AUTOMATIC MIXERS

Written by Tom Stuckman and Steve Marks

The use of sound reinforcement systems is common place in today's churches, boardrooms, meeting rooms, auditoriums and other facilities. In fact, it is hard to imagine what it was like before their wide spread use. With advances in the sound quality in the entertainment industry, television, CDs and home theatre for example, the expectation of quality sound reinforcement has risen as well.

When sound system requirements progress past the need of a single microphone, then mixing becomes an issue. Depending on how the system is operated, the performance of a sound system with multiple microphones can also change for the better, or worse. Automatic mixers are an important, and often misunderstood, tool of the sound system designer that can improve gain before feedback, improve intelligibility and reduce inter-microphone interference (comb filtering) in many multi-microphone systems.

To begin with, let us briefly examine why we use a sound reinforcement system. A sound reinforcement system is generally used to aid the listener by increasing the volume of a sound source sufficiently above the background noise so that it can be easily heard. Another benefit, when used in an acoustic environment, is to improve the clarity of the sound source at the listener. Stated another way, the goal is to make the sound source (talker) appear to be closer to the listener. There are several factors that influence how a sound system meets these goals, but we will concentrate on the microphone and microphone mixer and their use in an acoustic space.

The Acoustic Environment

Consider a light hanging in the center of a white room. The light radiates from the bulb until it strikes a wall, then reflects back into the room. The intensity of the light radiating from the bulb decreases with distance, but the intensity of the reflected light in the room is fairly even. The light bulb can be compared to the sound source, or loudspeaker system, and the reflected light to reverberant sound. There are however important distinctions that need to be made in this analogy. First, we are not using light to convey information as we are with sound waves. Second, sound travels at such a slow rate that we can perceive the reflections as an echo or reverberation. Reflected light in the room will make the room brighter and easier to see, but the slow moving sound waves can aid or hinder our ability to hear in a room. In the case of a large reverberant room, the volume may be louder, or fairly even, but the random reflections can restrict our understanding of what is being said.

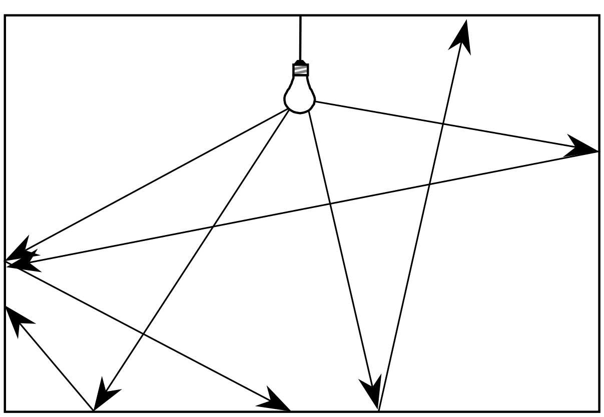

Light source in a reflective room.

Now let us look at how this affects our sound system requirements. In an environment without reflections, sound will decrease in volume from a point source at the rate of 6 dB every time the distance from the source is doubled. In a reverberant environment, the same holds true of the direct sound from the source.

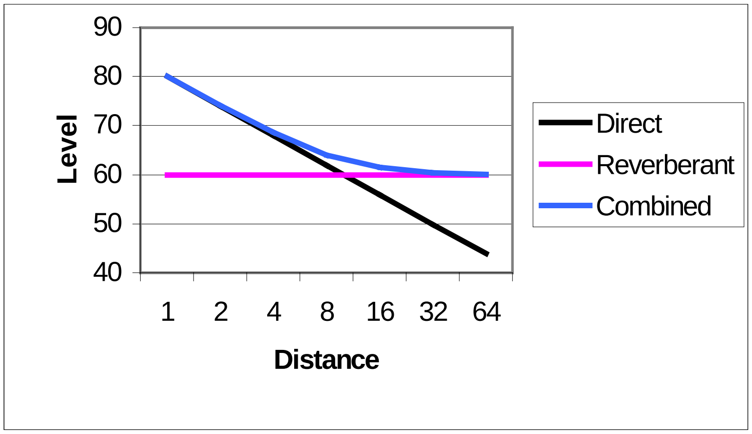

However, as we continue to move from the source, we reach a point called the critical distance where the direct and reverberant sound fields are equal. Moving beyond this point the volume remains constant due to even nature of the reverberant field. (Remember the reverberant field, like the reflected light in the room, is at a fairly even level in the room.) Therefore, the closer the listener is to the source, the louder the direct sound is relative to the reflected sound.

