Effects, Signal Processors & Noise Reduction
Dynamic Processors

By employing current modern analogue technology, it is possible to manufacture audio equipment with a dynamic range of up to 125dB. In contrast to analogue techniques, the dynamic range of digital equipment is approximately 25dB less. With conventional record and tape recorder technology, as well as broadcasting, this value is further reduced. Generally, dynamic restrictions are due to noisy storage in transmission media and also the maximum headroom of these systems.

Noise as a Physical Phenomenon
All electrical components produce a certain level of inherent noise. Current flowing through a conductor leads to uncontrolled random electron movements. This produces frequencies within the whole audio spectrum. Of these, currents are highly amplified, the result will be perceived as noise. Since all frequencies are equally affected, we term this *white noise (also known as thermal noise).*

The effect is similar when replaying tape. The *unidirectional* magnetic particles passing the tape playback head can also cause uncontrolled currents and voltages. The resulting sound of the various frequencies is heard as noise. Even the best possible tape biasing can 'only' provide signal to noise ratios of about 70 dB, which is not acceptable today since the demands of listeners have increased. Due to the laws of physics, improving the design of the magnetic carrier (tape) is impossible using conventional means.

What are audio dynamics
A remarkable feature of the human hearing system is that it can detect the most wide ranging amplitude changes – from the slightest whisper to the deafening roar of a jet plane. You would have learnt about this in *Fundamentals of Sound*. So, if one tried to record or reproduce this wide spectrum of sound with the help of amplifiers, cassette recorders, records, or even digital recorders (CD, DAT, etc), one would immediately be restricted by the physical limitations of electronic and electro-acoustic sound reproduction technology.

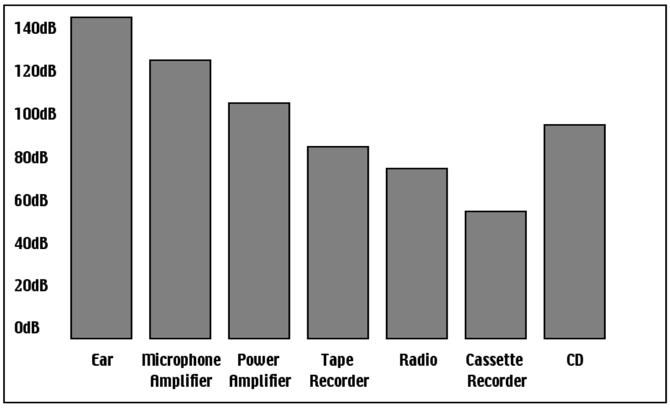

The dynamic range capabilities of various recording and reproduction devices

The usable dynamic range of electro-acoustic equipment is limited as much as at the low end, as at the high end. The thermal noise of the electrons in the components results in an audible basic noise floor and thus represents the bottom limit of the transmission range. The upper limit is determined by the levels of the internal operating voltages' if they exceed that, audible signal distortion is the result. Although in theory the usable dynamic range sits between these two limits, it is considerably smaller in practice. Since a certain upper reserve must be maintained to avoid distortion of the audio signal if sudden level peaks occur, we refer to this reserve range as *headroom* – usually about 10dB to 20dB. A reduction of operating level would allow for a greater headroom (i.e.

the risk of the signal distorting due to level peaks would be reduced). However, this could also lead to a problem. Since the signal is getting closer to the noise floor, we would perceive much more noise when we increase our monitoring levels. It is therefore useful to keep the operating level **as high as possible** without risking signal distortion in order to achieve maximum transmission quality.

One possible way to further improve the transmission quality is by having the console operator raise up the levels of the fader during the soft part and bringing it down during the loud parts. This is called *manual gain riding* and is based on a neural feedback path – which means that the operator is constantly monitoring the program material via his hearing system and manually moving the fader when the level drops or increases, all the time trying to achieve a balancing act to keep the levels within the average operational range. Of course it is fairly obvious that this kind of manual control is rather restrictive and laborious; it is difficult to detect signal peaks and it is almost impossible to level them out consistently all the time. Manual control is simply not fast enough to be satisfactory.

The need therefore arises for a fast acting automatic gain control system which will constantly monitor the signals and which will always adjust gain to maximize (or optimize) the signal to noise ratio without incurring signal distortion. This device is called a compressor or limiter. Compressors, limiters, gates, expanders, and other hybrid audio processing components come under the category of **dynamic processors**.

The principle function used in these devices is dependent on an *automatic gain control* device which reduces the amplitude of loud passages in program music and therefore restricts the original dynamics to a desired range. This application is especially useful in microphone recording techniques – to compensate for level changes which are caused by the singer varying his/her distance in front of the microphone during the record process. A *limiter* continuously monitors the signal and intervenes as soon as the level exceeds a certain point (which can be set by the user according to the needs). This level is called *threshold*. Any signal exceeding this threshold will be immediately reduced back by a factor of gain reduction set by the user.

A ***compressor*** also monitors the program material continuously and has a certain threshold level above which it will activate the gain reduction process. However, in contrast to the limiter, signals exceeding the threshold are not reduced abruptly but **gradually**. Above the threshold, the signal is reduced in level relative to the amount the signal exceeds this point. Although compressors and limiters perform similar tasks, one essential point makes them different :

Limiters abruptly limit the signal above a certain level with a very high gain reduction factor. Compressors on the other hand control the signal much more 'gently' over a wider range.

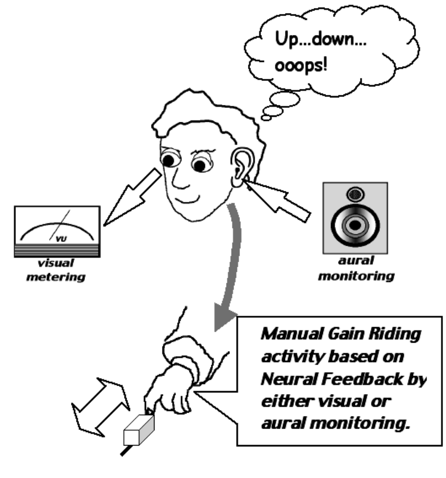

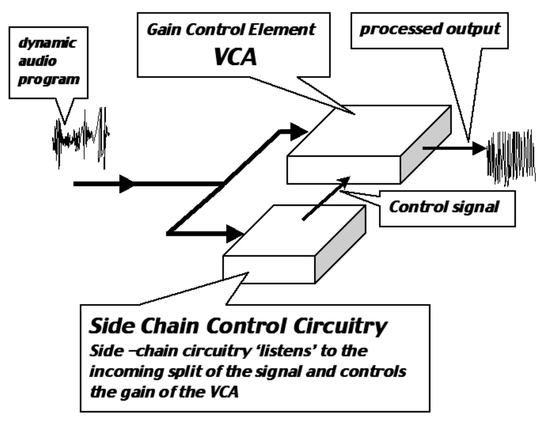

Compressors and Limiters

Unlike the human operator though, the compressor has no feel or intuition; it simply does what you set it up to do, which makes it very important that you understand what all the variable parameters do and how they affect the final sound. In order to react quickly enough, the compressor dispenses with the human ear and instead monitors the signal level by electronic means. A part of the circuit known as the 'side chain' follows the envelope of the signal, usually at the compressor's output, and, uses this to generate a control signal which is fed into the gain control circuit. When the output signal rises past an acceptable level, a control signal is generated and the gain is turned down. Figure 1 shows a simplified block diagram of a typical compressor circuit.

Compressor Settings and Characteristics

Threshold

With manual gain riding, the level above which the signal becomes unacceptably loud is determined by the engineer's discretion: if it sounds too loud to him, he turns it down. In the case of a compressor, we have to 'tell' it when to intervene, and this level is known as the *threshold*. In a conventional compressor, the threshold is varied via a knob calibrated in dBs, and a gain reduction meter is usually included so we can see how much the gain is being modified. If the signal level **falls short** of the threshold, **no processing takes place** and the gain reduction meter reads 0dB. Signals **exceeding** the threshold are reduced in level, and the amount of reduction is shown on the meter. This means the signal peaks are no longer as loud as they were, so in order to compensate, a further stage of 'make-up' gain is added after compression, to restore or 'make up' any lost gain.

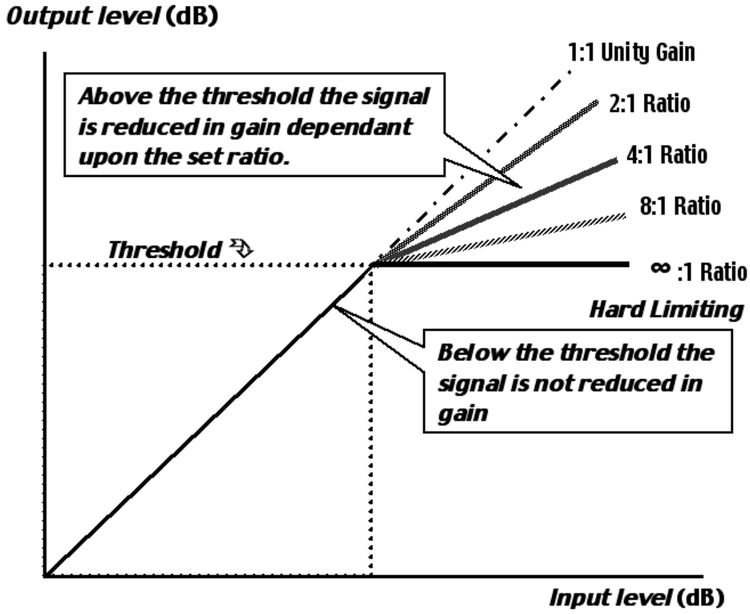

Diagram above shows the input versus output graph of a compressor at various gain reduction ratio settings

Ratio

When the input signal exceeds the threshold set by the operator, gain reduction is applied, but the actual amount of gain reduction depends on the **'Ratio'** setting. You will see the Ratio expressed in the form 4:1 or similar, and the range of a typical Ratio control is variable from 1:1 (no gain reduction all) to infinity (∞) :1, which means that the output level is never allowed to rise above the threshold setting. This latter condition is known as limiting, because the threshold, in effect, sets a limit which the signal is not allowed to exceed. Ratio is based on dBs, so if a compression ratio of 3:1 is set, an input signal exceeding the Threshold by 3dB will cause only a 1 dB increase in level at the output.

In practice, most compressors have sufficient Ratio range to allow them to function as both compressors and limiters, which is why they are sometimes known by both names. The relationship between Threshold and Ratio is shown in the diagram above, but if you're not comfortable with dBs or graphs, all you need to remember is that the larger the Ratio, the more gain reduction is applied to any signal exceeding the Threshold.

Attack

The attack time is *how long a compressor takes to pull the gain down once the input signal has reached or exceeded the Threshold level*. With a fast attack setting, the signal is controlled almost immediately, whereas a slower attack time will allow the start of a transient or percussive sound to pass through unchanged, before the compressor gets its act together and does something about it. Creating a deliberate overshoot by setting an attack time of several milliseconds is a much-used way of enhancing the percussive characteristics of instruments such as guitars or drums. For most musical uses, an initial attack setting of between 1ms and 20 ms is typical. However, when treating sound such as vocals, a fast attack time generally gives the best results, because it brings the level under control very quickly, producing a more natural sound. When compressing wide bandwidth material (for example the entire L-R mix), setting the wrong attack time with slow release time can cause a phenomenon known as *pumping*.

Release

The Release sets *how long it takes for the compressor's gain to come back up to normal (unity gain) once the input signal has fallen back below the Threshold*. If the release time is too fast, the signal level may 'pump'—in other words, you can hear the level of the signal going up and down. This is usually a bad thing, but again, it has its creative uses, especially in rock music. If the release time is too long, the gain may not have recovered by the time the next 'above Threshold' sound occurs. A good starting point for the release time is between 0.2 and 0.6 seconds.

Auto Attack/Release

Some models of compressor have an Auto mode, which adjusts the attack and release characteristics during operation to suit the dynamics of the music being processed. In the case of complex mixes or vocals where the dynamics are constantly changing, the Auto mode may do a better job than fixed manual settings. The release time is adjustable just as with the limiter, but the adjustment is much more critical, because the compressor is constantly changing gain whereas the limiter only cuts in occasionally. It is impossible to find a release time that is perfect for all situations. A very short release time will cause the circuit to attempt to trace the waveform of low frequency components of the signal, resulting in intermodulation distortion of high frequency components; whereas a long release time will leave the amplifier in the wrong mode if the signal changes quickly. That second error results in brief bursts of expansion, a very disconcerting effect. This effects is known as *breathing*.

Peak/RMS operation

Every compressor uses a circuit known as a side chain, and the side chain's job in life is to measure how big the signal is, so that it knows when it needs compressing. This information is then used to control the gain circuit, which may be based around a Voltage-controlled

Amplifier (VCA), a Field Effect Transistor (FET) or even a valve. The compressor will behave differently, depending on whether the side chain responds to average signal levels or to absolute signal peaks.

An RMS level detector works rather like the human ear, which pays less attention to short duration, loud sounds than to longer sounds of the same level. Though RMS offers the closest approximation to the way in which our ears respond to sound, many engineers prefer to work with Peak, possibly because it provides a greater degree of control. And though RMS provides a very natural-sounding dynamic control, short signal peaks will get through unnoticed, even if a fast attack time is set, which means the engineer has less control over the absolute peak signal levels. This can be a problem when making digital recordings, as clipping is to be avoided at all costs. The difference between Peak and RMS sensing tends to show up most on music that contains percussive sounds, where the Peak type of compressor will more accurately track the peak levels of the individual drum beats.

Another way to look at it is to say that the greater the difference between a signal's peak and average level, the more apparent the difference between RMS and peak compression/limiting will be. On a sustained pad sound with no peaks, there should be no appreciable difference. Peak sensing can sometimes sound over-controlled, unless the amount of compression used is slight. It's really down to personal choice, and all judgements should be based on listening tests.

Hold Time

A compressor's side chain follows the envelope of the signal being fed into it, but if the attack and release times are set to their fastest positions, it is likely that the compressor will attempt to respond not to the envelope of the input signal but to individual cycles of the input waveform. This is particularly significant when the input signal is from a bass instrument, as the individual cycles are relatively long, compared to higher frequencies. If compression of the individual waveform cycles is allowed to occur, very bad distortion is audible, as the waveform itself gets reshaped by the compression process. We could simply increase the release time of the compressor so that it becomes too slow to react to individual cycles, but sometimes it's useful to be able to set a very fast release time.

A better option is to use the Hold time control, if you have one. Hold introduces a slight delay before the release phase is initiated, which prevents the envelope shaper from going into release mode until the Hold time has elapsed. If the Hold time is set longer than the duration of a single cycle of the lowest audible frequency, the compressor will be forced to wait long enough for the next cycle to come along, thus avoiding distortion. A Hold time of 50ms will prevent this distortion mechanism causing problems down to 20Hz. If your compressor doesn't have a separate Hold time control, it may still have a built-in, preset amount of Hold time. A 50ms hold time isn't going to adversely affect any other aspect of the compressor's operation, and leaves the user with one less control to worry about.

Knee Control

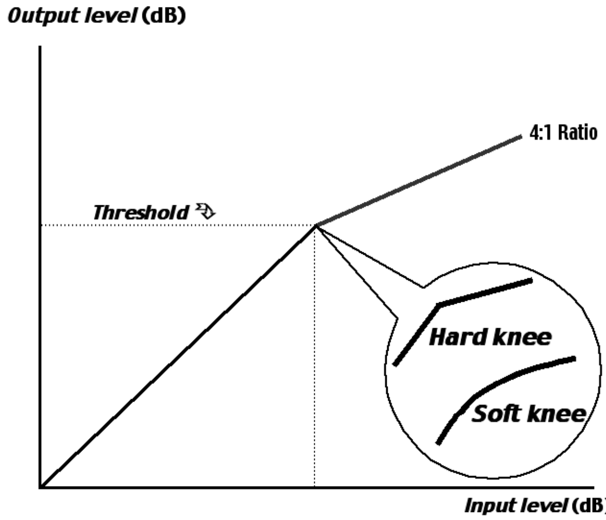

This is not a control or parameter, but rather a characteristic of certain designs of compressor. With a conventional compressor, nothing happens until the signal reaches the Threshold, but as soon as it does, the full quota of gain reduction is thrown at it, as determined by the Ratio control setting. This is known as **hard knee compression**, because a graph of input gain against output gain will show a clear change in slope (a sharp angle) at the Threshold level, as is evident in the diagram to your right.

Other types of compressor utilise a **soft knee characteristic**, where the gain reduction is brought in progressively over a range of 10dB or so. What happens is that when the signal comes within 10dB or so of the Threshold set by the user, the compressor starts to apply gain reduction, but with a very low Ratio setting, so there's very little effect. As the input level increases, the compression Ratio is automatically increased until at the Threshold level, the Ratio has increased to the amount set by the user on the Ratio control. This results in a gentler degree of control for signals that are hovering around the Threshold point, and the practical outcome is that the signal sounds less obviously processed. This attribute makes soft-knee models popular for processing complete mixes or other sounds that need subtle control. Hard knee compression can sometimes be heard working, and if a lot of gain reduction is being applied, they can sound quite heavy-handed. In some situations, it can make for an interesting sound—take Phil Collins' or Kate Bush's vocal sounds, for example.

Stereo Link

When processing stereo signals, it is important that both channels are treated equally, for the stereo image will wander if one channel receives more compression than the other. For example, if a loud sound occurs only in the left channel, then the left channel gain will be reduced, and everything else present in the left channel will also be turned down in the mix. This will result in an apparent movement towards the right channel, which is not undergoing so much gain reduction. The Stereo Link switch of a dual-channel compressor simply forces both channels to work together, based either on an average of the two input signals, or whichever is the highest in level at any one time. Of course, both channels must be set up exactly the same for this to work properly, but that's taken care of by the compressor. When the two channels are switched to stereo, one set of controls usually becomes the master for both channels—though some manufacturers opt for averaging the two channel's control settings, or for reacting to whichever channel's controls are set to the highest value.

All in the Ear

As you go through this course, you may discover during your sessions or at least read about, the fact that different makes of compressor sound different. But if all they're really doing is changing level, shouldn't they all sound exactly the same? As we've already learned, part of the reason is related to the shape of the attack and release curves of the compressor, and of course peak sensing will produce different results to RMS, but at least as important is the way in which a compressor distorts the signal. Technically perhaps, the best compressor is one that doesn't add any distortion, but most engineers seem to like the 'warm' sound of the older valve designs which, on paper, are blighted by high distortion levels. The truth is that low levels of distortion have a profound effect on the way in which we perceive sound, which is the principle on which aural exciters work. A very small amount of even-harmonic distortion can tighten up bass sounds, while making the top end seem brighter and cleaner. The best-sounding contemporary compressor designs include valve models with a degree of distortion built in, while others use FETs (Field Effect Transistors – a type of amplification circuit), which mimic the behaviour of valve circuits. As digital recorders and mixers are introduced into the signal chain, more people are becoming interested in equipment that can put the warmth back into what they perceive as an over-clinical sound.

Using Compressors

One problem newcomers to recording seem to have is deciding where in their system to patch the compressor. A compressor is a processor rather than an effect, so it should be used via an insert point or be patched in-line with a line-level signal.

Most engineers will normally add some compression to vocals while recording, and then add more if necessary while mixing. Working this way makes good use of the tape's dynamic range, while helping to prevent signal peaks from overloading the tape machine. It is best to use rather less compression than might ultimately be needed while recording, so that a little more can be added at the mixing stage if required. If too much compression is added at the beginning, there's little you can do to get rid of it afterwards. Similarly, if you have a compressor with a gate built-in, it might be better to leave this off when recording, and only use it while mixing. This will prevent a good take from being wrecked by an inappropriate gate setting. A further benefit of gating during the mix is that the gate will remove any tape hiss, along with the original recorded noise. If a gate is allowed to close too rapidly, it can chop off the ends of wanted sounds that have long decays, especially those with long reverb tails, so most gates (and expanders) fitted to compressors have either a switchable long/short release time, or a proper variable-release time control.

De-Essing

Another side chain-related process is the de-essing of sibilant vocal sounds. Sibilance is sometimes evident when people pronounce the letters 's' or 't', and is really a high-pitched whistling caused by air passing around the teeth. If a parametric equaliser is inserted into the side-chain signal path of a compressor and tuned to boost the offending frequency, the compressor will apply more gain reduction when sibilance is present than at other times.

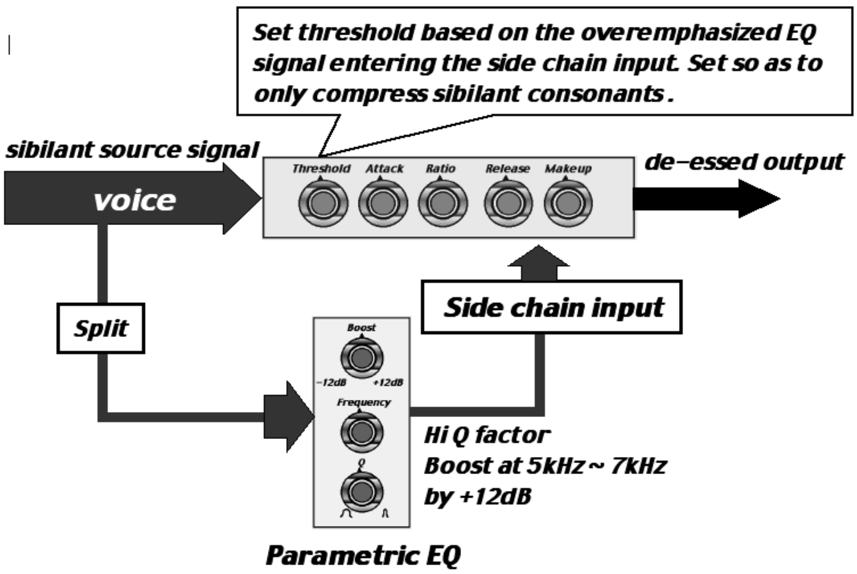

Most sibilance occurs in the 5kHz to7kHz region of the audio spectrum, so if the equaliser is tuned to this frequency range and set to give around I2dB of boost, then in the selected frequency range, compression will occur 1OdB before it does in the rest of the audio spectrum. The equaliser should be set up by listening to the equaliser output, and then tuning the frequency control until the sibilant part of the input signal is strongest. The diagram above shows how a compressor and equaliser may be used as a de-esser. Some compressors have a built-in sweep equaliser, to allow them to double as de-essers without the need for an external parametric equaliser.

Envelope Shaping

Not only can the compressor be used to limit the dynamic range, it can be also used creatively to shape the envelope of the sound passing through it. Being able to visualize how the envelope of various sound sources (instruments as well as voice) would be helpful in reaching to appropriate settings on a compressor to shape the envelope's attack, decay, and sustain. For example, a natural kick drum will sound 'flabby' (fast attack but long decay) due to the resonances inside the shell. It would be suitable for a jazz recording but not for pop music production which requires a 'thumping' and 'punchier' sounding kick. By gating the decay of the kick drum and applying compression to emphasize its attack, this can be achieved.

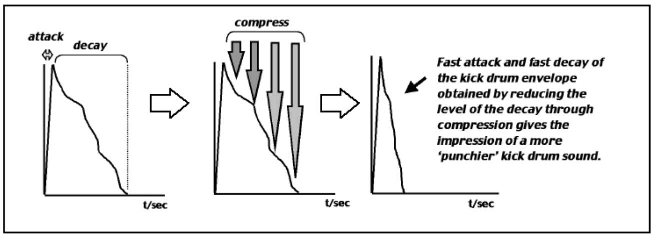

By using a gate to reduce the decay of a kick drum envelope and then appyling compression to emphasize the attack, the kick drum can be made to sound more percussive and 'punchier'.

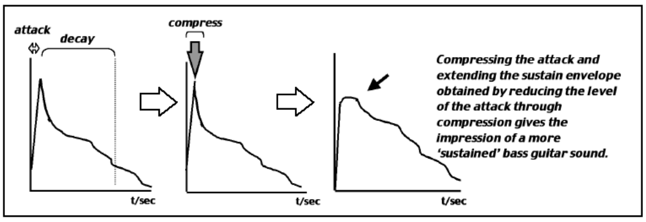

By using a compressor to reduce the transient attack of a bass guitar envelope and then applying makeup gain to the overall output, the decay would be apparently at the same loudness as the attack, allowing the listener to perceive a more sustained (longer) tone.

Another example of envelope shaping is applying compression on a direct injected (DI) bass guitar tone. In normal circumstances in which the bass guitar tone comes from a bass amp, the mass of the speaker cone itself makes it very difficult to produce the inherent attack of the bass guitar envelope. So when a bass guitar tone is heard through an amp, the speaker cone acts as a natural 'compressor', rounding off the transient attack to produce a much more rounder and sustained bass tone. However, in a direct injected guitar tone, the tone that would be produced will be inherently transient (pluck) and would have a shorter decay and sustain rate. By utilizing a compressor to roll off the transient attack of the envelope, we can emulate how a bass guitar amp reacts. By compressing the attack, the dynamics of the transient attack would be brought down to be almost as equal as the decay and sustain parts of the envelope. By raising the makeup gain, the overall decay of the bass guitar tone can be brought up to create a longer sustain effect.

General Guidelines

For some general advice on compression settings, take a look at the 'Useful Compressor Settings' box below.

