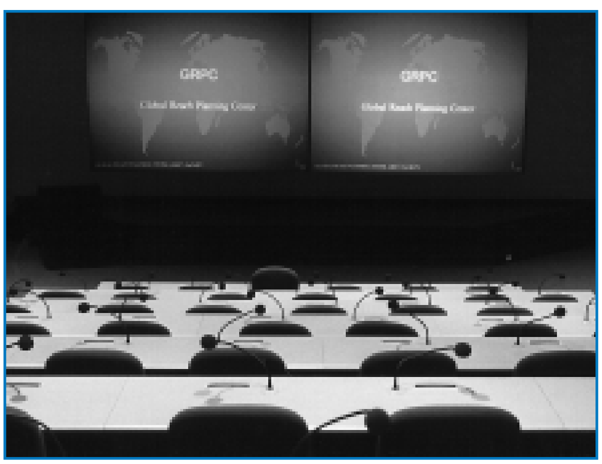

Scott Air Force Base

Summary of automatic mixer benefits

- The primary function of an automatic mixer is to keep unused microphones turned off and to instantaneously activate microphones when needed.
- Using an automatic mixer will:
  - Improve gain before feedback
  - Reduce audio degradation caused by superfluous open microphones
  - Control the build-up of background noise
- Keeping the number of open microphones to a minimum always improves overall audio and quality
- The additional control circuitry on automatic mixers provide a variety of additional functions like:
  - Audio privacy switches
  - Chairperson control of all microphones
  - Illuminated indicators of microphone status
  - Automatic video camera selection based on microphone activation

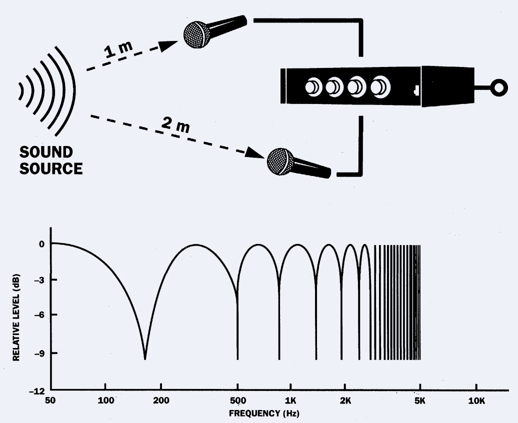

Figure 1-19: multi-mic comb filtering

Gated Automatic Mixers

The most basic form of automatic mixer functions as essentially a multi-channel noise gate. When the input signal surpasses a fixed threshold of some level, the channel activates. The input is attenuated when the level drops below the threshold. These mixers tend to either clip desired speech if the threshold is set too high, or trigger on undesired sounds if the threshold is set too low. Some designs only allow one talker at a time to prevent multiple microphones from gating on for a single source. A variable-threshold automatic mixer attempts to rectify these problems by using the signal from a remote microphone to provide a reference signal for setting the threshold. The desired talker must exceed this level by some preset amount to activate the channel. The remote microphone must be located such that it will not detect the program material, but only variations of room noise and reverberation. These levels, however, may not be identical to those at the talker's location. If the background noise levels are louder than those at the talker, the talker may not be speaking loudly enough to activate the channel. Some models use the sum of the outputs all the microphones to derive a threshold, rather than a remote microphone. This approach can work very well, because background noise is measured at the talker's location, and the talker will have a natural tendency to talk above the ambient level.

The noise-adaptive threshold automatic mixer employs a dynamic threshold unique to each input channel. Each input channel sets its own minimal threshold that continually changes over several seconds, based on variations in the microphone input signal. Sound that is constant in frequency and amplitude, like a ventilation fan, will not activate an input but will add to the noise-adaptive threshold. Sound that is rapidly changing in frequency and amplitude, like speech, will activate an input.

The mixer activates an input based on two criteria:

1. The instantaneous input signal level from the talker is greater than the channel's noise-adaptive threshold.
2. The input channel has the highest signal level for that talker.

This second criterion ensures that a very loud talker only activates one channel at a time.

Mixers of this type usually require a "last mic on" feature that keeps at least one microphone activated at all times to maintain a consistent level of background noise. A NOMA circuit is essential to keep the overall mixer output below the feedback threshold. An automatic mixer without NOMA is really nothing more than a multi-channel noise gate.

It should be noted that automatic mixers do not "mix" in the traditional sense. The gain adjustments made to individual channels are not continuously variable, but simply a transition from an "on" to an "off" state. Any balancing of signal levels must be accomplished by either a human operator or, to a limited extent, a dynamics processor such as an AGC. Consequently, automatic microphone mixers are not recommended for musical applications. Mixing for music is as much an art as a science, and artistic decisions are best left to a human being. Additionally, automatic mixers that use noise-adaptive threshold technology may be confused by musical program material, causing erratic gating.

Most automatic mixers share many of the same controls as manual mixers, including individual level adjustments, phantom power, basic equalization, etc. Several functions unique to automatic mixers are detailed below:

Input Channel Threshold: Determines the signal level where the mixer will pass the incoming microphone signal to the mixer's output.

Last Microphone Lock-On: Keeps the most recently activated microphone on until another channel is activated, maintaining room ambience when the automatic mixer is used to provide a feed for broadcast, recording, or to an assistive listening system.

Hold Time: Specifies the amount of time a channel stays activated after speech has ceased. The feature prevents the channel from gating off during the natural gaps that occur in speech patterns.

Off-Attenuation: Determines how much gain reduction is applied to an input channel when the channel is not active. The range of adjustment can vary from 3 dB to 70 dB, but 15 dB is a common value. Some mixers allow for a setting of infinity, or a true "off" setting.

Decay Time: Establishes the time required for an input to be lowered from the activated state to the attenuated state. As in a dynamics processor, decay time is always in addition to hold time.

CHAPTER 1
Types of Audio Processors

Gain Sharing Automatic Mixers

A gain-sharing automatic microphone mixer works from the premise that the sum of all the signal inputs from all microphones in the system must be below some maximum gain value that avoids feedback. The mixer maintains this level by distributing a constant total gain among the inputs, based on their relative levels. If nobody is speaking, the total available gain in the mixer is distributed equally to each input. When one person speaks, that channel has more signal than the others. Consequently, the mixer allocates more gain to that channel, and less gain to the others, roughly in proportion to the relative increase in signal level. The total gain of the system is the same as when no one is speaking.

For example, a 3 dB level increase at one microphone causes that channel gain to rise by 3 dB, while the gain of the other channels decreases by a total of 3 dB. When two talkers speak into separate microphones with levels that differ by 3 dB, they appear at the output of the system with a 6 dB difference. The microphone with the highest signal is given the most gain, while the microphone with the lowest signal is given the least. Since a gain-sharing automatic mixer increases the level difference between microphones, the key to transparent operation is fast action to prevent interruptions and overlaps in speech. Again, mixers of this type are not appropriate for music applications, where microphone signal levels should be balanced equally. Finally, microphones in this system are never turned "off", negating the need for last microphone hold or one-mic-per-talker circuits.

Feedback Reducers

As discussed previously, equalizers can be powerful tools for minimizing feedback problems in a sound system. The proper use of an equalizer for feedback control, however, requires a skilled operator with either a well-trained ear for identifying feedback frequencies or analysis tools to identify the problems. A *feedback reducer* (eliminator, suppressor, destroyer) accomplishes the same function automatically. These devices are basically adaptive equalizers. The equalizer employs a digital algorithm that can identify the characteristic build-up of a particular frequency that is feeding back, and places an extremely narrow filter at that frequency. The bandwidth of a feedback reducer filter typically ranges from 1/10 to 1/70 of an octave. The depth of the filter is usually dependent on the level of the offending frequency. Most feedback reducers will only cut the frequency as much as necessary. It is usually desirable that the filter width narrow as the depth increases, to prevent unwanted attenuation of adjacent frequencies. An effective feedback reducer should react quickly, with negligible effect on the overall sound quality of the audio system. The net effect of the feedback reducer should be to flatten the overall system response by using adaptive filters to reduce peaks. (See Figure 1-20.)

