Vacuum Tubes #20

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

My first studio used a collection of tube mic preamps that I had acquired over a few years of searching. Many were made by RCA, and most of those came from WPEN, where I worked while I was in high school and for a few years beyond. That radio station was entirely vacuum tube from microphone to antenna.

Most of the consoles at WPEN dated from the 1940s, and were really 1930s designs. They used large, octal tubes and all the amplifiers were rack-mounted. There were several types of preamps, and they were used as gain building blocks in the station. Many were configured as mic preamps but others were phono preamps, summing amplifiers, or level booster amps for the incoming broadcast lines.

Most were half-rack width, with a power supply, and plugged into a tray. Others were narrow and plugged into a different kind of rack tray. Those were the 500-series modules of the day. The power supply was a separate rack-mounted unit. Those trays held five or six preamp modules.

Some of the well-worn consoles were replaced with newer ones while I was there. I really wanted to buy the old consoles, but the station chief engineer wanted to keep them for spare parts, since the station had several consoles of that model still in use.

Same for the preamps and other amplifiers in the racks.

One day I did an inventory of all the preamps and told the chief engineer that we had three spare preamps for every preamp in use. Could he sell a couple of them to me? He was reluctant, but I told him if the station ever needed them, I would bring them back in.

I became the owner of about six of these RCA preamps. Some were the really old ones and a couple of them were the newer ones with their own power supply. For the little modules, I had to build my own power supply.

I also found some preamps that were made by GE, and I even used the electronics in Ampex 350-series tape machines as mic preamps.

Those preamps became the heart of my studio's console. Well, you could hardly call it a console by today's standards. It was really just a passive mixer, using a collection of Daven faders I found. A pair of the preamps were the summing amplifiers. There were no pan pots – for positioning, I simply patched a track into two mixer inputs and adjusted the stereo position with the level pots.

This was crude, but the studio was only 4-track and that did not take much of a console.

Still, I was looking forward to the day when I would upgrade to 8-track, and that would make my passive mixer obsolete. I needed a real console.

Back in the late 1960s, there were not many options for a console. Most studios used broadcast consoles, or built their own. Sigma Sound Studios in Philadelphia had an Electrodyne console with 12 inputs. It was a thing of beauty, with pan pots, echo sends, and some simple eq on each input. And it had a separate monitor section, which was really another console in itself.

But even the simplest Electrodyne console cost about $12,000, which would be about $90,000 in today's dollars. Not in my price range at that moment.

I found a company that was interested in adding a recording console to their product line and we signed a contract for a 12 input, 8 output console, with mic and line inputs, pan pots, eq on every input, effects sends, a separate monitor section, and even a compressor on several of the input modules.

The console took a few months to build. After wiring it up and checking it, I was excited to try it on a session. I asked several musician friends to come in to help me shake down the new console.

We had drums, bass, acoustic guitar, electric guitar, and lead vocal – a typical early 1970s session.

We did a take and sat back to listen.

I was shocked. It sounded nothing like what I was used to. I had a sense of that while we were recording, but actually listening to the sound was appalling. If the others were as disappointed in the sound as I was, they did not say so.

The next day, I ran some tests to measure the performance of the new console. Indeed, it met all the specifications. In fact, the measurements were a lot better than my previous mixer – quieter and with lower distortion. But it certainly did not sound better.

I talked to the manufacturer about it and was told that I just wasn't accustomed to "solid state perfection," after listening to tubes for so long.

I was skeptical of that. I knew what real music sounded like, and this was not it.

But I resigned myself to adjusting to this new "perfection." With all the inputs, eq, pan pots, etc. it was hard to resist the convenience and flexibility. It did give me a new creative tool to use in producing records, and it made the studio commercially viable.

Even after getting that console, and several others in the years that followed, I still often recorded individual microphones through one of my old tube preamps patched directly into the tape machine.

Since that day, nearly 50 years ago, I am still not convinced that solid state sounds all that good. The designs have gotten a lot better, and you cannot deny that most of our favorite records from the 1970s onward are recorded with solid-state equipment.

But there is something magical about the way that vacuum tubes handle music. This is something I have tried to understand for decades, and I think I know some of the explanation – but perhaps not all.

