John Curl interview

Intro:

John Curl is one of the most respected circuit designers of all time and the creative genius behind Parasound’s high end audio and home theater amplifiers. Since the mid 1970s he has left a trail of landmark hardware across the industry ... the classic Mark Levinson JC-2, the SOTA head amplifier, and his own Vendetta preamplifier. Each of the 200 or so Vendettas is a hand-built work of art that is treasured by its owners.

Curl's first job as an audio engineer was at Ampex where he worked on the design of tape recorders. Later he was involved in that company's pioneering research on video tape recorders. From there Curl moved into the rock and roll business, designing and building the sound systems for The Grateful Dead's road shows. As an independent consultant, Curl has worked dozens of projects both in pro audio and home audio; making master recorders, studio boards, microphone preamps, power amplifiers and many other products.

In 1989 John Curl was introduced to Parasound's founder, Richard Schram and since then, Curl has designed all of Parasound's high-power amplifiers and consulted on the design of the low-level circuits of many of its other components.

Curl earned the respect of everyone at Parasound by being somewhat conservative and uncompromising in his approach to circuit topology and component selection. He insists on using only the finest parts and balanced circuits. He avoids the use of capacitors and inductors in the circuit path. On the other hand, he has learned to avoid the excesses of design that can turn a great design into an overly-expensive design exercise. Curl and Parasound are dedicated to delivering products that offer very high value at a reasonable prices.

"He wouldn't last a week in a mass market receiver factory," says Schram. "The accountants would probably reject every part he picked because it cost too much. At the same time, he knows how to make a very, very good product at what we consider to be a reasonable price."

The following is an interview with John Curl recorded in August 1999.

Q: John, you’ve been designing amps and preamps for over 30 years. Is it still a challenge? Are you having fun?

JC: Yes, it’s still great. It never ends. We keep solving problems and then turning over new ones. It's very much like modern physics - it continually evolves. We have never yet been able to make a completely perfect amplifier at any price. And then there is the challenge of making very, very good amplifiers in a way that is cost effective.

Q: All your amplifier designs for Parasound use balanced, fully complementary circuits. Why do you do this?

JC: Well, it's inherently more linear. That type of design, all else being equal, always has less distortion. It's pointless to design something that is prone to distortion and then try to remove the distortion. You might as well start right and then work from there. I sometimes compare it to automobile engines, like a big Detroit V8 engine that's been made for the past 30 - 40 years. It's powerful, but basic design is relatively compromised when compared with something like a Porsche or a Ferrari. There is a reason for that -- it's not just a mark up, it's a matter of refinement. Porsche and Ferrari start with really sophisticated designs and then they refine them further.

Q: All the power amplifiers you have designed for Parasound run pure Class A bias at low-to-moderate signal levels. How important is Class A bias?

JC: All else being equal the more Class A you have the better. We don't need a full Class A amplifier, because people don't listen to full power levels all the time. They really need only 10 Watts or so because that is where most of our listening is done. However, getting that 10 Watts of Class A is really quite a feat considering of the all other constraints of having 200 Watts Class AB in reserve. The Class A portion is always smoother and better, while the Class AB is always a bit rougher because it is against the law of physics. It breaks the music into positive and negative cycles and then splices them back together. Only one side is working while the other side isn't... it's the nature of the beast. You can sometimes see the transition on a scope; it may not be very much but it will be there. It will also generate more high order harmonics. By definition, Class AB is not as linear as Class A, but it is certainly much more efficient. So it's better and easier to run as much Class A as possible.

Q: Is there anything else you can say about the Class A situation? One of the things you were talking about before was noise. You know, high order harmonics and things like that.

