29 VT-7 Compressor 10-7-20

I'm Doug Fearn and this is My Take On Music Recording.

Back in 2003, I was having dinner in a Nashville restaurant with my good friend Dave Hill, founder of Cranesong. Dave and I have been friends since 1993, around when we both started our professional audio products manufacturing companies.

It might seem strange that competitors would be good friends, but, really, there are only a handful of us in this high-end market, designing and manufacturing the tools used by the people producing the music we all enjoy. We might be competing in the same general market, but each of us has taken a slightly different tiny niche in a very small marketplace, based on our individual sense of what we can do to improve the quality of music recording.

So we have a lot more in common than you might think.

And really, who do we have to talk to about what we do? There are only a very small number of people who have the challenges we have, and most of us are dedicated to advancing the state of the music recording art.

My conversations with Dave usually start out with how business is. And then we move on to complaining about suppliers discontinuing parts we need, and the scramble to find a suitable replacement. We talk about how difficult it is to find vendors, like paint shops, that consistently provide us with the custom items we need, at consistent high quality.

And then we move on to the most interesting part – what new projects we are working on. Over the years, Dave has been very helpful to me with suggestions when I am struggling with a design, and I believe I have helped him, too.

I do not know anyone in our business who has a greater engineering insight into the technology we are using. Dave seems to get about ten times as much accomplished in a given time than I do. I am always impressed with what he is learning and doing.

So, on this evening we got to talking about new ideas and new products. I tell Dave that I just came across a new part, an FET switch array, that sounds like it would be perfect for a compressor. I talk about how it could be driven with a pulse-width modulator to provide extremely low distortion and transparent compression, something I have been trying to achieve for years. I must have a dozen prototype compressors in storage. Compressors that worked well, but never quite met my goal. And they didn't move the state of the art forward.

As I am talking about this, I can see that Dave is getting more animated. He is smiling, and I can see that he needs to say something. So, I stop talking and Dave begins to tell me he had the same reaction to the new part. And in fact, he already had a breadboard version of a compressor functioning. He was very excited about the potential of this design for a compressor.

I told him I had just received some samples of the part and I was currently researching pulse-width modulator designs to see what would be the best approach.

And then I said to Dave, if you have it working, perhaps I don't need to spend time on that part of the compressor design. Would you be interested in licensing your design to me?

Dave said sure. And we spent the next hour literally drawing schematic diagrams on napkins.

Dave sent me a prototype board with his design for the pulse-width modulator circuitry to drive the FET switch. To test it out, I used the amplifiers I had already designed for the VT-5 equalizer, with some modification for the specific needs of the compressor.

When I first listened to the prototype, I was initially disappointed. It didn't seem to be doing much, even though I knew it was functioning properly and providing a significant amount of compression. Then I switched the compression out and realized that it was doing almost exactly what I wanted. The compression seemed to disappear, leaving only the effect on the music that I wanted. This looked like it could be the sound I was looking for.

Of course, it didn't emerge from the prototype ready for market. There were still many months of development necessary. But that core sound is the basis for what became the VT-7 Compressor, a product we released in 2005 and it continues to be our most successful product in terms of the revenue it brings in. Dave's STC-8 uses the same core pulse-width modulator circuit, by beyond that commonality, the VT-7 and STC-8 are very different. That reflects the different direction that Dave and I took in bringing a product to market that meets our individual goals.

Now, you're probably wondering what all this about FET switches and pulse-width modulators means. And I'll get to that, with as little technical jargon as I can manage, a little later.

But first, I want to take a look back at the history of compression and limiting, to see how this indispensable recording tool came about.

The early 1930s were an amazing period in audio development. RCA came out with the first practical ribbon mic, and Neumann made the condenser mic a reality. Bell Labs figured out everything we needed for digital audio, even though the technology did not exist to fully implement it. And Bell Labs developed practical designs for high-quality amplifier and equalizers, among many other groundbreaking accomplishments.

Radio broadcasting had been around for about a decade and had become a universal force for communicating with the public and a primary source of entertainment. Radio back then was all AM broadcasting, using the same frequencies and technology we use today. Widespread FM radio did not come along until after World War II. And it wasn't until the 1970s that FM began to compete with AM.

AM radio has the potential to be a high-fidelity medium, but today our world is polluted with too much man-made noise, mostly from digital devices, that AM is no longer competitive. In fact, many new cars do not even include an AM radio, since the toxic noise environment in the modern car makes AM reception impossible. But back in the 1930s, AM was king, and it sounded great. And AM stations could cover vast areas of the country at night.