The Sound Reinforcement System

A typical sound system consists of four basic elements, although some may be combined or others added. These parts are the microphone, the microphone preamplifier or mixer, the power amplifier and the loudspeaker. Sound converted into an electrical signal in the microphone, is increased in amplitude in the mixer/preamplifier, is increased in power in the power amplifier and is then converted back into sound in the loudspeaker system.

When a sound system is introduced into an acoustic environment, inevitably feedback becomes an issue. Acoustic feedback occurs when sound from the loudspeaker system, both direct and reflected, arrives at the microphone at a volume greater than the original sound that entered the mic. This process generally occurs at one frequency before others, creating the howling sound we hear as feedback. Even if the sound arrives slightly lower in volume at the mic, it still recirculates through the system, reducing slightly in level each time, creating the sound we know as ringing.

To keep acoustic feedback from being a problem, we need to insure that the loudest sound from the loudspeaker system arrives at the microphone lower than the original sound. And to be safe, we generally allow a 6 dB margin. That is, the loudest sound from the loudspeaker system at the mic should be 6 dB lower than the original. In sound system design there are many techniques that affect this. But in a reverberant room, where the mic and listener are in the reverberant field, the gain must be set so that the sound level from the talker at the microphone is 6 dB louder than the reverberant sound from the loudspeaker system.

Why Multiple Microphones?

To this point we have been discussing a system with a single microphone. From the previous discussion, we can conclude that the volume of the talker at the microphone has a direct bearing on the maximum volume at a listener in the reverberant field. We can also conclude that halving the distance of the talker to the microphone increases the volume of the talker at the mic 6 dB without increasing the pickup of reverb and noise. The volume of the talker at the mic and the volume of the reverb and noise at the mic both affect how well we meet our original goal of allowing the listener to easily hear.

If there is only one talker using the sound reinforcement system, then our job is relatively easy. But as we add additional sound sources (talkers) that need to be reinforced, the difficulty increases. The trick is to position our single microphone so that we have a balanced level from each talker and still achieve adequate volume for the listener. Now let us consider the addition of a second microphone.

We begin by looking at a situation where two talkers are sitting at a boardroom table 3 feet apart. If we use one microphone, centered between the talkers and two feet in front, the distance from each talker to the mic is 30 inches. We will also assume that the level of each of their voices reaches the microphone at a sound pressure level of 70 dBA. Let us also assume this level is 6 dB above the level of the reverberant sound field from the loudspeaker system, which is at (64 dBA). Under these conditions we have a stable system; however, the best we can achieve is a level of 64 dBA for a listener also in the reverberant field.

Two talkers, each 30 inches from the microphone.

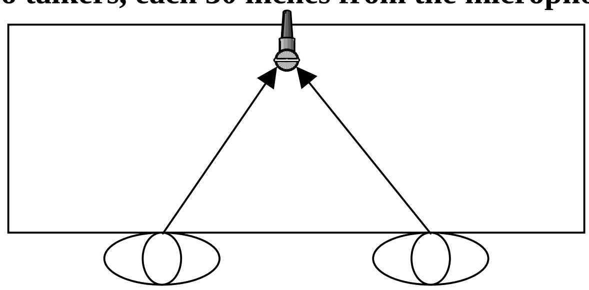

Now let us examine the same situation using two microphones, one positioned 15 inches in front of each talker. The level of each voice reaches their respective microphone 6 dB louder (76 dBA) than in the previous example because they are twice as close. The output from each microphone is added equally together at the mixer and fed to the sound system. If each microphone is given the same gain as our single microphone example, the level at the listener in the reverberant field would be 6 dB louder as well. But we now need to consider what happens to the reverb re-entering the microphones. Keeping the gain for each microphone to the loudspeaker system the same, the level arriving at the microphones from the loudspeaker is 6 dB less than the original sound. But this time the signal from the loudspeaker system re-enters both microphones. These random reverberant signals are then added together equally increasing the level of the reverb in the system by 3 dB. This leaves only 3 dB difference between the original voice signal and the reverb, reducing our feed back stability margin to 3 dB. To restore our 6 dB margin, we must reduce the gain of each mic by 3 dB. Even with this reduction in gain, the level of the talkers is still 3 dB louder than our original one mic example.