JC: Well, see this is the deal. When you're working with both sides of an output stage you actually have a true push / pull situation, you cancel out the even order harmonics. This is because you have one device turning on and the other device turning off and they are summing together. And actually the second harmonic just sums out - it cancels out. However, if only one side is on, the second harmonic cannot be canceled because there's nothing to cancel it with. But people will say, 'wait a minute I didn't measure any significant increase in the second harmonic. What happened?' What happened is the second harmonic turned into higher order harmonics, because the amp is only working on half the wave form. So the same nonlinearity, on half the wave form, makes third harmonic. So all of a sudden you don't see any second harmonics, but you will see a difference in the third. Of course if there is any fourth harmonic, which there usually is, it is fairly innocuous. You could have a lot a fourth harmonic and probably not even notice. But when it turns into fifth, you will notice it because it's more dissident. Fifth harmonic is out of tune. Third is tolerable, but fifth is pushing it -- it becomes dissonant. That's why people can tolerate these single ended amplifiers. They're all running pure Class A -- they have to by definition. They won't work any other way. They have lots of second harmonic; well nobody can easily hear second harmonic. It sounds sweet because it's still part of the music. But if it was fifth harmonic, trust me, nobody would want anything that had pure fifth harmonic.

Any time you go into Class AB you are creating at least some high order harmonics. It's converting the even orders that are naturally there and that would be naturally canceled out. But they can't be canceled anymore, so they have to be converted up. So they have to appear, but they don't appear as even orders anymore. They appear as odd order, so they add to odd orders that are there already because they can't be gotten rid of. All of a sudden you see this jump in the high order spectrum. Normal distortion meters won't measure this because they just measure the total harmonic distortion; not the individual harmonics. Unless you know what to look for with a spectrum analyzer you won't even know this is happening. Most people just see it as a little glitch on the meter. They don't know that the whole nature of the distortion has changed.

Q: What about output devices?

JC: Well, Parasound has a very good combination that's practical. They use bipolar output devices driven by a complementary FET driver stage. You get the best of both worlds. If you use a transistor driver stage, then the predriver stage -- the one that actually produces all the gain -- will get modulated by the loudspeaker load. In other words, the loudspeaker load reflects itself back as a changing impedance.

Lets say a loudspeaker momentarily changed from 8 ohms to 1 ohm. I have a speaker that does that: a pair of Wilson Audio Tiny Tots. They sound pretty good sometimes, but they are troublesome. They drop to less than half an Ohm at 2,000 cycles. In some amplifiers it can actually reflect that change of impedance back into the driver stage and change the dynamic gain of the amplifier. The FET driver tends to buffer the load change better than a transistor. And that's why I prefer FETs in that respect.

If FETs were perfect they probably would be better as output transistors than bipolar transistors, but it is almost impossible to build an all-FET amplifier. I can do it; however American FETs have all the qualities that you could ever want except that they're unreliable. They break if you look twice at them. You short the output for example, by accident and they invariably break.

O: Is reliability a major design goal for you?

JC: Absolutely. When I worked with Saul Marantz at Lineage 10 or more years ago, we built a all-FET power amplifier. I had read all the spec sheets and really thought that I had everything covered, but invariably, when we shorted the thing out, it would blow up. It gave me an appreciation for protection circuitry. Today, I believe in protection circuitry that is noninvasive. It will sense when you've shorted something out, and know that something's wrong and protect the amplifier.

Q: Well, one of the things about Parasound amps, they are very rugged and that's - is it an important part of your belief in amps?

JC: Well, it's not me that does that. Over the years Parasound has developed some very effective protection circuitry. It has an output relay that could be problematic in extreme circumstances. It's kind of like a rev limiter on a car. It doesn't do anything until you exceed the redline, and then it cuts power to keep the engine from blowing. It could reduce your performance sometimes if you like to push your engine beyond the redline. It depends on what you want. Do you want to be a hot rodder or do you want your engine to hold together? I have an old Porsche and it just can't breathe. It has a hard time just getting up to 6,000 rpm, which is it's red line. You're saying what am I doing here? I'd better shift because I'm not going anywhere. I also have an Acura and, man, that thing kicks in from 5,000 to 7000 rpm and I'm always going through the rev limiter at 7,000. If I didn't have a rev limiter, it'd be gone. Parasound's amps use a very sophisticated protection circuit that they've or developed over the years, and basically it's the equivalent of a rev limiter. It senses a number of things and when you're in trouble and it's too hot it fires this circuit to open the relay and protect the amplifier and the speakers.