But AM, or amplitude modulation, has one characteristic that made it a bit tricky to use. AM works by varying the strength of the pure radio wave. That's called the "carrier," and it's what you hear when there is no audio on the signal. The resting power ranges from 500 to 50,000 watts, depending on the station license. The audio signal acts like a power control for the radio station, pushing to power up and down with the sound. The effect is similar to the waveform you see on your computer screen in a digital format.

The problem arises when the audio pushes the signal down to zero. It's not possible to have negative power, so the result is undefined and the radio receiver does not know what to do with it. Sounds like digital overload, doesn't it? Different principle, but the same sort of effect. What you get is extreme distortion and noise.

And the transmitter doesn't like it, either. The transmitters of those days would usually fault and shut down if the modulation exceeded the maximum.

Early in my career, I worked for a radio station in Philadelphia that had both AM and FM outlets. But only the AM was profitable back in the 1960s. The main transmitter was from the late 1950s and it was built by RCA. The previous transmitter, from 1937 and made by Westinghouse, was used as the backup.

On occasion, the main transmitter would be down for routine maintenance and the backup transmitter would be used. On a couple of occasions, I was working at the AM transmitter site when the backup transmitter was in use. As sort of crude protection against damage should the modulation exceed 100 percent, there was a spark gap across the huge audio transformer that provided the modulating signal. Every once in a while, may once every few shifts, the spark gap would fire. This transformer was huge, so it was in a different room, along with the power supply. But even through a closed door, sounded like someone fired a gun, it was that loud.

Obviously, it was best not to knock the transmitter off the air with too much audio, so engineers had to keep the peak audio level well below 100%. That made the radio station difficult to hear, compared to a higher level of modulation. Like most things in life, it required a tradeoff. The balance was between loudness and reliability. And nothing strikes more panic at a broadcasting station than being off the air. And not being as loud as the competition was almost as bad.

Everything back then was made with vacuum tubes, of course. And one characteristic of a vacuum tube is its ability to act as a valve. A small signal can control a larger signal. That's the definition of an amplifier. That's why the British term for a vacuum tube is a "valve."

But vacuum tubes are not very linear. No amplifying device is, actually. But circuit designers realized that by placing a small negative voltage on the grid of the tube, the amplifier can be made much more linear, which means it has less distortion. And changing this negative voltage, called "bias," has an effect on the amount of gain the amplifier has.

There is actually only a small range of conditions that are optimum for the performance of a tube amplifier, and those parameters are what designers seek to optimize. But the gain can vary over a small range with a change in the bias.

The scientists and engineers who designed tubes were often brilliantly innovative, which is why there were thousands of different tube types during the golden age of vacuum tubes. And some clever designer determined that there was a way to make the grid, the control element in a tube, so that it could vary the gain of the tube over a wide range without introducing excessive distortion.

These innovative tubes were called “variable-mu,” mu being the Greek letter that was used to indicate the gain of a tube. The grid performed its usual function of a small audio voltage controlling a larger voltage, but it also could vary the overall gain as the bias voltage was changed.

By taking a sample of the audio signal and amplifying it in a separate amplifier section called the sidechain, and then converting that varying audio signal into a DC voltage, and then applying that voltage to the grid of a variable-mu tube, the input audio could increase substantially while the output audio remained relatively constant.

This is somewhat simplified, but the principle is easy to understand. The audio gets louder, but the “limiter” prevents it from getting too loud. If you set the maximum output to be close to the level that modulates the transmitter to 100%, you can safely run your program material at a much higher average level without fear of knocking the transmitter off the air. It was a brilliant solution and immediately improved the coverage area, and reliability, of radio stations.

There were other applications for this new technology, and one of them was in the design of the radio sets that people listened to at home.

Some people may have had a favorite radio station that they listened to exclusively, but many people had multiple stations that they liked to listen to. But not all the radio stations had equally strong signals. Some were a lot more powerful than others, and the distance to the station could be anywhere from less than a mile to thousands of miles.

That could result in a 20dB difference just in the strength of local stations. And if you wanted to listen to distant stations, the difference could exceed 60dB.

The answer was to adjust the gain of the receiver to match the level of the incoming signal, using a separate gain control, in addition to the volume control. That’s fine for technologically-adept people, but it was a lot to ask of the average listener.

