Good Dynamics Processing

• A COMPRESSOR
  • Threshold
  • Ratio
  • Attack time
  • Release time
  • Gain control
  • Hold time
  • Compression caused distortion
  • Signal measurement
  • What to look for when buying

• A LIMITER

• A DOWNWARD EXPANDER
  • What to look for in an expander

• A GATE

• HOW TO "KICK THE TIRES" OF A GATE

Ray Bennett
Rane Corporation

A COMPRESSOR

Producing a good, effective compressor is not a trivial task. It has to "rein in" excessive program dynamics without being unduly intrusive. It does this by reducing the gain in the audio path when the signal level exceeds a pre-determined threshold. Four important parameters need to be controlled. Refer to Figure 1 for the following definitions.
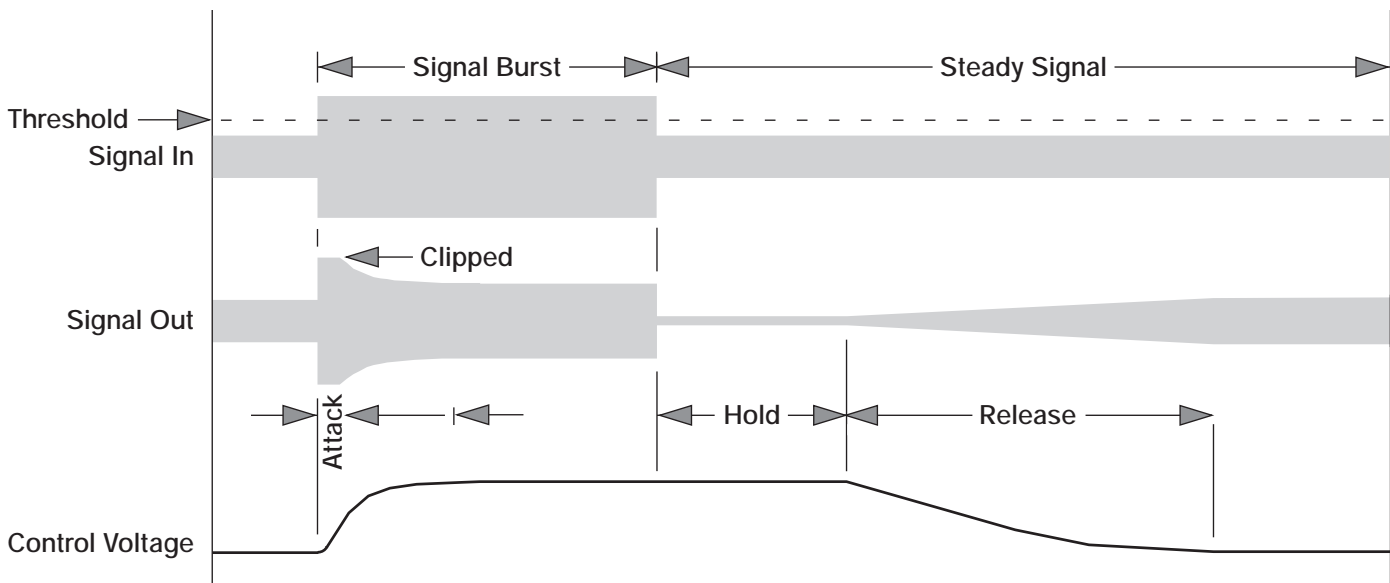
Figure 1. Short Segment of Signal Dynamics

Note that the signal increase is too fast to keep under control with the illustrated attack time. In this case, either the attack time must be decreased, or the high frequency content must be limited. Although not totally obvious, limiting the high frequency content also limits the rate-of-change of the signal level—that is, the sound can only get louder at some maximum rate.

Threshold

The signal level above which the compressor reduces system gain. This parameter is almost always adjustable. Note that a threshold that exceeds the system clip level is essentially a bypass function. A workable range for this control is –40 dBu to +20 dBu. Notice that if the threshold is set to a low level of –40 dBu, the compressor begins to look very much like a Leveler or Automatic Gain Control (AGC).

Ratio

A measure of how much the gain will be reduced under given conditions. A ratio of 2:1 means that once the signal level exceeds the threshold, the output signal level is allowed to increase by only 1 dB for every 2 dB of input increase. The desired setting during use should be the minimum that will provide the required overload protection. Normal ratio for a compressor ranges from 1.5:1 to 10:1. This parameter is almost always adjustable. Note that a ratio of 1:1 is analogous to a functional bypass.