Audible feedback must occur before the reducer can perform its task, hence, these devices are not "pre-emptive." A feedback reducer does not anticipate feedback, but reacts accordingly once feedback is detected. The faster the frequency detector algorithm works in a feedback reducer, the less chance that the audience will be annoyed by feedback. The speed of feedback detection is frequency-dependent, as well. For the detector to properly identify the

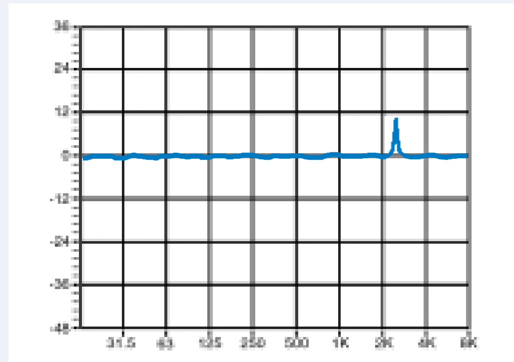

frequency response peak

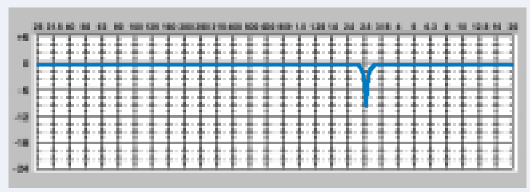

feedback reduction filter

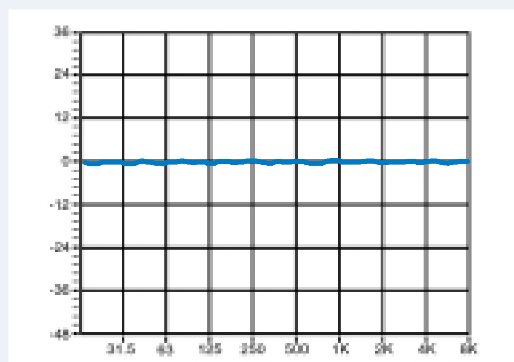

Figure 1-20: corrected frequency response

DFRs in rack

frequency, the sound wave must complete several cycles. The longer wavelengths associated with lower frequencies take more time to complete a cycle. Therefore, lower frequency feedback takes longer for the detector to properly identify. Assume that two frequencies begin to ring, one at 500 Hz and one 5000 Hz. A 500 Hz wave completes a full cycle in 1/500th of a second, or 2ms. The 5000 Hz wave will complete a cycle in 1/5000th of a second, or .2 ms. A feedback reducer should be able to identify the 5000 Hz feedback tone 10 times faster than the 500 Hz tone.

More importantly, feedback reducers are subject to the same limitations as manual equalizers. Foremost among these limitations, a feedback reducer cannot cause the sound system to achieve more gain-before-feedback than the levels dictated by the PAG equation (Appendix 2). Remember that adaptive equalization attempts to flatten frequency response anomalies in the sound system. Once this has occurred, no further benefit is possible from equalization. A feedback reducer can provide a maximum of 6 to 10 dB of additional gain. A feedback reducer is not a substitute for poor system design. Proper choice and placement of loudspeakers and microphones must be the first priority.

Applications Tip:* *Not Enough Feedback Filters?

Feedback reducers cannot deploy an unlimited number of filters. The number of possible filters is limited by the DSP capabilities of the device. Increasing DSP power makes it possible to deploy more filters, but if more than 10 filters are required, other problems with the sound system may need to be addressed. Instead of getting a feedback reducer that has more filters, investigate other alternatives to reducing feedback. (See Section 2.1.)

Acoustic Echo Cancellers

Echo cancellers reduce residual echo return in audio conferencing (teleconferencing) installations. Possible sources of echo in a teleconferencing system include: improperly balanced hybrids, signal coupling within the telephone lines, and satellite transmission links with long propagation delays. These types of echoes are electronic in nature and can be reduced by a line echo canceller. Acoustic echo occurs when audio received from the remote site reaches active microphones at the near site, and is transmitted back to the remote site with sound from both the near site talkers and acoustic echoes of the sound that originated at the remote site. This type of echo requires an acoustic echo canceller (AEC).

An AEC monitors the incoming audio signal from remote sites, and compares it to the signal that is about to be transmitted. If the echo canceller detects the presence of the incoming audio in the outgoing signal, it attempts to remove it electronically from the outgoing signal. This reduces the amount of echo, but does not completely "cancel" it. Notice that the echo canceller attempts to prevent the incoming audio from other sites from being sent back to them, but it does not do anything about echoes that other sites may be sending to your site. The AEC only improves audio for the remote site, not the one where the unit is installed. (See Figure 1-21.) Therefore, if one site on a network requires an echo canceller, all of the sites will probably need one.

More powerful processors and advanced cancellation algorithms have resulted in acoustic echo cancellers that are better and less expensive. It is now possible to have a separate echo canceller for each microphone input channel, which provides optimum echo reduction.

Acoustic echo cancellers are commonly believed to be capable of removing the hollow sound associated with a room that is too reflective. In fact, excess reflective sound makes it difficult for the echo canceller to work properly, and reduces its effectiveness. As with all audio processing, a room outfitted with proper acoustic treatment should come before attempts to fix a problem electronically.

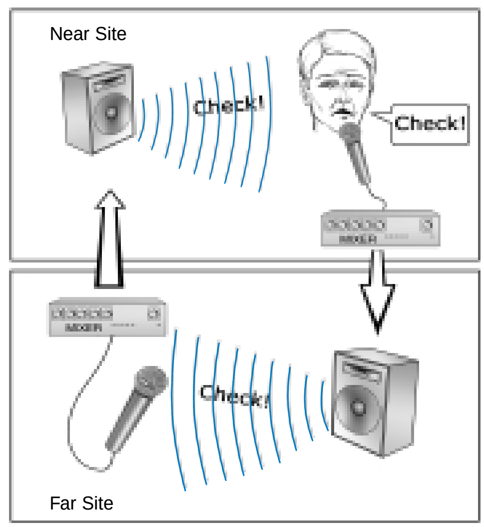

Teleconference system without AGC

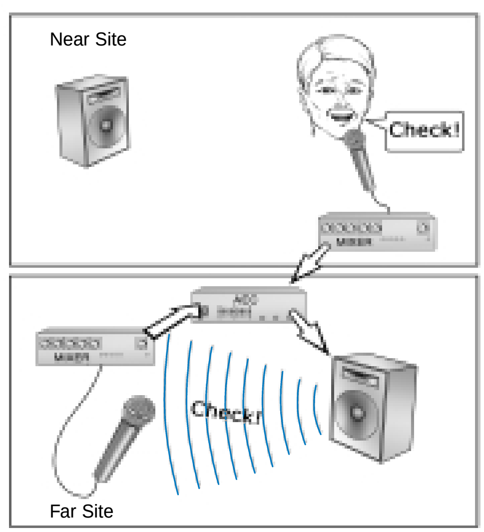

Teleconference system with AGC at far site

Figure 1-21

CHAPTER 2

Practical Applications for Audio Signal Processors

MAXIMIZING GAIN-BEFORE-FEEDBACK

If a sound reinforcement system cannot produce enough sound level at the audience position before it starts to feed back, intelligible audio and balanced frequency response are next to impossible. The first and most basic function of a sound system is to provide enough sound level at the audience position so that they can hear the performers at a comfortable level above the room's ambient noise.

Feedback occurs when the amplified sound from any loudspeaker reenters the sound system through any open microphone and is amplified again and again and again This is a physical phenomenon and audio processors cannot help a sound system obtain gain-before-feedback beyond the limits of physics. Most sound systems do not operate near their physical limit, yet they still experience feedback problems. Potential Acoustic Gain is the amount of level (in dB) that a sound system can produce before the onset of feedback.

The following list highlights the only possible solutions to feedback problems and what kind of improvement to expect. Note that some audio processors, like automatic mixers and feedback reducers, can help a system achieve the maximum amount of gain before feedback by optimizing some variables in the equation. Realize that the most helpful, yet least expensive options do not even involve an audio processor.

See Appendix 2 for a complete mathematical discussion of the Potential Acoustic Gain equation.

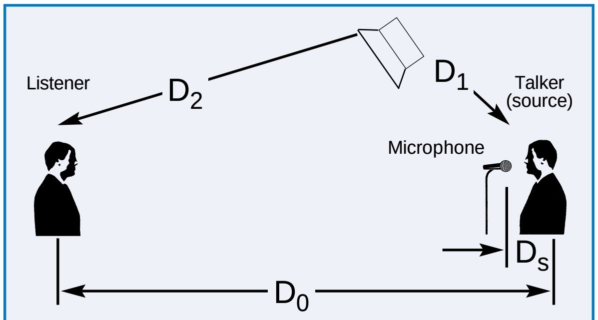

Figure 2-1: Minimize D_s and D_2

1. Move the microphone closer to the talker. This is the easiest, most effective, and sometimes most economical way of improving gain-before-feedback. Expect an increase in gain-before-feedback from 6 to 25 dB. Moving the microphone from a distance of 8 inches to 4 inches away from the talker provides a 6 dB increase in gain. Switching from lavalier microphones to headset microphones changes the distance to the talker from approximately 12 inches to less than one inch, which provides a 24 dB increase in potential acoustic gain.

2. Move the loudspeaker closer to the listener (i.e. away from the talker) or add a second loudspeaker for the rear part of the room. Expect an increase in gain-before-feedback from 3 to 15 dB. Installing a second loudspeaker set or satellite loudspeakers to provide sound coverage for the rear of the room allows the front of house (FOH) loudspeakers to be turned down, as they no longer have to project sound all the way to the back of the room. The second set of loudspeakers in the rear of the room effectively brings the loudspeaker much closer to the listener providing more gain-before-feedback. Every time the distance between the loudspeaker and the listener is cut in half, the potential acoustic gain of the system increases by 6 dB.