Two talkers, each 15 inches from their microphones.

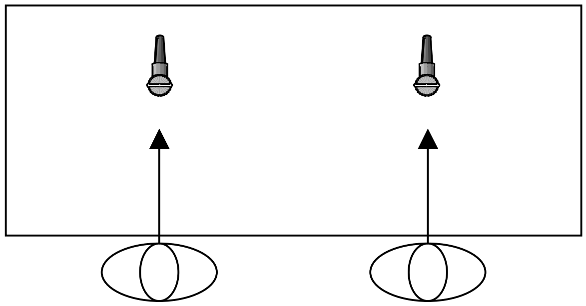

In our previous example we were able to increase the level of our talker 6 dB by moving the microphone closer to the talker. Continuing to move the microphone closer to each talker would at first appear to solve all of our problems of insufficient level at our listener. But moving the mic ever closer can create problems as well. As the distance from the mic to the talker decreases, small changes in distance can make a significant change is volume. A 3-inch movement from 3 to 6 inches, changes the volume in the same way as a 12 inch movement from 12 to 24 inches. The talker's movements are now restricted, and a small movement can quickly reduce the ability of the listener to hear.

Also, as we continue to add microphones to a system, we must reduce the gain of these microphones to prevent feedback. In fact, we must reduce the gain 3 dB every time the number of open mics is doubled. If we were able to control the microphone gain so that only one mic is on at a time, we could keep the 6 dB increase in volume first achieved by moving the mics closer without affecting our feedback stability margin.

Comb Filters

An additional problem the sound system designer must contend with affects the quality of the sound. This results whenever the sound from the talker arrives at one mic via two paths of different length or to two open mics that are different distances from the talker. This multi-path interference affects the quality of the sound source by emphasizing sound at some frequencies and canceling the sound at others. The effect of this interference is called comb filtering. If we look at the resulting frequency response, we see a series of bumps and notches that look like the teeth of a comb. Phasers and flangers intentionally create this effect, which is sometimes used in modern music.