Attack time

How long the compressor takes to control the signal after the actual overload occurs. In a good compressor with adjustable attack time, expect a range of adjustment from 500 microseconds (us) to 100 milliseconds (ms). Some manufacturers specify this time from the beginning of overload to some arbitrary gain reduction – perhaps –3 dB. Others specify it from the start of overload to the time the gain stops changing. Perhaps the most meaningful is to specify from start of overload to when the system gain is within, say, 3 dB of the final control point for 10 dB of gain reduction. Unfortunately, in the war of specsmanship, this measurement suffers by comparison. In a system where this parameter is adjustable, ridiculously short times are often spec'd just to look good. Some manufacturers state the attack time as from onset of a very large, very fast overload to the *beginning* of gain control. This is an artificially short time since it really doesn't reflect how long it takes to get the overload under control. The important consideration is "how long does it take to control an overload?" Equally important as speed, is the *shaping* of the attack function. If badly done, even a slow attack will sound abrupt and "clicky." Unfortunately, the buyer is not usually informed as to this critical part of the design. This is where a reviewer's article could pay off, as well as a carefully done listening test by the potential buyer. If during testing, the compressor sounds intrusive when the attack is reduced below 1 ms, try a different compressor.

Release time

How long the compressor takes to relinquish control once the overload passes. Same problems of specification as the attack time, but of less consequence. Normal adjustment range is from 100 ms to 3 sec or more. A short release time of 100 to 500 ms is a good starting range for spoken voice, while the longer times are better for instrumental music. This time inversely affects the distortion added by the compression process, as will be shown. Release time is usually adjustable. Here again, the buyer is at the mercy of the designer. Much has been done with release circuitry over the years to produce good compressors. Terms like "program dependant" and "dynamics dependant" abound. Some have genuine meaning, while some are hype to get the buyer's attention. Bottom line? How does it sound when compressing?
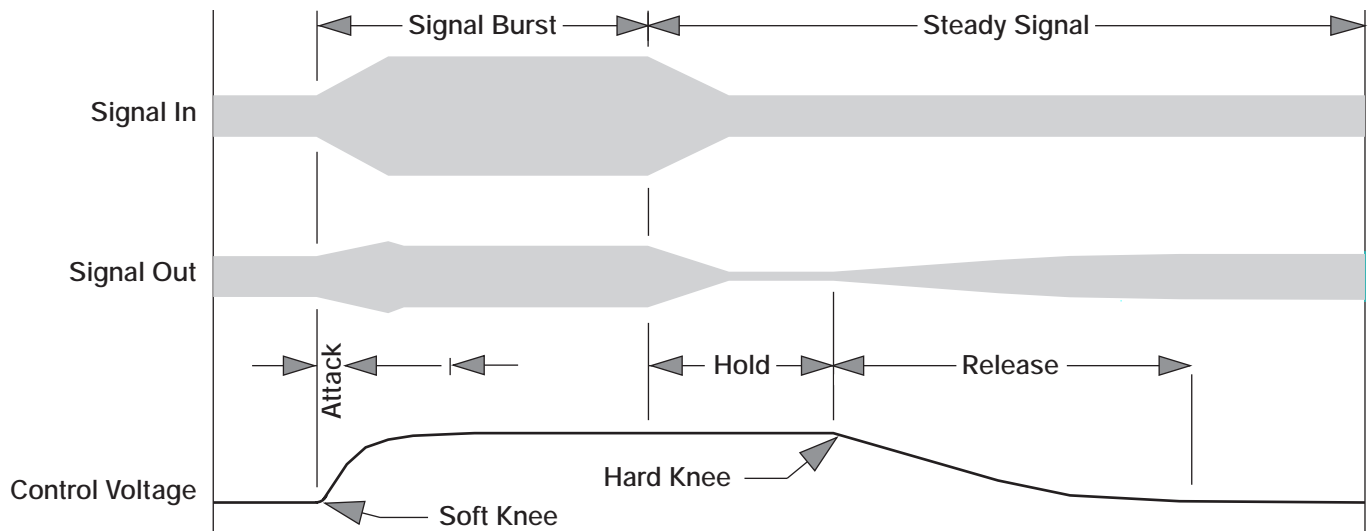
Signal Burst — Steady Signal

Signal In

Signal Out

Attack

Control Voltage

Hold — Release

Hard Knee

Soft Knee

Figure 2. Signal Dynamics with Limited Frequency Response

In this example, note that the signal attack time is slow enough to be kept under control. The hold time, if adjustable, could have been shorter.

Gain control

This is sometimes called "make up gain." It is used to adjust the output level to the desired optimum. This control usually has a range of plus and minus 15 to 20 dB. A center detent at 0 dB of gain is often provided. Beware of too much control, as it makes it difficult to adjust to the desired value. Some feel that ±10 dB is plenty. The author believes ±15 dB is a good range. If the user is employing lots of compression, then more gain range might be appropriate. However, if subtlety is the buyer's style, 10 dB may be better.

In addition, the following terms may be encountered:

Hold time

How long the compressor takes to *begin* the release process once the overload passes. Not normally spec'd in a compressor, but a valid parameter nonetheless.

If provided, the hold time allows the operator to maintain the important "tail" of cymbals, triangles, etc., without distorting the dynamics. In most compressors, this is accomplished by making the release time quite long. A hold function allows the operator to more accurately tailor the compressor function to the dynamics of the program material. If the adjustment is provided, it should be from 50 ms to 1 sec or more. Adjusting the hold time entails using the minimum setting that preserves the program dynamics. Speech would require minimum hold, while solo instruments such as vibraphone would probably sound better with a hold time of 1-2 seconds.