3. Reduce the number of open microphones. Expect an increase in gain-before-feedback from 3 to 12 dB. Every time the number of open microphones in a sound system doubles, the potential acoustic gain of the system is reduced by 3 dB. If 2 new microphones are turned on in a system that previously had 2 open microphones, the system will have to be turned down by 3 dB to maintain the same feedback stability margin. Adding more microphones can be a solution to feedback problems only if the microphones are being placed much closer to the talker than they were previously. For example, a few overhead, hanging microphones can be replaced with many lavalier microphones. In general, double the number of microphones only if the distance from the microphone to the talker is reduced by at least half. This should result in a minimum of a 3 dB increase in PAG with better coverage of the desired pick up area. Automatic microphone mixers greatly help sound systems with 4 or more open microphones by keeping microphones that are not being addressed turned down. This effectively reduces the number of open microphones and increases the potential acoustic gain of the system.

4. Use unidirectional microphones and loudspeakers. Expect an increase in gain-before-feedback from 3 to 8 dB. The proper use of unidirectional microphones, such as those with cardioid or supercardioid polar patterns, can help pick up more of the desired sources (the talkers) and less of the undesired sources (in this case, loudspeakers). They also help reject other undesired sources such as ambient noise and the room reverberation.

Loudspeakers with high directionality or narrow dispersion patterns are also available and can improve gain-before-feedback as well as intelligibility. They accomplish this by directing most of the sound energy to the audience. In doing so, less energy is sent to the reverberant field or back to the stage to cause feedback problems. The latter approach usually requires a complete redesign of the sound system.

5. Move the loudspeaker further from the microphones. Expect an increase in gain-before-feedback from 2 to 9 dB. Doubling the distance between the microphone and the closest loudspeaker results in a 6 dB increase in gain-before-feedback. However, moving the loudspeaker to twice the distance from the microphones is a less realistic option than the ones previously discussed. This approach usually results in inappropriate coverage for the audience in the front of the room, and there may be a space limitation that does not permit moving the loudspeakers.

If stage monitors are being employed, they should be used only for the monitoring of music, effects, cues, and playback. The signal from lavalier, boundary or overhead microphones on stage intended to pick up the performer's voice should never be routed to these monitors as it will severely handicap the amount of level the system can provide for the audience. If performers on stage need to hear each other or themselves, they must wear in ear monitors or consider using handheld, headset, or other microphone designs where the distance between the talker's mouth and the microphone capsule is extremely small (less than 1 inch.)

6. Reduce gain at feedback frequencies using notch filters (narrow equalizer filters). Expect an increase in gain-before-feedback from 3 to 9 dB. Narrow peaks in the overall frequency response of the sound system are the first to cause feedback. These peaks rise 2 to 10 dB above the overall system response and prevent the system from reaching its maximum potential acoustic gain. (See Figure 2-2.) This technique can be done with a manual equalizer (and the appropriate measurement tools) or a feedback reducer. A digital feedback reducer can detect feedback and insert a notch filter at the exact offending frequency, which effectively flattens the sound system's frequency response and allows it to reach its maximum possible gain-before-feedback.

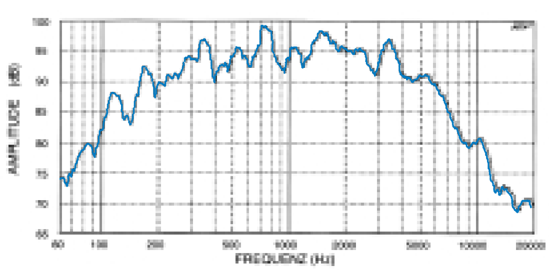

Figure 2-2: overall system response with peaks

7. Improve room acoustics with acoustic treatment. Expect an increase in gain-before-feedback from 3 to 6 dB. An acoustical consultant, working in conjunction with the architect, should design a venue with good acoustics before construction of the building even begins. Once the room has been built with inappropriate geometry, it is very difficult to fix acoustical problems. Covering the walls, floor, and ceiling with sound absorbent materials is at best a fair solution. While more expensive than all the other options discussed above, it can help reduce problems like long reverberation or decay times that affect intelligibility, and standing waves and reflections that affect gain-before-feedback and system frequency response. Keep in mind that, as a rule of thumb, to make a noticeable change you must treat 50% of the room's surfaces with sound absorbent materials. In some cases, a single wall or surface, such as the back wall, could be causing most of the feedback problems. Treating this surface alone could produce a good, noticeable improvement in gain-before-feedback, even though it will not dramatically improve intelligibility or reduce reverberation time.

There are no other solutions! These guidelines provide you with the ONLY options available to increase the potential acoustic gain of a sound system.

IMPROVING SPEECH INTELLIGIBILITY

Speech intelligibility is among the most difficult goals to achieve for any medium-to-large indoor sound reinforcement system. Some of the factors that play major roles in obtaining good intelligibility are beyond the control of the sound system or signal processing. These factors include the acoustic characteristics of the space (in particular its geometrical shape), its background noise level, and its reverberation time.

There are two basic ways that audio processing can improve the speech intelligibility of a sound system. The first, and most effective, is by reducing the number of open microphones. This approach involves using an automatic mixer to keep the microphones of participants who are not talking turned down. The second method employs an equalizer to limit the frequency response of speech microphones to the speech frequency range only. A bandpass equalizer is typically the appropriate tool for this job.

Reducing the Number of Open Microphones:

Automatic microphone mixers are typically the easiest audio processor to implement, since most designs require very little setup on the part of the user. In the majority of applications, microphones are attached directly to the mixer. Common applications include boardrooms, courtrooms, houses of worship, and theater.

Boardrooms/Meeting Rooms/Council Chambers: Any meeting facility that uses more than three microphones should consider an automatic microphone mixer. Even if the talkers are using push-to-talk microphones to keep the number of open microphones to a minimum, they often forget to activate the microphone,

CHAPTER 2

Practical Applications for Audio Signal Processors

leading to missed speech. Or, in the case of push-to-mute microphones, the delegate forgets to turn the microphone off. Momentary mute (or cough) switches are usually desirable, since the automatic mixer cannot distinguish between a cough and speech. A mixer with microphone logic capabilities can provide additional functionality for chairman microphone override, remote LED indication, and automatic camera-switching.

Houses of Worship: As above, use an automatic mixer if there are more than three microphones. Additionally, for worship leaders who use a lavalier microphone as well as a gooseneck microphone at the lectern, the automatic mixer will only activate one of the microphones, preventing comb filtering. The same applies to lecterns with two microphones. While logic dictates that two microphones provide better coverage for roaming talkers, the trade-off in comb filtering often creates more problems than it solves. Using an automatic mixer prevents comb filtering while providing a wider coverage area.

As mentioned previously, automatic microphone mixers are not recommended for music sources. Since most house of worship applications combine music and speech, both a manual and an automatic mixer should be used. The simplest setup could use the automatic mixer to submix the speech microphones into one channel of the manual mixer. Alternately, if using a sound system processor that has a matrix mixer, the outputs of the automatic mixer and manual mixer can be combined and routed by the processor. Either way, speech and music sources are handled independently. If the application only has an automatic mixer, use the logic functions to "force" the music microphones on so they will not mute. Note that for mixers with a NOMA circuit, this approach will reduce the output of the mixer, and any additional noise picked up by the music microphones will always be present unless muted by a human operator or traditional noise gate.

Theater: In theater applications, where the sound system operator requires complete control over the performer's audio, the preferred way of employing automatic microphone mixers is in the form of speech gates. In this scenario, the automatic mixer is connected to the mixing console on a per-channel basis via the insert jacks for each input channel. The operator has full control of each microphone's level when it is in use and retains all the functionality of the mixing console. The automatic mixer keeps only the microphones of performers that are talking turned up.

Equalizing for Speech Intelligibility:

Using equalization in sound reinforcement takes on two forms: the objective and the subjective. Objective equalization entails the use of corrective equalization to compensate for frequency response anomalies in the sound system components and room resonances that cannot (for financial or logistical reasons) be cured by acoustical means. Proper objective equalization requires the use of measurement devices to obtain a theoretically flat frequency response. Flat frequency response, while desirable as a good starting point, may not produce the most audibly pleasing result. Here is where subjective EQ enters the picture. Subjective equalization is more art than science, and requires a skilled operator with a trained ear to obtain optimal results. "Sounds good" cannot necessarily be quantified in measurable terms. However, some general guidelines can help with regard to enhancing intelligibility.

Reproducing intelligible speech demands a minimal frequency response from a sound system equal to that of a telephone system - about 300 Hz to 3 kHz. A wider frequency response can enhance the tonal quality of the reproduction but can also degrade intelligibility by emphasizing pops, rumble, hiss, room acoustics, and other noises that are extraneous to speech and would not be present in a normal conversation. Wider frequency response also permits more sound energy to unnecessarily contribute to the reverberant field of the room. This makes the system more prone to feedback and less intelligible.

Equalization can noticeably, but not dramatically, improve the naturalness or intelligibility of a sound reinforcement system by emphasizing the frequency ranges most critical for speech.

