

THE BIG SQUEEZE



By Michael Cooper

A comprehensive guide to compression and compressors.

Of all the processes used in modern music production, compression is perhaps the least understood. One reason is compression's sonic results are often subtle and thus hard to hear—especially for budding engineers. Another hurdle is presented by the various and differing compressor control parameters; those, too, are typically subtle in their individual sonic effects, and they work together interactively, further complicating the stew. Then there's the confusion that lies in the bewildering array of product types and models the engineer must choose from before even reaching for a control knob. For example, for a given application, should you select a VCA-based compressor or one controlled by an opto-electrical element? A solid-state or tube design (or a hybrid of the two)? Analog or digital compression? A hardware compressor or one that is software based? And so on.

With so many variables, it's no wonder compressors and compression remain a mystery for many users. Yet, if you want to master the arts of recording and mixing, learning compression's intricacies is imperative. After all, the production processes for most of today's popular music forms—with the notable exceptions of classical and some jazz—rely heavily on compression. Simply put, if you're not compressing properly, you're not getting the best sounds possible.

This article will guide you through the maze of compressor options and explain practical compression applications in plain English. I'll start with the basics of compression, citing examples of various production techniques and the theories behind them. I'll also tell you which features to look for in a compressor and why they're important. Finally, I will survey specific types and designs of compressors, describe some models, and offer opinions about which models do the best jobs on which instruments.

Improvement Plan

Compression falls under the broader category of dynamics processing. The term "dynamics" refers to changes in loudness level, so dynamic range is the difference between the softest and loudest sounds that a source produces, or that a track contains. A dynamics processor's purpose is simply to increase or decrease a signal's dynamic range, which alters how the levels fluctuate within

that range. Types of dynamics processors include gates, expanders, limiters, levelers, and compressors.

A compressor is a type of dynamics processor that “squeezes” a signal’s dynamic range—that is, it reduces the difference in volume, or level, between the loudest and softest parts of a performance. The process of reducing volume is called gain reduction. Properly applied, gain reduction makes a performance sound more consistent from beginning to end. For that reason, compression is a great remedy for a performance in which the levels fluctuate too widely.

By reducing dynamic range, a compressor also allows for the processed signal’s overall level to be raised—that is, become “hotter”—resulting in increased loudness without pushing the signal’s loudest parts into distortion. Bringing up the overall level has the additional benefit of making lower-level sounds louder than they were before compression. The result is that subtle nuances such as mouth sounds and ghosted notes—as well as burps, string buzzes, and snare rattles—are louder, clearer, and easier to hear.

Of course, you may not want to make burps, string buzzes, and other incidental performance sounds more audible. Therefore, apply compression only when musically appropriate—when the end result will sound better than what you started with.

You can always add compression after a track is recorded (during mixdown), but sometimes it is desirable to use compression during the recording process. That approach has several potential benefits. For one, a compressor makes it easier to capture usable tracks when recording an instrument with a wide dynamic range. Moreover, solving level-fluctuation problems during tracking frees you from having to solve them at mixdown. That, in turn, leaves more time and brain power—not to mention gear—for focusing on the mix’s creative aspects.

For those recording to any digital medium, using a compressor during tracking ensures that sounds are encoded at a higher level. Because more bits are used, better bit resolution results. Furthermore, by putting a lid on peaks, the compressor also helps avoid digital clipping on extraloud notes. For those recording to analog tape, compressing during tracking allows the signal level to be raised higher above the noise floor, which results in an improved signal-to-noise ratio.

Tricks of the Trade

In addition to problem solving—smoothing out rough performances, improving digital resolution and signal-to-noise ratio, avoiding digital clipping, and the like—you can also employ compressors in numerous creative applications. For example, a compressor can dramatically change the envelope of a sound in much the same way an envelope generator works in a synthesizer. That and other compression tricks can give a vicious attack to a lackluster snare drum, add crunchy edge and sustain to a mild-mannered electric guitar, make a lead vocal sound so urgent that listeners will dial 911, or pump up an entire mix until the band sounds like it’s exploding out of the speakers.



You can make the lead vocal sound so urgent that listeners dial 911.

In simplest terms, think of a compressor as an automatic volume controller. Indeed, before compressors were invented, engineers typically had to “ride gain” on a channel to maintain consistent volume levels. (Then again, many engineers still ride gain, even when using compressors.) However, a compressor controls levels with a speed and accuracy that is impossible to achieve manually—sort of like a magic genie adjusting the track’s fader with lightning-fast reflexes. The compressor’s control settings determine when and how much that fader moves.