Of course, the radio could not reproduce that 60dB range. The result would be that when you tuned from a weak signal to a strong one, the audio would suddenly increase dramatically, scaring all the people and pets in the house. And it wasn’t just way too loud, it was pure distortion. Ugly.

But by using the variable-mu tube in the early stages of the receiver, where the signal was still at radio frequencies and the audio had not yet been extracted, you could adjust the gain to match the strength of the signal.

Note that this does not change the dynamics of the audio. This is strictly to compensate for the differences in signal strength as you tuned across the AM dial.

One innovation that was used in radio receivers was to apply the gain control bias to more than one amplifier stage. This technique has been tried for audio compressors, but it doesn’t seem to work too well in that application.

This concept, called “automatic volume control,” is still used today, although some receivers implement their automatic volume control in the digital domain.

The limiter also found application in disc cutting. The 78RPM records of that era were intrinsically very noisy, so it was a huge advantage to keep the overall level up as much as possible. But you can only go so high in level before the cutting stylus moves side-to-side so far that it impinges on the groove before or after it. At its worst, this will cause the playback needle to jump right out of the groove. A limiter was a good solution to this problem, and it continues to this day in disc mastering.

The vacuum tube compressor-limiter was the only way to achieve dynamic control for a couple of decades until some other approaches were tried. Engineers at Teletronix came up with a novel way of level control using a photoresistor. This device changes its electrical resistance as the amount of light shining on it varies. This is the technology used for decades in the sound component of film projection, until a magnetic stripe replaced it. The result of this approach was the Teletronix LA-2.

This so-called optical limiter relied on an electroluminescent light panel, like you might find in a nightlight. That light source, fed from the sidechain of the compressor, varies in intensity with the changing audio level. The electroluminescent panel reacts much more quickly to a changing voltage than an incandescent light source does, which is necessary for quick response to peak levels.

The photoresistor is also critical to the operation. There are hundreds of different photo resistors available, all with different characteristics. What was needed was one that varied its resistance in a manner similar to how we perceive sound. It is that perfect combination of light source and photoresistor that makes the LA-2 sound so good.

You can also use a field-effect transistor as a quasi-substitute for a variable-mu tube. It doesn’t sound the same, but it works well. The 1176 is a good example of an FET-based compressor.

Another approach that came about in the solid-state era of the 1960s was the Voltage Controlled Amplifier, or VCA. This worked in a manner similar to the variable-mu tube, but instead of a fixed degree of level reduction, the VCA could be made to do a lot of different things – even acting as an expander. An expander is the opposite of a compressor. It increases the difference between the loudest and softest levels. The VCA could also be used as a remote volume control, which permitted automation and track grouping in a console.

Personally, I was never very impressed with the sound of the VCA compressor. Or fader, for that matter. But the technology was widely used in compressors and consoles for many years.

There are a few other, more obscure, ways of reducing the dynamic range of an audio signal, but I want to discuss the approach that is at the heart of the D.W. Fearn VT-7 Compressor, and also Dave Hill’s Cranesong STC-8: the pulse-width modulator.

By the way, over the years I have build compressors using all these techniques, just to understand what could be done with them. I learned a lot building all the circuits and experimenting with them, but the PWM and the variable-mu tube are still my favorites, and probably the only things I would use on my sessions.

So what is this pulse-width modulator anyway? Well, it’s not a new technology. I can no longer find the reference, but I believe the concept goes back to the 1930s.

There are probably dozens of pulse-width modulators in your home, varying the speed of motors in your HVAC, washing machine, etc. And a related concept is used in solid-state light dimmers.

The principle is simple: an AC waveform, which could be audio, or in this example, the 50 or 60Hz power used in your home, is a sinewave, or reasonably close to a sinewave. If you turn on a light switch, for example, the lighting device immediately sees the full 120 or 240 volts and lights to full brilliance.

But what if you rapidly turned the switch up and down, as fast as you could. The flashing light would be very annoying, but the total amount of light would be cut in half if the switch was on and off for equal amounts of time. Now imagine that you could turn that switch off and on much more rapidly. For a light source, it only has to be about 50 times a second before you can no longer detect the rapid off and on. The light just appears to be dimmer.

That 50Hz switching rate would work OK with a light, but it would be a terrible thing to do to audio.

But what if you could switch the audio off and on at a rate far above our hearing range? It would simply reduce the volume in half.