It's a lot better than fuses, because fuses are non - linear. They can't help it. They are designed to have to heat up before they blow up. This heating process changes it's resistance. The resistance invariably rises as the temperature goes up --that's pretty much the laws of physics-- and of course it's going to modulate. If you have a lot of current flow, even though you're below the limit of the fuse, the darn fuse is going to modulate. Of course, this is almost impossible to measure statically. Once you put a resistive load on and a certain amount of power the fuse is going to up to it's temperature and stay there and the resistance is not going to change very quickly, so you never know that the fuse is distorting. It only distorts getting up there and then going back down. It doesn't necessarily distort by the time you're ready to make the measurements. That's one of our problems with static measurements. You just don't know everything that you need to know.

Q: What is your philosophy on power supplies? Should they be tightly regulated?

JC: Actually, I don't use tightly regulated supplies by comparison to Krell or Levinson standards. They use tightly regulated power supplies and rightly so. If you want a very exotic, heavy and expensive amplifier you should super regulate the power supply. But I feel there is too little to be gained from this. It's just too exotic and too expensive, and also too costly as far as heat production is concerned. I'd rather save the heat sink space for the output devices, and we have plenty of those. But we don't have an infinite amount of space. We're pushing 100 pounds on our biggest amp, but we're not going go to 200 or 300 pounds like many of others do at ten times the price. We make a well regulated design in the sense that it has a lot of transformer and capacitor capability and will maintain itself well under dynamic conditions. Compared to a Levinson or Krell it will actually sag a little if you sustain the drive indefinitely. However, that is nothing compared to the power supply limitations of many of the budget amplifiers and receivers.

Q: How does Parasound's approach differ from Carver's?

JC: Carver's philosophy was a little bit different in the sense that he didn't use very much capacitance but he has used a very high voltage capacitor. His idea was to allow a lot of voltage swing -- more than most amplifiers. Carver was making 350 Watts channel amplifiers way back in the early 70's. However, when it comes to current drive he had almost no current drive. And that's because current drive is a function of the output devices and the amount of power supply capacitance. For a given size capacitor, the higher the voltage rating, the less capacitance there is available. So if you have a 150 volt capacitor, it may only have 10,000 microfarads capacitance. I would much rather use a 100 volt capacitor at 25,000 microfarads capacitance. The difference is that I can push out current 2 and half times longer and I'll only sag a little bit. The Carver's design might sag a lot, but since it might start at 130 or 140 volts, it will drop down to 100 volts. So, while the Carver design would have more dynamic voltage swing, the Parasound would have more dynamic current output.

Q: Why is this important?

JC: Loudspeakers have certain operating conditions where they demand more current than you would ever imagine -- 50 or 60 amperes can be demanded dynamically by a loudspeaker for a very short period of time. It doesn't have to be sustained; maybe only for 5 milliseconds or so. If the amplifier can't sustain this current, the capacitors are sagging and it is eternally clipping, or the protection circuits are firing. Something is happening that isn't necessarily right and I think this makes a difference. Since so many loudspeakers have these high peak current demands, we design our amplifiers to meet this requirement.

Q: What can you tell us about the transformers? Are they very important?

JC: Absolutely, but when we talk about transformers we have to separate power amps and preamps. While they are not perfect, Toroidal transformers are the logical choice for power amplifiers because they are very efficient, they tend to have a fairly low hum field, and they're readily available in large power ratings.

For preamplifiers and other line-level components, the old type EI transformers or what's called a D-core or split C- core transformer is actually better than a toroid. First of all, they tend to be more compact, and second, and perhaps more importantly, they have very low capacitance between the windings.

This can be a problem when Toroids are used in low signal level applications; the windings are on top of each other so they talk to each other. It used to not be so bad but today the AC power is so dirty. Harmonics that are created by high frequency fluorescent lights, fax machines, computers, you name it. All this new stuff, that's only been around for maybe the last 10 years, tends to get into the power supply through the transformer and then in to he grounding system and ultimately into our sound system. So then people have to of course use expensive power conditioners to repair the problem but if you fix it in the first place then its not so important.