Depending on how its controls are set, a compressor reduces either transient peaks—the short-lived, attack portions of a sound—or the average-level portions of the sound, and sometimes both. Examples of transient peaks include the stick strike on a drum head and guitar-string plucks. A sound’s average-level portions include a snare drum shell’s ringing and the sustain of a guitar note after it is plucked. Certain instruments—a wood block, for instance—produce mostly transients and very little sustain. Others, such as vocals and organs, typically produce mild transients that barely peak above their average levels.

The number of controls on compressors varies greatly, depending on design, cost, and other factors. Units that employ voltage-control amplifiers (VCAs), for example, typically have at least five controls: threshold, ratio, attack time, release time, and output level. Full-featured VCA models may offer more than twice that many controls, whereas some expensive opto-electrical compressors may provide only two control knobs.

Note that units with fewer controls are not necessarily less capable; rather, they typically provide automatic control of parameters such as attack and release time, or they “gang” two parameters (threshold and ratio, for example) on to one knob. I’ll discuss those types of compressors in more detail later. First, I’ll analyze the five controls common to most VCA-based compressors.

High Five

Threshold is the level at which compression kicks in and starts to reduce the signal’s level, or gain; the threshold control lets you set that level. With threshold at 0 dB, for example, all signals at or above 0 dB get compressed, while those that fall below 0 dB are unaffected. Therefore, to control peaks, set the threshold to a level below the level of the peaks but above the average level of the signal. That way, peaks that exceed the threshold get attenuated while the average levels pass unaffected through the unit.

Clearly, a proper threshold setting is critical to a compressor’s performance: if the threshold is set too high, the unit will not process any of the signal; if the threshold is set too low, the unit will react to—that is, attenuate—every portion of the signal.

Ratio expresses the difference between signal increases (volume) at the compressor’s input and increases at its output; the number on the left refers to input and the right to output. Therefore, the ratio control determines how much the signal will be




FIG. 1: A compressor set to a 2:1 ratio with a threshold of 0 dB produces an equal increase in output level respective to input level below the threshold, assuming that make-up gain is kept at unity. Above the threshold, output level rises only 1 dB for every 2 dB increase in input level. The compression curve shown is hard knee.

attenuated once it exceeds the threshold. For example, a 2:1 ratio will let a signal increase in level only 1 dB for every 2 dB it exceeds the threshold (see Fig. 1). Likewise, if the signal exceeds the threshold by 6 dB at a 2:1 ratio, the compressor attenuates the signal by 3 dB, a net gain increase of only 3 dB. In that case, the compressor's gain-reduction meter (if it has one) will show 3 dB of gain reduction.

Typically, different instruments and performances call for different compression ratios. For example, to compress a ballad's near-perfect vocal track, a mild 2:1 ratio would probably suffice; at that ratio, and with the appropriate threshold dialed in, the compressor tightens up the performance enough to ensure quiet phrases are not lost in the mix and higher levels are not overbearing. At the other extreme, a bass guitar track that alternates between mellow finger-pad technique and aggressive pop 'n' slap can easily have a huge dynamic range. To yield consistent levels from that type of performance, a higher ratio such as 10:1 may be in order.

Note that threshold and ratio work together to affect a signal's output level. The lower the ratio, the less control the compressor has on the signal; the lower the threshold, the lower the signal level subject to compression. The relationship between the two controls affords flexibility and sonic variation. There are, for example, two different-sounding ways to get the same amount of gain reduction out of a compressor—low threshold and low ratio or high threshold and high ratio.

Attack time is how long it takes—measured in milliseconds (ms) or microseconds (μ)—for the compressor to kick in once the signal exceeds the threshold. A slow attack time lets inherently fast transient signals pass threshold before compressing the rest of the signal; a fast attack catches transients, but may diminish high-frequency content.

One thing worth noting is that manufacturers sometimes measure attack times differently. Some specify attack time as the time it takes for the compressor to react after the threshold is exceeded, and others specify attack time as how long it takes for the compressor to reach, say, 67 or 90 percent of the maximum gain-reduction level it will ultimately achieve. Fortunately, the exact definition is of little importance, as typically attack time is set by ear. Depending on what kind of effect you're going for, simply decrease the attack time until unruly peaks are tamed or increase it until average levels are lowered and desirable peaks get through unscathed. If you're having trouble hearing your settings' effect, watching a downstream peak-level meter (that is, one that monitors the levels after the process—the compressor's output-level meter, for example) will let you visually confirm what portion of the sound is attenuated.

Release time is how long—measured in seconds or hundredths of a second—it takes for the compressor to return the signal to unity gain (its unprocessed state) after the signal falls back below threshold. That is, once the release time passes, the compressor lets the signal pass through unaffected. In general, slower release times result in a more natural sound.