Equalization cannot make a poorly designed sound system work satisfactorily or improve intelligibility problems caused by reflections, mechanical vibration, and high background noise levels. It cannot improve intelligibility problems caused by the talker being too far from the microphone, improve the performance of substandard audio components, or eliminate distortion and noise problems caused by mismatched audio levels between system components.

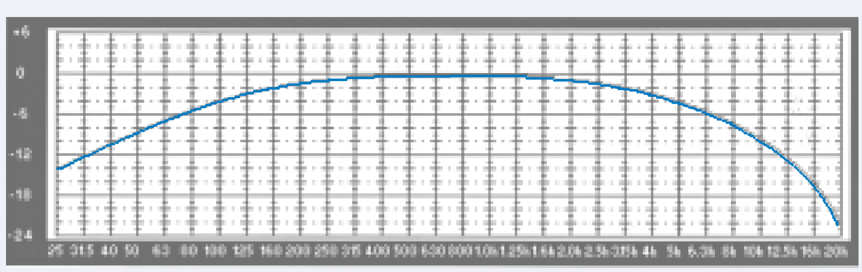

Figure 2-3: speech EQ curve

A hi-cut/low-cut (or band pass) equalizer is the most basic tool needed to equalize speech microphones for optimum intelligibility. Perception research and studies of human hearing suggest the following EQ curve as a good starting point. It maintains good, natural voice tonality while attenuating all unnecessary frequencies.

- **Low Cut Filter (LC) set to 125 Hz, 6 dB per octave.**
- **High Cut Filter (HC) set to 4000 Hz, 6 dB per octave. (See Figure 2-3.)**

Increasing the response bandwidth, for example from 80 Hz to 8000 Hz, would provide a slight improvement in tonal quality. Decreasing the bandwidth slightly from the low end should improve intelligibility. The minimum response should never be narrower than 400 Hz to 2.5 kHz and the filter slopes should not exceed 12 dB per octave. Note that the human voice contains very little energy below 100 Hz. While adding response below this point may sound impressive, the effect on intelligibility is more detrimental than helpful.

In addition to bandpass filters, a parametric equalizer can be used to boost a selective frequency range. Using a parametric filter to help intelligibility is mostly an experimental exercise and the exact frequency, bandwidth, and boost will vary from system to system. The idea is to boost a set of frequencies that are most essential to speech to overcome interference from the acoustical environment. This frequency is typically between 1 and 4kHz. The typical boost is 3 to 5 dB. The width of the filter can vary from 1/6 octave to 1 octave.

In general, approach equalization slowly. After every adjustment, listen carefully to the resulting sound. Most changes are not perceived as good sounding immediately. Listen for at least 3 minutes to each change to allow your ear to adapt. If the equalizer has a bypass button, use it often to provide a reference point. When the system is clear enough, stop equalizing.

When listening to live microphones, have someone else talk, never try to equalize to your own voice. When using recorded material to equalize, choose a recording that you are familiar with and have listened to many times in different sound systems.

SOUND SYSTEM GAIN STRUCTURE

Setting gain structure in a sound system concerns the proper calibration of signal levels between devices in the audio chain to achieve good signal-to-noise ratio and adequate headroom. Poor signal-to-noise ratio results in a high level of background noise (hiss) that, at best, is annoying for the listener, and, at worst, obscures intelligibility. Objectionable background noise usually results in a system with excessive headroom, where the desired audio signal level is close to the noise floor. In contrast, low headroom, where system noise is quiet but the audio signal is close to clipping, can lead to overload conditions that could cause distortion or loudspeaker failure. If every piece of audio equipment clipped (started to audibly distort) at the same level and had a similar dynamic range, then audio systems would be "plug-and-play." Unfortunately, this is not the case. (See Figure 2-4.)

Novice sound technicians commonly mistake the input sensitivity control on a power amplifier for a "volume" knob, often rotating the control to maximum in an attempt to get the highest possible level out of the sound system. Unfortunately, the end result is usually additional noise. The input sensitivity knob should be set just high enough to ensure maximum output from the amplifier. This point is determined by the setting at which the amplifier input sensitivity indicators begin to show clipping. Any additional boost beyond this point only adds noise. Maintaining the highest possible signal levels throughout the various components of the sound system in the easiest way to realize maximum output with minimal noise. If the power amplifier controls are indiscriminately placed at maximum, the sound technician must operate the mixer and other audio components in the signal chain at lower levels. Consequently, the program material is close in level to the noise floor of the mixer. Using the amplifier's input sensitivity control to compensate for low levels from the mixer only exacerbates the noise problem by raising the noise floor of the mixer as well as the program material. If sound levels in the room are too loud, the input sensitivity of the amplifier, rather than the level control on the mixer, should be reduced to maintain good signal-to-noise. In any case, amplifiers should be turned down, or off, until good gain structure is achieved in all components prior to the amplifiers.

This section introduces two methods of setting system gain structure, the unity method and the optimized method. Both methods rely on strong signal levels throughout the sound system, but differ in approach.

Chapter 2

Practical Applications for Audio Signal Processors

Clip	+24 dBu	+20 dBu	+18 dBu
Headroom	20 dB	16 dB	14 dB
Output level	+4 dBu	+4 dBu	+4 dBu

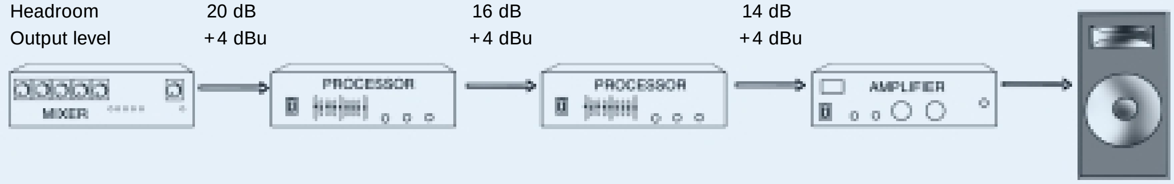

Figure 2-5: gain structure, unity method

The Unity Method

The historically conventional way to set sound system gain structure, this method relies on unity amplification, that is, every component after the mixer should produce an output voltage equal to the voltage at its input. If we assume typical line level, +4 dBu, each device in the system should be calibrated to produce this level at its output, ultimately resulting in +4 dBu at the amplifier input. The amplifier's input sensitivity control is used to set the desired sound level in the room. Advantages to this approach include:

1. Easy calibration
2. Easy to substitute components
3. Fast implementation

However, there are several significant disadvantages to the unity method. While operating levels throughout the system are consistent, headroom is not. The likelihood of clipping components post-mixer is the single biggest drawback. Consider a mixer with an output clipping level of +24 dBu. (See Figure 2-5.) Assuming that mixing at meter "0" produces +4 dBu output level, the mixer has 20 dB of headroom. If the output of the mixer is connected to an equalizer with a clipping level of +20 dBu, the equalizer only has 16 dB of headroom. Therefore, a waveform that contains transients well within the headroom of the mixer could potentially cause distortion at the equalizer. Mixing below meter "0" results in lower output voltage, which could help maintain 20 dB of headroom, but most likely will prove confusing for system operators unfamiliar with this sound system. Optimally, all components in a system should clip at the same point.