USEFUL COMPRESSOR SETTINGS

SOURCE	ATTACK	RELEASE	RATIO	HARD/SOFT	GAIN RED
Vocal	Fast	0.5s/Auto	2:1 - 8:1	Soft	3 - 8dB
Rock vocal	Fast	0.3s	4:1 - 10:1	Hard	5 - 15dB
Acc guitar	5 - 10ms	0.5s/Auto	5 - 10:1	Soft/Hard	5 - 12dB
Elec guitar	2 - 5ms	0.5s/Auto	8:1	Hard	5 - 15dB
Kick and snare	1 - 5ms	0.2s/Auto	5 - 10:1	Hard	5 - 15 dB
Bass	2 - 10ms	0.5s/Auto	4 - 12:1	Hard	5 - 15dB
Brass	1 - 5ms	0.3s/Auto	6 - 15:1	Hard	8 - 15dB
Mixes	Fast	0.4s/Auto	2 - 6:1	Soft	2 - 10dB (Stereo Link On)
General	Fast	0.5s/Auto	5:1	Soft	10dB

It should be should stress that these are just to get you started—the ideal settings vary from compressor to compressor. The more gain reduction is used, the higher the level of background noise, so never use more gain reduction than is necessary. Virtually all recorded pop music has a deliberately restricted dynamic range, to make it sound loud and powerful when played over the radio. The more a signal is compressed, the higher its average energy level. In addition to compressing the individual tracks during recording or mixing, the engineer may well have applied further compression to the overall mix. This can be very effective, but don't choke the life out of a mix by over-compressing it either. When it comes to individual tracks, it is pretty much routine to compress vocals, bass guitars, acoustic guitars and occasionally electric guitars, though overdriven guitar sounds tend to be self compressing anyway!

The most important of these to get right is the lead vocal, because even modest dips in level can make the lyrics difficult to hear over the backing. Sequenced instruments are less likely to need compression, because you can control the dynamics by manipulating the MIDI data in the sequencer. Most seasoned engineers try to avoid compression (or any other form of treatment) unless it's absolutely necessary. Even with vocals, if somebody gives a perfectly controlled vocal take, you wouldn't want to compress it just because compressing vocals is the in thing. Compression is a very valuable studio tool, but like all tools, it is just a means to an end—not an end in itself.

Side Effects of Compression

Most of the sound energy in a typical piece of music occupies the low end of the audio spectrum, which is why your VU meters always seem to respond to the bass drum and bass guitar. High frequency sounds tend to be much lower in level and so rarely need compressing, but even so, high-frequency sounds in the mix are still brought down in level whenever the compressor reacts to loud bass sounds. For example, a quiet hi-hat occurring at the same time as a loud bass drum beat will be reduced in level. One technique to reduce the severity of this effect is to set a slightly longer attack time on the compressor, to allow the attack of the hi-hat to get through before the gain reduction occurs. This is only a partial

solution, and if heavy compression is applied to a full mix, the overall sound can become dull, as the high-frequency detail is reduced in level.
Going back to the subjective effect of subtle harmonic distortion for a moment, some compressor designs make use of harmonic distortion or dynamic equalisation to provide an increase in high-frequency level whenever heavy compression is taking place. This helps offset the dulling of high-frequency detail, and can make a great subjective difference, but it isn't a perfect solution. More elaborate compressors have been designed which split the signal into two or more frequency bands and compress these separately. This is known as *multiband compression*. This neatly avoids the bass end causing the high end to be needlessly compressed, but it can introduce other problems related to phase, unless the design is extremely well thought-out.

DANGERS OF COMPRESSION IN LIVE SOUND

Compressors are wonderful allies, but there are two or three points you need to watch. If you're using a compressor to even out vocal levels, you're effectively reducing the peak signal level, so to get back to the same peak signal level you had before, you're going to have to turn the mic gain up. More mic gain means a greater likelihood of feedback problems, so rather than expecting too much of the compressor, it's probably best to set the threshold high with a fairly high ratio and use it just to tame the singer's worst excesses of level. In large concert situations where the PA is further away from the mics, this may be less of a problem, but if your setup is more 'Claire's brother' than Clair Brothers, it's something to keep in mind. By the same token, using compression to increase the average signal level means that low level noise will also be increased in level, so a compressor with a built-in gate or expander can help.

Finally, simply because compression does increase average signal level, it means your PA speakers are going to be handling a higher average power level. Although your limiters should protect the drivers from the results of clipping, a high average power means a lot of heat to dissipate. In extreme circumstances, running highly compressed signals at high levels can cause thermal failure of components, so if you do like a dense, compressed sound, make sure your speaker system is up to it.

Summary

For the benefit of those who are still a little unsure as to what a compressor does, it simply reduces the difference between the loudest and quietest parts of a piece of music by automatically turning down the gain when the signal gets past a predetermined level. In this respect, it does a similar job to the human hand on the fader—but it reacts much faster and with greater precision, allowing it to bring excessive level deviations under control almost instantaneously.

Expanders

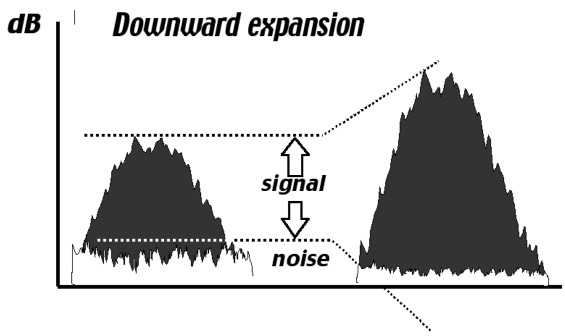

Expansion is the opposite of compression. It is accomplished with the same circuit as compression, but the rules are changed so that *any signal with amplitude above the threshold will be amplified and any signal with amplitude lower than the threshold will be attenuated. Expanders* are signal processing units used to increase (*expand*) the dynamic range of the signal passing through it. However, modern expanders operate only *below the set threshold point*, that is, they operate only on low-level audio. Operating in this manner they make the *quiet parts quieter*. The term *downward expander* or *downward expansion* evolved to describe this type of application. The most common use is noise reduction. For example, say, an expander's threshold level is set to be just below the quietest vocal level being recorded, and the ratio control is set for 2:1. What happens is this: when the vocals stop, the signal level drops below the set point down to the noise floor. There has been a step decrease from the smallest signal level down to the noise floor. If that step change is, say, -10 dB, then the expander's output attenuates 20 dB (i.e., due to the 2:1 ratio, a 10 dB decrease becomes a 20 dB decrease), thus resulting in a noise reduction improvement of 10 dB. It's now 10 dB quieter than it would have been without the expander.

Noise Gates

Sometimes a program is encountered which has proper dynamic range but which has objectionable background noise. Very little can be done to reduce the noise without affecting the signal, but the spaces where there is no signal can often be cleaned up with a *noise gate*. Noise gates (or *gates*) are expanders with fixed "infinite" downward expansion ratios. They are used extensively for controlling unwanted noise, such as preventing "open" microphones and "hot" instrument pick-ups from introducing extraneous sounds into your system. When the incoming audio signal drops below the *threshold* point, the gate prevents further output by reducing the gain to "zero." Typically, this means attenuating all signals by about 80 dB. Therefore once audio drops below the threshold, the output level basically becomes the residual noise of the gate. Common terminology refers to the gate "opening" and "closing." A gate is the extreme case of downward expansion.

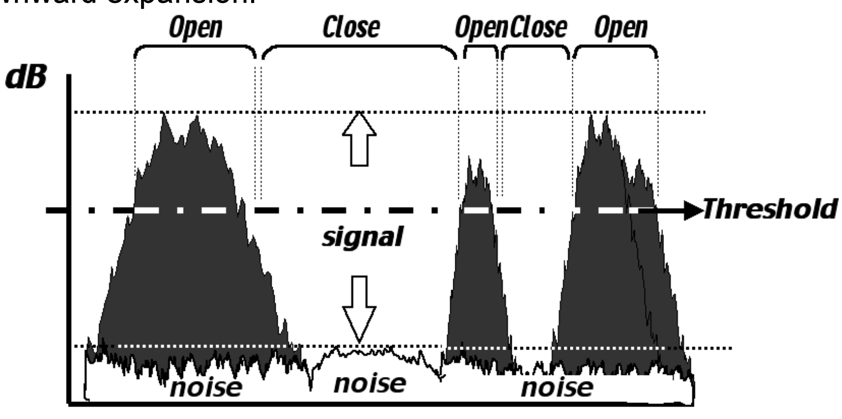

The action of a gate opening to let signal pass through when it goes above the threshold and closing when the signal goes below the threshold.

Just as poorly designed limiters can cause pumping, poorly designed gates can cause *breathing*. The term *breathing* is used to describe an audible problem caused by being able to hear the noise floor of a product rise and lower, sounding a lot like the unit was "breathing." It takes careful design to get all the dynamic timing exactly right so breathing does not occur. Another popular application for noise gates is to enhance musical instrument sounds, especially percussion instruments. Correctly setting a noise gate's *attack* (turn-on) and *release* (turn-off) adds "punch," or "tightens" the percussive sound, making it more pronounced — this is how Phil Collins gets his cool snare sound, for instance.

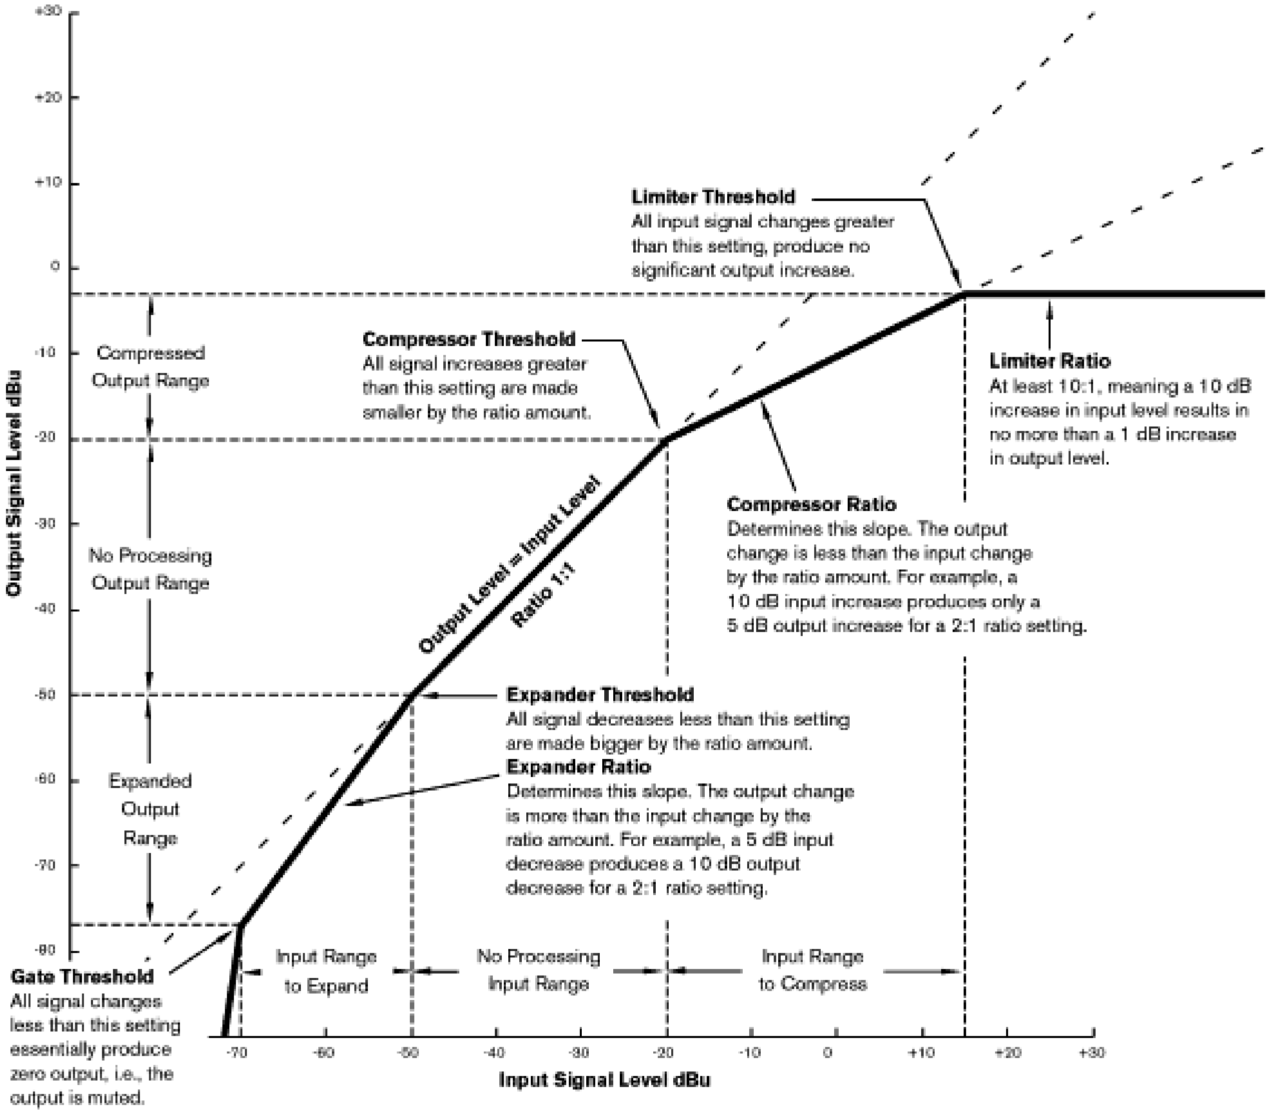

Graph Displaying Comparison between Gate/Expander/Compressor/Limiter action

Compression Techniques

Look Ahead Compression (Anticipated Compression)

Case 1

The dynamic transients in an instrument or program music are too fast for the attack time on the compressor to act upon.

Solution

Delay the input signal while supplying an undelayed signal to the external side chain (key input). This enables the compressor to ‘see’ beforehand the signal that is to come into its circuit. With experimentation, this process can create a ‘zero’ attack time at a given frequency.

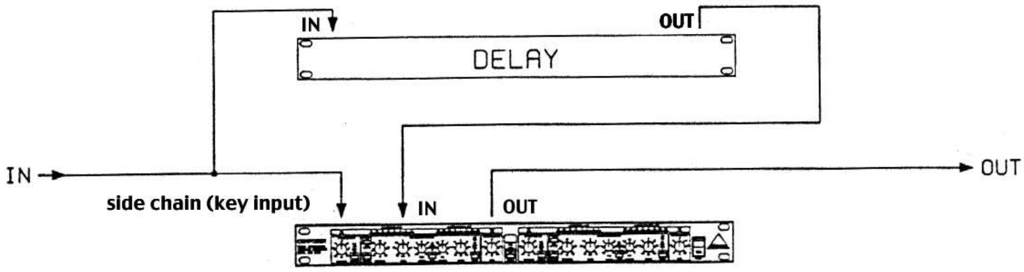

Voice Over Compression (Ducking)

Case 2

You are a disc jockey and require the music level to be automatically brought down to a lower level whenever you speak and brought up again when you are not speaking.

Solution

By routing a split of the disc jockey microphone output to the external side chain (key input), the VCA can be manipulated to activate compression whenever the external side chain input signal exceeds a set threshold. In this configuration, whenever the disc jockey speaks, the compressor will reduce the gain of the incoming source by a predetermined ratio set by the user.

Program Music

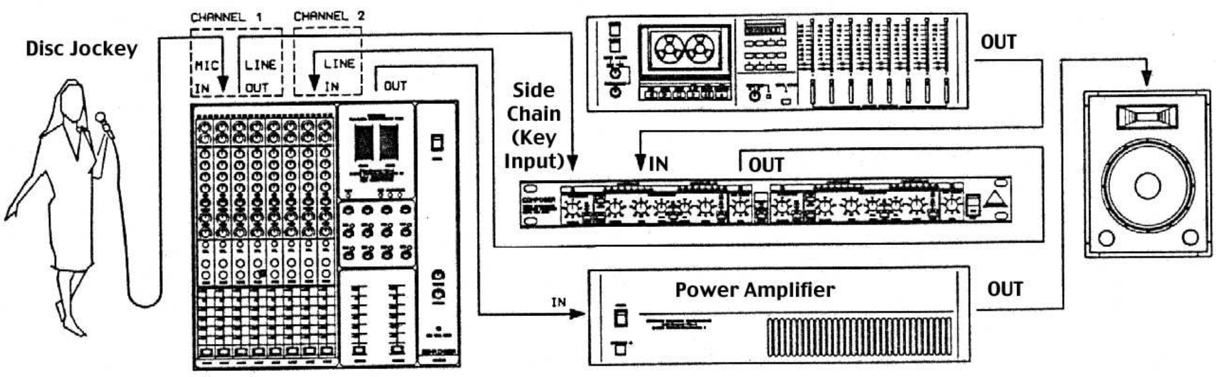

Blanking

Case 3

The kick drum that you recorded sound flat. You want to add a bit of body to the existing kick drum track to give the rhythm track a bit more ‘punch’ by adding synchronous impulses of a low frequency sine wave tone, in time with the kick, and recording it to another track. To do this :-

- ☐ Send an oscillator tone to the audio input of the compressor.
- ☐ Send a split of the kick track via auxiliaries to the side chain of the compressor

By routing a split of the drum track signal to the external side chain (key input), the VCA can be manipulated to open whenever the external side chain input signal exceeds a set threshold. By having a very fast gate close rate, the sine tone signal entering the gate input will be impulsive based on the opening and closing action of the gate triggered by the side chain. The audio out put of the gate can then be routed to another track and recorded. During the mixdown process, the kick drum and the sine tone pulses would be heard together simultaneously, reinforcing the weak kick drum.

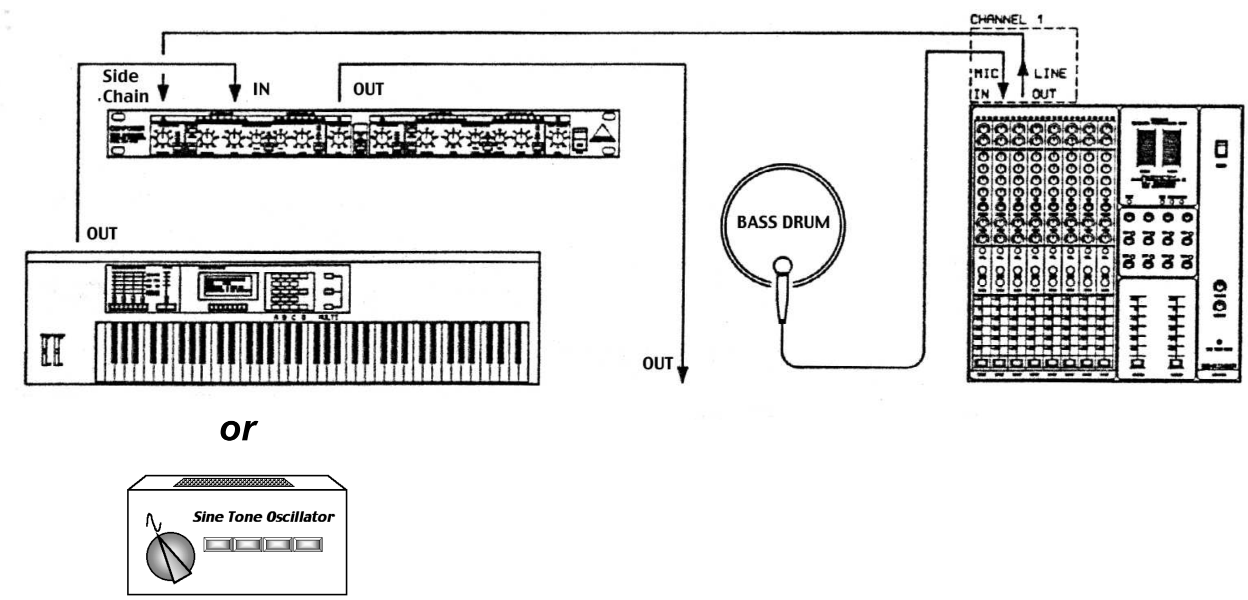

Gated Reverb

By gating reverb and setting a fast gate close time, the decay of the reverb can be 'chopped off' abruptly halfway through, creating an 'explosive' like effect. This is how Phil Collins gets his tom-toms and snare to sound larger than life.

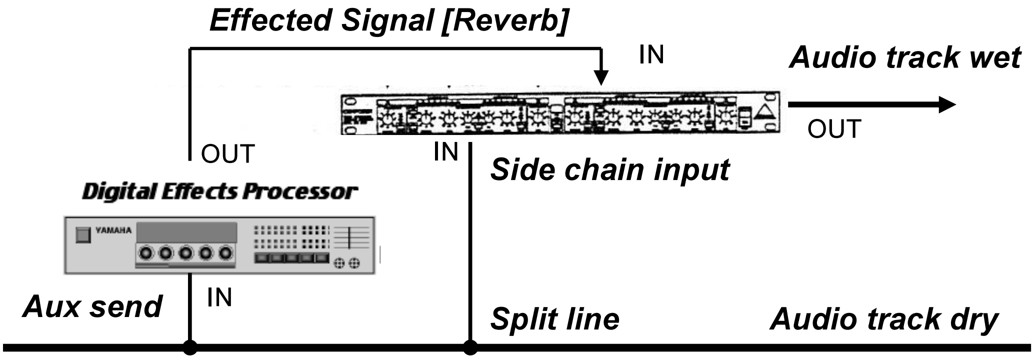

Effects, Signal Processors & Noise Reduction – Equalization

Signal Processing and Processors in Brief
The devices collectively known as audio processors were originally added to the recording studio to allow compensation for frequency response or dynamic range problems in the equipment. When carefully used, they can add to the fidelity of the sound, and occasionally improve on reality. If improperly used, they can seriously degrade the sound or even produce laughable results.

Processing Methodologies

There are two types of ways in which a signal can be process.

(i) *serial processing* *(ii)* *parallel processing*

the signal may be processed also by combining both serial and parallel methods the achieve the desired result (i.e. a gated reverb). The diagram below shows the signal flow of both processing methods.

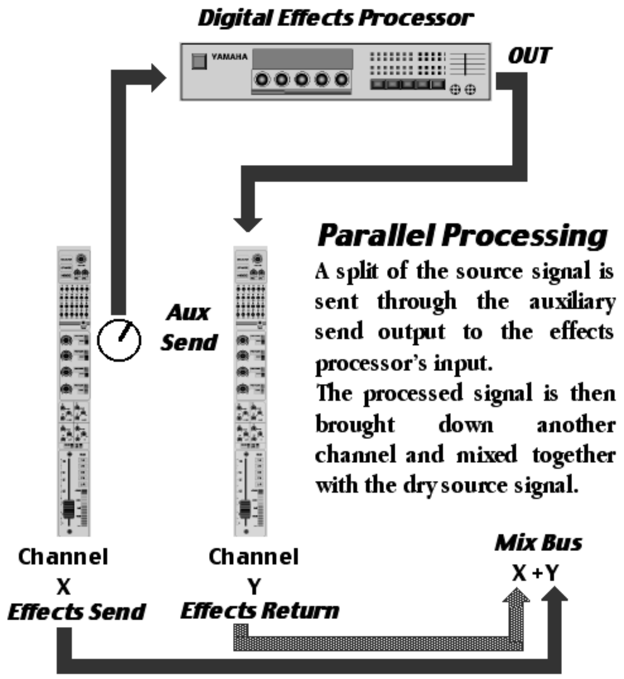

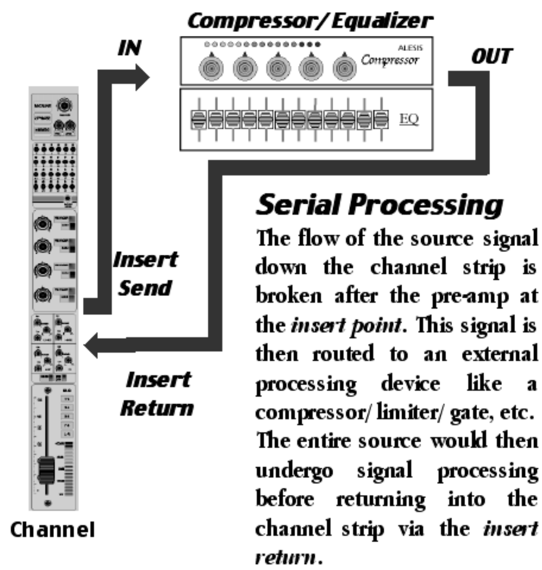

Defining Equalization (EQ)

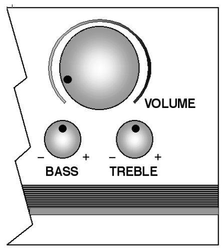

Simple Bass and Treble controls on consumer amplification systems

An equalizer can be defined as a frequency selective amplifier. It is a device that *can alter the spectral content of a signal that passes though it*. It is used to compensate for variations or discrepancies in frequencies present in audio signal. This can be done with any circuit that has an adjustable frequency response, the most familiar being the tone controls on a home stereo set. These tone controls typically affect the amplitude in two frequency regions, the treble and bass. This is sufficient for the minor changes the end user may wish to make in the program, but the recording engineer needs more flexibility and coverage of the entire audio spectrum. . Equalizers come in all different sizes and shapes, varying greatly in design and complexity.

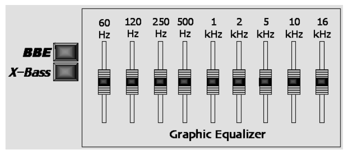

9 band graphic equalizer on consumer hifi mini component systems/ car audio systems.

At its simplest, the equalizer allows the engineer to accentuate or attenuate (cut or boost in simplified terms) any frequency or group of frequencies within the audio spectrum. By the same process, unwanted frequencies can be filtered out. This control over the spectrum of sound gives the engineer a good deal of creative freedom in changing timbre and harmonic balance. So, equalizers are used not only to correct a particular sound, but also as a *creative tool*. For example, you can use them during overdubs to match the sound of another instrument; or you can use them to control the relative balance of frequencies between different instruments in a mix to enable a better and cohesive blend between elements in a song without resorting to great level changes. Alternately, an equalizer can increase separation between instruments by helping to 'position' each instrument in the virtual three dimensional stereo image. So much said, the tasks of an equalizer range all the way from analyzing and improving a control room's acoustics, to getting that elusive sound which the producer can hear in his head, but never quite describe in words.