How fast does it have to switch? Well, basically, the faster the better. The rate has to be high enough that the segments cut out of the audio are much shorter than the duration of the audio wave. Otherwise, the audio will have some degree of distortion, since the output is no longer an exact replica of the input. In the VT-7 the switching rate is around 1.5Mhz. That’s one and a half million times a second – fast enough that the tiny slices taken out of each cycle of the audio are so short that the integrity of the waveform is maintained.

Just turning the audio off and on simply turns the volume down. But what if we leave the audio on most of the time, and just turn it off for a very short period of the 1.5MHz rate? The audio level would only drop a tiny bit.

Conversely, if it were off 90% of the time and on 10% of the time, the level would drop a whole lot more.

But what do we use for the switch? It can’t be a mechanical switch, which would only last a second or two at that switching rate, even if you could make it go that fast. Mechanical switching is possible, and was used for things like motor speed control many years ago.

Today we have solid-state switches that can operate at hundreds or thousands of MHz. This is the basis for all our digital equipment. Specialized transistors are the heart of all computers. And they only have to know if the circuit is on or off – a one or a zero in digital language.

The field-effect transistor is really good at this and forms the basis for all digital electronics, and even some analog circuits for audio. The FET makes an excellent, high-speed switch.

For audio, an FET switched at a high rate can instantly chop the audio into tiny off and on segments. But only works with FETs that are specifically designed for this precision chopping.

The next step is to vary the on-off time so we have continuous control of the audio level.

A pulse-width modulator works the same way to control the speed of the blower in your HVAC system. The system knows just how much motor speed you need for the desired amount of air movement, and varies the width of the pulse to provide continuous control of the motor speed.

In an application like that, it is likely that the high-speed switching will produce all sorts of extraneous electrical noise, which can get into any sensitive electronic device. Keeping that switching noise from ruining your radio reception, cell phone reception, or wifi can be a challenge. And it could be a problem in a high-quality audio device, too, if proper precautions are not taken to isolate the switching from the audio circuits.

In a PWM compressor or limiter, we use the same technique as the variable-mu vacuum tube. A sample of the audio is amplified separately from the main audio path and converted into a DC voltage. This control voltage is used to adjust the pulse width, instead of the bias on the tube.

It sounds simple, and the concept is pretty straightforward. The challenge is to make it sound good. The pulse has to have great precision and stability or it will produce ugly-sounding artifacts in your audio.

The key to all this is the FET switch, and that was the part that both Dave Hill and I came across that was perfect for a compressor.

I use the terms compression and limiting more-or-less interchangeably in my explanation because they are really the same phenomenon. The first level-control devices were called limiters because they set a hard ceiling on how loud the audio could get. As the signal got louder, nothing much happened until a certain threshold, at which point the limiting action kicked in and stopped the level from increasing. That was just what was needed to protect an AM transmitter… or these days, to prevent the audio from exceeding digital maximum level.

Compression, on the other hand, implies a simple reduction in the dynamic range. The quietest parts are much closer in level to the loudest parts. The dynamic range is compressed.

There is a lot of gray area in between these two extremes, and that makes the categories of limiting and compression difficult to specify.

I don't think there is any definition of limiting and compression that everyone agrees upon, and most people use the terms interchangeably. I tend to think of limiting as a ceiling on the level and compression as a reduction in the dynamic range. And that's how I use these terms in this discussion.

There are many more factors that affect the sound of a compressor, and I'll describe the most important ones.

The **ratio** defines how much the output will increase for a given increase in input level. For example, in a compressor with a 2 to 1 ratio, if the input increases by 2dB, the output will only go up by 1dB.

That's a pretty gentle, but noticeable, decrease in the dynamic range.

At the other extreme, a ratio of 20 to 1 will act more like the limiter I described. The input can increase by 20dB but the output only changes by 1dB. That's extreme, but it works well for protecting a sensitive device like a digital converter or a transmitter. You will hear this amount of compression for sure. And it makes a useful effect that can be used creatively in some recordings.

Another important set of parameters are the **attack** and **release** times. The attack time determines how long it takes for the compression to start after the level increases. A short attack time catches short peaks immediately, which is useful for protection. A longer attack time lets some of the audio get through before the compression starts.

Why would you want to wait for the compression to start? Well, a slower attack time lets the initial attack, or transient, get through without affect, preserving the impact of a percussive musical note or a drum hit.