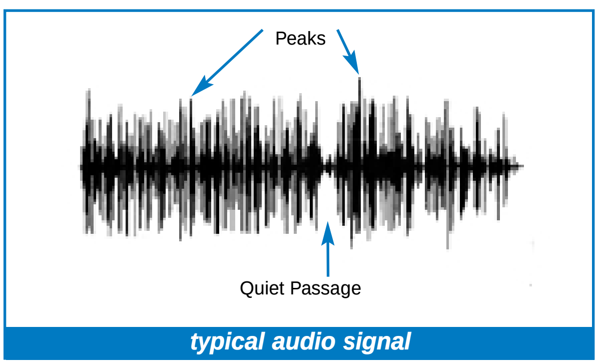

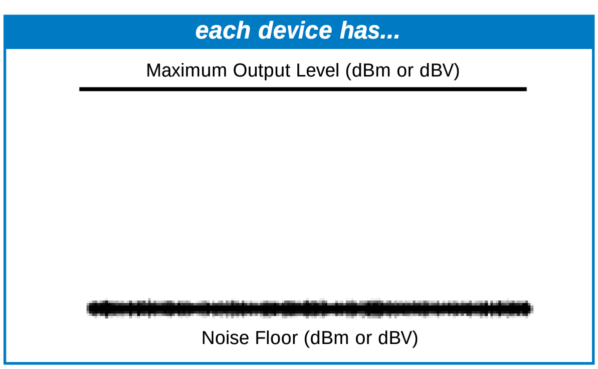

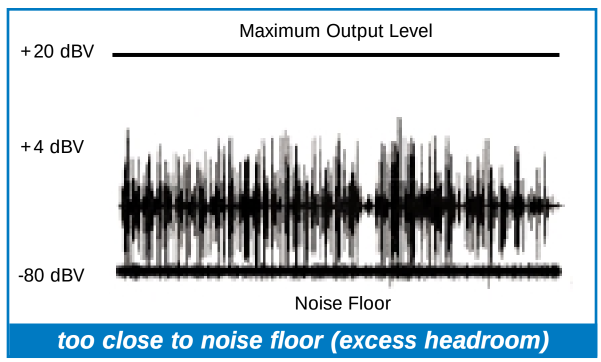

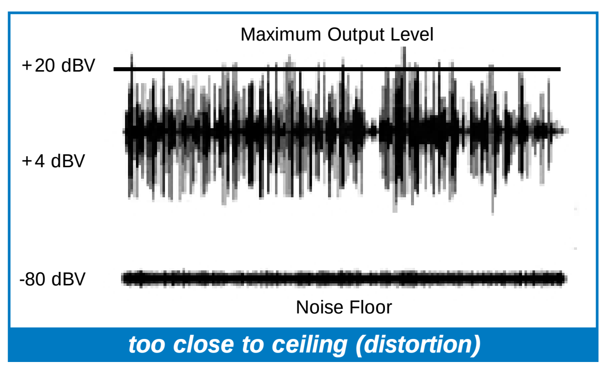

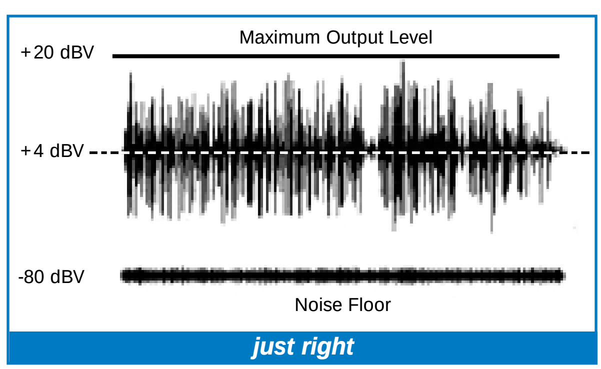

Figure 2-4

Clip	+24 dBu	+20 dBu	+18 dBu
Headroom	20 dB	20 dB	20 dB
Output level	+4 dBu	0 dBu	-2 dBu

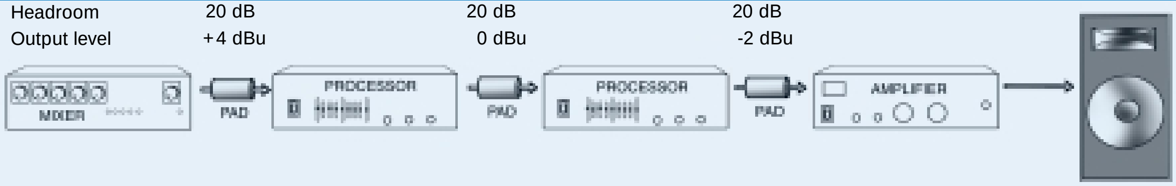

Figure 2-6: gain structure, optimized method

The Optimized Method

Establishing gain structure using the optimized method results in inconsistent operating level, but consistent headroom. With this approach, each device can output its maximum voltage, yet not overdrive the next component. This technique typically requires a resistive pad between components. Using the above example, the equalizer's clipping level is 4 dBu lower than the mixer. Therefore, the output signal from the mixer needs to be reduced by 4 dB before the input of the equalizer. (See Figure 2-6.) Occasionally, the attenuation can be achieved by lowering the input sensitivity control of the device. If not, a 4 dB attenuator should be placed between the mixer and the equalizer. The output signal from the mixer will be lowered to 0 dBu at the input of the equalizer, maintaining 20 dB of headroom. Advantages to the optimized method include:

1. Optimized signal-to-noise ratio throughout the system.
2. All components clip simultaneously. Mixing at meter zero results in the same headroom throughout the system.

Of course, this method requires more time and expertise on the part of the installer, and component substitution is more difficult since a replacement device may have a different clipping level.

A pad may be required before the input of the power amplifier if clipping occurs at a low gain setting. Otherwise, raise the input level control of the power amplifier until either the desired sound level is achieved for the audience, or the amplifier begins to indicate clipping. Realize that if clipping does occur before the desired sound level is achieved, a larger power amplifier (and consequently, loudspeakers that can handle the power) may be required.

DIGITAL SIGNAL PROCESSING

A digital signal processor (DSP) uses complex digital software algorithms to emulate the operation of analog signal processors in digital hardware. A DSP is nothing less than a specialized audio computer with its own operating system and software. Some models can be configured with front panel controls, but others need to be connected to a PC for setup. The latter requires a program called a Graphical User Interface (GUI) to control the DSP. (See Figure 2-7.)

While single-function DSP devices are available, the real advantage lies with multi-function devices. The majority of these products provide every type of processing required between the outputs of the mixer and the inputs of the power amplifiers and, in some cases, they can eliminate the need for a stand-alone mixer. Depending on the feature set, these devices can be classified as either loudspeaker processors or sound system processors. A loudspeaker processor tends to emphasize tools for protecting and aligning loudspeakers, such as crossovers,

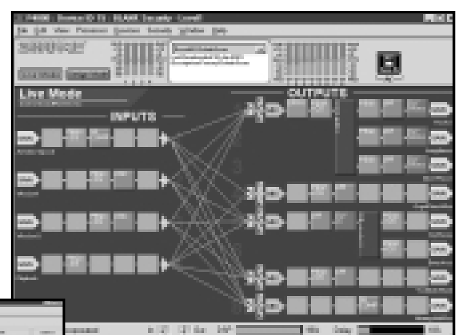

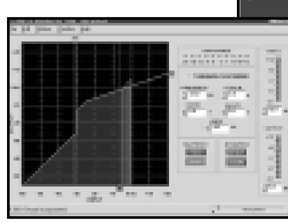

Figure 2-7: real time monitoring of Automatic Gain Control (AGC) functions

limiters, and delay. A sound system processor adds more front-end functionality, such as feedback reduction, echo cancellation, and more advanced matrix-mixing capability. Some processors even provide microphone inputs and automatic mixing. A key benefit of many DSPs is the ability to lock settings with password protection for installations in which a tamper-proof sound system is desired. Without a PC, the appropriate software, and the system password, access to parameters that could jeopardize the functionality of the sound system is eliminated. Other significant advantages to digital signal processors include:

Flexibility: While certain guidelines often dictate the order of components in the signal path, different situations may require a more flexible architecture. Some processors only provide a fixed signal path (for example, input-EQ-compressor-crossover-output). At the other extreme, some processors use a completely open architecture, where the designer is essentially given a blank page to design the sound system using a GUI that works like CAD (Computer Aided Design) drafting software. Lastly, a hybrid of the two methods offers a fixed number of place holders for processor modules, but gives the designer the ability to place the desired processing in any available place-holder, and route the signal as required.

Ease of programming: Using a computer for system setup should be intuitive and easy to learn. Hardware-based interfaces are typically more difficult to learn due to limited display area and multi-purpose controls. Adjusting a single parameter often requires searching through multiple layers of menus. Most digital processors that are programmed by computer present the user with GUI software that can make programming as simple as drawing lines or entering parameter values directly into the proper fields. The entire system layout and signal flow can be displayed on a single screen.

Work anywhere: The software for most processors does not require that the user by connected to the processor itself for design purposes. This functionality allows the installer to design the system anywhere there is access to a computer with the software, anytime it is convenient, and then load the design into the processor later. While certain parameters require on-site adjustment (such as equalization), signal flow, at the very least, can be planned in advance.

Control Options: Many digital signal processors offer control options for remote adjustment of certain processor parameters. These features are particularly useful for situations where the end-user needs some sound system control, but leaving behind a PC with the software could prove disastrous. Typical control options include preset selection, remote volume control, and remote muting of inputs or outputs.

Low noise and easy system connectivity: Gain structure is greatly simplified due to fewer physical components in the signal chain. Signal levels between functions within the processor do not need to be calibrated. Additionally, the noise floor of a single processor is significantly lower than that created by multiple devices.

Cost: A single multi-function digital signal processor typically costs far less than the equivalent amount of processing in several stand-alone devices. Also, if additional processing is required after the design phase, it is just a matter of reprogramming the software rather than re-laying out the equipment rack and purchasing another hardware device.

Time: It takes much less time to install a single DSP device compared to the time required to install, wire, and connect multiple processing components. The ease with which these processors can be programmed and implemented saves cost in installation and design time, as well.

The power and flexibility provided by digital signal processors gives sound system operators and installers all the necessary tools to provide an optimal auditory experience for the intended audience. As listener expectations continually get more and more sophisticated, a complete set of tools is required to meet those expectations: equalizers for tone shaping and feedback control, dynamics processors for increased audibility, and adaptive audio processors to automate control when possible. The combination of skilled design and proper application of the various audio processors results in superior sound quality for any venue.