It’s not that I have not tried to make transistors sound good. Way back around 1958 I saved up my allowance and bought a Raytheon CK722 transistor for $5. That would be about $26 in today’s dollars – a hefty purchase for a ten-year old. I used that CK722 in dozens of different experiments: as a radio receiver, an audio oscillator, an electronic switch, and as an audio amplifier. It was an amazing device, about the size of USB-C connector and colored a beautiful blue. I used it for years in various circuits until eventually one of the wire leads broke off.

But all the audio amplifiers I build were vacuum tube. Tubes were cheaper and I was much more comfortable with the circuitry.

Much later, I built solid-state mixers, and even an entire console. But I could never get any of those things to sound good to my ear.

Why is that?

Well, to explain the differences that I think contribute to the sound, we first have to understand a bit about how amplifiers work. This is not going to be overly technical. I’ll make it as simple as possible. Those of you that understand the theory will not find this rigorous, but I am trying to explain it for the non-technical listener.

What is an amplifier anyway? It’s simply a circuit that takes a small signal and makes it bigger, like a magnifying glass. The output of a microphone, for example, is very low in voltage – thousandths of volt. We can’t use that to drive a loudspeaker or headphones. We need to boost the level, and boost it quite a lot.

All amplifiers used in audio work on the principle of a small voltage controlling a much larger voltage.

In the UK, they call vacuum tubes “valves,” and that term is quite a bit more descriptive of the function. A valve uses a knob or lever to adjust the amount of water flowing. An amplifier does a similar thing to audio.

The “water” in electronics is electrons, tiny negatively charged sub-atomic particles. Electrons make our modern society possible, and they exist in everything.

To produce the electrons, we heat up a piece of metal by using a source of electricity to get a filament to glow – the same way an incandescent lightbulb works. But just heating the filament only produces a cloud of electrons with nowhere to go. To make them work for us, we have to create a path for them to actually do something.

In almost all vacuum tubes, the filament, the part you can see glowing dull red or orange, heats another metallic element surrounding it, called the cathode. It is the cathode that emits the electrons.

To get them to move, we need another metal piece inside the tube called the “anode,” which is more often called the “plate.” We put a relatively high positive voltage on the plate and the electrons scramble across the space and into the plate. Congratulations! We now have a stream of electrons flowing through our vacuum tube. That is not particularly useful, but this simple configuration, called a diode, can be used to convert alternating current into direct current – a rectifier – or as a way to detect radio signals.

But we want an amplifier, and the simple diode will not do that.

But if we put another metallic element in the path between the cathode and plate, we can use that as a valve to control the number of electrons that get through.

This is called the "grid," which is descriptive because it is more like a fence or gridwork. If it were solid, it would block most of the electrons, and that wouldn't be useful.

It's still not an amplifier but we're getting close. If we put a small voltage on the grid, it will influence how many electrons get through to the plate, and that suggests something beneficial can be done.

If the grid has a varying voltage, say from a microphone, it can control the number of electrons flowing through the tube. This will vary the much larger voltage that we put on the plate of the tube. A small voltage controlling a much larger voltage. Now instead of millivolts, we have an analog of our microphone signal that is many volts. And it tracks along "exactly" with the microphone voltage. It's an amplifier!

Transistors work in a similar fashion, but there are some significant differences. A vacuum tube is a voltage amplifier – a changing voltage on the grid causes a much larger voltage change on the plate. But transistors are current amplifiers. As the current on the "base," the transistor's equivalent of a grid, changes, the current on the collector, analogous to the plate in a tube, changes.

That's really a minor distinction, since both are amplifiers, but it does have implications for audio.

Transistors operate on low voltage – usually less than 48 volts – while vacuum tubes operate with much higher voltages – typically hundreds of volts.

And because of these differences, transistors are intrinsically low impedance devices while vacuum tubes are high impedance.

So, what has that got to do with how they sound? Theoretically, not much. But in the practical world of real circuits, those differences affect how the amplifier is implemented.

For one example, most amplifier designs require capacitors in the signal path. We won't go into why, but they are necessary. The value of capacitor required depends on the impedance of the circuit. A high impedance circuit needs small value capacitors, while a low impedance circuit requires much larger value capacitors. Why is that an issue?

Well, the quality of the capacitor has a significant effect on the sound, with some types of capacitors being very transparent-sounding, and others not so much. The low value capacitors in a vacuum tube circuit are small and we can use really high quality, sonically transparent capacitors. Those same capacitors for a transistor circuit would be gigantic – much too large to be practical.