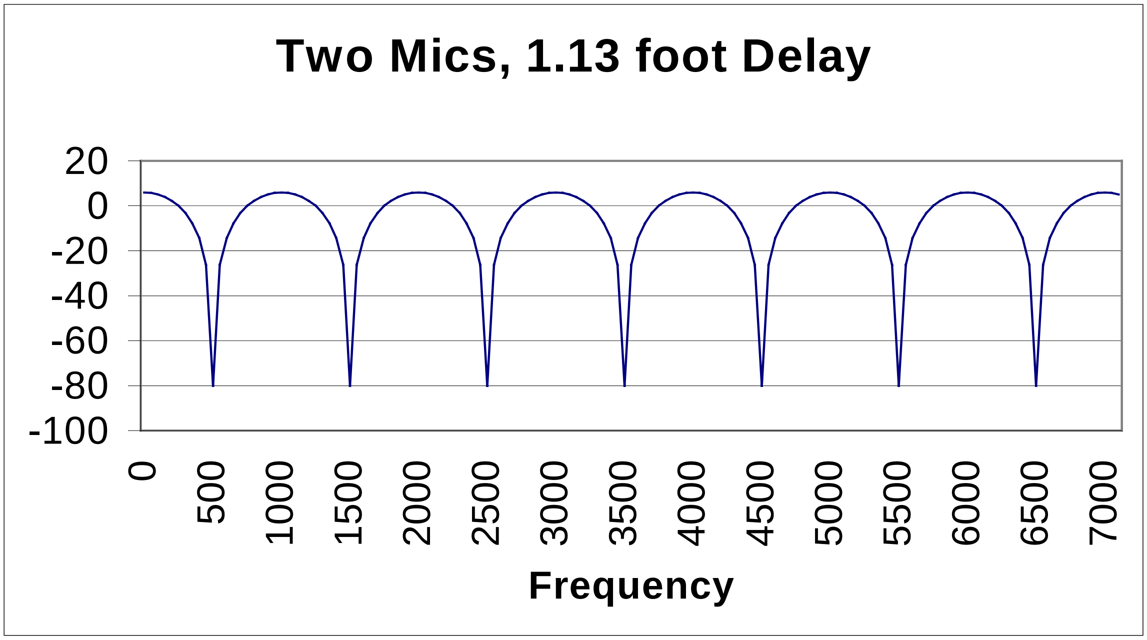

This effect can easily be demonstrated by making a recording of a talker sitting at a table using a single microphone. While the talker is speaking, lower the microphone from a position two feet above the far edge of the table down to the table's surface. The direct sound and reflected sound paths are now about the same length. The same demonstration can be performed with two mics mixed at an equal volume. Keep one mic two feet in front of a talker and slowly move the second mic from a distant position to the same position as the first. The effect of this interaction is quite dramatic and can be easily heard in the recording.

To prevent this annoying effect when multiple microphones are used, the level of the talker's voice should be 10 dB lower in the mix from microphones other than his own. There are two ways to achieve this goal. First, if all mics are open and set to the same gain, the distant mics should be 3 times the distance from the talker as his own mic. The second way of addressing this problem is to reduce the gain of microphones other than the talker's.

Automatic Mixing

A skilled sound system operator can greatly enhance the performance of a sound system. Gain before feedback can be increased, unwanted noise, reverberation, and comb filtering can be reduced as well. There are many situations, however, where even a skilled engineer cannot anticipate who will speak next and some speech is lost. A properly adjusted automatic mixer can greatly enhance the performance of a sound reinforcement system used in the church, boardroom, courtroom, theatre, council chambers or anywhere multiple mics are used for the spoken word. Next, consider the operation of these products.

An automatic microphone mixer, as its name implies, automatically mixes signals from multiple-microphones, without the need for a system operator. Automixers, as they are often referred, can significantly increase the gain before feedback in multiple microphone systems. This is accomplished by activating microphones only as needed and by adjusting the system gain to maintain system stability. Some automixers can also improve sound quality in multiple microphone systems by reducing the comb filtering created by close microphone spacing.

When a conventional microphone mixer adds individual microphone signals together, each doubling of the number of microphones feeding the system reduces the available gain before feedback by 3 dB. Thus a large system soon becomes cumbersome and ineffective if a system operator is not present to lower microphone levels or switch unused microphones off. Because acoustic gain is often marginal due to architectural room characteristics, the use of an automixer may be the only way to acceptably achieve high sound levels to the listener with an unattended sound system.

Decision Making

A viable automixer needs to decide which microphones need to be turned on and which ones need to be turned off. This may seem obvious but there are more issues to consider than may be readily identified. Several methods have been developed to make this decision.

Fixed Threshold

One of the first decision-making methodologies used in automixer design is the fixed threshold. The most familiar application is a VOX (voice operated switch). A detector in the microphone channel of the mixer simply switches the channel on when a signal is present and off again when the signal stops. To turn the channel on, the signal must be greater than a preset threshold for that channel. Although it may sound like a straightforward approach of controlling microphone switching, it has several drawbacks. First, there is the question of where to set the threshold to determine whether a signal is present. If the threshold is set too low, the microphone channel will easily respond to room noise and reflected sounds. Conversely, setting the threshold too high to avoid false triggering from room noise runs a high risk of chopping the signal and creating drop-outs. Ideally, the threshold needs to be set high enough to avoid turn on with random noises, but low enough to react to program signals. These are often incompatible conditions. Note in the figure below, the fixed threshold system cannot tell the difference between the intended “On Mic” sounds that originate near the microphone and “Off Mic” sounds. It only acts on the signal level. Furthermore when two or more channels are on at the same time, the total system gain increase may cause acoustic feedback.