Frequency Response

The term **frequency response** describes how a particular environment, whether it be acoustical or electrical, affects the signal within it or passing through it – that is whether it amplifies or attenuates particular frequencies. In discussing equalizers, a **frequency response plot** is other used. It displays the amount of energy (amplitude) at any given frequency compared to unity gain. The vertical (Y) axis of the plot displays the amplitude information, while the horizontal (X) axis displays the frequency bandwidth being measured.

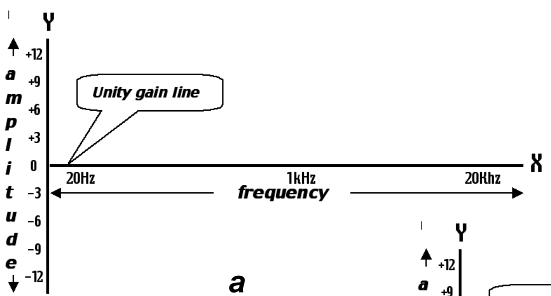

a

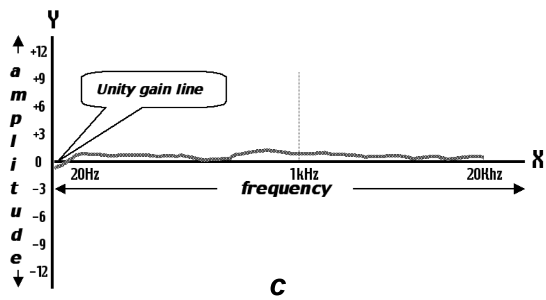

c

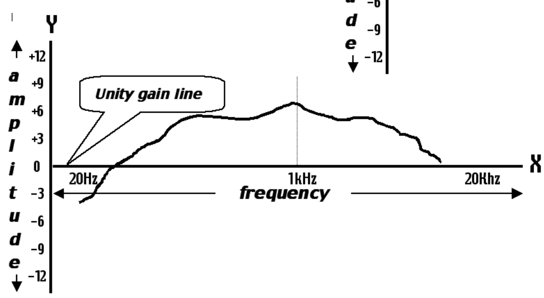

b

The straight horizontal line (*usually labelled zero*) shows the amplitude level of the incoming signal – unaffected by the EQ. It represents *unity gain*. As shown in diagram (a), to the far left of the X axis will be 20Hz (or lowest frequency to be displayed). To the far right will be 20kHz (or the highest frequency to be displayed). A graphical plot of frequency versus amplitude is then taken with reference to zero or unity gain and the resultant curve can then be plotted over the frequency -vs- amplitude graph as shown in diagram (b) and (c). This would show how the EQ is altering the signal away from the unity gain. If the curved line is above the zero line, then the EQ is said to be adding information. If the curved line is below the zero line, the EQ is then said to be subtracting information. Looking at diagram (b), we can see from the graph that 1kHz is boosted by +6dB. This will allow us to conclude that any 1kHz information entering this EQ unit would then be boosted by 9 dB at the output. Diagram (c) shows a much more moderate and almost flat response EQ which would affect the signal flowing through it to a lesser extent.

Frequency response can be applied to all types of electronic equipment, especially those dealing with the input and output of audio information such as amplifiers, speakers, microphones, tape recorders, and so on. A loudspeaker manufacturer may advertise that their product has a frequency response from 20Hz to 20kHz – which would be very good. But here's the catch… does it produce all the frequencies between 20Hz and 20kHz equally in intensity? Saying that the frequency response is 20Hz to 20kHz is not adequate; the important concern is how much variation exists between those two frequency extremes. To make the frequency values more specific and meaningful, a qualifier is added to specify the amount of variance from unity gain. For example a spcification might read 20Hz to 20kHz, ± 3dB. This tells us that this speaker would in fact reproduce the full audio bandwidth as claimed, never varying more that 3dB below or above the unity gain. Quite often, the term **linear** is used to describe a circuit that passes all signals at (or very close to) unity gain (refer to diagram C above). Linear devices tend to put out exactly what is put into them.

Real Time Analysis

A Realtime Analyzer (RTA) can be used to display a frequency response plot. A microphone with a very falt response is set up and connected to the RTA. Based on the signal from the microphone, the RTA then displays the amplitude of each of 31 different frequencies. This provides the engineer with a contstant display of frequency response. Any EQ changes made will be visible on the RTA display (sometimes an RTA is also called a *spectrum analyser*).

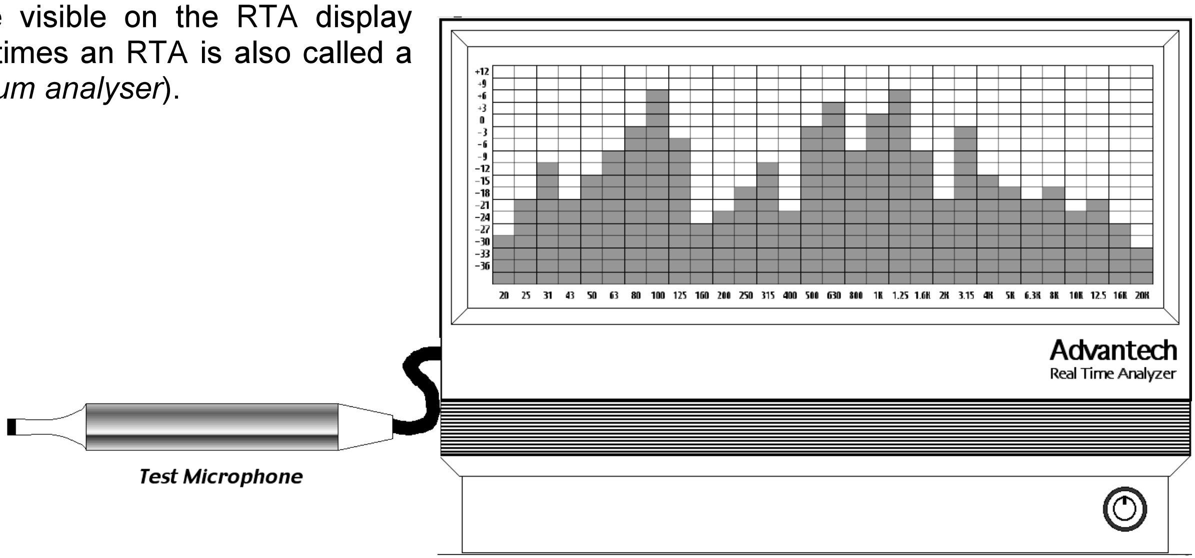

Equalizer Types and Characteristics

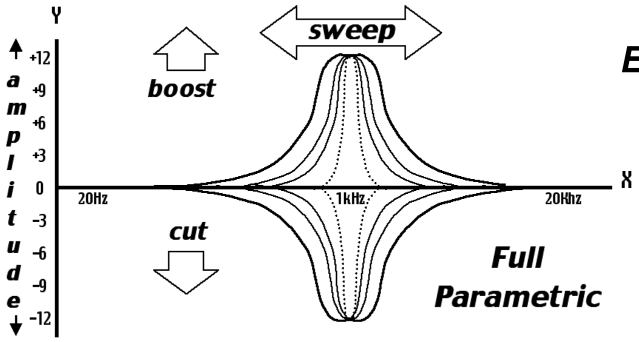

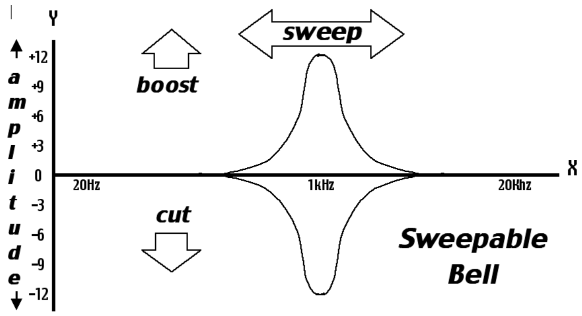

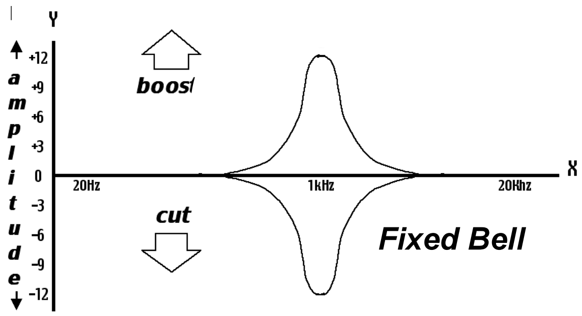

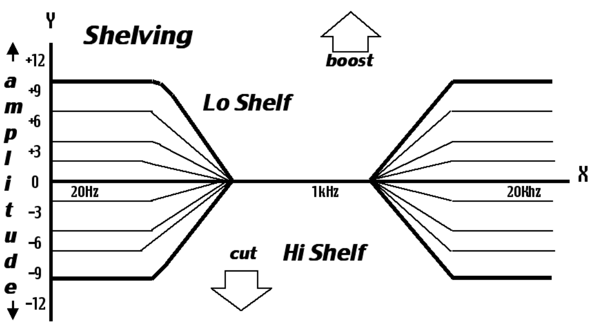

Bandwidth

The definition of **bandwidth** is the distance (or the frequency bandwidth) between the 3dB down points. Bandwidth is measured in octaves.

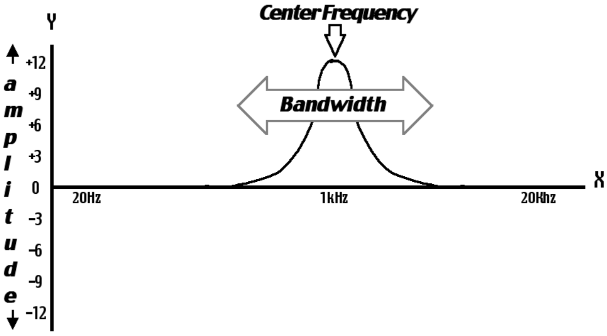

Some EQs have Q control instead of bandwidth control. This accomplishes the same thing, but is measured differently. Q is determined by dividing the centre frequency by difference between the upper 3dB down point and the lower 3dB down point or as shown in the formula on the below and right.

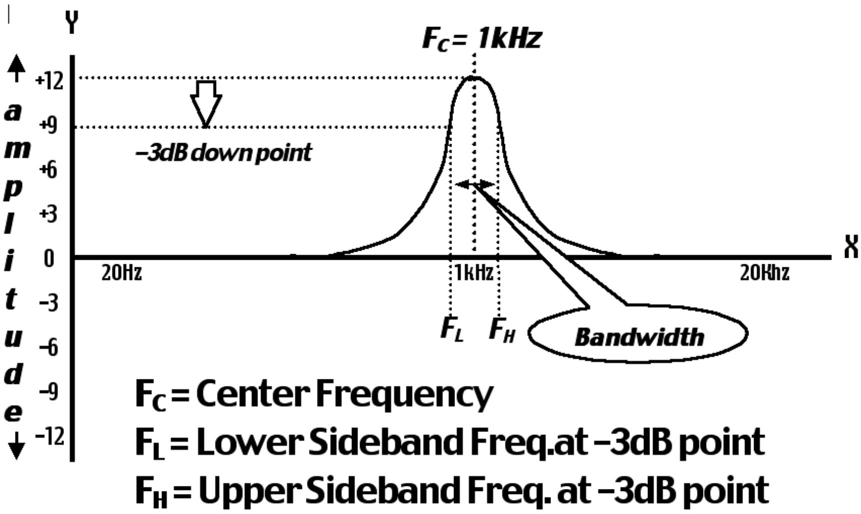

$$\mathbf{BW}_{\text{(Bandwidth)}} = F_H - F_L$$

$$\mathbf{Q} = \frac{F_C}{F_H - F_L}$$

Parametric Equalizers

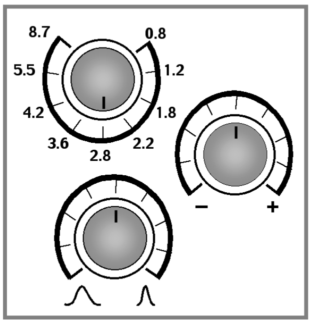

A parametric EQ is named for the user's ability to adjust all of the parameters of equalization. The parametric EQ allows manipulation of amplitude centre frequency, and bandwidth. In true parametric EQ, all four bands will have continuously variable amplitude (cut/boost) frequency, and bandwidth controls. Many manufacturers though create 'semi-parametric' or 'quasi parametric' where either one of these variables would be selectable between preset values.

Most parametric Eqs are divided into four bands: High Frequency (HF), High Mid Frequency (HMF), Lo Mid Frequency (LMF) and Low Frequency (LF).

Amplitude (Cut/Boost)
The amplitude controls provide a certain amount of cut or boost at the selected frequency range, usually affecting the centre frequency by ±12dB to ±18dB.

Frequency
The frequency control allows the engineer to tune in to a particular centre frequency. Each band provides a sweepable range that overlaps its neighboring bands.

Bandwidth
The bandwidth control determines what frequencies below and above the centre frequency will also be affected by the cutting or boosting process. A wide bandwidth setting will cause a wide range of sidebands (adjacent bands) to be also affected by any action on the centre frequency. A narrow bandwidth setting would affect the sidebands to a lesser extent. The general rule of thumb in adjusting bandwidth values is to use **wide range while boosting** and **moderately narrow when cutting** to obtain a natural frequency response.

Bell Curve (Peaking EQ)

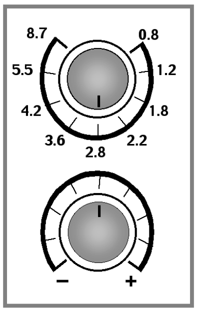

The bell curve characteristic (or peaking characteristic) has the most affect on the center frequency/. The information above or below the center frequency (sidebands) is affected according to the width of the bell. On the Mackie 8•Bus, the Low Midrange section has a sweepable bell characteristic – meaning that it offers fixed bandwidth setting but allows the center frequency to be swept across a specific range.

Shelving EQ

An EQ with a shelving characteristic works differently than a bell characteristic. A shelving EQ will treat all frequencies equally – an even amount of boost or cut for all frequencies beyond or before its knee frequency (cut-off frequency).\In a low frequency shelving EQ, all information below the selected frequency is affected equally. In a high frequency shelving EQ, all the information above the selected frequency is affected equally.

The complete studio will have complex EQ systems that might include 30 or more regions of control. Many of these machines have sliders to adjust the amplitude of each band, and those sliders are laid out in such a way that their positions visually indicate the frequency response. This feature gives rise to the name

Graphic Equalizer

The graphic equalizer derives its name from its front panel layout, which reflects its functions graphically by dividing the audio spectrum in to a number of separate frequency bands which can be modified independently. The 31 band graphic EQ is not commonly found on a mixing console as it takes up too much of space and its does not meet engineering need. A 31- band equalizer is going to be an expensive device, simply because of the sheer duplication of circuitry. Much of this circuitry is wasted most of the time because typical use only involves adjustment to two or three bands while the others are left alone. An equalizer with bands of adjustable frequency is a more economical approach to the problem, because only three to five circuits are necessary. Far and away, the simplest and most popular are the 1/3- and 2/3-octave graphics. They offer the best combination of control, complexity and cost.

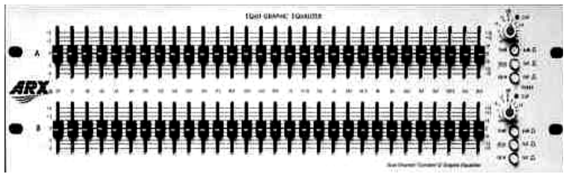

A dual channel 31 band graphic equalizer

In selecting graphic equalizers, the primary features to consider are the number of input/output channels, the number of boost/cut bands, the center-frequency spacing of each, and the *bandwidth behavior*. This last one may at first seem a bit odd, but it is perhaps the most important characteristic. Bandwidth behavior is either *constant-Q* or *variable-Q* (see diagrams). The *quality factor*, or *Q*, of a circuit relates to its bandwidth in an *inverse* manner. That is, narrow bandwidths result from high-Q circuits and wide bandwidths come from low-Q circuits. In the early `80s, Rane developed the first constant-Q designs to preserve the same shape (bandwidth) over the entire boost/cut range. In contrast, variable-Q designs have varying bandwidths (the shape changes) as a function of boost/cut amount. They start out very wide for small amplitude changes and become quite narrow for large changes. Rane's constant-Q design became the most popular, and changed the industry.

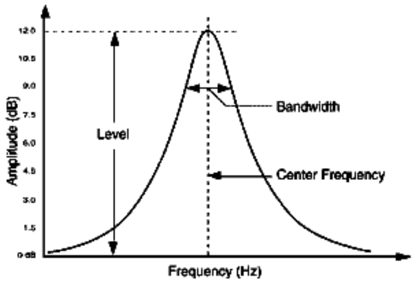

Bandpass Filter Parameters

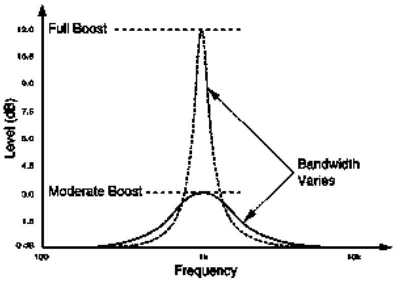

Variable –Q Graphic

Using Equalizers

Equalizers can do wonders for a sound system. Let's start with loudspeaker performance. An unfortunate truth regarding budget loudspeakers is they don't sound very good. Usually this is due to an uneven frequency response, or more correctly, a non flat *power response.* An ideal cabinet has a flat power response. say, 1 kHz as a reference signal, use it to drive the speaker with exactly one watt, measure the loudness, and sweep the generator over the speaker's entire frequency range, *all frequencies will measure equally loud.* Sadly, with all but the most expensive speaker systems, they will not. Equalizers can help these frequency deficiencies. By adding a little here and taking away a little there, pretty soon you create an

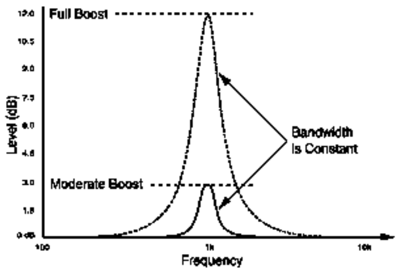

Constant – Q Graphic

acceptable power response — and a whole lot better sounding system. It's surprising how just a little equalization can change a poor sounding system into something quite decent.
The best way to deal with budget speakers — although it costs more — is to commit *one* equalizer channel for *each* cabinet. This becomes a marriage. The equalizer is set, a security cover is bolted-on, and forever more they are inseparable. (Use additional equalizers to assist with the room problems.) And now for the hard part, but the most important part: If you do your measurements *outside* (no reflections off walls or ceiling) and *up in the air* (no reflections off the ground) you can get a very accurate picture of just the loudspeaker's response, free from room effects. This gives you the *room-independent* response. This is really important, because *no matter where this box is used, it has these problems*. Of course, you must make sure the cost of the budget speaker plus the equalizer adds up to substantially less than buying a really flat speaker system to begin with. Luckily (or should this be *sadly*) this is usually the case. Again, the truth is that *most* cabinets are not flat. It is only the very expensive and well designed loudspeakers that have world-class responses.
The next thing you can do with equalizers is to improve the way each venue sounds. Every room sounds different -- fact of life -- fact of physics. Using exactly the same equipment, playing exactly the same music in exactly the same way, different rooms sound different -- guaranteed. Each enclosed space treats your sound differently.

Using Equalizers to Tune Room Response

Reflected sound causes the problems. What the audience hears is made up of the *direct* sound (what comes straight out of the loudspeaker directly to the listener) and *reflected* sound (it bounces off *everything* before getting to the listener). And if the room is big enough, then *reverberation* comes into play, which is all the reflected sound that has traveled so far, and for such a (relatively) long time that it arrives and *re-arrives* at the listener delayed enough to sound like a second and third source, or even an *echo* if the room is *really* big.

It's basically a geometry problem. Each room differs in its dimensions; not only in its basic length-by-width size, but in its ceiling height, the distance from you and your equipment to the audience, what's hung (or not hung), on the walls, how many windows and doors there are, and where. Every detail about the space affects your sound. And regretfully, there is very little you can do about any of it. Most of the factors affecting your sound you cannot change. You certainly can't change the dimensions, or alter the window and door locations. But there are a few things you can do, and equalization is one of them. But before you equalize you want to optimize *how* and *where* you place your speakers. This is probably the number one item to attend to. Keep your loudspeakers out of corners whenever possible. Remove all restrictions between your speakers and your audience, including banners, stage equipment, and performers.

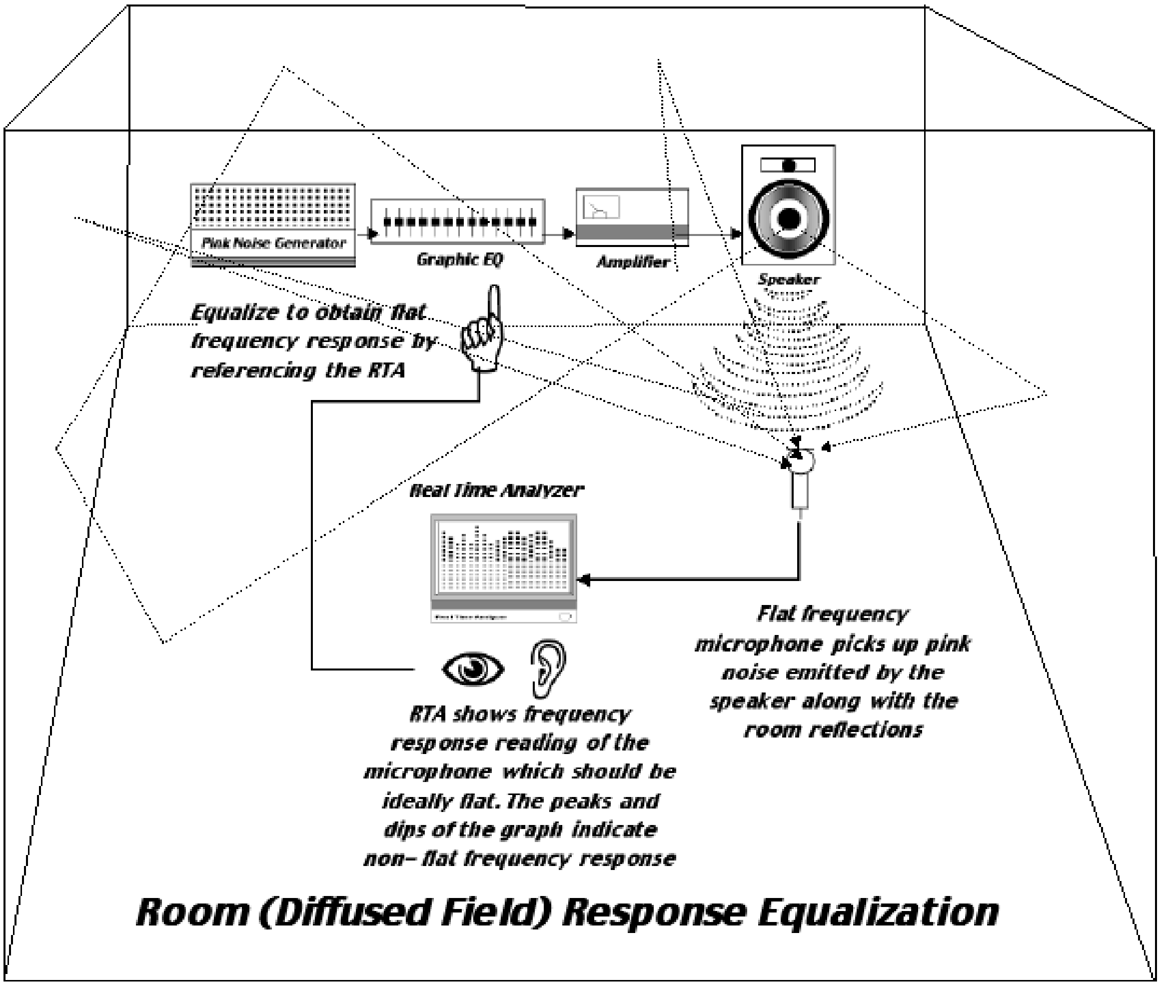

What you want is for most of the sound your audience hears to come directly from the speakers. You want to minimize all reflected sound. If you have done a good job in selecting and equalizing your loudspeakers, then you already know your direct sound is good. So what's left is to minimize the reflected sound.

Next use equalization to help with some of the room's more troublesome features. If the room is exceptionally bright you can beef up the low end to help offset it, or roll-off some of the highs. Or if the room tends to be boomy, you can tone-down the low end to reduce the resonance. Another way EQ is quite effective is in controlling troublesome *feedback* tones. Feedback is that terrible squeal or scream sound systems get when the audio from the loudspeaker gets picked-up by one of the stage microphones, re-amplified and pumped out the speaker, only to be picked-up again by the microphone, and re-amplified, and so on. Most often, this happens when the system is playing loud. Which makes sense, because for softer sounds, the signal either isn't big enough to make it to the microphone, or if it does, it is too small to build-up. The problem is one of an out-of-control, closed-loop, positive-feedback system building up until something breaks, or the audience leaves. Use your equalizer to cut those frequencies that want to howl; you not only stop the squeal, but you allow the system to play louder. The technical phrase for this is *maximizing system gain before feedback*.

It's important to understand at the beginning that you cannot *fix* room related sound problems with equalization, but you can *move* the trouble spots around. You can *rearrange* things sonically, which helps *tame* excesses. You win by making it sound better. Equalization helps.

Understanding EQ and its Effects on Signals

There are two areas of equalisation that to be touched upon. These two areas are vocals and music. We'll discuss the different effects of frequencies within audio signals. What do certain frequencies do for sound and how we understand those sounds. Why are some sound harsh? Why do things sound muddy? Why can't I understand the vocals?