In general, set fast attack and release times when you want the compressor to do its job and get out of the way quickly—for instance, when you want to put a lid on transient guitar plucks but allow the ringing notes to pass through unaffected. Conversely, a moderate attack time coupled with a long release is perfect for those David Gilmour-esque guitar solos in which you want notes to sustain forever. At two seconds or longer, the extended release time causes the compressor to slowly restore compressed levels to their original (higher) gain, just as the sustained notes start to naturally die off, which counteracts the

decay and makes the tails of the notes louder.

A compressor's last control stage is its output level. That control is also known as make-up gain because it is used to make up for the gain reduction caused by the compressor. The usual approach is to increase the processed signal's output level so it matches the unprocessed signal's level. That creates unity gain between the two signals, which makes it easier to compare them using the bypass switch and ensures appropriate levels when recording or mixing.

Path Not Taken

A compressor can degrade or ruin an audio signal as well as enhance it; therefore, one of the most useful features on any compressor is the bypass switch, which lets you compare processed and unprocessed signals. After using the output-level control to balance the levels of the processed and unprocessed signals—a critical step because louder signals sound brighter and fuller—you can judge whether your control settings actually improve the sound by switching the compressor in and out of the circuit.

Fortunately, most compressors provide a bypass (a notable exception is the Universal Audio Teletronix LA-2A). Typically, this is a switch that disables the compressor circuitry; ideally, it also disables the input- and output-level controls. A hardwire bypass is usually the best design because it routes the input directly to the output and bypasses all compressor circuitry, such as input and output amplifiers and gain-control devices.



Universal Audio reissued the legendary Teletronix LA-2A Leveling Amplifier—a favorite for processing vocals.

Actually, on some high-end units the signal is kept in circuit even in bypass mode—a viable design as long as the signal path is pristine enough not to color the original signal. The advantage of that design is it avoids the use of bypass relays and audio-path switches—elements that inevitably degrade during a period of time and compromise audio quality. The bypass on the Millennia TCL-2 Twincom Opto Compressor/Limiter (\$2,995), for example, defeats the unit's sidechain control only, thus preventing compression from taking place. (I'll discuss sidechains later.)

Don't Need Your Input

Many compressors also provide an input-level control, but those are often superfluous—if not undesirable. For one thing, a compressor with a wide-ranging threshold control can handle almost any input level. So an input-level control is necessary only if the threshold range is too high or too low to act on the input signal as is.

For example, if the threshold's highest setting is +2 dBV and you're feeding the compressor +12 dBV levels, you'll compress most or all of the time unless you somehow lower the level at the compressor's input. That's one instance in which an input-level control comes in handy. Conversely, if the threshold range's minimum setting doesn't go down very low, the compressor may not kick in when fed low levels. In that case, an input-level control is necessary to boost the input to an appropriate level.

The reason an input-level control can be thought of as undesirable is it adds yet another amplification/attenuation stage to the circuitry, thus degrading signal quality. For that reason, high-end, minimalist compressors, such as the Millennia TCL-2 Twincom, typically omit the input control to maintain a pristine

signal path.

Knee-jerk Reaction

In addition to the controls and parameters already discussed, several more-subtle parameters and design features often figure prominently into a compressor's performance or sound. One such parameter is the knee, which is related to the compressor's threshold control. The knee determines how quickly and smoothly the compressor will transition from no action to the full ratio of gain reduction set on the unit once the signal passes threshold. Generally, a compressor's knee is hard or soft, though some units—the Aphex 661 tube compressor/limiter (\$749), for example—provide switchable hard- and soft-knee compression.



The Aphex Systems 661 tube compressor/limiter features ultratransparent, VCA-based gain control and both manual and automatic operation.

In hard-knee compression, the unit processes the audio signal at the selected ratio once the input signal passes the threshold. Although useful for applications such as peak limiting and de-essing (discussed later), a hard knee can sound abrupt, especially with higher ratios.

A soft-knee compressor, or one set to soft-knee compression, begins to compress as the signal approaches the threshold level and gradually increases the ratio until the signal attains threshold, at which point it equals the selected ratio value. The gentler, logarithmic increase of soft-knee processing tends to sound more transparent (less noticeable) than hard-knee compression, and thus is usually preferable for most vocals and instruments.

In addition to manual controls for attack and release times, some compressors offer an automatic mode, called auto mode, that does some of the tweaking for you. That is often referred to as program-dependent or adaptive processing. In auto mode, the compressor's detector circuitry analyzes the program content (the audio-input signal) and dynamically adjusts the attack and release times accordingly. For example, if a guitarist starts picking harder, the unit automatically decreases and therefore quickens its attack time to catch the increased peaks. On the other hand, an increase in average levels typically prompts longer release times to avoid pumping while the compressor returns to unity gain.