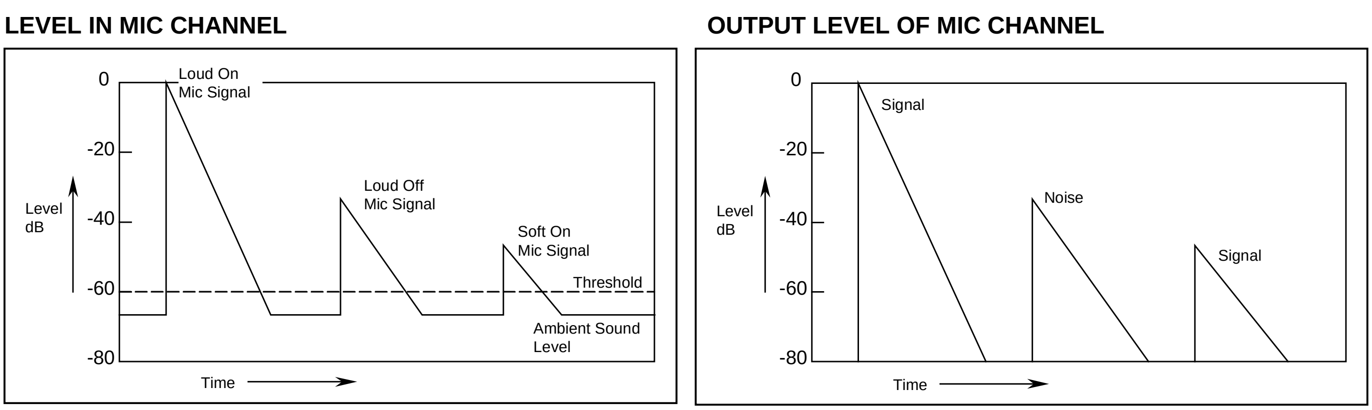

Performance of a fixed threshold system.

Adaptive Threshold

The adaptive threshold approach requires the automixer to automatically adjust its threshold level to the conditions of the room where the microphones are in use. For example, in a noisy room the automixer would increase the threshold level to prevent any of the microphone channels from being triggered on by the noise; therefore, a direct on-mic signal greater than the threshold level must be present before the channel is turned on. On the other hand in a quiet environment, the threshold level must be lowered to pick up a soft-spoken talker. However, if a loud noise is heard by all of the microphones (a door slam) the microphones channels must not all turn on. In the figure below, the adaptive threshold follows the ambient noise level and correctly identifies undesirable sounds beyond the critical distance from the microphone. If the sound source is beyond this critical distance or below the ambient noise level, it becomes impossible to amplify it in the sound system. Experienced mixer operators know that if the gain is increased trying to “reach” for such a sound source, the result will be feedback or unintelligibility. On the other hand, overall system gain and noise benefit whenever microphones not picking up any useful information are turned down. Accordingly, the adaptive threshold technique bases its decision making on whether the microphone hears any beneficial sound besides the general ambient noise in the room.

There are a number of methods that various manufacturers have developed to implement an adaptive threshold scheme in their automixers. Among these methods are ambience sensing, direction sensing, sweeping threshold, and gain-sharing. Any of these methods exhibit improved performance over the fixed threshold technique.