Q: So how would you fix it in the first place?

JC: By using a transformer that isolates the winding, which is important on low level circuits. The old style EI or the new C- or D-core is the ultimate in that respect -- as long as it is a dual-bobbin winding with physically separate bobbins for the primary and for the secondary. It makes a big difference in sound quality.

In power amplifiers, however, that isn't as important because the levels are just so much higher. If we had a choice, and if money was not involved, or weight, or anything else, we'd probably use an EI type or special type of transformer.

Q: You are very careful about the resistors you use. How can one low tolerance 47 Ohm resistor sound different from another one?

JC: There are a number of reasons, and we don’t know them all. People tend to think that a resistor is simply a measurement by an ohmmeter -that’s all resistance is. You measure it with an ohmmeter at 1 volt and you say 'okay, this is a 50,000 Ohm resistor.' Aren’t all 50,000 Ohm resistors measured this way all the same? Well, no; not at all. Any more than you can describe a car by saying it weights 2500 pounds. All cars that weigh 2500 pounds don’t perform he same, for example. That’s just one measurement.

The other resistor measurement is tolerance; how accurate is the resistance. Traditionally, people use 5% or 10% carbon resistors for most purposes. I don't think these are adequate for carefully balanced audio circuits. At room temperature a 10% resistor might be accurate within 1%, but its value will drift as it heats up. This can really throw a circuit out of balance, which can’t be very good. The newer, high quality 1% resistors hold their values as they heat up.

Resistors have other characteristics that effect the sound. One is the type of lead material used. The lead material of many inexpensive resistors is made of material that can be magnetized -- usually soft steel as far as I can tell. This is actually useful for automated parts insertion machines because they can use electromagnets to handle the resistor by it’s leads. Steel is also intrinsically rugged and it will stand a tremendous amount of vibration. The military, for example, doesn’t allow pure copper leads to be used except under very special conditions because, in a missile or something, it might break. They use a special alloy which, I think, has nickel in it. However, it is also magnetic and it doesn't sound as good as pure copper.

What we try to use is pure copper. Many of the precision resistors -- the really high quality ones, like those made for precision scales and instrumentation -- are made by Holco (which means all copper), Vishay and a number of others. These are all basically non-magnetic because the manufacturers know this can interfere. A stray electro magnetic field can actually be picked up by the circuit, so we try very hard to avoid it.

These things are very difficult to measure. When you measure something statically on a test bench, you test at one level and one frequency. That doesn't place much of a stress on a resistor, so the resistor acclimatizes itself to that condition and just does what it does. The problem is that music is continually changing so it's actually affecting the resistor much more than a static test will ever show. This is one of the problems with resistor testing because you have to use fairly sophisticated testing to actually see these problems. You can predict them by looking at the temperature coefficient of the resistor. If it's 50 parts a million, or 5 parts a million, or 100 parts a million, each one of these is going to be different. That means it is going to change its resistance for every degree of temperature change.

The size of the resistor matters too. If you use a very miniature resistor and put lots of dynamic power change across it, it is going to go from room temperature to maybe 50 degrees Centigrade and back. This rapid change can't be good. Normally, most resistors in a power amplifier or preamplifier don't do very much. They have current flowing through them continuously and they stay at a fairly regular temperature once they warm up. The feedback resistors, however, are very critical because they see the entire output signal and they swing back and forth depending on what the output signal is. They can be audibly effected by temperature changes.

Also there is noise. When you measure the resistance of a resistor, you can predict its lowest noise by its value. When you have current flowing through the resistor the noise will increase depending on the quality of the resistor. The cheaper carbon resistors can be extremely noisy, to the point where you put 5 or 10 volts across them and it's off the map. The higher quality resistors will change in almost no way, it's almost impossible to measure it.

Q: Okay, what about capacitors?