Compression caused distortion

This is largely ignored but is very significant. In one U.S. made compressor, which is generally considered to be quite good, operating with 10 dB of compression, minimum release time, ratio of 10:1, at a frequency of around 100 Hz, the distortion exceeds 5% THD (Total Harmonic Distortion)! This is not at all abnormal. As the release time is increased, this distortion decreases but never goes away completely. There is a technique that will eliminate it completely, but that is the subject of another paper. Incidentally, the author believes this distortion often causes the operator to adjust the release time longer than should be used. The assumption is that the sound improved because of the longer release time, but perhaps it actually sounded better because of the reduced distortion (to a point). The condition that saves the day is that this distortion is relatively difficult to hear, since it tends to be masked by the program dynamics that brought it about in the first place. The same effect masks occasional clipping. For example the cannons in the *1812 Overture* are almost totally clipped but still manage to sound reasonably good. Try treating a vocal that way and the singer will probably be unrecognizable.

The source of this compressor caused distortion is ripple remaining superimposed on the control voltage presented to the voltage-controlled amplifier (VCA). Because the rectification and log conversion is full-wave, the ripple is at twice the signal frequency. The release time acts as a single pole low-pass filter, which reduces this ripple at the rate of 6 dB per octave. That is, as the signal frequency goes up, so does the ripple frequency. As the ripple frequency increases one octave, it is reduced in amplitude by 6 dB (by one half) by the action of the release filter. Since the release filter is set to

around 0.5 seconds or more, its implied "rolloff frequency" (-3 dB point) is 0.32 Hz or less. (The formula is freq = 1/(2*pi*TC) where TC is the time constant – 0.5 sec in this example.) One would think that 0.32 Hz is so far below any signal of interest that there just can't be any ripple remaining. If the VCA didn't deal with it the way it does, that supposition would be true; however, the VCA functions by multiplying the control voltage times the signal voltage. This means that the ripple (at twice the signal frequency) multiplied times the input signal causes even-ordered harmonics. Often the designer has worked hard in the design process to eliminate the generation of these harmonics elsewhere in the audio path, only to generate them in the VCA.

Signal measurement

Much has been written concerning the proper rectification/log conversion process used to measure the incoming signal level. Some advocate peak logging, while others say that RMS log conversion is the only way to go. It must be kept in mind that, by definition, RMS log conversion requires averaging over several cycles of the signal. At most frequencies this makes the attack time unacceptably long. Often a compromise is used where the averaging time for the log conversion process is essentially a peak function at low frequencies and RMS at high signal frequencies. In- between it's anyone's guess. This sounds sloppy, but in practice, it works quite well. Ultimately, the release time filter converts the peak log to an averaged peak log. Since the logging process reduces the apparent extremes in the signal dynamics, this averaging is a reasonably good first-order approximation. One way to improve the ripple removal is to use a two-pole release filter. This two-pole approach reduces the ripple twice as fast – 12 dB per octave at the cost of more circuitry. Another advantage to using two-pole filtering (both attack and release) is the shaping of the transitions (the "knees") tends to be much smoother and less abrupt. The attack and release can be adjusted to be faster while still allowing the compression process to sound good. Unfortunately, unless the manufacturer "brags" about the method used, there really isn't any easy way to tell how the circuitry was designed. The bottom line remains the same – How does it sound? Keep in mind, how it sounds is the *only thing that will remain in your recording.* All the fancy lights, expensive extruded front panel, stainless steel cabinet, etc won't be in the credits.

What to look for when buying a compressor

A compressor is enough of an investment that, like a new car, it should be "test driven.". Look at the front panel. The function of each knob and switch should be clear and understandable. If the salesman has trouble explaining it so that the function makes sense, consider another brand or another salesman. Turn the knobs, feel the "action." It should be smooth, not "gritty." While this may seem silly for some, it is a good indication of how much effort the builder puts into buying quality parts. Listen to it carefully with program material similar to what will be used in the buyer's situation. Adjust for minimum attack, maximum ratio, mid-range release (perhaps 0.5 sec). Adjust the threshold until the gain reduction meter shows 10 dB or so of gain reduction with no

clipping during the loud music. No gain reduction meter or clip indicator? This is a good time to look at another brand. There should be at least a threshold light to help set the control levels. Once the unit is set up, listen very carefully to the transition from quiet material to the loud transients, such as drumbeats. There should be nothing added by the compressor except decreased dynamic range. If it sounds worse when compressing, look elsewhere. Understand that the rather extreme compression that was just set up will sound somewhat "squashed", but it shouldn't sound "clicky" or "nasty." The click sound usually comes from the attack not being shaped properly, while the nasty, distorted sound is simply the result of a poor compressor. This is the area that the buyer should try to get the best possible performance for the money being spent. Once this part is satisfied, the buyer can look for attractive graphics, lots of lights and extra functions. Be careful about paying extra for a function that isn't needed. Although, one extra function that is almost imperative is a limiter with separate controls. Often, a downward expander is also included. Depending on the intended use, this may be worthwhile.