Frequency	Range
Sub Bass	20Hz ~ 60Hz
Bass	60Hz ~ 300Hz
Low Mids	300Hz~2.5kHz
Hi Mids	2.5Khz~ 5kHz
Presence	5kHz ~ 7kHz
Highs	7kHz ~ 20kHz

The frequency ranges designated to the terms above are only a guideline. There may be variations by opinion.

Vocals

Roughly speaking, the speech spectrum may be divided into three main frequency bands corresponding to the speech components known as fundamentals (formants), vowels, and consonants.

Speech fundamentals occur over a fairly limited range between about 125Hz and 250Hz. The fundamental region is important in that it allows us to tell who is speaking, and its clear transmission is therefore essential as far as voice quality is concerned.

Vowels essentially contain the maximum energy and power of the voice, occurring over the range of 350Hz to 2000Hz. Consonants occurring over the range of 1500Hz to 4000Hz contain little energy but are essential to intelligibility.

For example, the frequency range from 63 Hz to 500Hz carries 60% of the power of the voice and yet contributes only 5% to the intelligibility. The 500Hz to 1KHz region produces 35% of the intelligibility, while the range from 1 to 8KHz produces just 5% of the power but 60% of the intelligibility.

By rolling off the low frequencies and accentuating the range from 1kHz to 5kHz, the intelligibility and clarity can be improved.

Here are some of the effect EQ can have in regards to intelligibility. Boosting the low frequencies from 100Hz to 250Hz makes a vocal boomy or chesty. A cut in the 150Hz to 500Hz area will make it boxy, hollow, or tube like. Dips around 500 to 1Khz produce hardness, while peaks about 1kHz and 3Khz produce a hard metallic nasal quality. Dips around 2 to 5kHz reduce intelligibilty and make vocals woolly and lifeless. Peaks in the 4kHz to 10kHz produce sibilance and a gritty quality.

Effects of Equalisation on Vocals

For the best control over any audio signal, fully parametric EQ's are the best way to go.

80Hz to 125Hz	Sense of power in some outstanding bass singers.
160Hz to 250Hz	Voice fundamentals
315Hz to 500Hz	Important to voice quality
630Hz to 1kHz	Important for a natural sound. Too much boost in the range produces a honky, telephone-like quality.
1.25kHz to 4kHz	Accentuation of vocals
5kHz to 8kHz	Important to vocal intelligibility.
5kHz to16kHz	Too much in this area can cause sibilance.

Too much boost between 2kHz and 4KHz can mask certain vocal sounds such as 'm', 'b', 'v'. Too much boost between 1kHz and 4KHz can produce 'listening fatigue'. Vocals can be highlighted at the 3KHz area and at the same time dipping the instruments at the same frequency.

Accentuation of vocals.

The range from 1.25kHz to 8kHz governs the clarity of vocals.

Equalizing Techniques for Instruments

Miking instruments is an art ... and equalizers can often times be used to help an engineer get the sound he/she is looking for. Many instruments have complex sounds with radiating patterns that make it almost impossible to capture when close miking. An equalizer can compensate for these imbalances by accenting some frequencies and rolling off others. The goal is to capture the sounds as natural as possible and use equalizers to straighten out any non-linear qualities to the tones.

Clarity of many instruments can be improved by boosting their harmonics. In fact, the ear in many cases actually fills in hard-to-hear fundamental notes of sounds, provided the harmonics are clear. Drums are one instrument that can be effectively lifted and cleaned up simply by rolling off the bass giving way to more harmonic tones.

Here are a few ideas on what different frequencies do to sounds and their effects on our ears.

31Hz to 50Hz
These frequencies give music a sense of power. If over emphasised they can make things muddy and dull. Will also cloudy up some harmonic content.

80Hz to 125Hz
Too much in this area produces excessive 'boom'. 160Hz to 250Hz This is the problem area of a lot of mixes. To much of this area can take away from the power of a mix but is still needed for warmth. 160Hz is a pet-peeve frequency of mine. Also, the fundamental of bass guitar and other bass instruments sit here. 300Hz to 500Hz Fundamentals of string and percussion instruments.

400Hz to 1kHz
Fundamentals and harmonics of strings, keyboards and percussion. This is probably the most important area when trying to control or shape to a natural sound. The 'voice' of an instrument is in the mids. To much in this area can make instruments sound horn-like.

800Hz to 4kHz This is a good range to accentuate instruments or warm them up. Too much in this area can produce 'listening fatigue'. Boosts in the 1K to 2K range can make instruments sound tinny. 4K to 10K Accentuation of percussion, cymbals, and snare drum.

Playing with **5kHz** makes the overall sound more distant or transparent.

8kHz to 20kHz
This area is often what defines the quality of a recording or mix. This area can also help define depth and 'air' to mix. Too much can take away from the natural sense of a mix by becoming shrill and brittle.

Here are a few other pin point frequencies to start with for different instruments. Remember…these aren't the answers to everything... its just a place to start at.

Kick Drum

Besides the usual cuts in the 200Hz to 400Hz area, some tighter (narrower) Q cuts at 160Hz, 800Hz and 1.3k may help. The point of these cuts makes for space for the fundamental tones of a bass guitar or stand up. It has been also found a high pass filter at 50Hz will help tighten up the kick along with giving your compressor a signal it can deal with musically. 5kHz to 7kHz for snap.

Snare Drum

The snare drum is an instrument that can really be clouded by having too much low end. Frequencies under about 150Hz are really unusable for modern mixing styles. A high pass filter would be suggested to suppress the lower frequency regions in this case. Most snares are out front enough so a few cuts might be all that is needed. Start with 400Hz, 800Hz, and some 1.3kHz. These are just frequencies to play with. It doesn't mean you will use all. If the snare is too transparent in the mix but you like the level it is at, a cut at 5kHz can give it a little more distance and that might mean a little boost at 10kHz to brighten it up.

High Hats

High hats have very little low end information. A high pass filter at 200Hz can clean up a lot of unusable mud in regards to microphone bleed (leakage of the other components of the drum kit into the hi-hat mic). The mid tones are the most important to a high hat. This will mean the 400Hz to 1kHz area but the 600Hz to 800Hz area to be the most effective. To brighten up high hats, a shelving filter at 12.5kHz does nicely.

Toms and Floor Toms:

Again, the focus here is control. Most toms could use a cut in the 300Hz to 800Hz area. And there is nothing real usable under 100Hz for a tom….unless you are going for a special effect. Too much low end cloud up harmonics and the natural tones of the instrument. Think color not big low end.

Over Heads

Drum over heads are the most important mics on a drum kit. They are the ones that really define the sound of the drums. That also give the kit some ambience and space. These mics usually need a cut in the 400Hz area and can use a good rolling off at about 150Hz. Again, they are not used for power.... these mics 'are' the color of your drum sound. Roll off anything that will mask harmonic content or make your drums sound dull. Cuts at 800Hz can bring more focus to these mics and a little boost of a shelving filter at 12.5K can bring some air to the tones as well.

Bass Guitar

Bass guitar puts out all the frequencies that you really don't want on every other instrument. The clarity of bass is defined a lot at 800Hz. Too much low end can mask the clearity of a bass line. A technique which most engineers use to shape the bass tone is to roll off everything below 150Hz, mold the mids into the tone you are looking for, then slowly roll the low end back in until the power and body is there you are looking for. If the bass isn't defined enough, there is probably too much low end and not enough mid range clarity. Think of sounds in a linear fashion, like on a graph. If there is too much bass and no clarity, you would see a bump in the low end masking the top end. The use of EQ can fix those abnormalities.

Guitar/piano/ etc.:

These instruments all have fundamentals in the mid range. Rolling off low end that is not needed or usable is a good idea. Even if you feel you can't really hear the low end, it still is doing something to the mix. Low end on these instruments give what is called support or *body*. The tone is in the mids. 400Hz and 800Hz are usually a point of interest as are the upper mids or 1kHz to 5kHz. Anything above that just adds brightness. Remember to look at perspective though. Is a kick brighter than a vocal? Is a piano bright than a vocal? Is a cymbal brighter than a vocal?

Filters

Equalizers are deliberately designed to create fairly minor changes in the signal. For more drastic effects, such as removing some region of the signal entirely, a FILTER is required. A filter is a circuit that sharply reduces the amplitude of signals of frequency outside of specified limits. The unaffected region is called the PASSBAND, and the filter type is named after the passband as low-pass, high-pass, or band-pass. The point where the signal attenuation just becomes noticeable (a reduction of 3 dB) is termed the CUTOFF FREQUENCY. A low-pass filter with a cutoff of 500Hz will attenuate any signal of frequency above 500hz.

The attenuation provided by a filter is never absolute. A graph of the frequency response of the filter shows that the amplitude of a signal decreases as the frequency moves beyond the cutoff. The actual rate of this decrease is a parameter of filter design called the SLOPE.

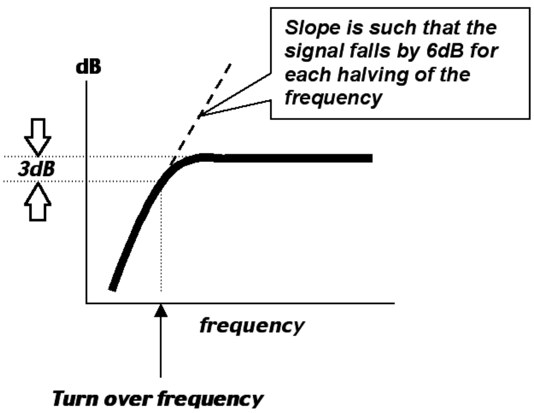

The slope is a general description of response from the cutoff. To begin with, the curve at which the signal begins to roll off is not a sharp and well defined frequency. It is a standard practice to take the frequency at which the signal level has fallen by 3dB as a useful criterion in describing the characteristics of the graph. This frequency is called, the *turnover frequency* or the *3dB down point*. Various circuits differ in the shape of the curve near cutoff, and the possible shapes have names such as Bessel and Chebychev, after the originators of the math formulas involved.

Another design parameter which affects the shape of the filter curve is known as "Q". The derivation of Q is too complex for this discussion, but it is useful to know that filters with a high value of Q have an amplitude bump near the cutoff frequency and have a tendency to oscillate at that frequency.

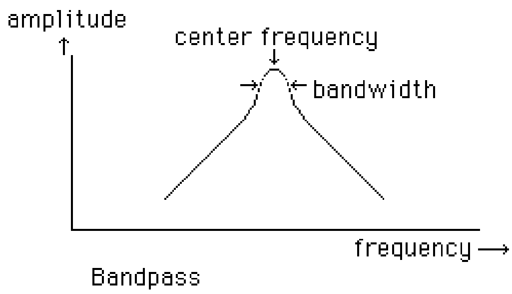

Bandpass response curve

Bandpass filters have two cutoff frequencies. The difference between these frequencies is the BANDWIDTH, and the mean of the two is the CENTER FREQUENCY. Bandpass filters are sometimes encountered in large groups of fixed frequency circuits similar to graphic equalizers. These filter banks are often called 1/2 octave or 1/3 octave filters after the spacing of the filter bands. Such devices used to be the mainstay of many tape music studios before stable synthesizer filters and parametric equalizers were available. The characteristic pitch associated with those machines is almost a trademark of early 60's tape music.

The 1/3 octave filter was originally designed as a research tool for use in the spectral analysis of sounds. The amplitude of the output of each filter band can be separately measured and those measurements plotted on a graph to show the spectral content of the signal. New versions of the spectral analyzer generate data suitable for feeding directly to a computer for measuring complex time/spectrum relationships.

The most complex filter system around is the VOCODER. This device contains a filter bank set up for spectral analysis of a signal and a similar filter bank set up to process a second signal. The measurements derived from each filter in the analysis bank are used to control the amplitude of the signal from the corresponding filter in the processing bank. This system will impose the spectral shape of the analyzed signal onto the processed signal. The effect produced is quite striking, especially if voice is used as the sound for analysis.

The most common slope specifications are :

6dB per octave – A gentle slope, often used in loudspeakers

12dB per octave – A very popular slope for consoles, Eqs, and loudspeakers

18 dB per octave – Used in consoled and Eqs, as well as many electronic crossovers used in live sound systems

24 dB per octave. Too dramatic for console use, but found in crossovers. A 24 dB per octave removal of information sounds somewhat unnatural on a single sound source

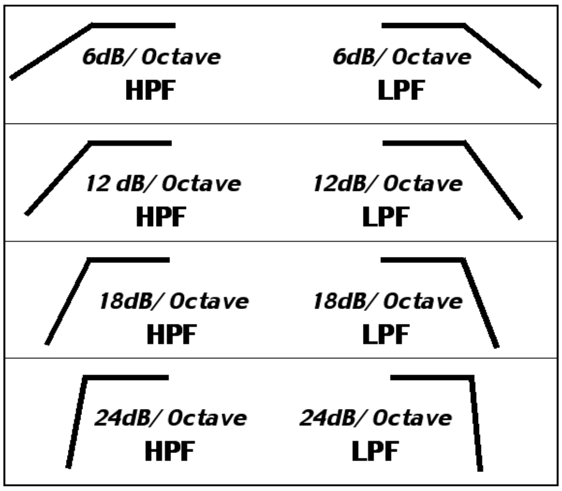

48dB per octave – These ‘brick wall’ filters are often found in inexpensive digital processors and known as anti-aliasing or anti-imaging filters. They are much too extreme to use in normal day to day operation.

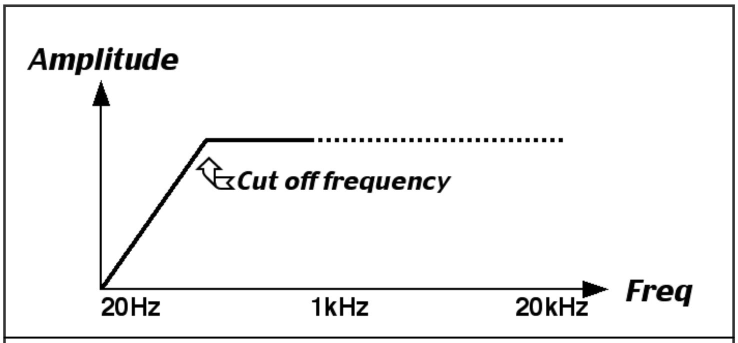

Hi Pass [Low Cut]
Attenuates frequencies below the cut off frequency. Suitable for rolling off low frequency rumble in the signal source.

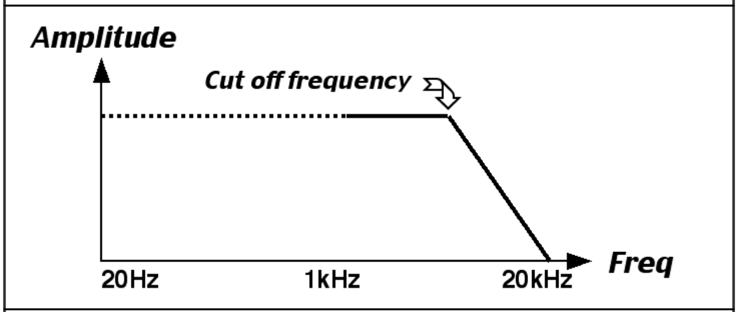

Low Pass [Hi Cut]
Attenuates frequencies above the cut off frequency. Suitable for rolling off higher frequency content of the signal which is unwanted (i.e. hi frequency rings, etc.)

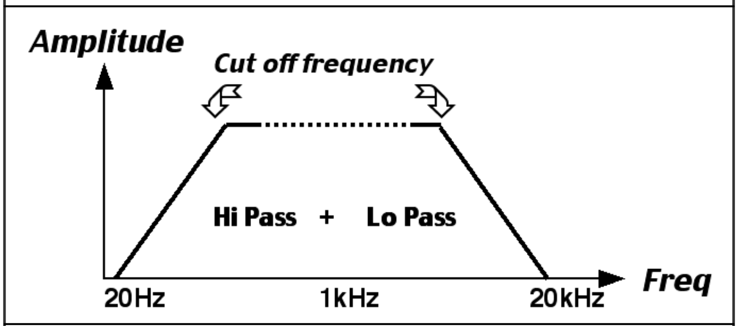

Band Pass
Attenuates frequencies below and above the cut off frequencies to create a bandwidth limiting filter This design can be achieved by combining a hi-pass and a lo pass filter.

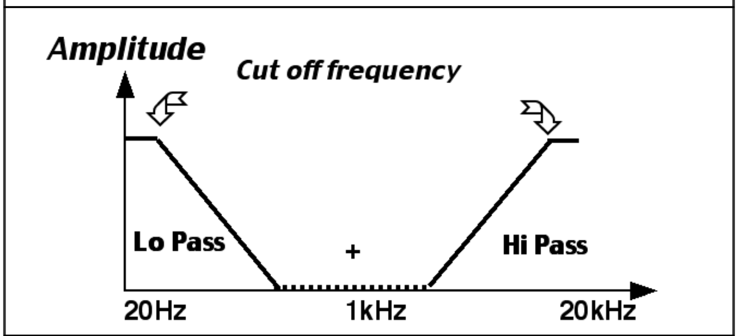

Band Reject [Band Stop]
Attenuates frequencies above and below the cut off frequencies This design can be achieved by combining a hi-pass and a lo pass filter. Allows a bandwidth of frequencies to be removed/ suppressed.

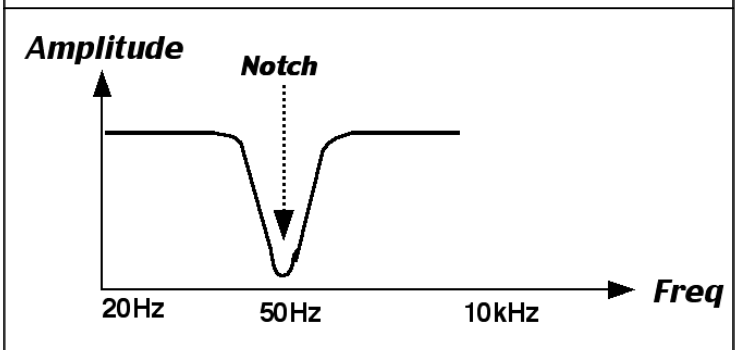

Notch
A very high Q attenuating filter which attenuates a narrow range of frequencies for a specific frequency centre. Examples of its use is to remove obtrusive electrical hums and ground loops.

Equalizer Terms

CUTOFF FREQUENCY

The frequency at which a high or low frequency EQ section starts to take effect, also referred to as turnover frequency.

SLOPE The rate at which a high or low frequency EQ section reduces the level above or below the cutoff frequency. Usually 6, 12, 18 or 24dB/octave.

PASS BAND
The frequency range that is allowed through.

STOP BAND
The frequency range that is attenuated.

FILTER
An EQ section of the following types:

HIGH PASS FILTER A filter section that reduces low frequencies.

LOW PASS FILTER A filter section that reduces high frequencies.

BAND PASS FILTER A filter section that reduces both high and low frequencies.

NOTCH FILTER A filter that cuts out a very narrow range of frequencies.

GAIN The amount of boost or cut applied by the equaliser.

Q How broad or narrow the range of frequencies that is affected.

SWEEP MID A middle frequency EQ section with controls for frequency and gain.

PARAMETRIC EQ An EQ section with controls for frequency, gain and Q.

GRAPHIC EQ An equaliser with a number of slider controls set on octave or third octave frequency centres.
BELL An EQ with a peak in its response.

SHELF A high or low frequency EQ where the response extends from the set or selected frequency to the highest or lowest frequency in the audio range.

HF High frequencies

LF Low frequencies

MID Midrange frequencies

[1]Effects, Signal Processors & Noise Reduction - Reverberation

Introduction

Reverberation (reverb for short) is probably one of the most heavily used effects in music. When you mention reverb to a musician, many will immediately think of a stomp box, signal processor, or the reverb knob on their amplifier. But many people don't realize how important reverberation is, and that we actually hear reverb every day, without any special processors.

Many of today's digital effects processors offer you considerable control over the creation of artificial ambiences for your music, and **if you know how reverberation works in real spaces**, you'll be better equipped for designing fake ones. Reverberation is something that few people are consciously aware of, yet it is one of the most fundamental aspects of a room's sound character. If you were to blindfold someone, take them to an unfamiliar building and lead them through a succession of rooms, clapping or shouting in each one, they would almost certainly be able to give a pretty accurate description of the size of each room. If they were being particularly perceptive, they might even be able to suggest where they were standing in each room and probably even give some idea of what there was in terms of wall coverings, curtains, soft furnishings and so on! In other words, it is reverberation that gives your brain most of the information it needs to create an aural picture of your immediate environment.

Every room has its own sound or 'acoustic', and part of the job of a recording engineer is to assess whether a room's characteristic sound is worth using in a recording. Some major recording studios spend thousands of dollars just to get their recording rooms sounding 'live' and right. If the dimensions, layout and fabric of the room enhance the recorded sound quality, all well and good, but if not, the microphone technique used should minimise any room sound so that an artificial acoustic can be added later from a reverb processor.

In this age of digital technology, artificial reverberation is not only more affordable than ever before, but can also be stunningly realistic and very controllable. With a good understanding of the physics of natural reverberation, and the fundamental operational principles of reverb processors, it is possible to quickly create the illusion of any acoustic environment you can imagine.

[1] http://www.harmony-central.com/Effects

The Physics behind Reverberation

Reverberation is the result of the many reflections of a sound that occur in a room. From any sound source, say a speaker of your stereo, there is a direct path that the sounds cover to reach our ears. But that's not the only way the sound can reach us. Sound waves can also take a slightly longer path by reflecting off a wall or the ceiling, before arriving at your ears, as shown in Figure 1. A reflected sound wave like this will arrive a little later than the direct sound, since it travels a longer distance, and is generally a little weaker, as the walls and other surfaces in the room will absorb some of the sound energy. Of course, these reflected waves can again bounce off another wall before arriving at your ears, and so on. This series of delayed and attenuated sound waves is what we call reverb and this is what creates the 'spaciousness' of a room.

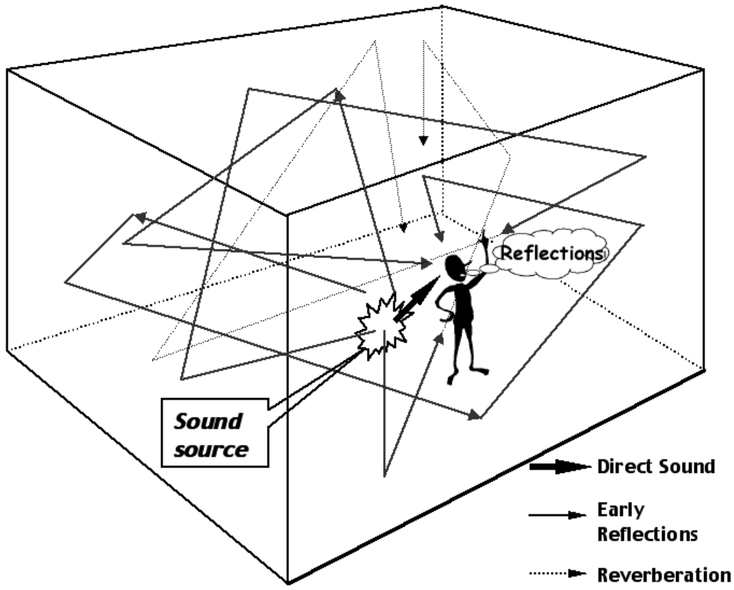

***Figure 1**: Sound waves travel many different paths before reaching your ears.*

It's very tempting to say that reverb a series of echoes, but this isn't quite correct. '**Echo**' generally implies a distinct, delayed version of a sound, as you would hear with a delay more than one or two-tenths of a second. With reverb, each delayed sound wave arrives in such a short period of time that we do not perceive each reflection as a copy of the original sound. Even though we can't discern every reflection, we still hear the effect that the entire series of reflections has.

So far, it sounds like a simple delay device with feedback might produce reverberation. Although a delay can add a similar effect, there is one very important feature that a simple delay unit will not produce - the rate of arriving reflections changes over time, whereas the delay can only simulate reflections with a fixed time interval between them. In reverb, for a short period after the direct sound, there is generally a set of well defined and directional reflections that are directly related to the shape and size of the room, as well as the position of the source and listener in the room. These are the **early reflections** (also called the '**early echoes**' despite the general meaning of the word 'echo'). After the early reflections, the rate of the arriving reflections increases greatly. These reflections are more random and difficult to relate to the physical characteristics of the room. This is called the **diffuse reverberation**, or the **late reflections**. It is believed that the diffuse reverberation is the primary factor establishing a room's 'size', and

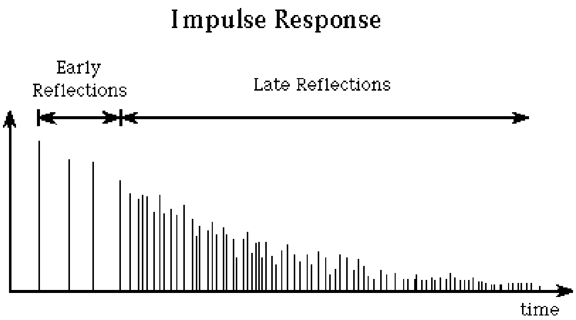

***Figure 2**: Impulse response of a room.*

it decays exponentially in good concert halls. A simple delay with feedback will only simulate reflections with a fixed time interval between reflections. An example impulse response for a room is depicted in Figure 2. (For those who are not sure what an impulse response is, think of it like this. If you consider a small piece of a sound, each vertical line marks when that same piece of sound is heard again, and the height of the columns is how loud the sound is at that time.)

One important point to remember is that as the number of reflections increase *exponentially* over time, the intensity of the reflections decrease *linearly* over the same period.