JC: I learned how to measure a capacitor for the first time about 25 years ago working with Tektronix. They were making a piece of equipment which measured capacitors and they were really worried about the values. They were doing it with one of their pieces of test equipment -- It's called a curve tracer. They modified it so it would measure the value of capacitors. They found that ceramic capacitors, for example, had another characteristic they had never been able to see on the screen before and it actually affected the measurement of the value of the cap. It showed a tremendous non-linearity. Interestingly enough, in this particular test and this method of measurement, only ceramic capacitors showed up to be really bad. We found that ceramic capacitors really were bad guys. Later we found that we could emulate the same problem indirectly using normal test equipment, but we had to operate the capacitor in some sort of real way. It couldn't just be sitting there with zero volts across it; it had to be working with some kind of a signal like rolling it off high frequencies, low frequencies or something.

I published a paper in 1978 and Audio magazine article in 1979 that showed this problem with ceramics, and we also found that Tantalum capacitors did almost the same thing. With this particular test (with normal test equipment) you could see the non - linearity of the Tantalum capacitors as well as the ceramics. This still allowed us, in theory, to use aluminum electrolytics -- we couldn't find any real problem with them as long as they were used properly -- or any kind of metal film capacitor.

A third type of distortion, which has been known for many years, but had been forgotten about, is called dielectric absorption. This particular problem used to be very important back in the 50s when people used to solve many engineering problems with analog computers. These analog computers would emulate mathematical equations with capacitors, resistors, or amplifiers. Music will also evoke dielectric absorption. Music tends to not be completely symmetrical at all times, and even though it averages out in the long run it isn't necessarily a test tone. If you put a symmetrical test tone through a mylar capacitor for example you won't find any real problem. However if you use an asymmetrical signal you'll find that it does have dielectric absorption. This is where the dielectric absorbs part of the signal and then spits it back later. Well this can't be good. Invariably you never get the musical peak, it cannot be completely passed by the capacitor because the capacitor has to take some of the energy from the musical peak. It stores it like a battery. Fortunately, this material property isn't shared by Polystyrene, Teflon, or Polypropylene, which is why we tend to use these caps instead of mylar. Tantalum, aluminum, and mylar are pretty bad in this area.

As a result of all this, we have to exclude many types of capacitors because they all have some problems to a greater or lesser degree. Ultimately we wind up with polystyrene, polypropylene, and Teflon. And that's why we tend to prefer these capacitors when we can. Except for the use of aluminum electrolytic for power supplies , the more capacitors we can eliminate the better it its.

Q: You've always been a proponent of trying to keep the signal path free of inductors and capacitors. Why is this so important?

JC: It's like this - it is easy enough today to design out capacitors between stages. It is rather redundant and wasteful to add capacitors between stages. First of all, they do not help the size of the unit. They're not very reliable. If anything is going to go bad, the capacitors will probably go bad first ... unless you have catastrophic failure. In short, they don't really do you any good so. The best capacitor is no capacitor...we don't need them anymore.

In the old days, when we didn't have complementary circuits, we needed capacitors. When you look at vacuum tubes there is no such thing as a complementary vacuum tube device. So you almost invariably needed transformers and capacitors. Then again, one of the advantages of the vacuum tubes is that they are very high impedance devices, so the capacitors could be small in value even though they might have to be high voltage.

Now, when you use capacitors in solid state transistor equipment, you generally need fairly large value capacitors, but their voltages don't have to be so high. These situations would seem perfect for aluminum or Tantalum electrolytics. However, these are the ones that are not very reliable and they have all these secondary distortion characteristics ...

dielectric absorption, nonlinear distortion, and that sort of thing. So, if you can eliminate these capacitors, why put them in in the first place?

Now, some people can say 'what about leakage or safety or something like that?' Well, of course, you have to be careful, and that is what modern protective circuitry is good at. It shuts down the amplifiers if they are behaving abnormally, yet it doesn't impact the signal when the amp is behaving normally. We also use servos, which are basically very precise well matched IC devices. In the factory, they laser trim them down to one or two millivolts and then we simply use these to compare the output to ground and then adjust very slowly to zero out any offset that might be inherent in the amplifier or preamplifier.