Auto mode's main benefit is it precludes the need to tweak attack and release settings on performances in which the dynamics change radically. It also lets you set up quickly yet still get good results when the pressure is on. The downside is you lose some control over the sound. For example, you may like those peaks when the guitarist picks harder—in which case you probably would not want to use auto mode.

Some compressors—such as the MindPrint T-Comp Stereo Tube Compressor (\$1,099)—offer a semiautomatic mode of operation. As the name suggests, semiautomatic mode lets the attack and release settings exert some influence on the adaptive processing.

Opto-electrical compressors may or may not offer an auto mode; however, even without one, those units provide something akin to automatic processing in that attack and release times—manually set or not—fluctuate based on program content. That is due to the inherent nature of opto-electrical compressors, which in general are slower and less exacting than VCA-based designs. Because the attack and release controls on optical compressors provide only approximate

response times, many manufacturers simply put “fast” and “slow” on either side of the knob, rather than hash marks indicating exact times. (More on optical compressors in a bit.)

Double Duty

Most dual-channel compressors offer stereo linking, a feature that lets you run two channels—for example, stereo acoustic guitar or even an entire mix—through the compressor and have each channel be attenuated the same amount. That keeps one side’s level from dipping more than the other, which would throw the stereo image out of whack.

True stereo linking works by having the channel that exhibits the most gain reduction determine the gain reduction for the other channel. Another form of linking establishes a master/slave relationship between the two channels in which one side (typically the left) is the predetermined master and the other follows its attenuation pattern.

It is commonly said that compression becomes limiting at ratios of 10:1 and higher, but that is not the entire story. Actually, the detector circuits in compressors and true limiters differ by design. A compressor’s detector circuit is usually designed to detect RMS, or average, levels rather than transient peaks. Therefore, transient peaks almost always overshoot a compressor’s threshold level, no matter how high the ratio and how fast the attack time is set. A true peak limiter, on the other hand, employs a detector circuit that responds to peak energy levels and thus reacts faster.

Whereas all true compressors use RMS-sensing detector circuits, detectors for different models can differ substantially in their reaction times. That means two different compressors set to the same attack, release, threshold, and ratio values may nevertheless respond quite differently to the same signal. (That is one of the many reasons it is difficult to recommend specific control settings for compressing various instruments.)

Chain, Chain, Chain

Every compressor has a sidechain detector circuit that “sees” when the threshold has been exceeded and tells the compressor’s gain-control element or amplifier to attenuate the signal. The sidechain is not in the audio path; it’s merely a traffic cop that tells the compressor when to attenuate the signal. The circuits for threshold, ratio, attack, and release are also found in the sidechain.

Full-featured compressors typically provide sidechain inserts on their rear panels. Think of a sidechain insert as an effects loop that patches into a compressor directly before the detector; like the rest of the sidechain, it is not in the audio path, so its effect isn’t directly heard. Sidechain inserts therefore let you process the compressor’s input signal before it reaches the detector. That permits de-essing and other frequency-conscious applications. Here’s an example of how to perform de-essing.

To de-ess a vocal, first patch the send and receive from the compressor’s insert into an equalizer’s input and output, respectively. Next, boost the equalizer’s high frequencies and cut its lows and mids. That causes the compressor’s detector to hear the vocal as having excessive highs. Whenever the whistling sound of sibilance raises its ugly head, the sensitized detector circuit hears it much louder than it really is, causing the circuit to vigorously reduce gain in the audio path. With attack time set to around 50 μ s and release time between 50 and 60 ms, the compressor can be made to quickly attenuate the sibilance and get out so the rest of the vocal is left unchanged. Of course, the compressor’s

threshold must also be set properly—above the vocal’s average levels—for that to work.

You can also use a sidechain insert to make the detector react to a signal entirely unrelated to the audio-input signal. The classic example here is “ducking”: a sidechain application in which an announcer’s voice is set to trigger a music bed’s attenuation. To set up this type of ducker, play stereo music tracks through a dual-channel compressor and patch the voice-over track (or channel) into the sidechain insert’s receive jack. Next, set the compressor threshold low enough that it responds to every vocal utterance. When the announcer speaks, the detector hears the voice and instructs the compressor to lower the music bed. You can also use that technique to automatically lower, say, guitar levels whenever a lead vocal comes back in. To accomplish this, patch a mult of the vocal into the insert receive jack of the guitar’s compressor channel.

Freq Show

The misconception that splitband compression is the same as frequency-conscious compression is common. A splitband compressor splits the audio signal into two or more frequency bands so each band can be processed by its own independent compressor circuitry (each with its own controls). That lets you compress, for example, a guitar’s bass frequencies differently from the highs.