A LIMITER

What makes a good limiter? One that is used infrequently but effectively. As will be shown, an effective limiter has constraints that preclude it sounding good. That may seem like a contradiction, but it really isn't. That's not to say that there aren't good limiters – there certainly are. There just aren't *good* limiters that make good *compressors.*

A Rane DC 24 is an example of a good feature set. It has two channels of compression, limiting, and downward expansion.

If we make it a condition that the limiter cannot pass any signal that exceeds some threshold, then we automatically impose certain restrictions as follows:

To prevent transient overload, it must have an attack time faster than the fastest rise time expected. See Figures 1 & 2. The fastest rise time may be defined as one divided by the highest system frequency. This frequency is estimated as four times the –3 dB rolloff of the system for a single processing device. "Four times" is an arbitrary and conservative constraint. It assumes that the highest frequency will be 12 dB lower in amplitude than the program material due to the 6 dB per octave rolloff attributable to the high frequency limitation. The 6 dB figure is correct only if the system rolloff is defined by a single filter. If several devices are in sequence in the audio path and they have similar rolloffs, the maximum frequency can be downgraded to two times the rolloff of the system. For the sake of discussion, let's assume there are two devices in the path and they each rolloff around 50 kHz at 6 dB per octave each. We estimate the highest frequency of interest at 100 kHz. Therefore, the attack time required is 1/100 kHz, or 10 us. An attack time this short represents a possible, but challenging goal.

A more common approach is to allow very short transients to pass through and be clipped by the system. Generally speaking this isn't too objectionable. In fact, it probably sounds better than allowing the attack time to be as fast as 10 us, for reasons to be explained. A common compromise

seems to be 100 us to 1ms for the attack time. Slower than 1 ms and the clipped peaks become wide enough to be quite audible. To avoid second-order effects, the fast attack transfer function must have "gentle transitions" known commonly as "soft knees." All this means is that as the change is made from no limiting to full limiting, the transition be performed smoothly rather than abruptly. If this isn't done, the engaging of the limiter introduces an audible click. This artifact is one of the main reasons operators "dial in" a relatively slow attack time in limiters with this parameter adjustable. *The operator finds the occasional clipping preferable to the click*. Consequently, if the click can be minimized, a faster attack time can be used. As the attack time becomes faster, this click becomes more difficult to avoid. Discussion abounds on the ideal shaping for this transfer function. Most common is an RC (exponential) charge with some diode action to "soften" the start of the limiting.

The limiter must have enough control ratio to guarantee the signal can't clip once it has begun to decrease system gain. Normally, the ratio of a limiter ranges from around 20:1 to "infinity." With an infinite ratio, once the input signal exceeds the limiter's threshold, the output is not allowed to increase at all. A 20:1 limiter is easily attainable in a "feedforward limiter" while a "feedback limiter" is required for an infinite ratio. A *feedforward* limiter is one, which measures the input level, then calculates the gain applied by other functions to obtain the assumed output level. It sounds complicated, but it isn't that bad. A *feedback* limiter simply measures the output level then generates whatever control voltage is necessary to prevent that level from exceeding the threshold. Why use one over the other? In a system where other dynamic gain control is used, such as a compressor, there is already a voltage available that indicates the level in dB of the input signal. In a feedback limiter, the rectification and logging (if any) must be added at the output. Also, all limiters add distortion to the audio. Some of the reasons have already been discussed. Another source of distortion is the ripple remaining on the control voltage. This ripple is inversely proportional to the release time and the signal frequency. If a half-wave rectifier is used (common if applied at the output), the ripple is at the signal frequency and double the amplitude compared to that left by full-wave rectification. With full-wave rectification, the ripple is twice the signal frequency. Also, if the rectified voltage is converted to a log function, it leaves less ripple than a simple, rectified voltage. This ripple causes distortion in the VCA as it did in the compressor, only worse. A serious problem with the feedback limiter is that the distortion caused by the limiter is already present in the signal being measured. In other words, the limiter reprocesses the distortion along with the signal and adds increasingly complex distortion. Fortunately, this "second-pass" distortion is much reduced in level. The ripple voltage may be minimized by use of a multiple-order release filter. However, as the order of the filter increases, it becomes more difficult to obtain reasonably short attack times. A second-order filter is a workable compromise. All this means is that the first release filter is followed by a second nearly identical filter. This reduces the ripple by 6 dB (by one half) at a given signal frequency compared to the standard first-

order release filter at the expense of additional circuitry. Nearly all limiters use a single-order filter for this function.