LEVEL IN MIC CHANNEL **OUTPUT LEVEL OF MIC CHANNEL**

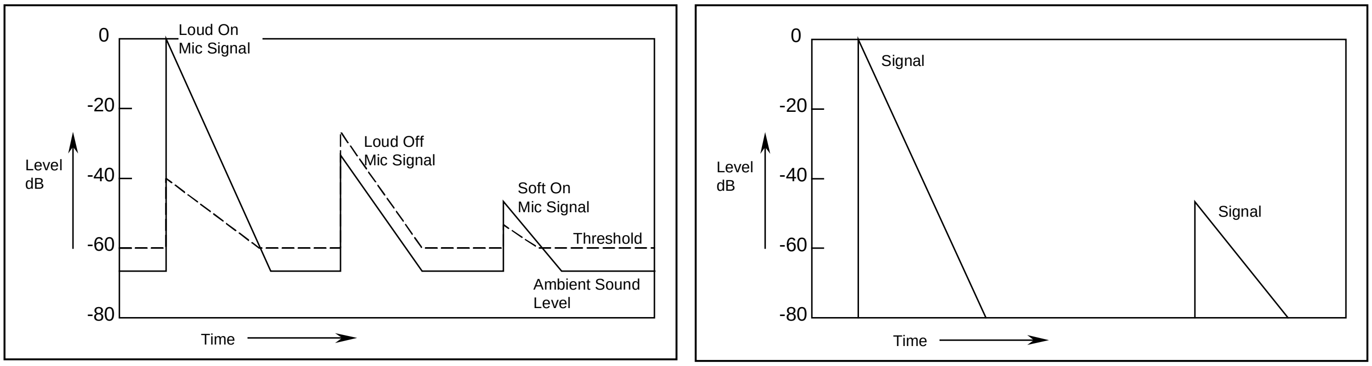

Performance of an adaptive threshold system.

Ambience Sensing

A straightforward method of implementing an adaptive threshold places a “dummy mic” in the room. This mic senses the ambient noise of the room and automatically adjusts the threshold level accordingly. However, there are some drawbacks to this approach. The placement of the mic proves to be critical. If the mic is placed near an air vent, it will pick up the air handling noise. This can prevent the mic channels from turning on. The opposite can also occur. If the mic is placed in a quiet area of the room away from such noise sources, the threshold level may be too low allowing the mic channels to trigger on falsely. Additionally, the dummy mic adds an extra microphone to the system, thus increasing the system cost without much advantage.

Direction Sensing

Another approach senses the direction from which the sound source arrives to the microphone. The automixer responds to signals having satisfactory levels within a predetermined space in front of the microphone. The mic channel then decides whether it should be “on,” based upon the relative signal levels at two back-to-back cardioid microphones within the single structure. As in the ambient sensing technique, an extra “direction” microphone is used to provide the sensing for each audio channel. An automixer that uses this approach will only work with special direction-sensing microphones. Therefore, a limited number of microphone choices are available to the system designer.

Scanning Threshold

Another practical approach involves scanning the level on all of the input channels and activating the channel with the highest level. This channel remains active while another scan begins. If the level of active channel is still higher than the other input channels, then it remains on and the process repeats. If the level of another channel is higher on the next scan, then that channel becomes active immediately. The channel that was previously on remains for a short period and then shuts off. This process happens very quickly so that dialog is possible.

Switching Method

Now that we have discussed some of the methods by which the decision is made to turn on individual microphone channels…how do we actually turn them on and off? Probably the most common method is a "gate." It is basically like a microphone on-off switch. However, a basic gating circuit is often undesirable due to its inherent tendency to "pop" when turned on and off. To overcome this switch noise problem a gating technology was developed that only permitted the gate to turn on or off when the amplitude of the switched signal was zero. This is known as a zero-crossing gate. The gate switching method is used with most of the decision-making schemes mentioned above excluding the gain-sharing principle. Other system use VCAs to switch the gain from "on" to "off". In some of these systems, the amount of "off" attenuation can be set.

NOM

It might appear that gating alone would provide a complete automix solution. But from our previous discussions, unless only one mic is allowed to be on at a time, we must reduce the system gain as the number of open microphones increases. Somehow the system gain must remain the same as it is with only one microphone on. This is known as NOM=1, or the Number of Open Microphones equals One. To perform this operation, many automixers add an NOM attenuator circuit that "counts" the number of microphones that are on in the system, and then it attenuates the output by a predetermined amount. For example if two microphones are on, the NOM circuit must attenuate the output by 3 dB to maintain NOM=1 and prevent acoustic feedback. The amount of output attenuation is seen in Table 1. Some systems allow the amount of NOM attenuation to be set for greater than 10log(NOM) attenuation.

Table 1.