APPENDIX 1

Sound

Sound is produced by vibrating objects. These include musical instruments, loudspeakers, and, of course, human vocal cords. The mechanical vibrations of these objects move the air which is immediately adjacent to them, alternately "pushing" and "pulling" the air from its resting state. Each back-and-forth vibration produces a corresponding pressure increase (compression) and pressure decrease (rarefaction) in the air. A complete pressure change, or cycle, occurs when the air pressure goes from rest, to maximum, to minimum, and back to rest again. These cyclic pressure changes travel outward from the vibrating object, forming a pattern called a sound wave. A sound wave is a series of pressure changes (cycles) moving through the air.

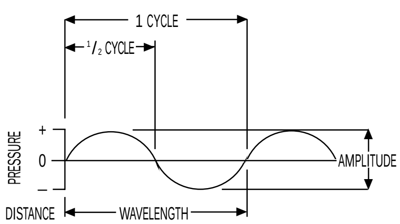

Schematic of Sound Wave

A simple sound wave can be described by its frequency and by its amplitude. The frequency of a sound wave is the rate at which the pressure changes occur. It is measured in Hertz (Hz), where 1 Hz is equal to 1 cycle-per-second. The range of frequencies audible to the human ear extends from a low of about 20 Hz to a high of about 20,000 Hz. In practice, a sound source such as a voice usually produces many frequencies simultaneously. In any such complex sound, the lowest frequency is called the fundamental and is responsible for the pitch of the sound. The higher frequencies are called harmonics and are responsible for the timbre or tone of the sound. Harmonics allow us to distinguish one source from another, such as a piano from a guitar, even when they are playing the same fundamental note. In the following chart, the solid section of each line indicates the range of fundamental frequencies, and the shaded section represents the range of the highest harmonics or overtones of the instrument.

The amplitude of a sound wave refers to the magnitude (strength) of the pressure changes and determines the "loudness" of the sound. Amplitude is measured in decibels (dB) of sound pressure level (SPL) and ranges from 0 dB SPL (the threshold of hearing), to above 120 dB SPL (the threshold of pain). The level of conversational speech is about 70 dB SPL. A change of 1 dB is about the smallest SPL difference that the human ear can detect, while 3 dB is a generally noticeable step, and an increase of 10 dB is perceived as a "doubling" of loudness.

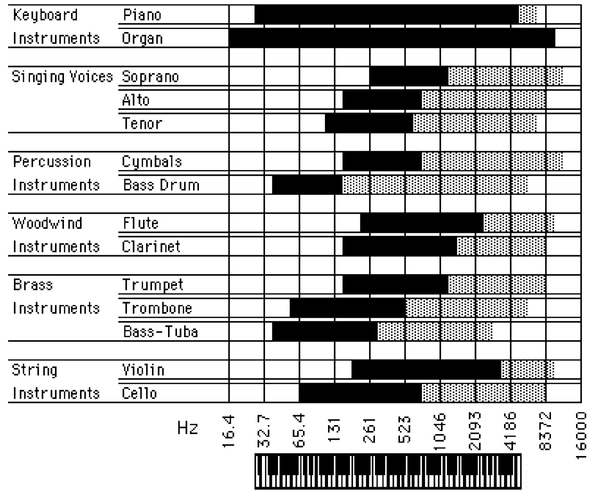

Instrument Frequency Ranges

Another characteristic of a sound wave related to frequency is wavelength. The wavelength of a sound wave is the physical distance from the start of one cycle to the start of the next cycle, as the wave moves through the air. Since each cycle is the same, the distance from any point in one cycle to the same point in the next cycle is also one wavelength: for example, the distance from one maximum pressure point to the next maximum pressure point.

Wavelength is related to frequency by the speed of sound. The speed of sound is the velocity at which a sound wave travels. The speed of sound is constant and is equal to about 1130 feet-per-second in air. It does not change with frequency or wavelength, but it is related to them in the following way: the frequency of a sound, multiplied by its wavelength always equals the speed of sound. Thus, the higher the frequency of sound, the shorter the wavelength, and the lower the frequency, the longer the wavelength. The wavelength of sound is responsible for many acoustic effects.

After it is produced, sound is transmitted through a "medium". Air is the typical medium, but sound can also be transmitted through solid or liquid materials. Generally, a sound wave will move in a straight line unless it is absorbed or reflected by physical surfaces or objects in its path. However, the transmission of the sound wave will be affected only if the size of the surface or object is

Sound Pressure Level of Typical Sources

APPENDIX 1
Sound

large compared to the wavelength of the sound. If the surface is small (compared to the wavelength) the sound will proceed as if the object were not there. High frequencies (short wavelengths) can be reflected or absorbed by small surfaces, but low frequencies (long wavelengths) can be reflected or absorbed only by very large surfaces or objects. For this reason it is easier to control high frequencies by acoustic means, while low frequency control requires massive (and expensive) techniques.

Once a sound has been produced and transmitted, it is received by the ear and, of course, by microphones. In the ear, the arriving pressure changes “push” and “pull” on the eardrum. The resulting motion of the eardrum is converted (by the inner ear) to nerve signals that are ultimately perceived as “sound”. In a microphone, the pressure changes act on a diaphragm. The resulting diaphragm motion is converted (by one of several mechanisms) into electrical signals which are sent to the sound system. For both “receivers”, the sound picked up is a combination of all pressure changes occurring just at the surface of the eardrum or diaphragm.

Sound can be classified by its acoustic behavior; for example, direct sound vs. indirect sound. Direct sound travels from the sound source to the listener in a straight line (the shortest path). Indirect sound is reflected by one or more surfaces before reaching the listener (a longer path). Since sound travels at a constant speed, it takes a longer time for the indirect sound to arrive, and it is said to be “delayed” relative to the direct sound. There are several kinds of indirect sound, depending on the “acoustic space” (room acoustics).

Echo occurs when an indirect sound is delayed long enough (by a distant reflecting surface) to be heard by the listener as a distinct repetition of the direct sound. If indirect sound is reflected many times from different surfaces it becomes “diffuse” or non-directional. This is called reverberation, and it is responsible for our auditory perception of the size of a room. Reverberant sound is a major component of ambient sound, which may include other non-directional sounds, such as wind noise or building vibrations. A certain amount of reverberant sound is desirable to add a sense of “space” to the sound, but an excess tends to make the sound muddy and unintelligible.

One additional form of indirect sound is known as a standing wave. This may occur when the wavelength of a sound is the same distance as some major dimension of a room, such as the distance between two opposite walls. If both surfaces are acoustically reflective, the frequency corresponding to that wavelength will be amplified, by addition of the incoming and outgoing waves, resulting in

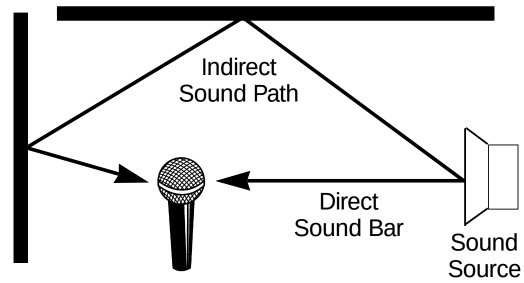

Direct vs. Indirect Sound

a strong, stationary wave pattern between the two surfaces. This happens primarily with low frequencies, which have long wavelengths and are not easily absorbed.

A very important property of direct sound is that it becomes weaker as it travels away from the sound source, at a rate governed by the inverse-square law. For example, when the distance increases by a factor of two (doubles), the sound level decreases by a factor of four (the square of two). This results in a drop of 6 dB in sound pressure level (SPL), a substantial decrease. Likewise, when the distance to the direct sound source is divided by two (cut in half), the sound level increases by 6 dB. In contrast, ambient sound, such as reverberation, has a relatively constant level. Therefore, at a given distance from a sound source, a listener (or a microphone) will pick up a certain proportion of direct sound vs. ambient sound. As the distance increases, the direct sound level decreases while the ambient sound level stays the same. A properly designed sound system should increase the amount of direct sound reaching the listener without increasing the ambient sound significantly.

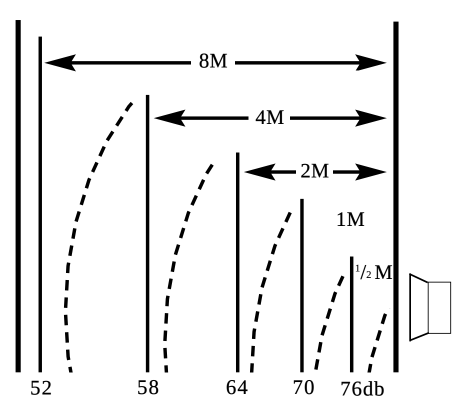

Inverse Square Law

APPENDIX 2

Potential Acoustic Gain and Needed Acoustic Gain

As previously discussed, there is a physical limitation to how much level a sound reinforcement system can achieve before uncontrollable feedback occurs. The available level, known as Potential Acoustic Gain (PAG), can be determined by a relatively simple equation. Before calculating acoustic gain, though, it is helpful to know how much gain is required to provide an adequate listening level for all members of the audience. The Needed Acoustic Gain (NAG) equation calculates the amplification necessary for the furthest listener to hear as well as nearest listener. This equation assumes the nearest listener is close enough to hear the sound source directly (without amplification).