Attack time is a creative decision for the engineer or producer to make.

Release time determines how long the compressor takes to "recover" to normal gain after the level decreases. As a general rule, a short release time will make the music, or voice, sound louder and more compelling, but with an obvious effect on the sound. That may be what you want.

A longer release time means the average level is more constant, but the dynamic range is not audibly decreased as much. The average energy level in the music is higher with a short release, and lower with a long release.

If the release time is too long, especially with a lot of compression, a sound that is louder than average can cause the gain to remain reduced for significant amount of time, maybe seconds. This is very obvious and annoying, although it can be used creatively on occasion, as an effect.

The optimum release time really depends on the type of music or other audio.

Some compressors and limiters have adjustable attack and release times, while other devices may have fixed times. Certainly a compressor with minimal adjustments makes using it simple. But with the added controls, much more creative uses may be found.

For device protection, a zero attack time would be best. But in the real world of analog circuits, nothing is instantaneous, so there will always be some slight lag before the limiting occurs. This could be measured in microseconds, millionths of a second, but it is still not zero and a digital converter will still be affected by that short initial peak.

Back in the 1970s, I use to ponder how to make an "anticipatory limiter," one that would know what was coming before the audio got there. The level could be reduced just in time to catch every peak.

I experimented with this using a tape machine running at 30 inches per second, which provided a delay of about 40 milliseconds or so. I fed the track I wanted to compress through the tape machine and into a limiter. But the control signal for the limiter came directly from the un-delayed audio.

It sounded awful. But it was an interesting experiment.

Later, when digital delays came along, the delay time could be set down to about 1 millisecond. I tried my anticipatory limiter with a digital delay, which sound a lot better, but it still was not a very useful sound.

And this was not a practical technique anyway, since the limited audio track would always be slightly delayed compared to the rest of the tracks. But I learned a lot from those experiments.

Today, we have digital limiters that can look ahead at the recorded digital data and apply limiting at exactly the right point to catch every peak. This is often called a "look-ahead" feature. I was skeptical of this, based on my experiments in the 1970s, but modern implementations of this anticipatory limiter actually work very well, and solve the problem of digital overload.

Another factor to consider is that a very short attack and/or release time may have a time delay comparable to spacing between cycles of a bass note. When this happens, the compressor attacks and releases on each cycle of the bass note. That does not sound good. It generates a lot of distortion.

If you were to graph the input level vs the output level, non-compressed signal would be a perfectly straight line at a 45-degree angle. This means that for every dB that the input goes up, the output also goes up by 1dB. But once the signal level reaches the threshold of compression, the line would immediately take a somewhat flattened out angle, indicating the output level is not going up as fast as the input level.

But in most limiters, the transition from no limiting to limiting does not work that ideally. The shape of the transition is rounded off, and this turns out to be a good thing for many times of compression applications. The graph of this sort of looks like a bent knee, so the term "**knee**" is often applied to this characteristic. Note that this term was not invented for audio. The knee characteristic goes back long before compressors, in many areas of physics.

The shape of that knee affects the sound of the compressor or limiter. Some compressors have a control to adjust the knee.

Ratio, attack time, release time, and knee are the main characteristics that define the sound of a compressor. There are others, more-subtle, but not significant for our discussion.

The designer is faced with a huge number of decisions on these parameters. And if they are controllable by the user, he must decide what range of adjustment is best.

This is one reason why most compressors sound different from other, similar compressors.

What I wanted for the VT-7 was compression that was transparent and free from distortion. Compressors using variable-mu vacuum tubes introduce a degree of harmonic distortion, which generally increases with the amount of gain reduction. This may not be entirely a bad thing, and it gives those compressors part of their characteristic sound. I like that sound, too, on certain instruments and voices. But that is a different compressor design, and not what I was looking for in the VT-7.

I wanted to make it work like the perfect engineer, who could turn a fader up and down at practically instantaneous speed, tastefully controlling the level.

I also wanted a compressor that, when used on a mix, could help make all the parts fit together more naturally and smoothly. Some engineers refer to this as “glue” for the mix, which I think is a pretty good description of the effect.

The pulse-width modulator, coupled with a precision FET switch, fulfilled the requirement for transparency and low distortion. In fact, the VT-7 does not exhibit any change in the amount of distortion when going from zero compression to the maximum available.