Another parameter is the release time. If too short, the action of the limiter sounds choppy; if too long, the limiter "breaths". That is, after the offending overload, the limiter holds the system gain down too long. As the gain is brought back to normal, the increase in level is audible – not very pleasant. Also, as mentioned above, as the release time is shortened, the distortion increases. For a limiter operating with a release time of perhaps 500 ms, at 500 Hz this distortion easily exceeds several percent THD. It is quite audible. This may be improved by using a second-order release filter as mentioned, but ultimately the best fix is to use the limiter as little as possible.

A serious limiter *must* use full-wave rectification to measure the signal level. Let's assume in a half-wave system that the positive peaks are measured. It is entirely possible to have a major overload on the negative half of the signal only. In such a case, the overload wouldn't be detected and there wouldn't be any limiting. Full wave rectification adds so little to the cost that the choice is trivial. Unfortunately, the buyer has no easy means to determine the method used. If the ratio is mentioned and it's greater than about 20:1, it is *probably* a feedback limiter. A feedback limiter "tacked on" to another function (usually a compressor) is *not likely* to sound good if used heavily.

Normally the limiter is part of some other function, such as a compressor, as in the Rane DC 24. A limiter by itself is difficult to evaluate. It should be listened to with a compressor ahead of the limiter in the sound processing. With the compressor set up and functioning normally, have the clip light come on *occasionally*. Adjust the threshold of the limiter to just light on the loudest music peaks, and then only briefly. The clip light *should not come on except very rarely*. The limiter should be considered "disaster prevention." With it operating as described, its action should be quite unobtrusive. If it can be heard working, it probably is either not set up properly or is a bad limiter. It is the author's opinion that there are a lot more bad limiters than good ones. Fortunately, if used carefully, even a bad one won't be a total disaster. However, it *must be fast enough* to avoid clipping on all but the most difficult program material. If it can't do this properly, it's no better than a really bad compressor.

A DOWNWARD EXPANDER
The function of a good downward expander is to increase the *apparent* dynamic range of the system by decreasing the gain during the relatively quiet times thereby moving the apparent noise floor downward. It does this by comparing the signal level to a threshold. When the signal level drops below this threshold, the expander decreases the system gain by some ratio. This ratio is the same as defined in the compressor, only opposite in sense. The difference between this ratio and that of the compressor is that the gain of the expander decreases when the input signal is *below* the threshold. Expanders usually operate with a rather gentle ratio – perhaps 1.5:1 to 2.5:1. Occasionally higher ratios are used, but the audio tends to take on an odd breathing sound as the background noise comes and goes due to the expander's action.

To be effective, the dynamics of the expander must have moderately fast attack and fairly slow release times. The attack is the time spent restoring the system gain to normal when the signal level increases. Times of 0.5 ms to 1.5ms work well. The release is the time spent reducing system gain when the signal decreases. For the release, a time of 0.5 sec to 1.5 sec is common. The waveform of the control voltage must avoid sudden changes. Transitions must happen smoothly with a decidedly "soft knee." This is especially important with the downward expander since the audio will be, by definition, rather quiet at the threshold. Without audio dynamics to mask the gain adjustment, it must be done smoothly and gently or the action adds objectionable artifacts to the sound. To recap, the five determinates of the dynamics of a good downward expander are:

Threshold. The signal level below which the expander begins to reduce the system gain. This parameter is almost always adjustable and typically ranges from –50dBu to 0 dBu.

Ratio. The rate at which the gain reduces expressed as a ratio of output level change to input level change, such as 2:1. The normal range of adjustment is 1:1 (bypass) to 5:1 or so.

Release Time. The time it takes the expander to reduce system gain after the signal has dropped below the threshold level. If this parameter is adjustable, it is set to fit the dynamics of the particular audio source. Speech typically requires a fairly short release time while music sounds better with a longer release time. Adjustment range should be from 100 ms to 3 sec or more.

Attack Time. The time it takes the expander to restore the system gain after the signal level rises above threshold. A variety of methods exist to define this time, but it seems to be less of a problem than with a Gate, Limiter, or Compressor. The most descriptive is to declare the time it takes the expander to increase the gain to within some percentage of the final level after the signal is suddenly restored. This time may depend on how much the input level exceeds the threshold. Attack time is usually fixed. If it is adjustable, it should range from 0.5 ms to 10 ms. If it is fixed, an attack time from 1 to 5 ms would be appropriate.

Depth of gain reduction. This refers to the limit of the gain reduction caused by the expander. As a practical matter, it's usually limited to perhaps 60 dB of gain reduction. Sometimes the amount of gain reduction is adjustable to keep the system from sounding "dead." If the depth is fairly modest, this becomes a "ducking" function. If the system gain is decreased to a very low level, the constraints placed on the control dynamics are relaxed. For example, if the system is not passing audio due to heavy downward expansion, the transition at the beginning of the attack time is not so critical. However, the apparent attack time appears longer because of the time it takes the control to transit through the relatively inaudible region.