NOM	Output Gain Adjustment (dB) 10log(NOM)
1	0
2	-3.0
3	-4.8
4	-6.0
5	-7.0
10	-10.0
20	-13.0
100	-20.0

Until this point, we have assumed that all of the microphones in our system are strictly within the reverberant field of the loudspeaker system. The means that the signal from loudspeaker reaching the microphones is random, and when added, the level would increase based only on the addition of power, or 3 dB per doubling of sources. At this point, however, we will consider a system where some correlation exists in the signal that arrives at the microphones. (Completely correlated signals at all mics would increase the level 6 dB per doubling of open microphones, but this is nearly impossible to achieve.) Consider how a system responds as the number of open mics increases. This situation may result from several people talking at once, or by a noise in the room. If the addition of signals from the loudspeaker system at the mics exceeds the 3 dB per NOM compensation, the system can become unstable. In systems that only provide 10log(NOM) compensation, the stability with all mics on must be tested. The overall gain may need to be reduced, to prevent the system from becoming unstable under this multiple microphone on condition.

Gain Sharing

Still another viable approach uses voltage-controlled amplifiers (VCAs) to vary the gain of each microphone channel instead of using a switch. This approach uses a special control bus to sum all of the input channels. The gain of each channel is then adjusted by comparing its level to the level of that sum. The gain is computed so that the combined system gain of all microphones remains constant. In this system, the microphones with the strongest signal are given the highest gain and those with low level signals have their gain reduced. Specifically, each microphone channel is attenuated (turned down) by the amount, in dB, equal to the difference between that microphone channel's level and the sum of all microphone channel levels.

The operation of this system is more easily understood if we look at some examples. If two talkers are each speaking on their own microphone channel at the same level, then the level of the sum of all channels would be 3 dB higher than the level from each of these mics (power addition). Each microphone would then be turned down 3 dB. If a person is being picked up equally by two microphones, the sum of these two signals would increase 6 dB. (The 6 dB increase results from the addition of two coherent signals.) The level in each mic channel would be 6 dB lower than the sum and the gain of each mic is reduced 6 dB. The resulting output would be the same as if only one mic were on. In a final example, a talker is speaking into his mic, but his voice also enters an adjacent microphone at a level 4 dB lower. In a gated system, this second mic could very easily be gated "on" creating comb filtering effects. In the gain sharing system, the gain of the second mic will be attenuated 4 dB. This makes the level from the second mic 8 dB lower than the first, which will greatly reduce the interference.

Last On

Remembering that the purpose of an automixer is to function the same as a manual mixer with an operator, the condition of the system when no microphones are on requires examination. In a manually controlled system, a good operator always leaves one microphone on in order to preserve the natural ambience of the room. This is especially important in applications where a recording is made, such as in a courtroom. In a gated automixer where none of the microphones are on, several models include a "Last On" feature. This keeps the microphone channel that was last used on until another microphone demands attention. This feature operates along with the NOM method of attenuating the output (see Table 1) to maintain system stability.

It seems apparent that the solution to this dilemma is to provide an automixer that allows for some of the microphones to be almost fully on some of the time and the others almost fully off some of the time, in the same way an operator might adjust the system. If the maximum system gain could be "programmed" into the automixer, the system would never experience acoustic feedback. The gain sharing automixer exhibits this behavior. As previously mentioned, each microphone channel's gain is controlled by its comparison to the entire system gain. Its gain is then "shared" with all of the other microphone channels regardless of their activity. Inherent in the gain sharing principle is room ambience because all of the microphones are always somewhat "on" maintaining NOM=1. This eliminates the need for a "Last On" feature. Additionally since an individual microphone's gain is determined by the sum of the signals of all microphones, the automixer will handle coherent inputs accordingly without acoustic feedback. There are also no thresholds to set, no special microphones to purchase, and no additional ambience microphones to install for the automixer to make mixing decisions. It can even reduce comb filtering by reducing the gain of microphones that may be gated on in other systems. Seemingly, this approach to automixer design would be most cost-effective and useful. Its operation is straightforward and most closely approaches the sound system that has an operator. After all, this is what the automixer should accomplish.