In its original form, the VT-7 prototype was capable of about 60dB of compression. That is well beyond what is needed in any audio application, so I throttled it back to about 20dB of maximum compression. That’s still a lot more than you would typically ever need, but it did make the Threshold control a lot easier to adjust.

The attack and release times are constrained to the range that is useful for just about any audio application. That means they can be set fast enough to introduce distortion on low bass frequencies if you’re not careful. This goes against another design principle I have, which is to make it nearly impossible to make any D.W. Fearn product sound bad. But in this case, I thought the ability to use very fast attack and release times could be useful in some situations. I just had to depend on the user to understand what trouble they could get into if they were not aware of this.

Let’s go through the controls on the VT-7 and see what they do.

First is the Threshold control, which essentially adjusts the amount of compression. With the control fully counter-clockwise, there is no compression at all. At full clockwise, the compression could be 20dB or more, depending on many other factors.

When we use compression, the overall level going through the compressor is reduced by the amount of average gain reduction. To restore the level to normal, the Gain control is used. More accurately, this could be labelled make-up gain, but that label didn’t fit very well on the panel, so it’s just labelled Gain.

The amount of make-up gain should match the maximum amount of compression available, and that is how the VT-7 is designed. You will rarely use the full capability of the gain reduction, so the Gain control will usually only be set near its lower limit.

By the way, turning the Gain control all the way down does not turn the audio off. It is set so that it is unity gain, so with no compression, the output level is the same as the input.

The Attack and Release controls set the times for those parameters, with clockwise rotation making the times longer. I never measured the exact time ranges for these controls. Like most of my designs, it was determined by ear.

Another control that is initially a bit confusing is the Harder/Softer control. This adjusts several parameters simultaneously, and it’s difficult to describe the effect with a single word on the panel.

The VT-7 actually has two side-chains. Remember, the sidechain is the separate section that converts a sample of the audio into a DC control voltage for the pulse-width modulator. The sidechain also has components and adjustments that determine the time constants. That's mostly the attack and release time ranges.

The two sidechains have different compression ratios, attack and release time ranges, and different shapes of the compression knee. The two sidechains also have different ranges of the Threshold, too.

The front panel control can be thought of as a pan-pot between the two sidechains. Rather than a hard switch to select one or the other, you can mix the two in various combinations.

On some material, the effect is very subtle and almost anywhere in the Harder-Softer range will sound similar. On other material, the difference will be much more obvious.

There are no rules for how to use this control. You have to use your ears and determine where in the range the VT-7 sounds best to you, for the material you are processing.

I find that I use the control around its middle position when the VT-7 is on the mix buss. For vocals, and some acoustic instruments, more towards the softer end sounds best. That's how I set it for this podcast recording.

Music that is more hard-hitting might benefit from the Harder end of the range. For the type of recording I do, which is mostly acoustic music, I rarely venture into the Harder end of the range. But it's available to you if you find it fits your style better.

The **Separate-Link-HPF** switch determines whether the two channels of the VT-7 are independent or coupled together for stereo. Generally, a stereo mix should have the compression linked, which keeps the stereo image stable, even when there is an imbalance in the left and right content. That imbalance could happen if there is a solo part that is panned towards the left or right. Linking the channels keeps the two compressors locked together, so no matter what happens in the stereo mix, the image will remain the same.

Some people like the sound of un-linked left and right channels, and that can work well, especially with mixes where most of the content is panned to the center. But, generally, the linked position is best for a stereo mix. Or for mastering.

The third position on this switch is labelled HPF, which stands for high-pass filter. Or you could think of it as a low-cut filter, which actually makes more intuitive sense than the electrical engineering term" high-pass filter."

This filter does not affect the audio going through the VT-7, so you will not hear any difference switching the HPF in or out if there is no compression. What it does is to reduce the sensitivity of the gain reduction to bass frequencies.

Why is that helpful? Well with many types of music, where the bass and bass drum are dominant in the mix, the overall amount of compression can be “captured” by the low frequencies. This will reduce the bass content, and may even cause the kick drum to control the overall volume. That could be a desired effect sometimes, but usually it is not what you want. So by placing the HPF in the sidechain, the low bass mostly passes through the VT-7 with little or no compression. This helps to maintain the proper balance between the instruments.

The HPF has a very gentle curve, starting around 200Hz on the high end, and becoming more aggressive at the lowest frequencies. This seems to work the best for all material.

Note that the HPF position has the two channels linked.