What we have to use in transistor circuits are electrolytic capacitors, which have the property of a lot of capacitance in a small package. Problem solved!

But what do electrolytic capacitors sound like when you put audio through them? Not as good as the type used in tube circuits, which are generically called "film" capacitors.

The problem with electrolytic capacitors is that they tend to be non-linear – in other words, what comes out of them is not quite the same a what went in. We call this "distortion," a general term for any modification of the signal.

We might not even hear the sound as distorted in our usual sense of the term, but there is usually a subtle change in the character of the sound. Going through one electrolytic capacitor may not degrade the sound all that much, but when you go through a lot of capacitors, that distortion adds up. Through a typical console, from mic to recorder, the audio might go through dozens of electrolytic capacitors. That is much more likely to be audible.

Also, electrolytic capacitors have a finite lifetime. Eventually they will have to be replaced. They gradually change characteristics, and as they do, the sound changes. That's why you have to "re-cap" consoles.

The capacitors used in the signal path of vacuum tube amplifiers rarely deteriorate. They will sound the same decades later. No one really knows how long they last, but it could be hundreds of years.

That's why you never hear about "re-capping" tube gear.

Just to be clear, all equipment uses electrolytic filter capacitors in their power supplies, and those will eventually have to be replaced. But aside from an increase in noise, power supply capacitors will not affect the sound in the same way that they will when they are in the signal path.

There is another possible difference, and that has to do with headroom.

I learned about this from an article in the May 1973 issue of the Journal of the Audio Engineering Society called, "Tubes vs Transistors: Is there an audible difference?" by Russ Hamm. Russ was an engineer at legendary A&R Studios in New York, and his research had a profound effect on me, starting with the day I read his article.

Remember my disappointment with my first solid-state console? Well, I wasn't alone in noticing that. And Russ Hamm was also puzzled by what he heard and decided to take a scientific look at what was going on.

By the way, about 15 years after his research, Russ and I became friends and we have discussed this and many other aspects of audio quality over the years.

One thing that Russ decided to look into was just what kind of peak level came out of the microphones at A&R, when used in actual recording situations. What we read on a VU meter does not tell the entire story. In fact, no meters of that day could accurately measure the short duration of the attack transient of many instruments.

But you could measure it with an oscilloscope, which responds instantly to the absolute peak level of a signal and displays it on a screen. Well, the old oscilloscopes that used a cathode ray tube, like the picture tube on an old TV set or computer monitor, could show this instantaneously. Modern digital oscilloscopes are marvels of engineering, but they are not as instantaneous as a CRT in displaying peak levels in real time.

Russ put a variety of condenser, dynamic, and ribbon mics on various instruments and measured the true level coming out of the mic. What he discovered was astounding to me. The level was much higher than I imagined. I found it difficult to believe, so I duplicated his experiment and found the same thing.

Some percussive instruments, like drums or a piano, could produce nearly 1 volt of output with some mics. That was about 20 to 30 dB higher than what we assumed the level was, judging from a VU meter reading.

When I tried the same experiment, I found even higher levels than Russ did. Maybe the players I was recording were hitting things harder, or I was mic'ing closer, but whatever the reason, this totally changed my understanding of recording.

How does this relate to tubes and transistors? Well, we have to go back to the fundamental differences and see that any amplifier cannot produce a clean output if the level exceeds the power supply voltage. At that point, there is nothing left to increase and the sound waveform is "clipped" at that level.

Tubes have an advantage here because we are dealing with hundreds of volts available, as opposed to the 30 to 36 volts used by most transistor amplifier stages.

So, are transistor amplifiers crippled by this restriction? Not at all. It just means the circuit designer has to allow a huge amount of margin in the gain structure to accommodate the peaks without exceeding the available voltage. This is not a difficult engineering challenge, but it must be done with an appreciation of the real-world levels or else the amplifier will go into clipping on the peaks.

And in all audio circuit designs, there is always a battle between noise and headroom. The more headroom you provide, the higher the noise level, and vice-versa. Determining the best compromise is a constant challenge for circuit designers.

The average level out of your microphone on a piano or a snare might be -20dBm, while the peak level could easily exceed 0dBm. But the duration of the transient peak at the beginning of the sound only lasts for a tiny fraction of a second. That's so short that you will not hear distortion on the transient peak. It just sort of rounds it off at the power supply voltage and acts almost like a peak limiter.