What to look for in an expander

Normally an expander is packaged as part of some other processing function, such as a compressor or a gate. A good listening test for an expander is to set the threshold *higher* than the threshold of the companion compressor. This would *not* be a normal setting, but the overall control function should still be smooth and without "odd anomalies." All this means is that it shouldn't "go wacky" even though the controls are set in a non-standard way. When operating by itself, it should be possible to set the controls so that the expander doesn't breath excessively nor sound choppy. An expander is particularly effective with spoken material. Subtle control should normally be the goal – avoid extremes. If the ratio is fixed, make sure it is gentle, perhaps 1.5:1 or so. Make sure there is a way to defeat the expander if the program material doesn't require it – a real bypass function is better than a very low threshold. If a mic is available for the listening test, set the threshold to substantially decrease the gain with no input to the mic. While listening with earphones, say something like "putt, putt, putt" into the mic. Make sure the expander "comes alive" quickly enough to include all of the "p" sound. If the release is adjustable, it should be possible to set it such that the gain decrease after the input ends is not overly noticeable, yet the softer "t" sound isn't chopped off. If used properly, the expander should do a good job of cleaning up a sound track during the quiet times without making it sound like the background noise is "doing pushups."

A GATE

An excellent gate is a very challenging balance of constraints. A gate is a device used to pass or not pass audio, based on the signal level. Refer to Figure 3. If the level exceeds an established threshold, the audio is passed (the gate is "open"). If the level is below the threshold, the audio is blocked (the gate is "closed"). It's never "partially open" except during transitional times. A deceptively simple definition. The easy approach is to use the equivalent of a simple switch – either on or off. Indeed, some are done that way. Unfortunately, they sound terrible.

By definition, sound is present when the switch is "thrown." The rapid change in level causes a serious "click" to be added to the audio. This happens both when the gate is opened and when closed. Various tricks are added to avoid these artifacts and that is where the design challenge happens. One method, which has some acceptance is to only open or close the gate when the actual signal waveform crosses zero volts (remember, it's alternating voltage). If the VCA is a reasonably good one, no output signal is produced at that instant. The problem is that most waveforms resulting from a percussion instrument have a very large initial transition followed by smaller "tremors." By the time the threshold is exceeded and the next zero crossing happens, the main event is over. The loud drum "rim shot" is reduced to a dull thud, for example.

Another approach is to use a downward expander with a very high ratio and a very fast attack. The problem with this method is that almost no downward expander is capable of the dramatic attack times required. The attack time is defined as the time required for the gate to open, or pass signal, after the signal level exceeds the established threshold. To be effective, this time must be less than 100 us. One gate measured, advertises 5 us, although in actual tests it didn't do anything before 100 us, and it took an additional 50 us to actually control the signal! Another gate has been measured at 20 us
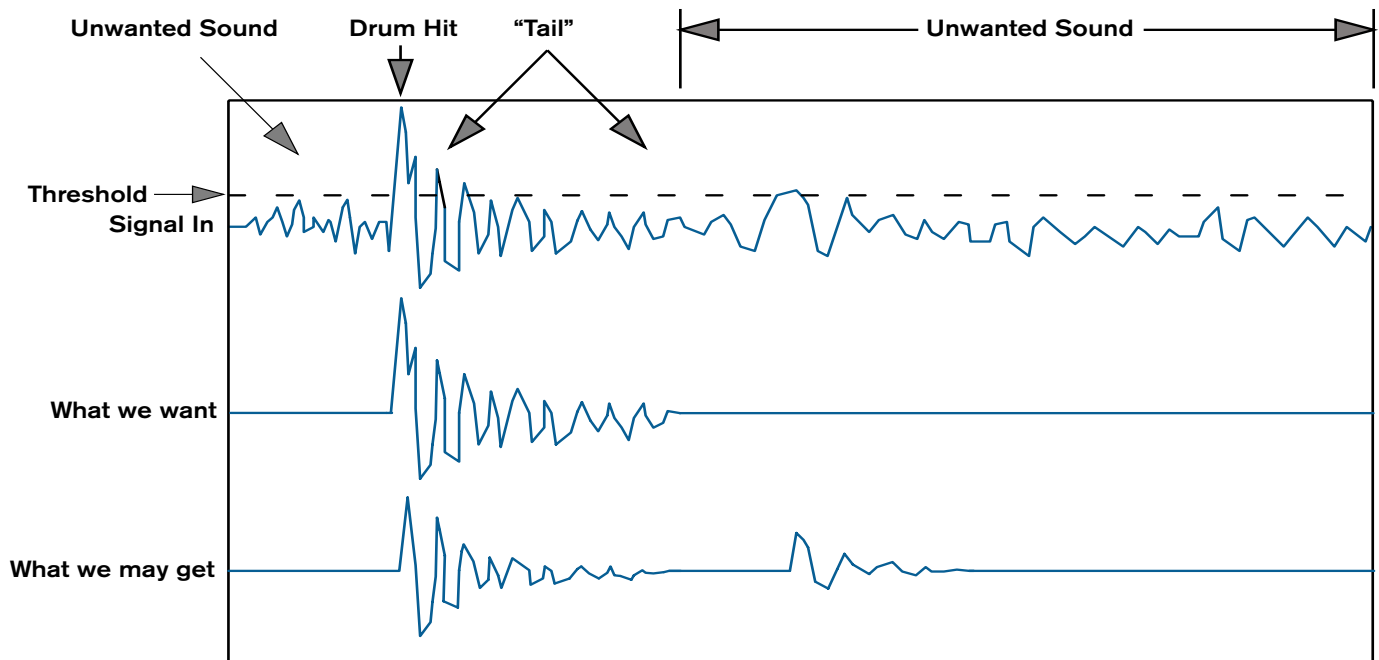
Figure 3. Typical Gate Dynamics