When the VT-7 is in the Link or HPF positions, only the controls on the left, or A, channel are operative. The right-hand controls do nothing. However, the Gain control is still active on the right channel, to adjust the amount of make-up gain.

The VT-7 has a VU meter for each channel, and that meter can be switched to read either the output level, which is labelled VU, or the amount of gain reduction, labelled GR. I find that I leave the meter in the GR position almost all the time, using the VU position only to check how the output level looks. That can be helpful in setting the Gain controls.

You will notice that there is no bypass switch on the VT-7. This is by design, since adding a bypass would also require an adjustment to the gain controls in order to maintain a constant level with the compression switched in or out. That’s not practical, unless there was a hard bypass that essentially just connects the input and output connectors on each channel. I have a VT-7 that I modified with that hard bypass. But I did not find it useful and I never used it.

If you look at all the classic compressors and limiters that engineers and producers love, like the LA2A, the Fairchild 660, or the UREI 1176, none of those had a bypass switch.

The VT-7 can be used on individual tracks, or on an entire mix, or on a sub-mix of, say, the drum tracks. It seems to have found a home on the mix bus at many studios, and it indeed works very well there. But I think people miss an opportunity for using the VT-7 on individual tracks, where it also shines. If you need a transparent compressor on a sound, the VT-7 may be just what you need.

The VT-7 is also an indispensable tool in mastering facilities. Many mastering engineers have told me they would not start a mastering project without their VT-7. Mixers tell me that, too.

In either mixing or mastering, most people use the VT-7 to provide that special processing that changes a typical recording into a record. You know the sound. It adds the final polish and cohesiveness and makes the resulting song competitive. For many engineers, the VT-7 is used in conjunction with the VT-5. In fact, we frequently sell a VT-5 and VT-7 together, since they make such a wonderful pair.

It does not take much compression in mixing or mastering. In fact, many mixing and mastering engineers tell me that the gain-reduction meter is barely moving, and that's all it takes to achieve the sound they want. A lot depends on if there is compression on individual tracks. I tend to only use compression on vocals and bass, with the VT-7 providing the bulk of the compression on the mix bus.

The VU meter that indicates the amount of gain reduction does not respond instantly to every peak that the VT-7 may see, so even if the needle is barely moving, there may be more compression going on than you think.

Settings for the controls are, of course, an artistic decision, and what you like may be very different from what I like. But as a general rule, I have the attack set around mid-point or slightly to the right of that. That's a fairly slow attack.

I tend to set the Release in the faster half of the dial, around 9 or 10 o'clock.

The Harder-Softer control is usually in the middle, or slightly to the right of straight up.

And, of course, the Threshold is set for the desired amount of compression. On a mix, that is usually about 1 to 3dB on the Gain Reduction meter. But sometimes it could be as much as 5dB, or as little as just a flicker on the highest peaks. It all depends on what my ears tell me.

For this podcast recording, I set the VT-7 with a bit more compression, usually 3 to 5dB. The Attack is set around 1 o'clock and the release around straight up. The Harder-Softer is at around 2 o'clock. I use the link position, although it doesn't matter because the end product is mono. I do not use the HPF.

For the podcast, I aim for a loudness of about -18dB loudness units full scale. Any more than that and I find the compression too obvious and annoying.

I also use a digital limiter, a Flux Elixer plug-in, after the VT-5 and VT-7, to catch the occasional peak that would otherwise set the maximum overall level. This brings up the loudness level about 2dB, since I no longer have to worry about digital overload. By the way, I set the output of the Flux Elixer to -2dB, and the podcasts stay at that level until they are encoded into mp3 by the podcaster provider. I find that much above -2 does not translate well into mp3.

I am always amazed at what the VT-7 does, and that reaction seems to be universal among users. Careful design of the input and output amplifiers, along with Dave Hill's masterful pulse-width modulator design, fit my notion of what a compressor should do.

And the VT-7s built after September 2020 do not have a fan. A return to the original amplifier design, first used in the VT-5, reduced the heat to about half of what it was, so we could eliminate the fan.

That's the VT-7 story. If you have questions, comments, or suggestions, I always appreciate hearing from you. Send me email at dwfearn@dwfearn.com

And if you have friends who would find this podcast interesting, please pass along the link to them.

Don’t forget that you can use a podcast app, like Apple Podcasts and many others, to automatically notify you when a new episode is available, or to automatically download each episode.

This is My Take On Music Recording. I’m Doug Fearn. See you next time.