But it does change the sound. The percussive attack lacks impact and the sound loses its punch.

Properly-designed solid-state amplifiers handle this very well, but in many designs, those peaks are going to get chopped off and nothing can really restore them. There is an abrupt point where the amplifier ceases to follow the input and clipping occurs. And you can hear it.

Tubes, on the other hand, have a couple of characteristics that mitigate this phenomenon. For one thing, theoretically they have a much greater dynamic range because they have so much headroom. But there is another factor, and that is that the transition from perfect reproduction to the clipped state happens much more gracefully, compared to the hard cut-off of the transistor. The onset of tube distortion is usually gradual, which makes it less audible. Even when the level does exceed the headroom, the output is much more acceptable to our ears.

And this leads us into what is probably the biggest factor in the difference in the sound of tubes.

Before we get into that, we have to first understand that there are several generic varieties of tubes, based on how many elements are inside them. The tube I described earlier, with the cathode, grid, and plate, is called a triode. Three elements.

But you can make tubes with additional grids – as many as five total. Those tubes are called tetrodes, pentodes, for the two and three grid versions. And each configuration has its advantages and disadvantages. Adding more grids allows the tube to have more gain in the same package, for one thing. That can reduce the number of tubes necessary to achieve the required gain. And that reduces the cost of the product. That's great for mass-produced consumer electronics, but do we really want to save a couple of dollars in the design of high-quality professional audio equipment? I don't think the trade-off is worth it, so I only use triode tubes in the products I design.

Triodes are the simplest amplifying tubes, and as a general rule, I like "simple" in my audio path.

The less stuff you put the audio through, the better it sounds.

And triodes have another interesting characteristic when it comes to distortion. When the tube begins to distort, because it can't keep up with the range of voltage being asked of it, the distortion products are predominantly even-order harmonics.

Even-order harmonics are simply octaves of the original sound. Even-order harmonics are musical. In fact, most pleasant-sounding instruments produce mostly even-order harmonics or overtones. We like that sound.

Triodes tend to emphasize the second harmonic when they begin to distort. That's an octave above the original sound. That adds "fullness" or "musicality" to the sound. It may not be a precise rendition of the original sound, but this slight modification is generally pleasing to our ears.

Sounds below the onset of distortion are reproduced very accurately.

And, you may have guessed, transistors emphasize the odd-order harmonics, which are discordant, strident, and annoying to our ears.

In fact, when a transistor amplifier goes into clipping, the resulting waveform becomes a square wave, which by definition is the fundamental plus an infinite array of odd-order harmonics. Ouch! Not easy on our ears. That sound is part of the appeal of a fuzztone pedal, which might sound great on an electric guitar part, but definitely sounds awful on almost everything else.

Unless you intentionally want to annoy your listener, you want to avoid odd-order harmonics. Of course, sometimes you do want an annoying sound, and you can use this characteristic of odd-order harmonics to do that.

So, why do solid state devices even sound acceptable to our ears? Again, it is careful design that avoids that realm of overload.

But in the case of a solid-state mic preamp, the headroom may be exceeded constantly from the attack transients that are many dB higher in level than the average level. The sound is modified in a way that can make it sound smaller and less exciting. Our ears are very good at detecting transients, since we need those for localization and in interpreting the meaning of a sound.

The tube mic preamp may have a similar problem when its headroom is exceeded, but the resulting distortion is much more acceptable to our ears. In fact, it often makes the instrument or voice sound bigger and stronger.

I use a mic preamp as an example because it is the most critical of the devices we have control over in our recording. But every stage of the recording process has amplifiers that can suffer from these problems. It is important to recognize this effect, whether you are a circuit designer or a recordist. Some of these characteristics you have control over, and some you do not.

There are other factors that explain the difference between tubes and transistors, but these are ones that I think are most significant.

And that's why I design my products using tubes. They just sound better to me.

A tube circuit can certainly sound bad if improperly designed. That is entirely possible, and you may have examples of that in your experience.

On the other hand, a guitar amplifier sounds good for its purpose. Tube guitar amps sound great, even though they are not good audio designs. The original goal in the 1930s was to increase the volume of a guitar so it would be comparable to the drums. And this new technology could not be outrageously expensive or no one could afford to buy it. Guitar amplifiers were stripped down to the bare minimum number of parts necessary to make them function. Every corner was cut to keep them simple and affordable.