Another very important characteristic of reverberation is the correlation of the signals that reach your ears. In order to give a listener a real feeling of the 'spaciousness' of a big room, the sounds at each ear should be somewhat incoherent. This is partly why concert halls have such high ceilings - with a low ceiling the first reflections that reach you would have bounced off of the ceiling, and reached your ears at the same time. By using a very high ceiling, the first reflections to reach the listener would generally be from the walls of the concert hall, and since the walls are generally different distances away, the sound arriving at each ear is different. This characteristic is important for stereo reverb design.

A measure that is used to characterize the reverberation in a room is the **reverberation time**. Technically speaking, the reverb time is the amount of time it takes for sound pressure level or intensity to decay to 1/1,000,000th (60 dB) of it's original value (or 1/1000th of it's original amplitude.) Longer reverberation times mean that the sound energy stays in the room longer before being absorbed. Reverberation time is associated with what we sometimes call the 'size' of the room. Concert halls have reverberation times of about 1.5 to 2 seconds.

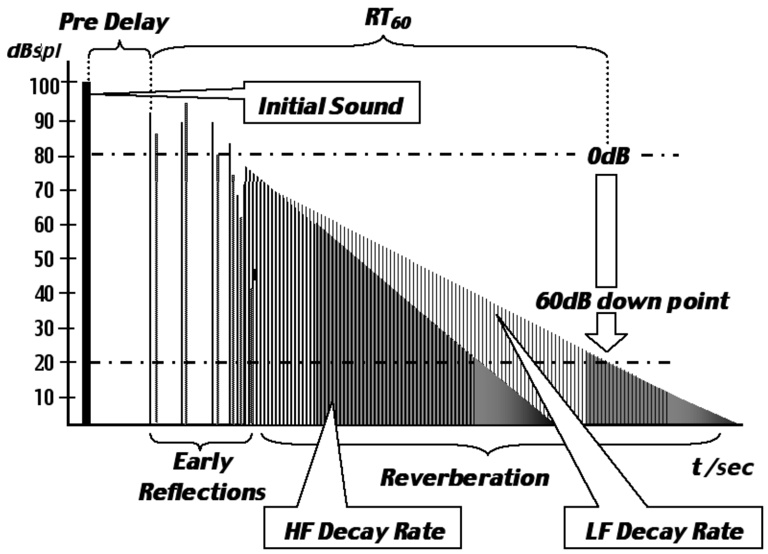

The reverberation time is controlled primarily by two factors - the surfaces in the room and how they absorb or reflect sound, and the size of the room. The surfaces of the room determine how much energy is lost in each reflection. Highly reflective materials, such as a concrete or tile floor, brick walls, and windows, will increase the reverb time as they are very rigid. Absorptive materials, such as curtains, heavy carpet, and people, reduce the reverberation time (and the absorption of most materials usually varies with frequency). You may be able to this notice difference on a gig. During the sound check, the room will sound 'bigger', but during the actual performance, the room may not sound as empty. People tend to absorb quite a bit of energy, reducing the reverberation time. Bigger rooms tend to have longer reverberation times since, on an average the sound waves travel a longer distance between reflections. The air in the room itself will also attenuate the sound waves, reducing the reverberation time. This attenuation varies with the humidity and temperature, and high frequencies are affected most. Because of this, many reverb processors incorporate lowpass filters which allow a more realistic emulation of natural reverb characteristics.

Since we are so accustomed to hearing reverberation, we often have to specifically listen for it in order to notice it. Probably the best way to notice reverb is to listen after short, impulsive sounds, while the sound is still bouncing around. If you want to test out the reverb in various rooms of your house or apartment, clapping your hands works pretty well. Try clapping in the classroom and then as you are coming down the stairs, clap in they stairway. You would be able to hear a greater amount of echoes and reverberation in the stairway due to the fact that it is a compartment which is four storey high!

Direct and Reverberant Sound Fields

(This section is mostly a background in acoustics, and not directly related to reverb effects design and usage)

In acoustics, we talk about the direct and reverberant sound fields in a room. If the direct sound from a source that reaches you is louder than the reflections, you are in the **direct field**. If, on the other hand, the sound pressure due to the reflected sounds is greater than the direct sound, then you are in the **reverberant field**. The point at which the direct field and reverberant field intensity are the same (50 % direct: 50% reflected) is called the **critical distance**.

The reverberant field is extremely important. In fact, most of the time you are in the reverberant field and without it, any performance or lecture would be very hard to follow. As you may know, trying to speak to a group of people outside requires that you speak louder than necessary when speaking in a room. The reverberation of a room helps to keep the sound energy localized in the room, raising the sound pressure level and distributing the sound throughout it. Outdoors, many of the reflective surfaces are missing, and much of the sound energy is lost.

The reverberant field is also important for music. First, it helps you to hear all the instruments in an ensemble, even though some of the performers away be further away then others. Also, many instruments, such as the violin, don't radiate all frequencies equally in all directions. In the direct field alone, the violin will sound quite different (and even unpleasant) as you move with respect to the violin. The reverberant field in the room helps to spread out the energy the instrument makes so it can reach your ears - it truly can enhance a performance. If you can get access to an anechoic chamber (a room designed to have almost no reflections), see if you can get someone to bring an instrument in and see what happens.

Of course, there can be too much of a good thing. As the reverberation time becomes very large, it can be very difficult or impossible to comprehend speech and follow lines of musical instruments. This can be noticed in many gymnasiums and large rooms or hallways with many windows.

Why use Reverb?

If reverb is always around us, why do we add reverb to recorded sounds? Well, many times we are listening to music, we are in environments with very little or poor reverb. The reverberation in a car for example, may not be sufficient to create the majestic sound of a symphony orchestra. And when using headphones, there is no reverberation added to the music. A very dry signal can sound quite unnatural. Since we can't always listen to music in a concert hall or other pleasing environments, we try to add reverberation to the recording itself.

To add reverb, one could make the recordings in a highly reverberant room such as a concert hall, but this is often impractical since such rooms may not be easy to access, be located far away, or too expensive to use. This has caused the development of a variety of ways to synthetically add reverb to recordings.

The Early Days of Artificial Reverberation

Reverberation Chambers

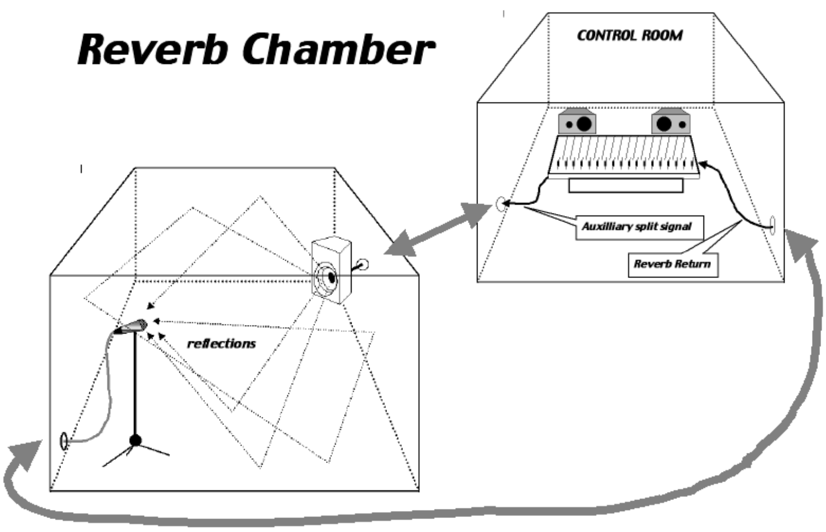

Although you may not have a concert hall at your disposal, studios have used rooms to add reverb to recordings. Elevator shafts and stairwells may work as highly reverberant environments. In most cases, it is literally a room, often tiled and full of ceramic sewer pipes to provide a wealth of reflective surfaces. A loudspeaker generated the direct sound, and one or more microphones collected the resulting reverberation. The echo room has the advantage that the reverb is naturally very complex, but it is also difficult to adjust, and requires a large and quiet room!. The amount of reverb generated by the room can be controlled by adding some absorptive materials (i.e. a studio can have curtains that can be drawn across the more reflective surfaces).

Spring Reverbs

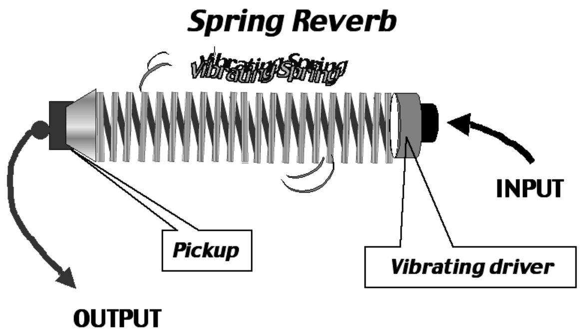

Spring reverbs provide a relatively simple and inexpensive method for creating reverb effects. Spring reverbs have been used in Hammond organs and will still find them in many guitar amplifiers. In amps, the spring reverbs are usually enclosed in a metal box, called the reverb pan, which is attached to the bottom of the amp. The pan takes an audio signal and produces a reverberated version which is then

mixed into the dry signal.
The operation of a spring reverb is pretty simple - the audio signal is coupled to one end of the spring by a **transducer** (a transducer is simply a device that converts energy in one form to another - in this case, electrical and mechanical energy. Some other familiar transducers are the pickups on a guitar, microphones, and speakers). This creates waves that travel through the spring. At the other end of the spring, there is another transducer that converts some of the motion in the string into an electrical signal, which is then added to the dry sound. When a wave arrives at an end of the spring, part of the wave's energy is reflected and stays in the spring. It is these reflections that create the reverb characteristic sound.

Often you will find several springs being used together in a reverb unit. Each spring can be of a different length or under a different tension to avoid the uniform behavior in a single spring, where all the reflections occur at fixed times. In a sense, it increases the 'randomness' of the echoes. However in most reverb units, the spring lengths and tensions are fixed in the design process, and not left to the user to control. The character of a particular spring reverb unit can be optimized for the sound source at the design stage by careful choice of the number, length, diameter and compliance of the spring(s).

Plate Reverbs

Plate reverbs are not used widely outside of the studio since they are expensive and rather bulky. The setup is similar to a spring reverb, but instead of being connected to the springs, the two (or more) transducers are connected to different points on a metal plate. These transducers send vibration waves throughout the plate, and reflections occur each time a wave reaches the edge of the plate. The reverberation can be controlled by adjusting the damping of the plate and the location of the transducers. The plate has a characteristic metallic, bright, sound quality which has become intimately associated with pop music.

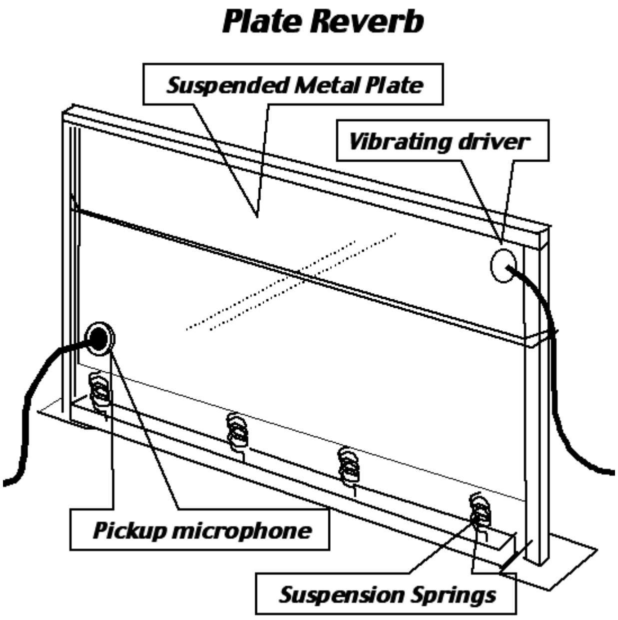

Other reverberators

There are many other ways to create reverb effects, although these are generally not used often. Before digital technology and memory was so inexpensive, there were some tape based reverb units. All manner of record-replay systems have been developed to provide a reverb effect, but none have survived the digital revolution. The earliest ideas simply

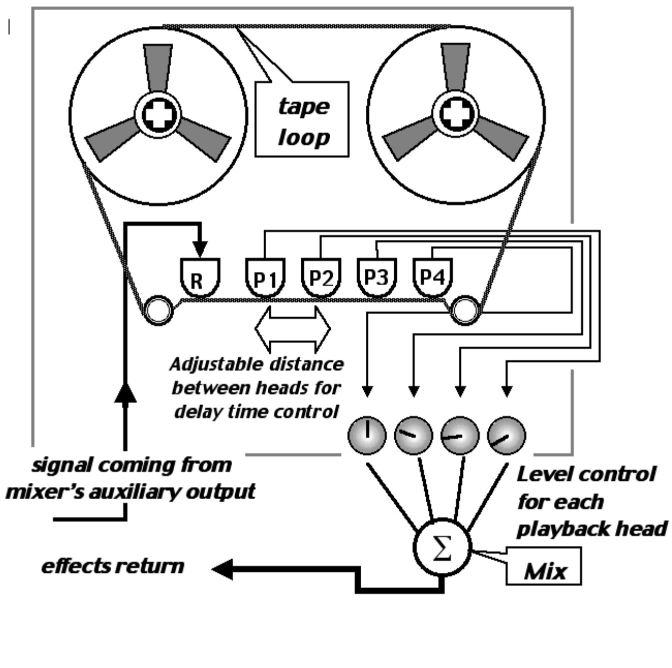

used a three-head tape machine, where the direct signal was recorded onto tape and the replay signal provided the reverberation.

The tape speed and head spacing determined the pre-delay and if some of the replay signal was mixed with the direct signal, a pseudo-reverb could be created. The results are hardly realistic, but the system was popular at a time when the alternatives were too expensive or impractical.

The record-replay theme was further developed into machines like the WEM Copycat and the Roland Space Echo, which used tape loops and multiple replay heads, with the ability to adjust the contribution and feedback of each head -- but then solid-state technology arrived...

Digital Signal Processing (DSP)

Digital Processing Techniques

The advent of digital technology really revolutionized artificial reverb, basically because the time-domain signal processing of digital audio lends itself very well to the kind of sound manipulation needed to create realistic reverb. Creating a pre-delay is simply a case of storing sound in a memory until the required time has passed. The early reflections are created by replaying the direct sound repeatedly at suitable moments, with level and equalization changes as necessary. The main body of the decay is created by cycling the direct sound through a complex set of short feed-back and feed-forward delays, configured to introduce the desired equalization characteristics.

Digital reverbs are available to suit all pockets from a large number of manufacturers, including Yamaha, Sony, Digitech, Klark Teknik, ART, Ensoniq and Alesis, to name but a few. However, if you ask any professional sound engineer to name their favourite reverb machine, chances are you'll hear one name above all others. Lexicon is probably the most popular manufacturer of digital reverbs, and their product line extends from push-button preset units to the most complex state-of-the art systems. The advances in digital hardware have made reverb processors available at inexpensive prices that are portable, and quite flexible.

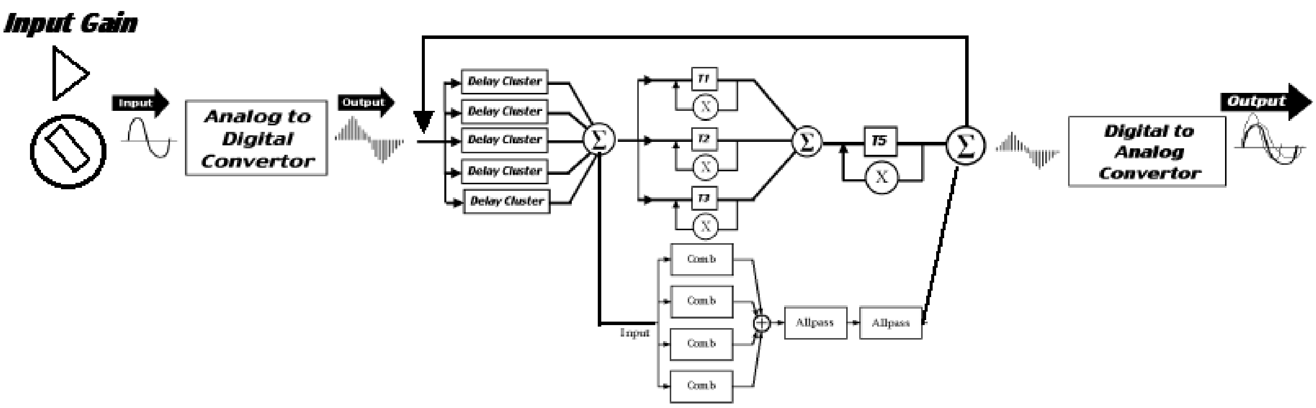

A DSP Reverb Processing Algorithm

Lexicon 480L Mainframe Unit and Large Alphanumeric Remote Control (LARC)

Lexicon 960L Mainframe Unit and Large Alphanumeric Remote Control (LARC)

LEXICON PCM 81/ 91

EMT 250 Digital Plate Reverb

AMS RMX16

SONY DRE-S777

Sony's latest weapon in the cutthroat market for high-spec reverb units is a cunning and powerful digital data-manipulation technique known as convolution, which can create stunningly realistic reverb derived from the measured (or 'sampled') acoustic signatures of real performance venues around the world. Essentially Sony's new DRE S777 imposes (convolves) the carefully recorded and measured reverberation characteristics of selected venues upon its audio inputs. Samples can be loaded via CD-ROMs. The result is a reverberant sound stage which, to all intents and purposes, sounds just as it would if the sources were actually recorded live in the selected environment!

ENSONIQ DP4+

The DP4+ has four processing engines and has the ability to process four different signals in parallel.

General Reverb Processor Control

Let's recap on the parameters that today's digital reverb processors are likely to offer for the simulation of real acoustic spaces.

Wet/Dry Control
Firstly, there is usually a means of balancing the direct sound against the reverb. The direct sound is often referred to as *'dry'* and the reverberation as *'wet'*, so a wet/dry control will normally be included on a effects processor

Input Level Meters
Some of you will soon know how unpleasant it sounds when an analogue-to-digital (A/D) converter is overloaded, so a critical control on digital reverbs is the input level control, and its associated headroom meter.

This will be followed by one or more parameters for controlling the number, timing, amplitude, and timbre of the other early reflections. Some machines provide controls called 'Pattern', 'Level' and 'Room Size'; others might simply offer preset venue simulations ('Hall', 'Chamber', 'Jazz Club', and so on). Finally, having skillfully set all of these parameters to create a wonderfully believable artificial acoustic (!), you can usually store your fake room in one of a number of user memories.

Input Level (in dB's)

Determines the level of the signal at the input stage of the unit most reverb devices are mono in and stereo out but generate different reflection patterns for the left and right channels. The two independent signals is assimilated by our stereo hearing mechanism enabling us to perceive spaciousness

The actual input level is program dependant but a good starting point is to set the input level of the signal processor at zero gain and increasing the auxiliary send level to maximum (+4dB, OVU). Then the input gain can be increased until the meter reaches nominal input level on the effects processor. This would enable us to obtain the best signal to noise ratio.

Pre Delay *(also known as Initial Delay)* **[1ms ~400ms]**
Delays the input signal by 1ms to around 400ms before it reaches the reverb cluster. It indicates the length of time between the initial sound source and the onset of reverb. In other words, it effectively defines the distance of the first reflective surface; this gives the effect of separating the dry source from the reverb. Longer pre-delay times, where there is a slight gap before the onset of reverb will add a greater sense of clarity as the reverb does not occur too quickly and cause the dry source to sound muddy

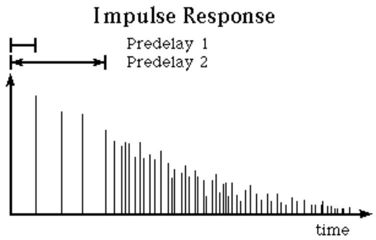

General pre-delay guideline

Up tempo Drums/ Percussion	25ms-50ms
Ballad drums	40ms-80ms
Acoustic instruments	40ms-80ms
Vocals	75ms-125ms
Strings	100ms-200ms
Brass	50ms-100ms

E/R Level [0% ~100%]

In simulating ambient spaces, there is a distinct echo prevalent at the beginning of the sound. These are the first reflections which result form the contact between the sound source and the closest reflective surfaces. Early reflections are important as it gives a sense of spaciousness. Appropriate levels of early reflections on an instrument can reinforce the sound of the instrument which will help project the perceived 'size' of the instrument.

Decay Time *(RT60)* **[0 sec ~12sec]**
Long Decay time simulate larger ambient spaces while shorter decay times put the instrument in a smaller environments

Density [0 ~10]
This parameter determines the density of the reverb reflection (average amount of time between reflections). Lower settings produce minimum reverb density and lead to a more spacious sound in which one can perceive the individual echoes between reflections. Higher settings produce dense reverb where the reflections are tightly spaced, causing it to sound closely packed and 'thicker' sounding. This is due to the fact that the gap between reflections is not distinct and seems to blend between one another.

Diffusion
Diffusion is the complexity of the reflections and how they spread out. If diffusion is set to minimum complexity, a cleaner reverb effect is produced. As the diffusion value is increased, the complexity of the reflections increases, producing a larger more diffused field off reverb

Feedback/ Regeneration [0%~100%]
This is a variable control that sends the output of the delay unit back into the input. It will create multiple repeats of the reverb. This parameter should be used moderately as an excessive amount will result in an uncontrollable feedback and creates a squalling loop.

Brightness
Brightness is also an important control. When reverberation occurs naturally, frequency response of the reverberation pattern is not linear. Frequencies at the top of the spectrum tend to reflect with much less energy that the lower frequencies. In fact reflection hardly occurs at all above 6kHz. So its will sound quite false if an artificial reverb is too bright, and can be very distracting to hear sibilant frequencies reproduced loudly in reverb. The brightness control acts as a filter to cut down the very high frequencies.

HPF, LPF
Included in most reverb processor, the high pass and Low pass filter help the user change the lower and high frequency response of the input signal. More advanced units have high and low frequency damping so that these frequencies will still be present, but their undesirable effects on the reverb signal will not be noticeably audible.

Reverb Decay
The reverb decay indicates how you how long the reverb can be heard after the input stops. The actual measure of what can be 'heard' can vary among manufacturers. The reverb decay is typically in terms of milliseconds. It's important to note that the reverberation tail lasts for different durations at different frequencies. High-frequency sound waves have a lot of trouble persuading air molecules to vibrate quickly enough to pass the sound energy onwards. Consequently, high-frequency sounds tend to die away, as they travel, much faster than mid-frequency sounds. On top of that, high-frequency sounds are absorbed by soft furnishings (which includes people and even wallpaper!). On the other hand, high frequencies reflect strongly from a wide range of surfaces, such as windows, sound desks, equipment racks, and so on. At the other frequency extreme, low frequencies are only reflected by large and very solid objects, so there may be little LF in the reverberation at all in some circumstances, but a definite bass 'bloom' in cave-like rooms!

To help provide this level of realism, most reverb units allow you to adjust the reverberation time for high (and sometimes low) frequencies relative to middle frequencies, and introduce some kind of overall equalization to the reverberation tail.

Gate Time
This parameter applies to gated reverbs. The gate time is simply the length of time that the reverb is allowed to sound. This may also refer to the length of a reverse reverb.

Gate Decay Time
Some units with gated reverbs will also provide this parameter, which controls how the gate is actually applied or 'closed'. A very short gate time means that the reverb is cutoff rapidly, such as shown in Figure 3. Longer decay times means that the reverb is given some time to fade away gradually.

Gate Threshold
Rather than apply a gated reverb to an entire signal, you could very well only gate the reverb depending on signal levels. Typically, the gate on a reverb will be kept open (the impulse response is not truncated) when signals are above this value, but as when the signal drops below the threshold, the gate closes and the number of reflections is reduced. The gate will open again when the signal rises back above the threshold. Some gated reverbs may use a threshold that is not user programmable.

Other Reverb Types

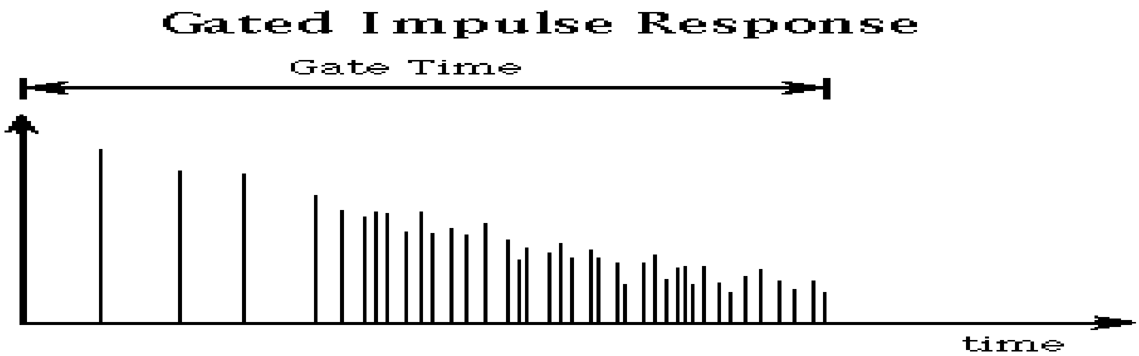

Gated Reverb
A gated reverb is created by simply truncating the impulse response of a reverberator or alternatively, only allowing a sound to make a certain number of reflections. The amount of time before the response is cut off is called the **gate time**, as labeled in the figure on your left. Some reverb units may allow a more gradual decay of the sound, rather than an abrupt silence. Gated reverbs are most commonly implemented with digital processing and are commonly used on drums.