In this example, what we want isn't going to sound like what we may get! First, part of the drum hit is lost. Second, the end of the tail is attenuated too soon. Third, unwanted sound is picked up. In a really good gate, the drum will sound exactly the same as the original, only the unwanted sound will be gone. Incidentally, part of the solution to this problem may be to reposition the microphone to remove some of the unwanted sound. No gate can perform magic, though with the right settings it can seem like it.

(start of +20 dBu burst of 20 kHz to *beginning* change of control voltage). In addition, this gate uses an analog delay to "hold back" the audio path to give the gate time to open. It appears to open before the increase in signal level happens. These are rather expensive but effective tricks just to make a "switch" sound good.

The following parameters are required in a superior gate:

Threshold. The signal level above which the gate opens to pass audio. The operator almost always adjusts this parameter. A normal control range is from –50 dBu to +20 dBu. A threshold setting *below* the noise floor becomes a bypass –gate always open.

Attack Time. The time from when the signal exceeds the threshold to when the gate opens. Unfortunately, it is often defined as when the gate just begins to open in order to "improve the spec's." This is probably the first specification a prospective buyer should look at when evaluating a gate. When an attack time of less than 100 us is specified, be suspicious of how it was measured. This parameter is some-times adjustable, often by a 2-position switch (fast and faster). If adjusted by potentiometer, it needs to be set fast enough to include all of the percussive rise time of the audio. If it won't open that fast, go buy a better gate. The attack time should be adjustable, or fixed, to 500 us or *less*. A reasonable maximum is 1.5 ms.

Release Time. The time from when the signal drops below the threshold to when the gate closes. This parameter is usually adjustable. It is set to include the "tail" of the sound before the switch completely closes. For a kick drum, this time is fairly short, while for a triangle it may be as long as 4 seconds. The range of adjustment should be from 100 ms to 4 sec or more.

The following parameters are optional:

Hold Time. Companion to the release control. The time from when the signal drops below the threshold to when the gate just *begins* to close. If present, this parameter is normally adjustable by the operator and ranges from 10 ms to 4 seconds. This parameter is used to preserve the dynamics of the "tail" of a sound, such as a cymbal or bell. If used thoughtfully with the release time the over all quality of the transition is much improved.

Depth of Cut. The amount of gain reduction when the gate is "switched off." Since a VCA is usually used for the actual gain control element, a total switch function is not practical. Also, in some situations, total gain reduction is not wanted. If total silence isn't necessary while the gate is closed, the attack time is improved by limiting the gain reduction. Therefore, this parameter is sometimes adjustable. It should range from 90 dB reduction to 0 dB (effective bypass). A depth of 60 dB to 40 dB would be typical.
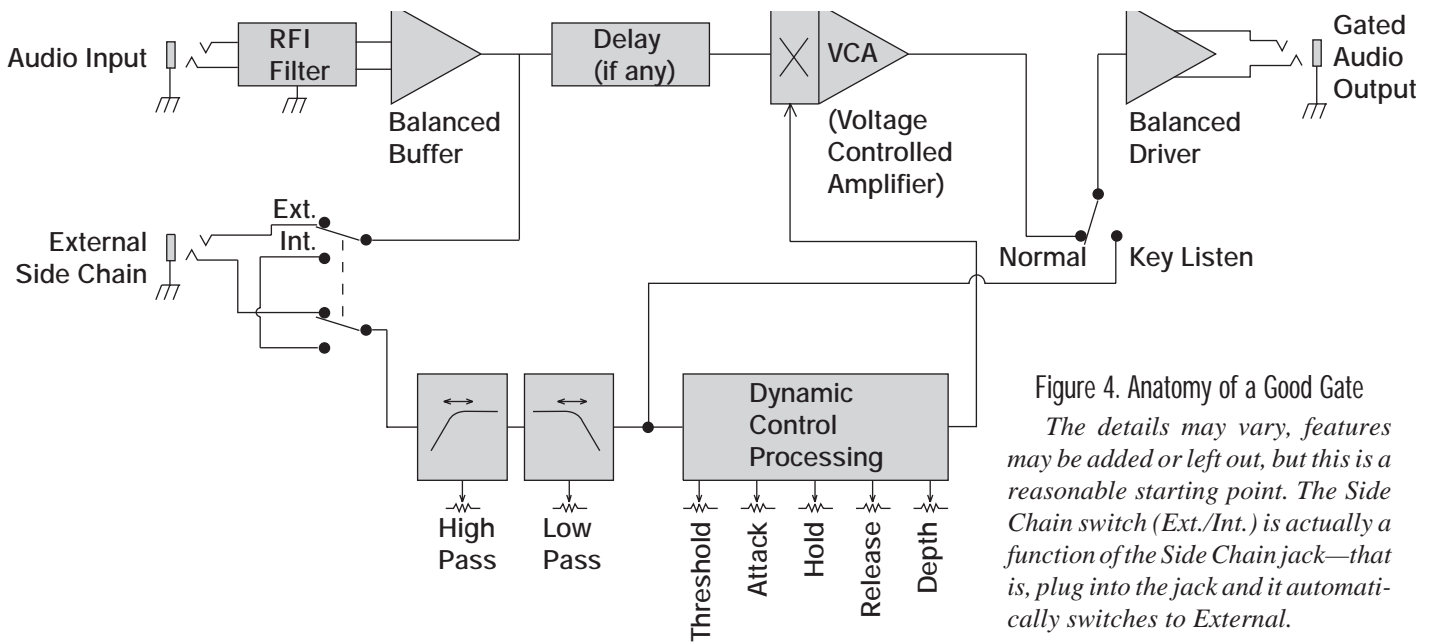
Figure 4. Anatomy of a Good Gate