NAG = 20 x log (Df/Dn)

Where: Df = distance from sound source to furthest listener
Dn = distance from sound source to nearest listener
Log = logarithm to base 10

The NAG equation is based on the inverse-square law, which states that sound level decreases by 6 dB for each doubling of distance from the sound source. For example, the front row of an audience (10 feet from the stage) may experience a comfortable level (without a sound system) of 85 dB. The last row, which is 80 feet from the stage, will only experience 67 dB; 18 dB less than the front row. Therefore, the sound system needs to provide 18 dB of gain to the last row of the audience, so it will experience the same listening level as the front row. Using the equation:

NAG = 20 x log (80/10)
NAG = 20 x log 8
NAG = 20 x 0.9
NAG = 18

Potential acoustic gain (PAG) is calculated from the distances between various components in the sound system, the number of open microphones, and other variables. The sound system is sufficient if PAG is equal to or greater than the Needed Acoustic Gain (NAG). While it appears somewhat complex, the equation is easily solved with a scientific calculator:

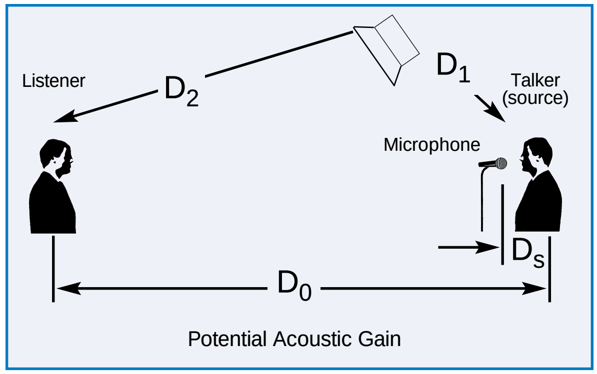

PAG = 20 (log D1 – log D2 + log D0 – log Ds) – 10 log NOM – 6

Where: PAG = Potential Acoustic Gain (in dB)

D_s = distance from sound source to microphone

D_0 = distance from sound source to furthest listener

D_1 = distance from microphone to nearest loudspeaker

D_2 = distance from loudspeaker to furthest listener

NOM = number of open microphones

-6 = a 6 dB feedback stability margin

log = logarithm to base 10

The 6 dB feedback stability margin is required to provide a small amount of "headroom" below the feedback threshold, even when NAG and PAG are equal. The NOM term reflects the fact that gain-before-feedback reduces by 3 dB every time the number of open microphones doubles. For example, if a system has a PAG of 20 dB with one open microphone, adding a second microphone will cause a 3 dB decrease to 17 dB. Doubling the number of open microphones again, to four, drops PAG to 14 dB. Consequently, the number of open microphones should always be kept to a minimum. Unused microphones should be turned off or attenuated, either manually (by a human operator) or electronically (by an automatic mixer). In fact, using an automatic microphone mixer with a NOMA (Number of Open Microphones Attenuator) circuit removes the NOM component from the equation, since NOMA ensures that the overall output of mixer will always be equivalent to one open microphone.

APPENDIX 2

Potential Acoustic Gain and Needed Acoustic Gain

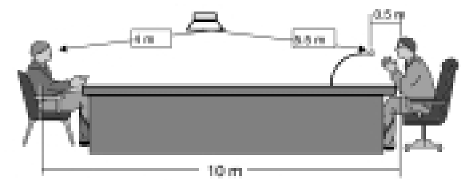

$D_1 = 5.5$ m $D_2 = 4$ m $D_s = 0.5$ m $D_0 = 10$ m NOM = 1

PAG= 20 log D_1 - 20 log D_2 + 20 log D_0 - 20 log D_s - 10 log NOM - 6
PAG= (20 log 5.5)- (20 log 4) + (20 log 10) - (20 log 0.5) - 10 log 1 - 6
PAG= 14.8 - 12 + 20 -(-6.0) - 0 - 6
PAG= 22.8 dB (NAG = 20 log 10/1 = 20 dB)

System will work: PAG>NAG

To provide maximum gain-before-feedback , the following rules should be observed:

1. Place the microphone as close to the sound source as practical.
2. Keep the microphone as far away from the loudspeaker as practical.
3. Place the loudspeaker as close to the audience as practical.
4. Keep the number of open microphones to a minimum.

Achieving noticeable results when making changes to a sound system requires a level difference of at least 6 dB. Due to the logarithmic nature of the PAG equation, a 6 dB change requires a doubling or halving of the corresponding distances. For example, if a microphone is placed 1 ft. from a sound source, moving it back to 2 ft. away will decrease gain-before-feedback by 6 dB. Moving it to 4 ft. away will cause a 12 dB decrease. Conversely, moving it to 6 inches away increase gain-before-feedback by 6 dB, and moving it to 3 inches away will increase it by 12 dB. The single most significant (and inexpensive) way to maximize gain-before-feedback is to place the microphone as close as possible to the sound source.

The PAG equation allows the performance of a sound system to be evaluated solely on the basis of the relative location of sources, microphones, loudspeakers, and audience, as well as the number of microphones, but without regard to the actual type of component. Note that the equation also assumes omni-directional components. As discussed previously, using directional microphones and loudspeakers may increase PAG. Component characteristics notwithstanding, the results provided by this relatively simple equation still provide a useful, best-case estimate.

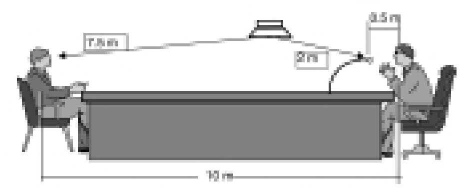

$D_1 = 2$ $D_2 = 7.5$ $D_s = 0.5$ $D_0 = 10$ NOM = 1

PAG= 20 log D_1 - 20 log D_2 + 20 log D_0 - 20 log D_s - 10 log NOM - 6
PAG= (20 log 2) - (20 log 7.5) + (20 log 10) - (20 log 1) - 6
PAG= 6 - 17.5 + 20 - 0 - 6
PAG= 2.5 dB (NAG = 20 log 10/1 = 20 dB)

System will not work: PAG<NAG

GLOSSARY

Active
A device that requires power to operate.

Acoustic Echo Canceller (AEC)
A processor that attempts to remove acoustic echoes in a teleconferencing system.

Ambience
Room acoustics or natural reverberation.

Amplitude
Magnitude of strength of signal or wave.

Audio Chain
The series of interconnected audio equipment used for recording or reinforcement.

Automatic Gain Control (AGC)
A signal processor that attempts to compensate for the differences in level between different sound sources.

Band Pass Filter
A filter that only allows a certain range of frequencies to pass.

Band Reject Filter
A filter that reduces a range of frequencies.

Bandwidth
The range of frequencies that a filter affects.

Cardioid Microphone
A unidirectional microphone with moderately wide front pickup (131 deg.). Angle of best rejection is 180 deg. from the front of the microphone, that is, directly at the rear.

Clipping Level
The maximum electrical signal level that a device can produce or accept before distortion occurs.

Comb Filtering
The variations in frequency response caused when a single sound source travels multiple paths to the listener's ear, causing a "hollow" sound quality. The resultant frequency response graph resembles a comb. Can also occur electronically with multiple microphones picking up the same sound source.

Compressor
A signal processor that reduces the level of incoming audio signals as they exceed a given threshold. The amount of reduction is usually defined by the user.

Crossover
A processor that divides the audio signal into two or more frequency bands.

Decade
The distance between two frequencies that are multiples or divisions of ten (e.g. 200 Hz – 2000 Hz).

Decibe
A number used to express relative output sensitivity. It is a logarithmic ratio.

Dynamic Range
The range of amplitude of a sound source. Also, the range of level between the noise floor and clipping level of a device.

Echo
Reflection of sound that is delayed long enough (more than about 50 msec.) to be heard as a distinct repetition of the original sound.

Equalizer
A signal processor that allows the user to boost or cut selected frequencies. Used for tone shaping and limited feedback control. Variations include graphic or parametric.

Expander
A signal processor that expands the dynamic range of an audio signal.

Feedback
In a PA system consisting of a microphone, amplifier, and loudspeaker, feedback is the ringing or howling sound caused by the amplified sound from the loudspeaker entering the microphone and being re-amplified.

Fidelity
A subjective term that refers to perceived sound quality.

Filter
A processor that cuts or boosts a specific frequency or frequency range.

Frequency
The rate of repetition of a cyclic phenomenon such as a sound source.

Frequency Response
Variations in amplitude of a signal over a range of frequencies. A frequency response graph is a plot of electrical output (in decibels) vs. frequency (in Hertz).

GLOSSARY

Gain
Amplification of sound level or voltage.

Gain-Before-Feedback
The amount of gain that can be achieved in a sound system before feedback or ringing occurs.

Gate (Noise Gate)
A signal processor that mutes the audio signal when it drops below a given threshold.

Headroom
The difference between the nominal operating level of a device and the point at which the device clips.

High Pass (Low Cut) Filter
A filter that attenuates low frequencies below a certain frequency.