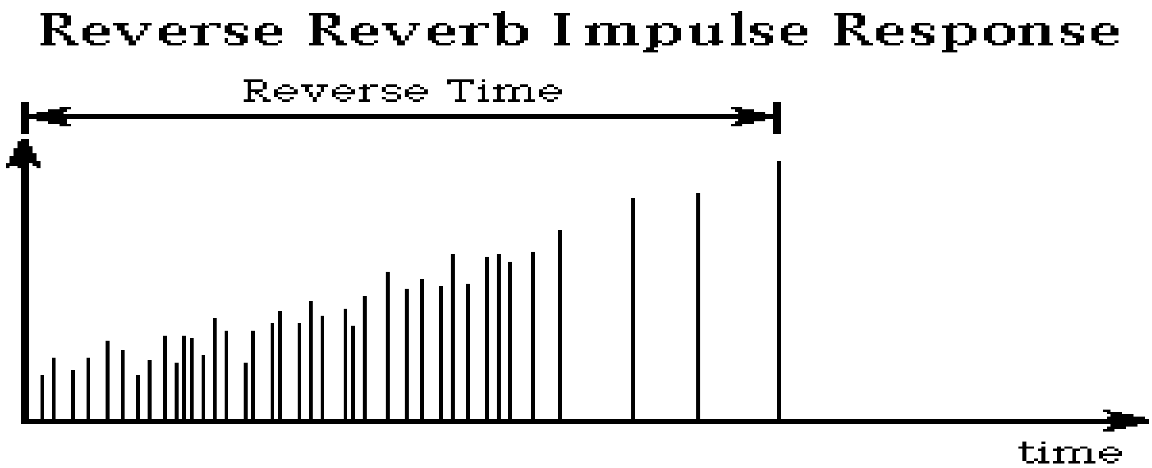

Reverse Reverb
The reverse reverb puts a little twist on reverb responses we looked at above. Instead of simulating reflections that become quieter and gradually fade away, the reverse reverb simulates reflections where the sound gets louder over time, and then abruptly cuts off. The length of time the sound builds up is often referred to as the **reverse time**, or the gate time, as it resembles a gated reverb simply reversed in time. The figure on your right shows an example of what a reverse reverb impulse response would look like.

Techniques behind the Use of Artificial Reverberation

Creating Spaces by learning to Listen

You cannot create an artificial acoustic space if you don't know what a real one sounds like to start with, and the only way you can find out is to actively listen to sounds in as wide a range of environments and circumstances as possible. Everyone has a very detailed subconscious knowledge of how different rooms sound, and although few are able to analyze the reverberation structure, most spot incongruities in artificial reverberation very easily.

It's very revealing and informative to consciously listen to the sound of different rooms as you move around in a building -- even in places with which you thought you were familiar. Try to analyze in your own mind what sort of pre-delay, decay time, early reflections and high-/low-frequency decays naturally occur to create the 'sound' of that room. Don't just listen to indoor reverberations, either -- try to assess the reverberant features of the local high street, the great outdoors, a wood or forest, or wherever you happen to be. You will find reverberation in places you didn't expect it, and may be surprised to discover that places you assumed to be reverberant actually are not!

Mono and Stereo

The very nature of genuine reverberation is that it tends to come at you from all over the place, but particularly from the sides of the room. This has significant effects on the compatibility between stereo and mono versions of your mix, since the mono listener is effectively denied any information from the sides of the stereo image.

To see how this happens, consider a simple M&S (Middle & Side) stereo microphone technique (you will learn this in the next subject in this Unit), being used to record something in a reverberant room: the stereo listener hears the full acoustic in all its glory, but the mono listener hears only the forward-facing 'M' microphone, not the sideways-facing 'S' microphone -- and guess which one picks up the bulk of the room sound? This absence of reverberation in mono afflicts artificial reverb processors as well as natural acoustics. In practice, the amount of reverberation heard in mono may be substantially less than that in the stereo balance, and if mono listeners are likely to be an important part of your music audience, always check for mono compatibility. In general, you almost always have to compromise the balance in some way because either the mono will be too dry, or the stereo will be too wet!

Tricks worth trying include reducing the stereo width of the reverb (turn the pan-pots in a bit towards the centre instead of having the reverb returns running out to full left and right), or mix in a small amount of reverb from another reverb processor, panned centrally. The extra reverb should be set up with the same parameter values as the stereo reverb, although a slightly shorter pre-delay and longer HF decay time often work well. The balance between the dry sound, the mono reverb, and the stereo reverb needs to be adjusted carefully, while you continuously switch between mono and stereo listening to find the most uniform results in the two modes.

In matrix surround systems (such as Dolby Surround), real or artificial stereo reverb tends to spread across the rear channel quite naturally as a result of the way in which the rear-channel information is encoded and decoded. Altering its stereo width controls the front-back balance, narrowing the reverb pulls it to the front, and increasing the width pushes more to the sides and rear.

Many stereo digital reverb units have a single input and a stereo output, and this often causes people to wonder how the reverberation can be 'true stereo' with only a mono input. The answer is simple if you consider the real situation of a sound source within a reverberant space.

If someone claps, there's only one sound source, yet the reverberation will come from all directions and could be captured by a simple stereo microphone array -- a mono input to the room and a stereo output from it. Of course, in a more complex situation with, say, a string quartet in the room, there are multiple sound sources and each will have slightly different pre-delays and early reflection patterns, but this is usually a very subtle distinction, and in practice the mono-in, stereo-out system of most digital reverb units works perfectly adequately.

Something few people ever check is the line-up of a stereo reverb unit. However, it is a stereo source and should be treated in just the same way as any other stereo signal, which means making sure that the left and right reverb outputs have the same gain and equalization through your mixer. I find that a quick, easy and reliable method of doing this is to simply dial up a 3- or 4-second decay time and send a brief burst of signal into the machine. Listen carefully to the dying reverb tail: it should decay centrally, possibly even becoming narrower in width as it goes (although this depends on the particular algorithm). If the reverb tail appears to collapse towards one side or the other, your return channels have different gains and should be adjusted.

Choosing and Using Reverb

In general, two reverb units will meet the needs of pretty much every recording situation. One machine would normally be set for a short, bright sound (perhaps a plate setting) for percussive sounds, whilst the other would be set to a longer, warmer patch, providing a 'lush' quality for vocals and solo instruments. You could also try passing some instruments through both reverbs (percussive one first) for a third alternative.

Some engineers like to use several reverbs to create a layering effect, but you will generally find that this approach causes a loss of definition and adds confusion to the overall sound (unless it is something intentionally required like in trance, new age, electronica music production genres). Going back to the idea that the artificial reverb is merely replacing the poor acoustics of a less-than-ideal recording venue, it could be argued that there should only be one reverberation sound for everything, as would be the case if the musicians all played live in the same reverberant room!

The next issue to address is how much reverb to use. The classic mistakes of the novice are using too much reverb return on everything and allowing reverb tails to be too long. Reverb generally needs to be subtle, and ideally only the loudest musical peaks should cause obvious tails. Even the biggest halls rarely have a reverb time in excess of four seconds, and often a two-second decay time is easily long enough.

The choice of reverb parameters is dependent on both personal taste and the nature of the programme material, so it's impossible to give specific recommendations, but try to create life-like environments wherever possible. Most reverb units offer a number of special effects, such as gated or reverse reverbs, and these are best used sparingly, so that they keep their impact. While we're on the subject of special effects, it's worth trying out the pseudo-reverberation programs too. Algorithms such as 'Ambience' or 'Alive' can often add extra definition and life to dull vocals, or spice up closely-mic'd solo instruments without your having to resort to using exciters.

Normally, reverb sends are taken post-fader, so that direct signal level adjustments are reflected in their reverb returns. However, *it can often be useful to send pre-fader*, and not allow any direct signal into the final mix at all. This is particularly effective with sustained keyboard string sounds and the like, where it helps to make less-than-ideal synth sounds blend a lot more smoothly.

Another useful trick is to set up a reverb specifically for the keyboard sounds, and route the reverb returns through a chorus unit. This provides a completely different kind of sound to chorusing the keyboards directly and adds an interesting 'swirling' quality which can be very effective if used discreetly.

Although reverb processors are most important during mixdown, they're also vital during recording, especially when recording vocals. Many singers have enormous trouble pitching properly without reverberation and it's essential to have the ability to route reverb returns to the headphone monitor mix. The reverb setting for the cue monitor is not particularly critical to the performance (provided it is broadly appropriate) and need not be recorded, although some engineers do like to record voice and reverb together (occasionally as a complete mix but more usually on adjacent channels on the multitrack machine). This is particularly useful if the reverb plays a part in the performance (through timing or percussive vocal effects, for example).

As artificial reverb becomes more and more elaborate, there's a trap which many engineers find themselves falling into. It's possible to become so engrossed in adjusting each parameter minutely, trying every possible combination along the way, that you lose sight of the original idea. The best way of getting the sound you want, quickly, is to understand the nature of real reverberation and apply that knowledge to creating the acoustic space you've imagined. It's far better to think for a minute or two, and then dial the right numbers in, than to sequentially try every preset on the machine, hoping to stumble across something that sounds OK.

One Algorithm or Two?

A word of warning -- not every reverb processor is as flexible as it might seem. Particularly with multi-effects units, it is quite common to find that there is actually only one reverberation algorithm. The wide range of supplied preset environments (Hall, Room, Plate, and so on) is actually composed of variations in the delay, decay and EQ settings of a single algorithm. In these cases, you'll find that no matter how you adjust the reverb parameters, all settings sound very similar: the overall character of the room does not seem to change, and this is because the pattern of the early reflections remains fixed. The better machines have a number of different algorithms and a variety of early-reflection patterns, which allow a larger range of different room types to be created, each with distinct and individual sonic characters.

Fortunately, there is an easy way to find out which category a particular machine falls into. Select two, theoretically diverse, programs -- perhaps a Hall and Plate. Set the delay, decay, EQ and any other parameters to identical values and store the new settings in a couple of user memories so that they can be recalled easily. Next, listen critically to the quality of the reverberation while switching between the two presets. There should be an obvious difference in the character of the room acoustic if the machine uses different algorithms, with different early reflection patterns. (Try closing your eyes and imagining the dimensions and furnishings of the fake room.) If you cannot spot any differences, the chances are that the machine uses the same algorithm for all its reverb programs.

In the early days of reverb technology, some engineers got into the habit of layering a number of reverb and delay effects in order to generate a richer reverb than would otherwise have been possible. However, while this particular technique generates a characteristic sound which some may wish to recreate, it is not nearly as necessary with today's digital units.

This does not mean that most commercial records only ever use one reverb nowadays — some *do* only use a single program in order to ensure uniformity of ambience, but these projects are still in the minority. Many modern records use bizarre reverb treatments on certain instruments as a special effect, but when different algorithms aren't being used in this way, they still fulfil a definite production function. Placing instruments in different environments can increase the contrast between one and the other — if you put a vocal in a small bright room and a guitar in a large warm-sounding hall, the voice would stand out, albeit at the expense of a uniform acoustic.

Effects, Signal Processors & Noise Reduction | Time Based Processors

Delay

The Basic Delay

Simply put, a delay takes an audio signal, and plays it back after the **delay time**. The delay time can range from several milliseconds to several seconds. Figure 1 presents the basic delay in a flow-graph form. This only produces a single copy of the input, and thus is often referred to as an echo device.

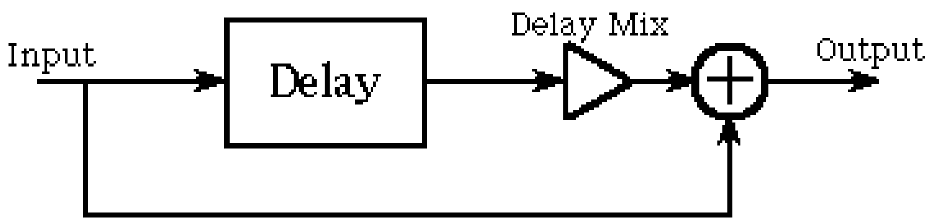

Figure 1: *Diagram of the basic delay unit, or an echo device.*

Just having a single echo effect is rather limiting, so most delays also have a **feedback** control (sometimes called **regeneration**) which takes the output of the delay, and sends it back to the input, as shown in Figure 2. Now you have the ability to repeat the sound over and over, and it becomes quieter each time it plays back (*assuming that the feedback gain is less than one. Most delay devices restrict it to be less than one for stability*). With the feedback, the sound is theoretically repeated forever *(at least until you turn the unit off*), but after some point, it will become so quiet that it will be below the ambient noise in the system and inaudible.

Figure 2: *Diagram of the basic delay unit with feedback.*

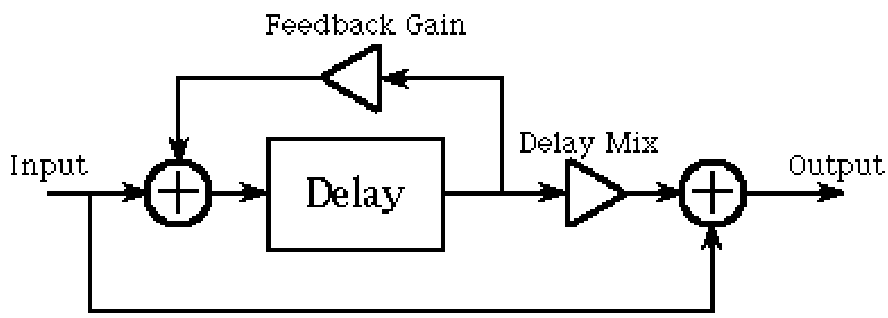

Delays are very useful for filling out an instrument's sound. Playing through a delay unit with a short echo, say 50 to 100 milliseconds, creates a doubling effect, as though two instruments were being played in unison. Using several delays together with feedback can be used to create a reverb-like sound, though a typical reverb unit will create a more complex sound pattern.

As you increase delay times beyond 100 milliseconds or so, the delay no longer a subtle effect. One interesting possibility is to match the delay time to the tempo of a song so that the delayed copies of the sound fall on the beat. Extending to very long delay times close to a second or more gives you a chance to play over yourself and develop harmonies even though you may only be playing one note at a time.

Looping and sampling are just a short jump away. Instead of repeating everything you play, you can record a segment of your playing, say a chord progression, and then loop it - play the recorded audio over and over. This lets you go a step further so you can actually solo over yourself when you don't have a rhythm player at your command. Some delay pedals include sampling capability, though the length of the sample may be limited to two seconds or less. For serious looping, you will need devices with longer recording times, such as Lexicon's *JamMan* and the Oberheim *Echoplex*. are some of the popular units on the market for looping, and they offer other capabilities over straight looping, such as recording additional sounds onto the sample, playing the loop backwards, as well as recording several different loops while being able to

Delays are also very important when building a mix of instruments in a stereo environment. It can enhance stereo placement of instruments, and making the mix sound 'bigger'. A little delay can be more effective than panning for spreading tracks out in the stereo field. Just a simple delay on the order of 20 milliseconds can make a big difference.

Other Delay Types

Slapback

A slapback delay is not a new algorithm of any sort. It is the same as the basic delay without feedback discussed in the opening of this article. A delay is called a slapback delay if the delay time is very short - say between 40 and 120 milliseconds. A longer delay is often referred to as an echo, rather than a slapback delay.

Multi-Tap Delay

In some cases, you might want more flexibility in a delay unit, and a multi-tap delay offers gives you just this. Multi-tap delays are interesting because they allow you to create more complex patterns that can add a rhythmic quality to the instrument.

In the delay units discussed above, the output is taken after the signal has been delayed for the total delay time. But you can also take outputs such that the signal has only been delayed a portion of the total delay time. Taking outputs from points within the delay line is referred to 'tapping' the delay line, much like a tap in a water pipe allows you to get water at points long the pipe. Units are usually labeled with the number of available taps - a 3-tap delay has three taps to use, a 4-tap has 4, etc. Unwanted taps can be removed by setting the tap output level to zero. The amount of delay between the various taps can be different. A diagram for a multi-tap delay is shown in Figure 3.

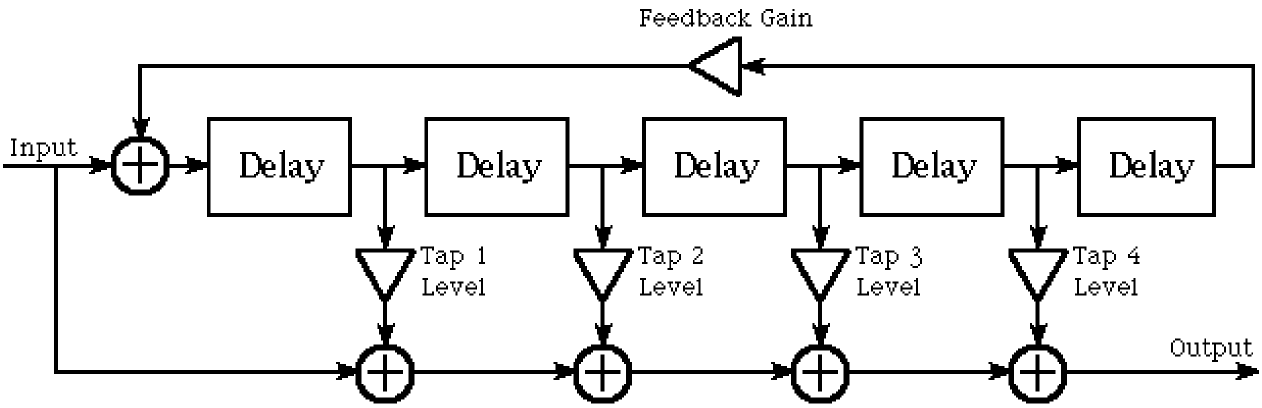

Figure 3: - *Flow diagram of a 4-Tap delay. If the last delay value is zero, and only the fourth tap is used, the system is equivalent to the basic delay.*

The multi-tap delay is really a more general case of the basic delay design. If you set all but one of the tap gains to zero, and place the remaining tap at the end of the delay-line, you then have the basic delay discussed above. The multi-tap delay could be generalized even further, by allowing feedback from the tap outputs to the beginning of the delay line as well. You're are unlikely to find this kind of control in products though, primarily because it becomes very easy to create an unstable system.

It can take a while to understand the multi-tap delay. Notice that if you look only at the output of a single tap, the sound repeats according to the total delay time. The input sound will appear at the tap output before the total delay time (assuming that the tap isn't located at the far right in the diagram).

Ping-Pong Delay

As the name implies, the ping-pong delay produces a bouncing sound, where it's typically bouncing between the left and right channels of a stereo signal. The ping-pong delay uses two distinct delay lines, each driven by an input (the inputs could be the same signal if desired). Rather than feeding back on themselves however, the output of the delay lines feedback into the other delay line, as shown in Figure 4. This setup produces two output signals, which when panned hard left and right can create the classic 'bouncing' sound.

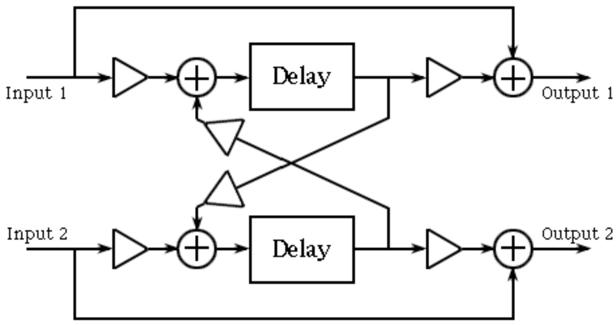

Figure 4: - *Flow diagram of a Ping-Pong delay unit.*

Delay Time Calculations

For TIME in milliseconds,

MilliSeconds = (BARS x 240,000) / BPM

So to work out the delay-time for 3 sixteenths of a bar would be,

((240,000 x 3) / 16) / BPM

which equals to 45,000 / BPM. Assuming the tempo is 120bpm, the delay time for 3/16 of a bar would be

45,000 / 120 = 375 milliseconds.

The table below lists the various delay times in milliseconds for 16ths of a bar and 12th

Time (ms)	Sixteenths of a Bar								Twelfths of a Bar			
BPM	1	2	3	4	5	6	7	8	1	2	4	5
60	250	500	750	1000	1250	1500	1750	2000	333	667	1333	1667
65	231	462	692	923	1154	1385	1615	1846	308	615	1231	1538
70	214	429	643	857	1071	1286	1500	1714	286	571	1143	1429
75	200	400	600	800	1000	1200	1400	1600	267	533	1067	1333
80	188	375	563	750	938	1125	1313	1500	250	500	1000	1250
85	176	353	529	706	882	1059	1235	1412	235	471	941	1176
90	167	333	500	667	833	1000	1167	1333	222	444	889	1111
95	158	316	474	632	789	947	1105	1263	211	421	842	1053
100	150	300	450	600	750	900	1050	1200	200	400	800	1000
105	143	286	429	571	714	857	1000	1143	190	381	762	952
110	136	273	409	545	682	818	955	1091	182	364	727	909
115	130	261	391	522	652	783	913	1043	174	348	696	870

BPM	1/16	1/8	3/16	1/4	5/16	3/8	7/16	1/2	1/12	1/6	1/3	5/12
120	125	250	375	500	625	750	875	1000	167	333	667	833
125	120	240	360	480	600	720	840	960	160	320	640	800
130	115	231	346	462	577	692	808	923	154	308	615	769
135	111	222	333	444	556	667	778	889	148	296	593	741
140	107	214	321	429	536	643	750	857	143	286	571	714
145	103	207	310	414	517	621	724	828	138	276	552	690
150	100	200	300	400	500	600	700	800	133	267	533	667
155	97	194	290	387	484	581	677	774	129	258	516	645
160	94	188	281	375	469	563	656	750	125	250	500	625
165	91	182	273	364	455	545	636	727	121	242	485	606
170	88	176	265	353	441	529	618	706	118	235	471	588
175	86	171	257	343	429	514	600	686	114	229	457	571
180	83	167	250	333	417	500	583	667	111	222	444	556

Chorus

Introduction

Just as a chorus is a group of singers, the chorus effect can make a single instrument sound like there are actually several instruments being played. It adds some thickness to the sound, and is often described as 'lush' or 'rich'.

How it Works

The algorithm behind the chorus effect isn't a spectacular or amazing trick - it's actually fairly simple. What happens when two people play instruments in unison? Well they are not always playing in precise synchronization, so there is some delay between the sounds they produce. In addition, the pitch of the two instruments can deviate somewhat, despite careful tuning. These are the functions that your chorus effect is reproducing.

This slight delay can be easily implemented with a delay line. Creating the detuning effect may not seem very simple at first, but it can be achieved by transforming the simple delay line into a variable length delay line. The 'variable length' part just means that the delay time changes over time, though it's effect on the pitch may not be very clear at first.

To understand how the pitch is changed, picture the delay as a recording device. It is storing an exact copy of the input signal as is arrives, much like a cassette recorder, and it then outputs that a little later, at the same rate. To increase the amount delay, you want a longer segment of the signal to be stored in the delay before it is played back. To do this, you want to read out of the delay line at a slower rate than it's being written (the recording rate is unchanged, so more of the signal is being stored). Reading back at a slower rate is just like dragging your fingers on the wheel of the cassette, which we know lowers the pitch. Similarly, to reduce the delay time, we can just read back faster, analogous to speeding up a playing cassette, which increases the pitch - the 'munchkin effect.'

So now, by mixing this delayed and pitch modulated copy of the input together with the original, we have the chorus effect. A diagram for the chorus effect is given in Figure 1.

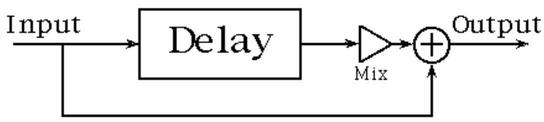

Figure 1: *Diagram of the chorus effects. The delay changes with time.*

This structure may look very familiar to you - it's basically our other counterpart the flanger. The chorus differs in only a couple of ways. One difference is the amount of delay that is used. The delay times in a chorus are larger than in a flanger, usually somewhere between 20 ms. and 30 ms. (the flanger's delay usually ranges from 1 ms. to 10 ms.) This longer delay doesn't produce the characteristic sweeping sound of the flanger. The chorus also differs from the flanger in that there is generally no feedback used.

The only remaining point to discuss is the manner in which the delay time actually changes. In general, some periodic waveform, such as a sine wave, is used. This waveform changes slowly (say than 3 Hz and below.) and is referred to as a **LFO (Low Frequency Oscillator)**. You can control the chorus sound by changing the waveform's frequency, its amplitude, and its shape. We make a simple change to our diagram of the chorus to denote this LFO dependence as in Figure 2.

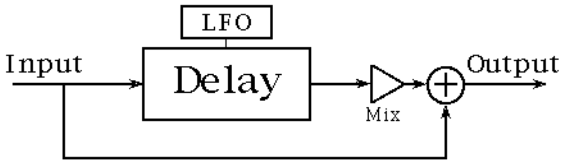

Figure 2: *The flow diagram for the chorus effect including its LFO dependence.*

Other variations on the chorus effect are also possible. For example, rather than using an LFO, you could use a randomly changing delay time, which might model musicians playing in unison a little better. Also, when playing in unison, there will be some loudness differences between the players, so we could also vary the amplitude of delayed signal. This amplitude parameter could then be controlled by another LFO.

Common Parameters

Delay

The delay parameter simply controls the amount of delay used. More specifically, it actually controls the minimum delay time that is used. As the delay becomes very small, the chorus will act as a flanger. Typical delay times range between 20 and 30 ms.

Sweep Depth/ Width

The sweep depth controls how much the total delay time changes over time. It is usually expressed in milliseconds, and the sum of the sweep depth and delay parameters is the maximum delay used in processing the signal. Alternatively, you can think of the sweep depth as the amplitude of the LFO. The relationship between the delay and sweep depth parameters is depicted in Figure 3.