The details may vary, features may be added or left out, but this is a reasonable starting point. The Side Chain switch (Ext./Int.) is actually a function of the Side Chain jack—that is, plug into the jack and it automatically switches to External.

High-Pass and Low-Pass Filters. These are applied in the control signal path (referred to as the "side chain") **not** the actual audio signal path. Usually there is a switch called a "key listen switch" that allows listening to the control signal to judge the affect of the filtration. Often there is a back panel jack called the "side chain jack" that allows this control signal path to be interrupted and modified externally then returned to control the gating. Any delay in this signal path degrades the attack time of the gate, so filters should be used sparingly since **all** filters introduce delay (we'll save the semantics of "lead" and "lag" for another discussion).

An analog delay in the audio path. If used, it is essential that the curve of the delay-versus-frequency be smooth and free of sudden changes. Ideally, it would seem that a constant delay would be best, however, a delay that varies **smoothly** with frequency works well. In any case, its function is to delay the audio path slightly to give the gate time to react before too much of the signal transition has occurred. This delay should range from 10 us to 100 us or more. If it is too long it will cause other problems. This delay is difficult to evaluate in testing. The only practical approach is to measure phase shift at a rather low frequency. Incidentally, this delay thoroughly "trashes" a square wave, or any other complex wave when viewed on an oscilloscope, due to phase dispersal. Fortunately, the signal sounds unchanged, but it certainly doesn't look unchanged.

Notice that there is no reference to "ratio." For those who insist on thinking in terms of a downward expander, the ratio of a gate is infinite. That is, above the threshold it's on and below the threshold it's off. Nothing in between other than transitional shaping.

Note that, because of the very fast attack time required, shaping of the attack function is very important. This alone will make or ruin an otherwise good gate. The real secret is a very fast turn on done carefully and smoothly – not a trivial task. Most, but not all, gates that are fast enough to be effective have a serious "click" at the turn on-transition.

HOW TO "KICK THE TIRES" OF A GATE

The best way is to play a drum track recording which has picked up low levels of some other instrument. The drum should be a snappy snare, preferably with lots of "rim shots." The sound to look for is a sharp percussive attack with some trailing sound. With the audio signal passing through the gate, adjust the threshold to trigger on the drum but not quite trigger on the background sound. Set the attack time for the minimum. There should be no "click" added even at minimum attack time. Listen very carefully to the drum sound while enabling and disabling the gate action. The sound of the drum *should not change* while the background sound between drum hits should be gone. Adjust the release and hold controls for the best result. This is the essential purpose of a gate. If it can't do it well, find a gate that will. Take heart – good gates do exist.

In summation, let it be said – buyer, be careful. Don't expect top performance if you're shopping for the lowest price. By the same token, a high price doesn't ensure the best sound. If you don't need "gimmicks", don't be talked into them. The beautifully done, extruded front panel may look really sharp, but it doesn't make a recording sound one bit better. If the cabinetry is tinny and cheap looking, the circuit design *may* also be minimal and of poor quality, but this isn't necessarily true. Review articles are a good starting point, but be careful of the reviewer's personal prejudices – they may not coincide with those of the buyer. Listen to the opinions of colleagues, again with a view to personal prejudice. When all is said and done – listen to the equipment. Make the test as representative of the buyer's actual situation as possible. Listen carefully, be critical. If something sounded funny, don't necessarily believe the salesman's claim that there's "something wrong with the recording." If necessary, the buyer should provide a recording for the listening. Make sure the equipment can be returned if it doesn't work out. After all, the ultimate test is how the potential hit recording sounds. Will it go platinum or just be another "not bad" recording?