A compressor that offers—or is set up to provide—frequency-conscious compression is still a full-band device acting on the entire signal. The difference between it and normal compression is simply that the detector is set to be called into action by the prevalence of specific, user-selected frequencies. Frequency-conscious compression has dozens of useful applications, but space limitations dictate I save that vast subject for a future article.

Down and Dirty

Now I’ll discuss various types of compressors on the market and which designs and models are best for different recording and mixing applications. I’ll start with analog compressors, which can be subdivided into four categories based on the type of gain-control element they use: opto-electrical, Variable-Mu, FET, or VCA. Each design has benefits and drawbacks.

To increase the usefulness of this article, I tested a sampling of compressors chosen to represent the various design and feature sets available. Bear in mind that this sampling is not meant to be comprehensive, nor is it intended to overlook or denigrate units not mentioned here.

An opto-electrical compressor (opto, for short), uses a special opto-electrical cell in its sidechain that consists of either an electro-luminescent panel—basically a night-light—or an LED that shines on a light-dependent resistor (LDR). In simple terms, the light panel, or LED, shines with increasing intensity on the LDR as the audio-input signal gets louder, and the LDR causes a corresponding increase in compression of the audio-input signal. Because the LDR has an inherent memory effect, it releases slower when the light is brighter or has been shining for a while. In practical terms, that means heavy or near-continuous compression results in longer release times.

Vintage-style optos generally have only two control knobs, typically labeled gain reduction and gain. Turning up the gain-reduction knob feeds more signal to the opto cell, effectively lowering the threshold and causing more compression. The gain knob sets post-compression output level, or make-up gain.

As mentioned earlier, opto elements have a natural lag time in their attack response; indeed, all the time constants are inherently adaptive. Optos also have, by nature, a soft knee. For those reasons, opto compressors tend to have a natural-sounding attack and release. The downside is they are usually not quick enough to catch fast transients, so substantial overshoot is not uncommon. Typical applications for opto compressors include vocals, bass, and electric guitar. However, one thing I've learned is not to try to pigeonhole equipment based on general design characteristics.

A case in point is the Joemeek C2 stereo opto compressor (\$399). A solid-state unit, the C2 sounds great on snare drum, serving up an outstanding power-pop snare tone with a dark yet explosive attack (see the sidebar "Dialing in Hot Sounds"). The ability to get a defined but warm edge also makes the C2 a good choice for creating crunchy electric-guitar sounds. The C2 also does a surprisingly good job of smoothing out levels on inconsistent kick drum tracks. Whereas some optos respond to such transient material with inconsistent attack and release responses, the C2 is rock steady. You get a bit of transient overshoot, but overall the control is excellent. Overall, the Joemeek C2 is—at least for certain applications—the best compressor I've heard in its price class.



The Joemeek C2 stereo-optical compressor offers outstanding performance at a buyer-friendly price.

Another dual-channel opto compressor that does a good job of reining in kick drum tracks is the Bellari RP583 (\$650). This hybrid unit employs a tube circuit for the internal processing and solid-state op amps for the input and output circuits.

An even better-sounding unit on kick drum is Joemeek's SC2.2 stereo optical compressor (\$799). The SC2.2 offers outstanding level control yet doesn't thin out the drum sound very much—a universal problem when heavily compressing that instrument. It is one of the best compressors I've heard on kick drum.

In general, opto compressors are not my first choice for compressing arpeggiated acoustic guitar tracks. That's because optos tend to pump in this application; you can clearly hear the level dipping and then recovering. Pumping occurs when the bass frequencies in a plucked note trigger heavy compression, which in turn attenuates highly noticeable mid and high frequencies. That said, the SC2.2 delivers a surprisingly transparent sound on acoustic guitars.

The C2 and SC2.2 also sound smooth, natural, and beautifully warm on vocal tracks. But the most venerable of vocal compressors is the single-channel Universal Audio Teletronix LA-2A Leveling Amplifier (\$3,495), a faithful reproduction of the highly coveted, '60s-era LA-2A tube compressor. The LA-2A has an uncanny ability to warm up a piercing or thin vocal, and it smooths the most unruly vocal dynamics with a transparency hard to match in other compressors. The LA-2A also sounds great on bass guitar, kick drum, and snare drum.

Another unit that sounds absolutely gorgeous on vocals is the aforementioned Millennia TCL-2 Twincom, a dual-channel unit with switchable tube and solid-state audio paths for each channel. If you're looking for a compressor with stunning clarity, nuance, and depth, look no further than the TCL-2.