Inverse Square Law
States that direct sound levels increase (or decrease) by an amount proportional to the square of the change in distance.

Limiter
A signal processor that prevents signals levels from exceeding a certain threshold.

Low Pass (High Cut) Filter
A filter that attenuates high frequencies above a certain frequency.

Mixer
A device which allows the combination, manipulation, and routing of various audio input signals.

NAG
Needed Acoustic Gain is the amount of gain that a sound system must provide for a distant listener to hear as if he or she was close to the unamplified sound source.

Noise
Unwanted electrical or acoustic energy.

Noise Gate
A signal processor that mutes the audio when the signal level drops below a certain threshold.

NOM
Number of Open Microphones in a sound system. Decreases gain-before-feedback by 3 dB every time the number of open microphones doubles.

Octave
The distance between two frequencies that is either double or half the first frequency (e.g. 500 Hz to 1000Hz).

Omnidirectional Microphone
A microphone that picks up sound equally well from all directions.

Q
Quality Factor. Indicates how tightly a filter is focused near the center frequency.

PAG
Potential Acoustic Gain is the calculated gain that a sound system can achieve at or just below the point of feedback.

Passive
A device that does not require power to operate.

Phantom Power
A method of providing power to the electronics of a condenser microphone through the microphone cable.

Reverberation
The reflection of sound a sufficient number of times that it becomes non-directional and persists for some time after the source has stopped. The amount of reverberation depends on the relative amount of sound reflection and absorption in the room.

Shelving Equalizer
Reduces (or raises) the frequencies below (or above) a certain frequency to a fixed level. The response when viewed on a frequency response graph resembles a shelf.

Signal to Noise Ratio
A measurement of the noise of device expressed as a ratio between the desired signal level (dBV) and the noise floor.

Sound Reinforcement
Amplification of live sound sources.

Speed of Sound
The speed of sound waves, about 1130 feet per second in air.

Supercardioid Microphone
A unidirectional microphone with tighter front pickup angle (115 deg.) than a cardioid, but with some rear pickup. Angle of best rejection is 126 deg. from the front of the microphone.

Voltage
The potential difference in an electrical circuit. Analogous to the pressure on fluid flowing in a pipe.

PRODUCT SELECTION CHART

Mixers+Amplifiers

Model >>	FP16	M367	SCM262	SCM268	SCM410	SCM800	SCM810
Features:							
Transformer-balanced input	●	●		●			
Active-balanced input			●		●	●	●
Transformer-balanced output	●	●		●			
Active-balanced output			●		●	●	●
Low-Z mic-level input	●	●	●	●	●	●	●
Line level input	●	●	●		1	●	●
Aux level input			●	●		●	●
Mic level output	●	●	●	●	●	●	●
Line level output	●	●	●	●	●	●	●
Phono jack aux level output			●	●	●		
Headphone output		●				●	●
Phantom power	●	●	●	●	●	●	●
48 V phantom power		●				●	●
VU meter		●					
Peak meter				●	●	●	●
EQ		●	●		●	●	●
Tone oscillator		●					
Linkable	●	●			●	●	●
Slate mic + tone		●					
Limiter		●			●	●	●
Stereo operation		●	●				
AC operation	●	●	●	●	●	●	●
Battery operation	●	●					

1
Internal modification or optional accessory.

Integrated Signal Processors

	DSPs			
Models >>	DFR11EQ	DP11EQ	DFR22	P4800
Features:				
Inputs x outputs	1x1	1x1	2x2	4x8
Connectors	XLR & 1/4”	XLR & 1/4”	XLR & Phoenix	Phoenix
Rack space	1/2 rack	1/2 rack	1 rack	1 rack
Audio specs	Dynamic range > 104 dBA	Dynamic range > 104 dBA	Dynamic range > 110 dBA	Dynamic range > 100 dBA
Matrix Mixer	Signal goes straight through	Signal goes straight through	Full matrix mixer	Full matrix mixer
Front panel controls	Scene selector for 3 scenes. Controls for DFR parameters.	Bypass	Preset selector for 16 presets. Controls for DFR parameters	No front panel controls
Front panel audio metering	Single-LED signal strength indicator	Single-LED signal strength indicator	Mute, 20 dB, 0 dB, Clip LEDs for each input and output	Full string metering for each input and output
Automatic feedback reduction	10-band DFR	None	Drag and drop blocks for 5-, 10-, and 16-band single channel and stereo DFR	Drag and drop blocks for 5-, and 10-band single channel DFR
DFR filter removal	Hold mode	N/A	Auto clear	Hold mode
Additional processing	GEQ or PEQ, limiter, delay	PEQ, gate, downward expander, comp, limiter, delay. Option for AGC & peak stop limiter	Drag and drop blocks for GEQ, PEQ, cut/shelf, delay, single channel and stereo compressors and limiters, peak stop limiter, AGC, gate, downward expander, ducker, crossover	
External control options	Compatible with DRS-10, AMX or Crestron control	N/A	DRS-10 & serial commands (AMX or Crestron); contact closures and potentiometers for preset, volume and mute.	
Control pin inputs	None	None	4	8
Logic outputs	None	None	None	8
Security	Front panel lockout	Front panel lockout	Front panel lockout with password protected multi-level security	password protected multi-level security
Shure link	Yes	Yes	Yes	Yes

Reference Information

Bibliography & Additional References:

Bohn, Dennis
"Linkwitz-Riley Crossovers," Pro Audio Reference, Rane Corporation, Mukitelo, WA

Bohn, Dennis, and Pennington, Terry
"Constant-Q Graphic Equalizers," Pro Audio Reference, Rane Corporation, Mukitelo, WA

Brown, Pat
"System Gain Structure," Handbook for Sound Engineers, 3rd Edition, Focal Press, Boston, MA

Davis, Gary D., and Jones, Ralph
Sound Reinforcement Handbook. Hal Leonard Publishing Co., Milwaukee, WI

Lyons, Christopher
Audio for Distance Learning, Shure Incorporated, Niles, IL

McMannus, Steven
"Filters and Equalizers," Handbook for Sound Engineers, 3rd Edition, Focal Press, Boston, MA

Vear, Tim
Audio Systems Guide for Houses of Worship, Shure Incorporated, Niles, IL

Whitlock, Bill, and Pettersen, Michael
"Preamplifiers and Mixers," Handbook for Sound Engineers, 3rd Edition, Focal Press, Boston, MA

Acknowledgements:

The following individuals contributed to this publication, either with their words or their editing skills:

Bob Rieder
Luis Guerra
Tim Vear
Michael Pettersen
Cris Tapia

Biography: **Gino Sigismondi**

Gino Sigismondi, a Chicago native and Shure Applications Engineer, has been active in the music and audio industry for nearly ten years. In addition to his work as a live sound and recording engineer, Gino's experience also includes performing and composing. Gino earned his BS degree in Music Business from Elmhurst College, where he was a member of the Jazz Band, as both guitar player and sound technician. After college, he spent several years working for Chicago area sound companies, local acts, and night clubs, and currently mixes for Chicago's "Standing Room Only Orchestra." As a member of Applications Engineering, Gino brings his years of practical experience to the product training seminars he conducts for Shure customers, dealers, distribution centers, and internal staff. He is the author of the Shure educational publications "Selection and Operation of Personal Monitors" and "Audio Systems Guide for Music Educators", and has written for the Shure Web site. He recently contributed a chapter on in-ear monitoring to the Handbook for Sound Engineers (Focal Press). Gino continues to remain active as a sound engineer, consults musicians on transitioning to in-ear monitors, and dabbles in sound design for modern dance.

Additional Shure Publications Available:

These guides are available free of charge. To request your complimentary copies, call one of the phone numbers listed below.

- Selection and Operation of Personal Monitor Systems
- Audio Systems Guide for Video Production
- Audio Systems Guide for Houses of Worship
- Audio Systems Guide for Meeting Facilities
- Microphone Techniques for Studio Recording
- Microphone Techniques for Live Sound Reinforcement

Our Dedication to Quality Products

Shure offers a complete line of wireless systems for everyone from first-time users to the biggest names in the industry— for nearly every possible application.

For over seven decades, the Shure name has been synonymous with quality audio. All Shure products are designed to provide consistent, high-quality performance under the most extreme real-life operating conditions.

SHURE®
It's Your Sound®

Shure Incorporated
5800 West Touhy Avenue
Niles, IL 60714-4608 USA

Phone: 847-600-2000
Fax: 847-600-1212
E-mail: info@shure.com

Shure Europe GmbH
Wannenäcker Str. 28,
74078 Heilbronn, Germany

Phone: 49-7131-72140
Fax: 49-7131-721414
E-mail: info@shure.de

Shure Asia Limited
Unit 301, 3rd Floor
Citicorp Centre
18, Whitfield Road
Causeway Bay, Hong Kong

Phone: 852-2893-4290
Fax: 852-2893-4055
E-mail: info@shure.com.hk

Elsewhere
Phone: 847-600-2000
Fax: 847-600-6336
E-mail: international@shure.com

www.shure.com

AL1517 Printed in U.S.A. 15K 8/03