The result was a pretty terrible audio amplifier, but it was exactly what was needed to amplify a guitar. Those deficiencies became the definition of the electric guitar sound. And that's why it is nearly impossible to recreate that with solid state devices.

The products I design do not have the traditional "tube sound," which is largely based on cheap designs. Good designers have always known how to make tubes sound wonderful, and I am fortunate to have their legacy to guide me.

Good design also usually equates to expensive products, unfortunately. But every product I designed is something that fills a need I have in the studio. It has to sound and operate better than anything else available to me or else it is not going to be released as a product.

Still, there are situations where tubes are impractical to use and transistor devices are necessary. One example is in A to D converters. Although the input analog amplifier could be a vacuum tube, the next stage, which interfaces with the actual converter chip, requires characteristics that would be very difficult to do with tubes. It might even be impossible, or at least extremely complex. So, we are stuck with solid-state for that function.

Also, in a console, the summing amplifiers that extract signals from a mixing buss have some tricky requirements. Solid-state integrated circuit opamps are perfect for this.

The first op amp was invented by Bell Labs, not surprisingly, in the late 1940s, and it used vacuum tubes. The concept was later adopted to use discrete transistors, and eventually became the integrated circuit devices we know today.

And speaking of consoles, the complexity of a full-featured console using vacuum tubes is certainly possible, but it would outrageously large and expensive. Solid-state integrated circuits place a huge number of transistors in a tiny package and make the modern console practical.

These days, most studio operations revolve around computers in one form or another, and here again solid-state makes this practical. You would need a building the size of an aircraft hangar to house a tube computer with the capability of even the simplest laptop or tablet.

Want an example?

Well, when I was around 12 years old, a neighbor asked me if I wanted to see the Univac computer he used as an engineer at General Electric. Of course I did!

This was a vacuum tube computer, obsolescent even at that time, as transistors replaced tubes in all digital computers thereafter. This computer was already a piece of history.

The computer took up a very large room. It was noisy and generated a lot of heat, so there was a massive air conditioning system in the building just for this computer. The front panels contained an array of flashing lights, which was very impressive, plus many reel-to-reel tape transports, which were used for data storage. These were beautiful multitrack machines running one-inch tape, and they were constantly starting and stopping, and reversing direction.

My neighbor asked if I wanted to go inside the computer. Inside? Wow.

We walked around to the side of this massive device to a glass and metal door that looked like something on a submarine. It had big latches and a thick rubber gasket all around the opening. The doorway was rounded on the top and bottom. When he opened the door, a huge amount of air was sucked in, and I was almost knocked off my feet.

We went inside and closed the door behind us. There were no lights inside, but it was still pretty bright from the seemingly infinite number of vacuum tubes with their orange filaments glowing. The tubes looked like the typical dual triode tubes I was familiar with – and like the tubes I still use in my equipment designs.

It was like an oven inside, even with the hurricane of air blowing through. Actually, it was sort of an infrared oven. The noise was deafening. It was impossible to hear each other. My friend just pointed at things and I nodded my head. The explanation would have to wait until we got outside the machine. We exited through a door on the opposite side of the computer.

Outside in the relative comfort of the room, he showed me the core memory of the computer. Random access memory, or RAM, as we know it today had not yet been invented. This computer used a three-dimensional matrix of tiny metallic donuts with wires going through them. When a specific core was addressed by applying a voltage to the two wires that intersected a particular donut, it became magnetized and stored a single on-off bit. To read the data back, a similar process was used to ask the donut if it were magnetized or not – a one or a zero.

I don't remember exactly, but I think the memory core was cable of storing less than one kilobyte of data. The engineers were very proud of that massive capability.

The computer program was on punched cards.

Permanent storage was on magnetic tape, and they had a large room with floor to ceiling boxes of 3M one-inch tape.

Needless to say, this made a big impression on me.

So, no, I don’t want a vacuum tube computer. Or really any tube device that isn’t for audio. Transistors are much better at that than tubes.

But for me, recordings made with vacuum tube audio equipment sounds much more pleasing and realistic to my ear. There is no “tube sound.” There is only the harshness of audio solid-state devices. Tubes sound like music.

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