The Anthony DeMaria Labs ADL 1500 (\$2,995) is another great opto/tube compressor for recording and processing vocals. It's a dual-channel,

vintage-style (two-knob) opto with a fat sound and a soft top end.

Although out of production, the Groove Tubes CL1S dual-channel opto/tube compressor is worth mentioning in case you can find one used (the original price was \$2,295). The CL1S is one of the most lush-sounding compressors ever made; it imparts rich yet present textures to everything it touches. The unit's only downside is it is quite finicky; it pumps readily if not set up with painstaking care. But for dialing in huge, velvety vocals, thundering bass, or slammin' electric guitar, it's the compressor to get—if you can find one. The CL1S also sounds great on kick and snare.

Continent of Mu

The first compressors ever made incorporated a Variable-Mu design. Those use a vacuum tube—such as a dual triode or pentode—for the gain-control element. Technically oriented readers will be interested to know that this is a fully differential, push-pull design in which one side of the tube handles the positive waveform phase and the other side the negative phase.



**The
Manley
can really
fatten up a
mix.**

Variable-Mu compressors do not offer an adjustable ratio control. What people love about the Variable-Mu compressors is that they continuously increase their ratio the harder they're pushed (the higher above threshold that the input signal rises), resulting in an increasing desensitification of sound. Though Vari-Mu compressors offer faster attack and release times than optos, they are not as fast as VCA designs, and therefore they're not as effective at handling peaks as VCA-based units. Also, as a class, Vari-Mu compressors cannot produce as much gain reduction as other types of compressors because the employed tube typically runs out of dynamic range sooner than other types of gain-control elements. A Vari-Mu usually gets 12 to 15 dB of gain reduction and sometimes considerably more.

Few Variable-Mu models are currently on the market. I am aware of high-end units made by boutique manufacturers such as Manley Labs and Pendulum Audio. Because of design differences, Pendulum Audio's Variable-Mu compressors typically offer faster attack times than Manley's.



**The Manley Stereo
Variable-Mu
Limiter/Compressor features
an all-tube audio path that is
optimized for pro (+4 dBu
nominal) levels.**

The Manley Stereo Variable-Mu

Limiter/Compressor (\$4,000) sounds awesome on bass and electric guitar; it imparts a fat, lush tone with plenty of presence and clarity. When set to limit mode, the Manley also sounds great on snare drum. However, the unit is intended primarily as a stereo-bus compressor for processing an entire mix. In this application, I had to keep the attack time near its slowest setting to avoid pumping. But properly dialed in, the Manley Stereo Variable-Mu can really fatten up a mix.

FET in the Cap

FET-based compressors are scarcely available now. They use an FET (Field Effect Transistor) as the gain-control element, which has advantages and disadvantages. On the plus side, FETs offer attack and release times that are much faster than optos can provide, and even faster than many VCAs. The downside is FETs have a somewhat limited dynamic range. In traditional FET designs, very hot levels can cause amplitude-modulation artifacts (for example,

pumping) and waveform distortion.

Crane Song's single-channel Trakker (\$2,550) and dual-channel STC-8 (\$4,450) compressors use a unique Pulse Width Modulator (PWM) FET design to solve the distortion problem. Unlike traditional FET designs, a PWM FET does not modulate the gain reduction, so there is less distortion at high gain. In technical terms, the PWM is essentially a switch that turns the audio signal on and off at a 1.1 MHz frequency, letting only a percentage of its energy through at any given time to control the gain. A filter rids the signal of switching artifacts. Because the signal is either on or off, it doesn't modulate the gain-control element's resistance or resulting gain reduction. The result is that the distortion content is linear as gain reduction increases. In addition, attack time can be as fast as a few microseconds. Unfortunately, I have not worked with any Crane Song compressors in my studio.



The single-channel Universal Audio 1176LN Limiting Amplifier (\$2,295) is a faithful reproduction of the late '60s-era, FET-based compressor prized for its lightning-fast attack and crystalline sound.

Universal Audio's reissue of the vintage (late '60s era), solid-state 1176LN Limiting Amplifier (\$2,295) is one of the best compressors I've heard on kick drum, electric guitar, and especially snare drum. Many of the explosive snare sounds heard on past and present hit records were processed with an 1176LN. The single-channel unit also dishes up crystalline vocals and a burpy bass-guitar tone. Whereas the LA-2A shines on thin, piercing vocals, the 1176LN sounds particularly flattering on woolly vocals by improving clarity and intelligibility as it reins in levels.

Gimme a V, Gimme a C . . .