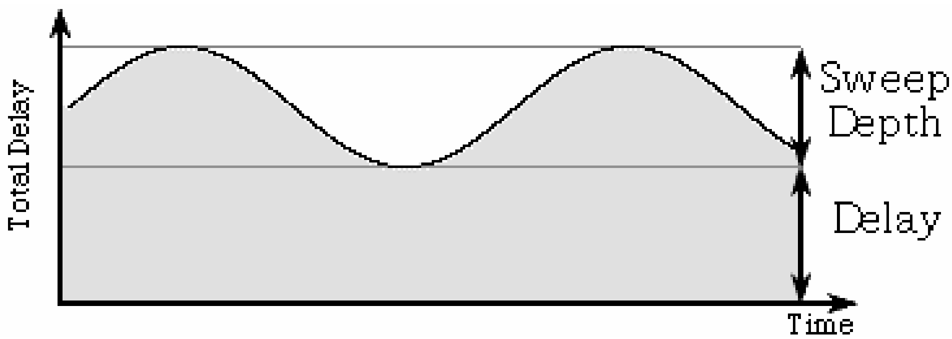

Figure 3: *The delay used in the chorus is the sum of the delay and sweep depth parameters, where the latter changes over time.*

The sweep depth also increases the pitch modulation introduced by the time-varying delay line. This happens because you have to read even faster and slower to cover the total change in time. Large sweep depths will create a 'warble.'

LFO Waveform The LFO waveform shows how the delay changes over time. When the waveform reaches a maximum, then the delay is at its largest value. When the waveform (and total delay time) is increasing, the pitch becomes lower. Some commonly used LFO waveforms are shown in Figure 4.

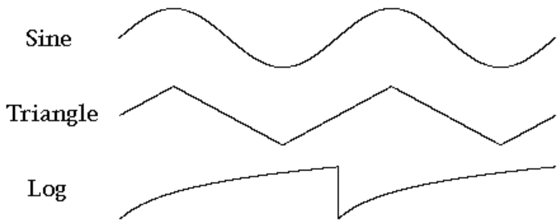

Figure 4: *Three commonly used LFO waveforms.*

The amount of pitch modulation introduced by the chorus is related to how quickly the LFO waveform changes - the steepest portions on the waveform produce a large amount of pitch modulation, while the relatively flat portions have very little or no effect on the pitch. We can use this view to understand how the sweep depth varies the pitch. If you increase the sweep depth, you are effectively stretching the waveform vertically, which makes it steeper, and thus, the pitch is altered more.

The sine is a very smooth function, and is always changing so the pitch is constantly changing as well. The triangle on the other hand only produces two pitches because the slope only takes on two different values, and the change between the pitches is sudden. The log waveform is smooth over it's cycle, but there is a jump at the end of each cycle. Since the slope at the beginning and end of the waveform are different, there is an abrupt change in the pitch as well. The sweep depth has been made very large in this example to stress the LFO shape's effect. This is generally not how the chorus is applied.

Speed/Rate

The speed control is pretty straightforward. This parameter refers to the rate at which the LFO waveform repeats itself, and it also is a factor in the pitch modulation. Using the view from above, increasing the rate is equivalent to compressing the LFO waveform in time, which makes it steeper, resulting in more pitch modulation.

Number of Voices

Up to this point, we have only covered what is referred to as a single voice chorus, meaning that there is only one copy of the input. But there's no reason why you can't have multiple copies of the sound, modeling a situation with more than just two instruments are being played. Some chorus units allow you to do just this, and may even allow you to choose how many voices to use.

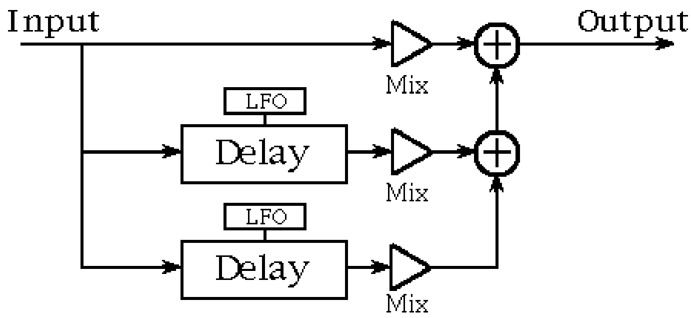

Typically a multi-voice chorus uses a single LFO waveform for all the voices, but each voice has a different phase. This means that at any point in time, each voice is at a different point along the waveform, so they have different delay times. If all the voices were in phase, it would have the same effect as a single voice chorus with an increased level. Of course, it's possible to build a chorus such that each voice use it's own LFO shape, and even rate.

Figure 5: *A dual-voice chorus diagram.*

Other Notes

Stereo Chorus

A stereo chorus is generally constructed by running two monophonic chorus run in **quadrature phase**. This simply means that the LFO in each monophonic flanger differ in phase by ninety degrees (or one-quarter of a wavelength). This technique creates a 'wider' sound because the sound arriving at each of your ears is different

Son of Flanger

As mentioned above, the chorus is essentially the same algorithm as a flanger, but the flanger is characterized by its smaller delay lengths (from about 1 ms. to 10 ms.) and a feedback path that routes the delay line output back into the input.

Flanger

Introduction
Flanging has a very characteristic sound that many people refer to as a "whooshing" sound, or a sound similar to the sound of a jet plane flying overhead.

Flanging is generally considered a particular type of phasing (another popular effect). As will be shown below, flanging creates a set of equally spaced notches in the audio spectrum. Phasing uses a set of notches as well, but the spacing of them can be arbitrary and the notches in a phaser are usually created using allpass filters.

How it Works
Flanging is created by mixing a signal with a slightly delayed copy of itself, where the length of the delay is constantly changing. This isn't difficult to produce with standard audio equipment, and it is believed that flanging was actually "discovered" by accident. Legend says it originated while the Beatles were producing an album. A tape machine was being used for a delay and someone touched the rim of a tape reel, changing the pitch. With some more tinkering and mixing of signals, that characteristic flanging sound was created. The rim of the reel is also known as the 'flange', hence the name 'flanging'. Most modern day flangers let you shape the sound by allowing you to control how much of the delayed signal is added to the original, which is usually referred to as a '**depth**' control (or 'mix'). Figure 1 is a diagram of a simple flanger with this depth control.

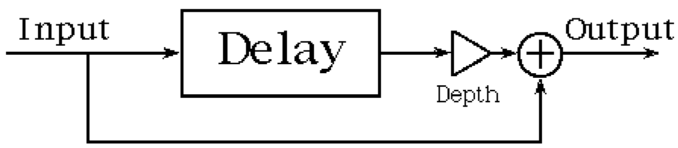

Figure 1: *Diagram of a simple flanger. The delay changes with time.*

When we listen to a flanged signal, we don't hear an echo because the delay is so short. In a flanger, the typical delay times are from 1 to 10 milliseconds (the human ear will perceive an echo if the delay is more than 50-70 milliseconds or so). Instead of creating an echo, the delay has a filtering effect on the signal, and this effect creates a series of notches in the frequency response, as shown in Figure 2. Points at which the frequency response goes to zero means that sounds of that frequency are eliminated, while other frequencies are passed

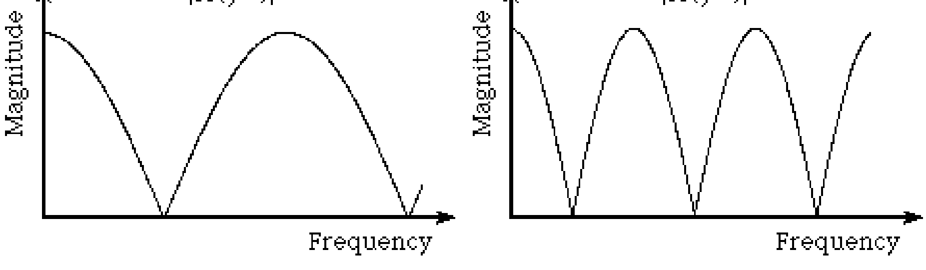

with some amplitude change. This frequency response is sometimes called a **comb filter**, as its notches resemble the teeth on a comb.

Figure 2: *The frequency response of a simple flanger with two different delay times (both with a depth of 1). The plot on the left would be for a flanger with a smaller delay than that on the right.*

These notches in the frequency response are created by destructive interference. Picture a perfect tone - a sine wave. If you delay that signal and then add it to the original, the sum of the two signals may look quite different. At one extreme, where the delay is such that the signals are perfectly out of phase, as one signal increases, the other decreases the same amount, so the entire signal will disappear at the output. Of course, the two signals could still remain in phase after the delay, doubling the magnitude of that frequency (constructive interference). For any given amount of delay, some frequencies will be eliminated while others are passed through. In the flanger, you can control how deep these notches go by using the depth control. When the depth is at zero, the frequency response is flat, but as you increase the depth, the notches begin to appear and extend downward, reaching zero when the depth is one. Even if the notches do not extend quite all the way to zero, they will still have an audible effect.

The characteristic sound of a flanger results when these notches sweep up and down the frequency axis over time. Picture the notches compressing and expanding like a spring between the two plots in Figure 2. The sweeping action of the notches is achieved by continuously changing the amount of delay used. As the delay increases, the notches slide further down into the lower frequencies. The manner in which the delay change is determined by the **LFO (Low Frequency Oscillator)** waveform (discussed below).

This changing of the delay in the flanger creates some **pitch modulation** - the perceived pitch 'warbles'. This happens because you have to 'read faster' or 'read slower' from the delayed signal. Picture a flanger created by two tape reels running the same audio signal. To increase the delay between the two signals, you have to slow one of the reels down. As you may know from experience, as you slow down a tape, the pitch drops. Now to decrease the delay, you have to catch up - sort of like fast forwarding, which increases the pitch (also known as the 'munchkin effect'). Of course only the delayed copy of the sound has this pitch change, which is then mixed in with the unaltered signal.

Common Parameters

Depth (Mix)
This is the depth parameter referred to above. The larger the depth, the more pronounced the notches in the flanger. In multi-effects units, the depth may only be controllable in the mixer section, and not available within the flanging processor. Some people use the term 'mix' interchangeably with 'depth'.

Delay
The delay parameter specifies the minimum delay used on the copy of the input signal - as the delay changes, this will be the smallest delay. Looking at the frequency response, this value determines how high the first notch will go. As the delay is increased, the first notch drops down. In some cases, the delay parameter can be set to zero, in which case the notches will sweep the uppermost frequency range, and essentially disappear momentarily. In other cases, you may not be able to control delay parameter.

Sweep Depth (Width)
The sweep depth determines how wide the sweep is in terms of delay time - essentially the width of the LFO. This sweep depth is the maximum additional delay that is added to the signal in addition to the delay in the delay parameter. It determines how low the first notch in the frequency response will reach. A small value for the depth will keep the variance in the delay time small, and a large value will cause the notches in the frequency response to sweep over a larger area. Figure 3 shows how the delay and sweep depth parameters are related to the LFO. The minimum delay applied to the signal is given by the delay parameter, and the maximum delay is the sum of the delay and sweep depth parameters.

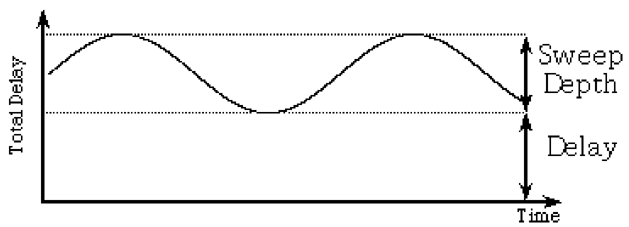

Figure 3: *The relationship between the sweep depth and delay parameters.*

As the sweep depth is increased, the pitch modulation effect mentioned above will become more noticeable. The flanger needs to read even faster or slower to change the delay in the same amount of time.

Note that when you vary the delay parameter, both the upper and lower limits of the first notch are changed, but when you adjust the depth, only the lower limit is affected. So when you are setting up a flanger to sweep over a particular range, first set the delay so that the high point of the sweep is where you want it, and then adjust the sweep depth to set the low point of the sweep. (When I refer to the sweep here, I'm talking about the notches, not the delay time. Remember that the parameters you set are controlling the delay, and that the notches result from this delay. As the delay increases, the notch frequencies decrease.)

LFO Waveform
Some flangers will allow you to choose the LFO waveform. This waveform determines how the delay in the flanger varies in time. Figure 4 show some common LFO's. The triangle is probably the more common choice in flangers.

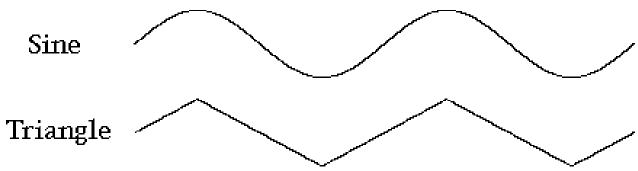

Figure 4: *Two common LFO waveforms.*

Feedback/ Regeneration

Some units will give you an option for taking a portion of the flanger's output and routing it to the input. In some cases, you can also specify whether to add or subtract the feedback signal. A large amount of feedback can create a very 'metallic' and 'intense' sound. A diagram for the flanger with feedback is shown in Figure 5. Of course as the feedback gain approaches one, the system can be come unstable, possibly resulting in overflow or clipping.

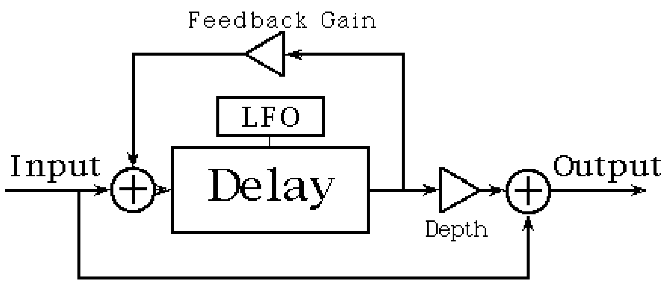

Figure 5: *The more complex flanger including a feedback path and LFO control over the delay.*

Speed/Rate

The speed control is pretty straightforward. This parameter refers to the rate at which the LFO waveform repeats itself, or equivalently, how many times per second the notches sweep up and down. The speed also affects the amount of pitch modulation. By increasing the speed, the flanger will have to sweep through the depth in less time.

Phaser

Introduction
The phase shifter (or phaser) achieves its distinctive sound by creating one or more notches in the frequency domain that eliminate sounds at the notch frequencies (the flanger also makes use of notches, and it is actually one specific type of phasing). The notches are created by simply filtering the signal, and mixing the filter output with the input signal. The filters can be designed so that we can independently control the location of each notch, the number of notches, and even control the width of the notches. This can lead to many interesting sonic possibilities .

How it Works
The notches needed for phase shifting (or simply called phasing) are most often implemented using a special group of filters called **allpass filters**. As the name implies, the allpass filter passes all frequencies - that is it allows all frequencies to appear in the output with no attenuation or amplification. So if you were to put any sine wave into the allpass filter, you would see a sine wave at the output, with the same amplitude it was at the input. To complete the phase shifter, we just add the filter output to the input signal, as in Figure 1. The amount of the filtered signal that appears in the output is set by the **depth** (also called the 'mix' or 'level') control.

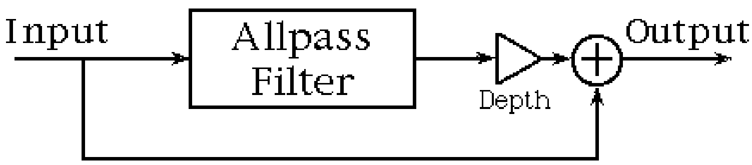

Figure 1

Well that's interesting, but if the filter passes all frequencies equally, how does it alter the sound, and where do the notches come from? Well there's one other characteristic of the filter we haven't mentioned yet, and that is the filter's phase response.

The **phase response** is kind of difficult to explain, so let's keep things relatively simple here. Let's say you're given a block box, and you don't know what circuit is inside of it. You can test it by putting a sine wave generator on the input and then look at the output and the input together on an oscilloscope. There are two important characteristics that can be observed. The first is the relative amplitudes of the two signals (which is referred to as the magnitude response). Again, for the allpass filter, these amplitudes are equal, so the magnitude response is one for all frequencies. The other important characteristic is the relative alignment of the two signals in time, i.e. do they both cross zero or reach the maximum values at the same time. The difference between the two signals is the **phase lag**, or the phase response. The black box is altering, or shifting, the phase of the input signal - hence the name 'phase shifting.'

The phase lag is usually measured in fractions of a wavelength rather than an amount of time. A single cycle of the wave is 360 degrees (regardless of the period or frequency of the sine wave), so 90 degrees is a quarter of a cycle, 180 degrees is half a cycle, and so on. Figure 2 shows some examples of a phase lag with a sine wave. The horizontal axis is time, so a delay moves things to the right.

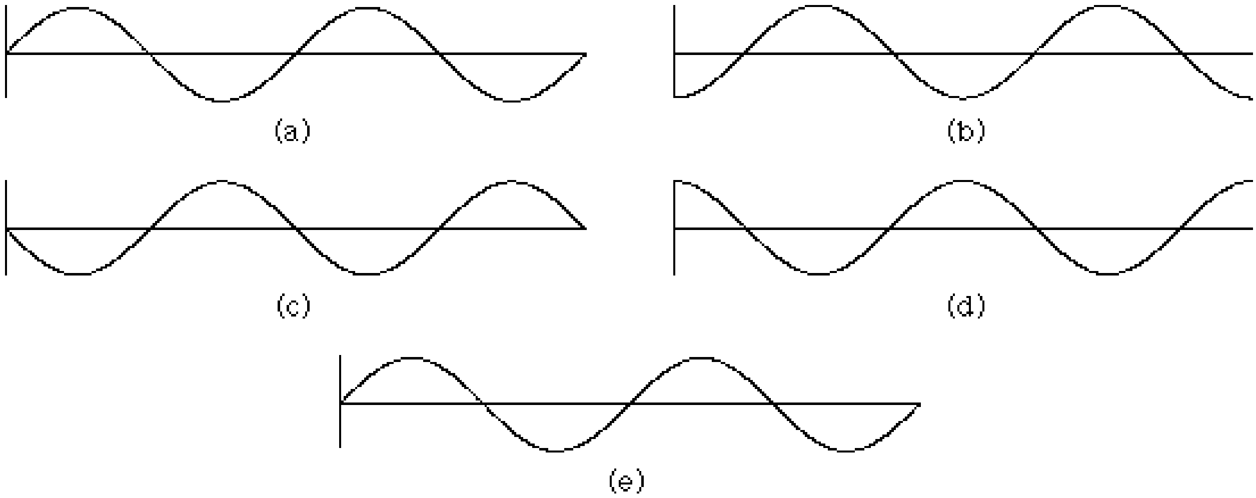

Figure 2: *A simple sine wave (a), and the same signal with a phase lag of (b) 90 degrees, (c) 180 degrees, (d) 270 degrees, and (e) 360 degrees (one cycle of the waveform). Note that when (a) and (c) are added together, the result is zero - complete cancellation.*

All practical filters have a phase response that changes with frequency, and Figure 2 is a plot of the magnitude and phase response of one possible allpass filter. One interesting case is a linear phase response (i.e. the phase response is a straight line, angled downwards). In this case, doubling the frequency means doubling the phase lag. The wavelength of the doubled frequency is half that of the original. This basically keeps all the frequency components aligned in time - it delays the signal. So the pure delay (no feedback or mixing with the original signal) is one type of an allpass filter. (On some audio products, there is a specification 'Deviation from linear phase.' Linear phase is desirable in some applications so that all frequencies are kept together in time and the sound isn't colored by 'phase distortion.')

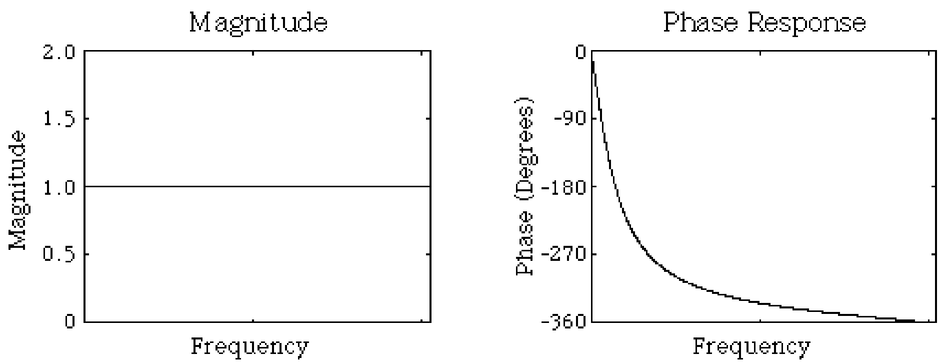

Figure 3: *The magnitude and frequency response for a particular allpass filter (of second order). Only the phase varies, while all frequencies are passed with a gain of one.*

Now we are ready to understand where the notches come from. To create the notches, we only need to mix the allpass filters output with the input. Why? At some frequencies, the phase lag introduced by the filters will be 180 degrees, which is equivalent to taking the negative of the input. When you mix this signal with the input, those frequencies that experience 180 degrees of phase lag will exactly cancel with those frequency components in the input, and that is the notch. Frequencies near the notch will also be attenuated somewhat. Essentially, a filter with a non-linear phase response, though it may be hard to believe, is delaying the signal, but not all frequencies are delayed by the same amount.

Going back to the linear phase case (the pure delay), the phase response hits values of -180 degrees minus multiples of 360 degrees (-180, -540, -900, -1260, etc.) at equally spaced frequencies. So when we mix a delayed copy of the signal with the original, there will be notches at equally spaced frequencies. This is exactly how the flanger operates, and thus, the flanger is merely one type of a phase shifter.

Common Parameters

Depth (Mix/Level)
The depth parameter controls the amount of the filter output that is added to the sound. It's called the depth control because as it increases, the depth of the notches increases as well. When the depth is set to 1 (or 100%), then the notches reach all the way to zero. (In many cases, people use 'depth' to describe how wide a range the notches sweep across, which is referred to as the 'sweep depth' in this article.) In some multi-effects processors, the depth or level control may only be controllable in the mixer section.

Sweep Depth (Range)
This parameter is used to control how far the notches sweep up and down in frequency. In some cases, you may be able to select actual frequency values, and in other cases, the base frequency may be set to some value and you can only control how far from that frequency the sweep will go.

Feedback/ Regeneration
The phase shifting effects can be made more intense by using feedback - adding part of the filter output to the input, as in Figure 4. Some units may also allow you to have negative feedback gains, which is equivalent to subtracting the output from the input. When working with several stages of allpass filters, you can conceivably use feedback on each individual stage as well as the entire structure.

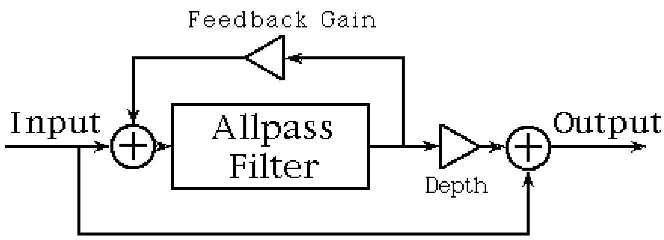

Figure 4: *The phase shifter with an added feedback path. The feedback gain can be negative as well as positive.*

Speed/Rate
This parameter simply controls how quickly the notches sweep up and down over the frequency range. The rate sets how many times the notches sweep up and back down per second only. The speed at which the actual notches are moving is determined by this rate control, the sweep depth, and the sweep pattern.

Other Notes

Stereo Phase Shifter
A stereo phaser could be constructed along the lines of a stereo flanger/chorus unit by using two filters with notches falling at different frequencies. It's also possible to create more complex sounds by selectively mixing the outputs of the two filters as well which can create additional notches that can exist on a single output. This general structure is depicted in Figure 7. If the feed-across gains are set to zero and each output fed to a speaker, you will

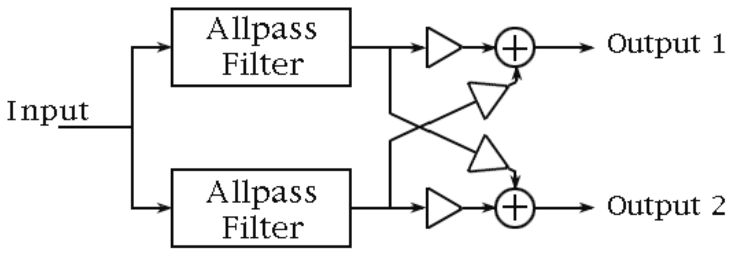

have a 'spatial phaser' where the notches exist only in the air and can be heard by simply moving about the room.