VCA-based compressors can be lightning fast and therefore generally offer better peaks control than opto and Variable-Mu units. They can also attain absurdly high gain-reduction levels—useful when you really need to smash down a signal. The downside of VCA designs is the tendency of lesser-quality units to dull high frequencies at high gain-reduction levels. In addition, some engineers don't like the coloration VCAs can impart to the audio path. However, VCA-based compressors' quality ranges widely, and some units sound far better than others.

I frequently use a pair of stereo-strapped Aphex Expressors (the solid-state Model 651s, which have been discontinued) for compressing an entire mix because, simply put, they are more transparent than any other full-band (as opposed to splitband) analog compressors I've heard. By transparent I mean that, at reasonable settings, the Expressors exhibit virtually no discernible amplitude-modulation artifacts (pumping and breathing, for example).

The acid test for any compressor is how well it can handle broadband, percussive material—that is, material with a lot of sharp transients and extreme bass and high frequencies mixed together. A good example of broadband, percussive material is a full-ensemble stereo mix complete with drum set, bass, and arpeggiated acoustic guitar. Most full-band compressors on the market do a mediocre if not terrible job handling that kind of material. When transients or heavy bass content cause abrupt compression, full-band compressors tend to pump, sucking down easily audible mid and high frequencies.

A pair of stereo-linked Expressors can really beef up a mix without pumping.

The Expressor is also extremely stable on acoustic guitar, and it is the best compressor I've heard on kick drum. A serviceable workhorse on bass, vocals, and electric guitar, the Expressor's only downside is it can be a tad noisy and thin sounding. Because compressing a full mix usually makes bass content perceptually louder, the latter consideration is not much of an issue—I've heard tube compressors that offer a much fatter sound but are not nearly as stable. (If you're interested in the Expressor note that the original, solid-state model 651 was replaced by the Aphex 661 tube compressor/limiter, which adds a tube in the audio path and an auto mode. The 661 features the same ultratransparent VCA 1001 gain-control element, but the audio path sounds a bit veiled compared to that offered in the original solid-state Expressor.)

Another really great VCA-based compressor is the Empirical Labs Distressor (\$1,499 for the single-channel version, \$2,899 for the dual-channel). The Distressor is a digitally controlled analog compressor featuring a custom Class A VCA and standard Class A/B op amps in the signal path. The digital circuitry actually switches the unit between four different and independent solid-state compressors—all in one box—making for one of the most versatile compressors on the market.



The extremely versatile Empirical Labs Distressor features digital-control circuitry that switches the unit between four VCA-based analog compressors inside the same box.

If you buy a Distressor, order the new British mode option, which costs an additional \$100. Simply put, British mode “kills.” With the right settings, you can make a Distressor in British mode sound a lot like a vintage 1176LN; it serves up unbelievably savage power-pop snare sounds, crunchy guitars with beautifully long sustains, and in-your-face vocals with crystalline highs. (See the sidebar “Dialing in Hot Sounds” for sample control settings.) The Distressor also delivers fat, burpy electric bass-guitar tracks (with British mode turned off) that sound quite similar to what a great tube compressor would produce.

The MindPrint T-Comp Stereo Tube Compressor (\$1,099) is a tube- and solid-state hybrid, soft-knee compressor that performs transparently on stereo mixes and acoustic guitar. However, it's quite noisy and lacks the headroom to handle really hot pro levels. Just the same, the T-Comp is a good workhorse compressor for vocals, bass, and snare-drum tracks.

At the bargain end of the VCA price scale is the PreSonus Bluemax Smart Compressor/Limiter (\$199). This fixed-stereo unit provides numerous compression presets, and it also offers a manual setting with which you can dial in attack and release times to taste. Despite its low cost, the Bluemax is one of the best compressors I've heard on kick drum. It also provides excellent control of acoustic guitar tracks.

Rubik's Tube

Many engineers lust after the warm, round tones that great tube gear can deliver. Tubes—especially when driven hard—generally produce more even-order harmonics than solid-state devices and they also tend to saturate in a more gradual and pleasing way. On the other hand, the most pristine solid-state designs tend to offer a more focused sound and slightly better transient response (detail). Both topologies have their place.

Quite a few hybrid devices are available for purchase. Hybrid designs employ both tube and solid-state devices in their audio paths. The dual-topology Millennia TCL-2 Twincom lets you switch between completely independent

all-tube and all-solid-state audio paths in the same box.

It's helpful to know a compressor can be marketed as a tube processor and yet have a solid-state device—for example, a VCA or opto cell—in its audio path. Some purists insist a compressor is not “all tube” unless the gain control element is also tube-based (as in Variable-Mu designs), but that is not a practical concern. The gain-control element affects the compression curve's envelope, which in turn can affect the sound's timbre. But a well-designed tube output amplification stage can certainly give the lush timbral coloration and depth most folks are looking for in a tube device. For example, the Groove Tubes CL1S combines an opto element with a tube output stage, and that box has incredible warmth and depth.

On the other hand, just because a compressor has a tube in it doesn't mean it's going to give a fat sound. Just by listening to the Bellari RP583, for example, I would never know it was a tube compressor. Also, I've heard other tube units that sound more like distortion pedals than pieces of studio equipment. Generally speaking, you get what you pay for—and the best-sounding tube gear tends to cost a lot.

Interestingly, there are also quite a few solid-state compressors that offer wonderfully warm tones (the Empirical Labs Distressor and Joemeek units immediately come to mind). My advice is to judge each compressor, tube or solid state, on its sonic merits and try not to get caught up in the hype.

Ones and Zeros

One advantage of digital compressors is most of them offer “look ahead” circuitry. (Interestingly, dbx also offers an analog compressor with that feature.) Because the compression algorithm is in software, the compressor can analyze what it is about to process and place the attack time right at the onset of—or even before—the sound, resulting in a zero attack time. However, while a super-quick (or zero) attack time is great for catching transients, it doesn't always sound the best. Therefore, use such power judiciously; the crack of a snare drum without any attack just doesn't sound right.

In addition, digital compressors usually offer incremental control of every parameter imaginable, as well as the ability to store settings for later recall. Perhaps the biggest benefit of working with digital compressors is the ability to stay in the digital domain. If you're working with a digital audio workstation or digital mixer, there are strong arguments for not re-entering the analog circuits. Most importantly, by staying in the digital domain, you avoid the signal degradation and distortion caused by multiple conversions.

If you're considering buying a hardware digital compressor, make sure it has great-sounding A/D and D/A converters. It's also helpful if the software is upgradeable through user-installable EPROM, CD-ROM, or some other user-friendly method. In addition, you should insist on a box with a word clock input. Without word clock inputs, you will be limited to using only one digital compressor at a time.

The main reasons to buy a software digital compressor are easy upgrade capability and less hassles with dither and jitter—key considerations in yet another vast and controversial subject that I don't have space here to discuss. Suffice it to say that working within a self-contained digital environment (such as a DAW)—as opposed to patching together peripheral digital devices—can simplify things a great deal.

The biggest concern with all digital compressors, of course—both software and hardware based—is audio quality. Quite frankly, until just a few years ago, most of the digital compressors on the market sounded terrible. They would just suck the air out of any audio they processed and then spit out dull, lifeless tracks with no top end or depth. That, however, is changing rapidly. Although there are still relatively few digital compressors to rave about, some truly singular ones stand out.



**The jack-of-all-trades
Arboretum Systems Ionizer
is a cost-effective plug-in for
digital audio workstations
that offers killer splitband
compression.**

Because I usually work with digital compressors only when mastering, I can speak authoritatively in that context only. The compressors in the Alesis Masterlink ML-9600 Master Disk Recorder (\$1,699) sound outstanding. In DAW, the Waves Renaissance Compressor (sold as part of the Renaissance Collection bundle—\$300 for the Native version and \$600 for TDM) and Arboretum Systems Ionizer (\$499 for the MAS version) are extremely transparent compressors. Ionizer offers great splitband compression and other dynamics processing, as well as EQ and noise reduction.

Fresh Squeezed

Clearly, it's important to choose the right compressor for the job at hand. One thing I learned from years of working with various compressors is it's not so much the design but the execution of the design that makes a compressor good or bad for a specific application. Be wary of any generalizations about compressors. We've all heard, for example, that opto-electrical compressors provide transparent and natural-sounding compression—as if that were a given. But the fact is, some optos do and some don't. As always in audio, it's the sound that counts, not the hype.

Hopefully, this article not only opened your eyes to the wide, wonderful world of compression but also has helped you steer an easier course through the labyrinth of compressor types, features, applications, and models. Deft and artful compression takes time and experience to master, of course. Naturally, as with any other discipline, practice makes perfect. So go forth and squeeze!

Michael Cooper is the owner of Michael Cooper Recording, located outside the beautiful resort town of Sisters at the base of the Oregon Cascades.

NEXT

Reprinted with permission from  Magazine, February, 2001
© 2001, Intertec Publishing, A Primedia Company All Rights Reserved

[an error occurred while processing this directive]

Interactions



[an error occurred while processing this directive]

404

[an error
occurred while
processing this
directive]

We recently redesigned our website, replacing many of our pages, updating information and making it more convenient for you to find what you are looking for. It seems that the page you are looking for is one that has been moved or replaced.

Please visit [AEC Newsroom](#) or use the links on this page to find your way around our website.