Figure 7: *A generalized diagram of a stereo phaser.*

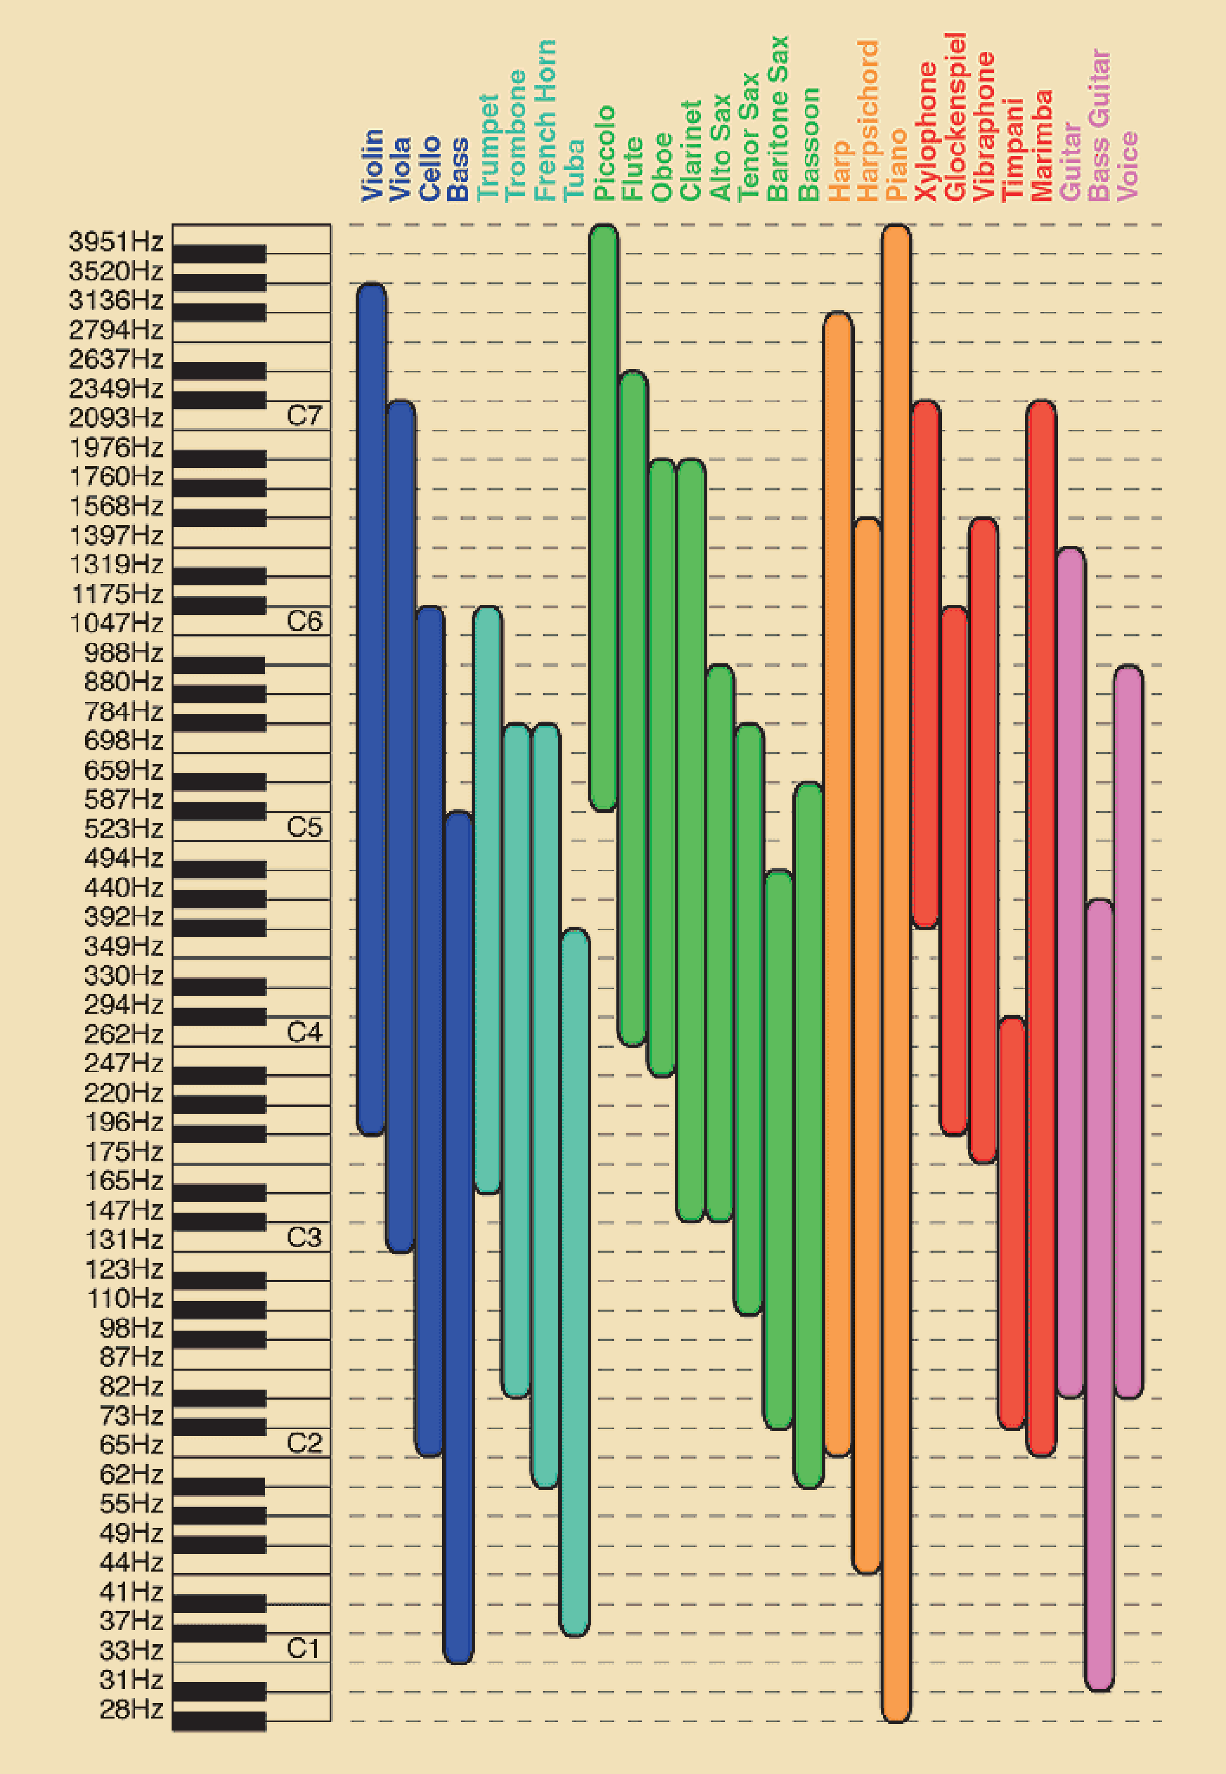
Violin
Viola
Cello
Bass
Trumpet
Trombone
French Horn
Tuba
Piccolo
Flute
Oboe
Clarinet
Alto Sax
Tenor Sax
Baritone Sax
Bassoon
Harp
Harpsichord
Piano
Xylophone
Glockenspiel
Vibraphone
Timpani
Marimba
Guitar
Bass Guitar
Voice
3951Hz
3520Hz
3136Hz
2794Hz
2637Hz
2349Hz
2093Hz
C7
1976Hz
1760Hz
1568Hz
1397Hz
1319Hz
1175Hz
1047Hz
C6
988Hz
880Hz
784Hz
698Hz
659Hz
587Hz
523Hz
C5
494Hz
440Hz
392Hz
349Hz
330Hz
294Hz
262Hz
C4
247Hz
220Hz
196Hz
175Hz
165Hz
147Hz
131Hz
C3
123Hz
110Hz
98Hz
87Hz
82Hz
73Hz
65Hz
C2
62Hz
55Hz
49Hz
44Hz
41Hz
37Hz
33Hz
C1
31Hz
28Hz

OVERVIEW OF EFFECTS
MidiVerb 4

Reverb Effects

Reverb is made up of a large number of distinct echoes, called reflections. In a natural acoustic space, each reflection's amplitude and brightness decays over time. This decaying action is influenced by the room size, the location of the sound source in the room, the hardness of the walls, and other factors. The MidiVerb 4 offers the following types of reverberation:

Concert Hall
This is a simulation of a large concert hall. Halls tend to be large rooms with lots of reflective surfaces, where sounds can swim around, changing timbre over time. This is a classic reverb which sounds good on just about anything. Try it on vocals, drums, acoustic, electric, or orchestral instruments.

Real Room
This algorithm gives you the sound of a medium size studio room. This algorithm uses a lot of processing power for a rich sound and smooth decay. It compares favorably to high end studio reverbs for its rich sound. The attack is also more reflective. It sounds good on drums, keyboards and guitars.

Realroom & Room
These are less processor intensive versions of the Real Room effect, used in Multi Chain and Dual Mono Configurations.

Ambience
This algorithm simulates a very small room. It can be used when just a slight amount of ambient character is needed to augment a sound. For example, if playing a solid body guitar, use the Ambience effect to simulate the sound of an acoustic guitar's hollow body.

Plate Reverb
This is a simulation of a classic echo plate, a 4' by 8' suspended sheet of metal with transducers at either end used to produce reverb. Popular in the 1970's, it still prized for its transparent sound, particularly on vocals and guitars. This algorithm uses the most processing available for a truly realistic reverb plate simulation. It works well for a lush lead vocal, piano, or guitar, especially when looking for a classic rock and roll sound.

Nonlinear
This reverb effect's direction can be set either forwards or backwards. Selecting the forward direction provides a classic "Gated" digital reverb sound. Selecting the reverse direction gives you a backwards reverb sound. A popular trick in the 80's was to record the reverb with the tape flipped over, so it would play backwards in the mix. The reverse reverb is a useful effect for drums and other percussive sounds –— adding space without washing out the instrument.

Reverb Parameters

Most of the reverb effects in the MidiVerb 4 operate under the same set of control parameters, which are listed and described in this section. However, reverbs which use more processing power (i.e. the Single Configuration reverbs) provide more parameters which take advantage of their extra processing power; parameters which are not found in the other, smaller reverb algorithms. For example, Reverberation Swirl is a parameter found only in the Single Configuration reverb types. Here are the reverb parameters:

Decay
The Reverb Decay determines how long the Reverb will sound before it dies away. When using the Reverse Reverb effect type, the Reverb Decay parameter controls the Reverse Time.

Low Pass Filter
The lowpass filter that can be set between 059 Hz and 36.2 kHz, and attenuates all frequencies above this value by 6dB per octave. The lower the setting, the less high frequencies of the input are allowed to pass thru to the reverb effect.

Pre-delay
All the reverb effects have pre-delay parameters. Pre-delay slightly delays the reverb itself up to 175 ms, so that the dry signal more easily stands out from the reverb. A bit of pre-delay can make certain instruments (such as snare drums) sound bigger.

Pre-delay Mix
This allows you to balance the amount of Pre-delay to Direct Signal as a percentage of each. This gives you the ability to hear a bit of the Reverb before the loudest part of the Reverb (the Pre-Delayed Reverb) sounds, and makes for a bigger, smoother sounding Reverb.

Density
Density controls how the first reflection of the reverb effect will appear. When set to 0, the first reflection is heard alone without any other reflections. When set to 99, the first reflection appears to “fade-in” and then “fade-out”. This is because a number of reflections will occur just before and just after the first reflection, in addition to the remaining reflections heard after the first reflection. Thus, the reverb sounds more “dense”.

Diffusion
Diffusion determines the “thickness” of the reverb sound by adding more reflections to the reverb’s decay. With lower diffusion settings, you may be able to actually hear the individual echoes that make up the overall reverb sound. With higher diffusion settings, the echoes increase in number and blend together, washing out the reverb’s decay. Greater diffusion works better with percussive sounds, whereas less amounts of diffusion work well with vocals and other sustained sounds.

Frequency Damping – Low & High

These two parameters allow you to control the equalization of the reverb's decay separately for both the low and high frequencies. This means that you have control over the tonal shape of the Reverb itself, being able to cut the high frequencies if the effect is too bright, and being able to cut the lows if the effect is too boomy. These parameters allow you to simulate different surfaces of a room or hall, with softer surfaces absorbing more high frequencies and smaller rooms having less low frequencies. *Example:* If a room has lots of drapes hung, the high frequencies will decay faster than the lower frequencies.

Reverberation Swirl
This parameter is very useful for smoothing the decay of the reverb when set at a low value. When set to a high value, it creates a more dramatic detuning effect as the reverb decays.

Gating
Gating is the process which abruptly cuts off the reverb's decay for a more "choppy" sound. This effect is very popular on drums because it makes them sound HUGE. It is achieved by dropping the level of the signal very rapidly after the initial attack making a short, sharp sound.
In all the Single Configuration reverb effect types and most of the Double and Multi Chain Configurations which use the Realroom effect type, there are three gating parameters available. These include: Gate, Hold Time, and Release Time. The Gate controls the level of the reverb signal after the gate closes, and can be set between 001 and 100%. In other words, if Gate is set to 100%, then no reverb will sound after the gate turns it off. If Gate is set to 50%, then some reverb signal will still be present even after the gate turns off the main reverb signal. Alternatively, the Gate parameter can be set to "OFF" when you do not wish to use the gating effect. The Hold Time determines how long the gate will be held open before it begins to turn off; this can be set from 0 to 500 ms. The rate at which the gate closes is determined by the Release Time, which can be set from 0 to 500 ms.
In the case of the Chorus->Room, Flange->Room and Room->Flange Configurations, only one parameter is available: Gate. This can be set between 10 and 500 ms, and controls both the hold and release times of the gate effect. Alternatively, the Gate parameter can be set to "OFF" when you do not wish to use the gating effect.

Delay Effects

Delay provides a discrete repetition of a signal. By adding feedback within the effect, the delayed signal can repeat many times, with each successive decay softer than its predecessor. Each of the Delay types allow you to adjust their delay time in milliseconds, however, the BPM Delay effect will display the equivalent musical tempo in BPM (beats per minute). MidiVerb 4 offers the following types of delay:

Mono Delay
This Single Configuration provides delay of signal up to 1299 ms. The delay time can be adjusted separately by 100ths, 10ths and 1 ms increments. Feedback is also available to increase the complexity of the signal. You also have high and low frequency cutting, which gives you the ability to equalize the effect's decay. This can help emulate an old tape-style echo where each succesive echo is darker than the previous one.

Stereo Delay
This Single Configuration provides two separate delays which can be individually adjusted for delay time, feedback and high and low cutting. The delay time can be adjusted separately by 100ths, 10ths and 1 ms increments.

Ping Pong Delay
So called because the output bounces from left to right in stereo with the speed determined by the delay time. Again, low and high frequency cut is available. The delay time can be adjusted separately by 100ths, 10ths and 1 ms increments.

MultiTap Delay
This is like having three delays at once. Each of the 3 "taps" have individual delay, level, panning and feedback controls. By adjusting the delay time of each tap, you can create sophisticated rhythms.

BPM Mono Delay
This is a mono delay which can have its delay time parameter set to a specific tempo or BPM (beats per minute) value. This allows you to reference the delay time to the tempo of the music you are playing, rather than searching for the correct delay time in milliseconds.
An additional parameter, called Note, is used to determine what beat value your tempo represents. For example, if you set the note to 4, then you can set the tempo using quarter-note beats to establish delay time. If instead you set the Note to 4t, the same delay tempo setting will play faster because it is simulating quarter-note triplets in relation to the selected tempo. You can also choose dotted-note values, such as 4d or 8d for different rhythms relative to the selected tempo.

Delay & DLY
These effects are mono, less processor-intensive versions of the Stereo Delay effect, used in the Multi Chain Configurations Delay->Realroom, Chorus->Dly->Room and Flange->Dly->Room; the Double Configuration Realroom+Delay; and the Dual Mono Configurations Delay:Delay, Chorus:Delay and Flange:Delay. They provide only high frequency cutting ability with no control over the low frequencies.

Setting Delay Time Using Tap Tempo
You can adjust the delay time using a technique called “tap tempo”. By tapping the button which corresponds to the Tap parameter, you can have the MidiVerb 4 follow your tapping and adjust its delay time to match the tempo you are using. If the Footswitch parameter (UTILity mode) is set to *Control*, you can tap your delay time by repeatedly pressing down on the footswitch.
You can also adjust the delay time using tap tempo from the audio source being routed to the MidiVerb 4’s input(s). This can be done in two ways:
• Hold the button which corresponds to the Tap parameter; or
• Hold down the footswitch (if the Footswitch parameter is set to the Control function).
While using either of these methods, feed signal to the MidiVerb 4. This could be done by hitting a drum, plucking notes on a guitar or keyboard, or by singing some “doot doots” into a microphone (depending on what is connected). *Note:* When the Footswitch parameter is set to the Control function, you can control tap tempo as described above while in either Program mode ([PROG] button lit) or Edit mode ([EDIT/PAGE] button lit), unlike when using the front panel for tap tempo which requires that you be in Edit mode. For more information on connecting a footswitch and selecting the Footswitch parameter’s function, see Chapter 2.

Pitch Effects

The Pitch effects alter the pitch of a signal in various ways to produce “layered” timbres that are more complex than the original signal. Although some of these effects can sound similar to one another depending on the parameter settings, each is achieved differently and can be quite dramatic under the right circumstances. Pitch effects are achieved by splitting the signal into at least two parts, effecting the pitch of one of the parts, then mixing them back together. This eventual mixing is essential since the overall sound of the effect is achieved by the actual difference between the dry, uneffected signal and the effects signal. The various types of Pitch change are:

Stereo Chorus
The Chorus effect is achieved by splitting the signal into three parts with a dry signal and a separate Detuning section for both left and right channels. When the left channel is detuned sharp, the right is detuned flat, and vice versa. The detuning is further effected by being modulated by an LFO (low frequency oscillator) which causes the detuning to vary. Many variables are available in this scheme: the Predelay can be varied, the LFO depth can be varied, the LFO speed can be varied, and a portion of the detuned signal can be fed back to the input to increase the effect. Finally, the waveform shape of the LFO can be changed from a smooth sinewave, to a more abrupt squarewave to make the pitch detuning more pronounced.

Quad Chorus
Quad Chorus modulates four delayed signals, each with its phase offset by 90°. Each of the four signals has a separate Predelay variable, allowing you to change the “rhythm” of the phasing.

Chorus
This is a mono, less processor-intensive version of the Stereo Chorus effect, used in the Multi Chain Configurations Chorus->Realroom and Chorus->Dly->Room, the Double Configuration Realroom+Chorus, and the Dual Mono Configurations Chorus:Chorus and Chorus:Delay.

Stereo Flange
First used in the 1960s, “flanging” was achieved by the use of two tape recorders that would record and play back the same program in synchronization. By slowing down one tape machine, and then letting it catch up with the other, different phase cancellations would occur at different frequencies. Since the slowing down of the tape machines was done by hand pressure against the flanges of the tape supply reels, the term “flanging” came into being.
Flanging is similar to chorusing , but modulates the delayed signal over a much shorter delay range (typically 0-12 ms). This produces a “jet airplane”-like sound. The flange modulation sweep can be triggered by the audio input (either the left or right input, or both), in order to sync up with the rhythm of your playing. You can adjust the attack and release threshold of this audio triggering function.
In the case of the Stereo Flange, the signal is split into three parts with a dry signal and a separate Delay section for both left and right channels with one channel flanging up while the other channel flanges down. Once again, this causes the effect to become more pronounced and dramatic.
When flanging was done using two tape machines, it was possible for one to be behind the other, catch up and then go past the other. This is called passing “through zero”. The “zero” point is when both signals were in perfect synchronization. Since the MidiVerb 4 is digitally simulating the flanging effect, it normally cannot provide the through zero effect. Instead, it delays the effected signal to a point, then brings it back to the zero point, and repeats this over and over. The “Thru0” parameter found in the MidiVerb 4’s flanging effects lets you create the appearance of the effected signal passing through the zero point. It does this by actually delaying the uneffected signal by as mush as 12 milliseconds (an amount virtually undetectable to the human ear). This allows the wet signal to move “behind” the dry signal as it cycles.

Flange
This is a mono, less processor intensive version of the Stereo Flange, used in the Multi Chain Configurations Flange->Realroom, Realroom->Flange and Flange->Dly->Room; the Double Configuration Realroom+Flange; and the Dual Mono Configurations Flange:Flange and Flange:Delay. The effect of mono flanging is achieved by splitting and slightly delaying one part of the signal, then varying the time delay, with an LFO. The delayed signal is then mixed back with the original sound to produce the “swishing” or “tunneling” sound.

Lezlie
With the Lezlie effect (found in the Lezlie->Room Configuration), the pitch change block becomes a rotating speaker simulator. This effect was extremely popular during the 1960’s and was achieved by mechanically rotating the speakers to produce complex timbral changes. The Lezlie speaker system is most often used with rock organs, but is occasionally used for guitar amplification as well. Parameters include: Motor on/off, Speed, which can be slow or fast; and High Rotor Level, which lets

you attenuate the volume of the high frequencies. When switching the Lezlie effect on and off, or when changing the speed between fast and slow, the effect will ramp rather than change abruptly, just as a true Lezlie speaker system would do. By raising the High Rotor Level, you can really make this effect scream. *Tip:* Try modulating the Motor or Speed with aftertouch.

Stereo Pitch Shifter
The Pitch Shifter effect transposes the pitch of the incoming signal by a fixed amount. It is useful for creating parallel harmonies, detuning, chorusing, and special effects. The Semi parameter shifts the pitch in increments of one half step, with a range of up or down one octave. The Fine parameter detunes the signal in very fine increments, with a range of up or down one half step. Also available are Delay and Feedback parameters. The Delay parameter delays the shifted signal up to 250 ms, or can be used with the Feedback parameter to produce decaying arpeggio effects. There are also Low Cut and High Cut filters in the feedback loop which can be used to alter the timbre of the sound as it repeats. The Stereo Pitch Shift configuration provides two discrete pitch shifters, each with their own Pan and Level control.

Pitch
This is a less processor-intensive version of the Pitch Shifter effect, used in the Multi Chain Configuration Pitch:Delay, and the Dual Mono Configuration Pitch:Pitch.

Auto Pan
The Auto Pan effect alternates the loudness of the signal in opposite channels at a definable rate. Low and high frequency cutting is available, and (like the flange effects) can be triggered by the input signal (either the left or right input, or both).

The Direction parameter determines which direction the panning will start in when triggered (this has no effect if the Trigger parameter is turned off). You can adjust the attack threshold of this audio triggering function, and adjust the hold time (how long the Auto Pan will wait before it can be retriggered). The Direction parameter can be set to alternate; that is, the Auto Pan will change its start direction with each succesive trigger.

THE SPX990 SYSTEM

EFFECT CONFIGURATION

The effect programs of the SPX990 consist of three separate effects – pre-effect, main effect and post-effect.
The unit offers 36 main effects, 4 pre-effects and 3 post-effects to choose from. The pre- and post-effects provide equalizing and dynamic control of the main effects.
In addition, each of these effects is controlled by a number of parameters for almost unlimited possibilities in the creation of effect programs.

4 Pre-effects

3-band-PEQ:	P.EQ
Compressor:	Comp.
Aural Exciter:	AuralX
Compressor, Distortion & EQ:	Dist.

36 Main Effects

REVERB EFFECTS

Reverb:	Reverb
Filtered Reverb:	FiltRev
Stereo Reverb:	St.Rev
Echo Room:	EchRoom

ER (EARLY REFLECTION) EFFECTS

Thin Early Reflection:	ThinER
Fat Early Reflection:	FatER
Gate Reverb:	GateRev
Reverse Gate:	Reverse

DELAY, ECHO, EFFECTS

Delay L, C, R:	Dly-LCR
Echo:	Echo
Multi Tap Delay:	Mlt.Tap
Stereo Echo:	St.Echo
Tempo Mono Echo:	TmpEch 1
Tempo Stereo Echo:	TmpEch 2
Tempo Quad Echo:	TmpEch 4

MODULATION EFFECTS

Flange:	Flanger
Dual Flange:	DualFlg
FM Chorus:	FM.Cho
AM Chorus:	AM.Cho
Phaser:	Phaser
Symphonic:	Symphon

PITCH CHANGE EFFECTS

Mono Pitch Change:	MonoPit
Dual Pitch Change:	DualPit
Triple Pitch Change:	TripPit
Stereo Pitch Change:	StPitch

PAN EFFECTS

Auto Pan:	AutoPan
Triggered Pan:	TrigPan

COMPLEX EFFECTS

Chorus & Reverb:	Cho & Rev
Symphonic & Reverb:	Sym & Rev
Flanger & Reverb:	Flg & Rev
Reverb (L)/Reverb (R):	Rev/Rev
ER (L)/Reverb (R):	ER/Rev
Echo (L)/Reverb (R):	Ech/Rev
Chorus (L)/Reverb (R):	Cho/Rev
Pan (L)/Pan (R):	Pan/Pan

FREEZE EFFECT

Freeze:	Freeze

3 Post Effects

3-band-PEQ:	P.EQ
Compressor:	Comp
Aural Exciter:	AuralX

MODULATION EFFECTS

Mixing sounds with slightly varied delay times results in tone alterations due to mutual phase interference. In addition, the delay time and delay sound level can be modulated using the LFO, resulting in further tone alteration over time.

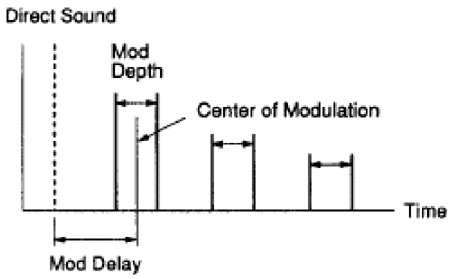

Flange (Flanger)	**2in/2out**
Dual Flange (DualFlg)	**2in/2out**

The flanging effect is produced by varying the delay between 2 identical signals, thus producing a complex varying "Comb Filter" effect.

Parameters

① ModFrq (Modulation Frequency: 0.05Hz ~ 40.0Hz)
Sets the speed of modulation, and hence the rate at which the effect varies.

② Depth (1, 2)(Modulation Depth: 0% ~ 100%)
Sets the amount of delay time variation, thus adjusting the depth of the effect. A bigger value will give at deeper modulation.

③ Delay (1, 2)(Modulation Delay Time: 0.1msec ~ 100.0msec)
Sets the basic delay time from the initial direct sound to the beginning of the flange effect. A setting of 1.0msec and below causes interference in the high frequency range.

④ Phase (Phase: -180.0deg ~ +180.0deg)
Sets the phase between Modulation Delay 1 and 2.

⑤ FbGain (Feed Back Gain: 0% ~ 99%)
Determines the amount of flange signal fed back to the input of the processor for further modulation. More feedback increases the overall complexity, "strength" and decay time of the effect.

⑥ InMode (Input Mode: Mix, Stereo)
Selects the input mode for processing of a mixed left and right channel signal (Mix) or separate left and right channel signals (Stereo).

⑦ HPF (high Pass Filter Frequency: THRU, 32Hz ~ 1.0kHz)
Permits rolling off the low frequency content of the reverb signal below the set frequency. The HPF is OFF when set to THRU.

⑧ LPF (Low Pass Filter Frequency: 1.0kHz ~ 16kHz, THRU)
Permits the rolling off the high frequency content of the reverb signal above the set frequency. The LPF is OFF when set at THRU.