It’s easy to do these things now. Thirty years ago it wasn’t easy because we didn’t have FET input ICs, much less very well matched FET input ICs. For example, the JC -2 didn’t have servos because they weren't practical in 1973 when it was designed. Maybe the military could’ve done it, but the real world had to wait until about 1978 or so.

Also, we couldn’t use mylar capacitors, which are fairly efficient coupling capacitors. While mylars are fairly efficient from a size and cost point of view, we realized they have problems with dielectric absorption. I didn’t believe it at first. I was working with Noel Lee and a company called Symmetry. We designed this crossover and I specified these one microfarad Mylar caps. Noel kept saying he could 'hear the caps' and I thought he was crazy. Its performance was better than aluminum or tantalum electrolytics, and I couldn’t measure anything wrong with my Sound Technology distortion analyzer. So what was I to complain about?

Finally I stopped measuring and started listening, and I realized that the capacitor did have a fundamental flaw. This is were the ear has it all over test equipment. The test equipment is almost always brought on line to actually measure problems the ear hears. So we're always working in reverse. If we do hear something and we can't measure it then we try to find ways to measure what we hear. In the end we invariably find a measurement that matches what the ear hears and it becomes very obvious to everybody.

Years ago, there was a time when people used to think you could have a two- or four-foot path difference between loudspeaker components; like the Klipschorn, for example. Everyone said this time difference was inaudible, and it didn't really matter because Bell Labs' research, Ohm's law of acoustics, Helmholtz and all these other people believed that the ear was completely insensitive to phase. So it didn't matter how you built the speaker as long as it sort of averaged out sort of okay in the room. You could take five microphones and measure them all together, if that measured out okay within a few DB's then heck with it. Well, that really isn't true and of course when stereo came along all of a sudden you had these big Klipschorns and they wouldn't image for anything. At least that was my personal experience. I owned them and I was a believer too.

Then I started measuring them and I said 'oh my goodness, this is a problem.' The late Richard Heyser tried to tell people that a two foot path difference might be audible. People were going crazy and saying this was impossible and it was a big controversy. Now, of course, no fool would design a speaker with a two- or four-foot path difference. John Dunlavy was very outspoken on the Internet this week, criticizing a loudspeaker that wasn't completely phase aligned to within one inch.

See how we change. I don't disagree with John Dunlavy, although I do think he is overstating his case in this particular one. But, there was a time when we didn't. The same thing happens with capacitors. There was a time when we didn't know better and we just used any old capacitor as long as it had the right values.

Q: Many companies use cooling fans for their most powerful amplifiers. How do you feel about them?

JC Well, in principle I'm not against cooling fans. We try to put as much heat sink on as possible and if I can find, in fact if I was asked to find I could probably do this, a fairly quiet cooling fan and if it was mounted properly internally or on the bottom it would be very helpful. In fact sometimes I use a cooling fan on my own Parasound, just by on the side of the heat sink. I have one heat sink that's slightly miss-calibrated and it tends to go into over temp, so I put a cooling fan on it. It works great but most cooling fans are rather noisy. You always have to have a certain amount on noise but you can go below 20 SPL -- that ought to be quiet enough for almost anyone. Not everybody in the known world, but at least those in the home theater and that sort of thing. My personal feeling is there is nothing really inherently wrong with fans it's just that they are often audible, especially cheap ones, because they're just a cheap add on. But if you think about, you could use good dampening materials to isolate the fans. One thing I used to do with my Dyna kits, was to put a whisper fan on top blowing down on the tubes. Now this is not the most efficient way, but it's safer on the fan because the heat's not passing through it. I would use a rubberized damping material that would fit around the fan so it didn't vibrate. That was a pretty efficient way of doing things and it extended the life of the vacuum tubes.

Q: That's about all the time we have. Is there anything else that you would like to share?

JC: That’s about it. We keep working on trying to make amplifiers faster, use less feedback and that sort of thing. I’m working on a new pre amp of my own manufacture.

Actually this sort of thing is more like a Formula One car, where the designers push the limits of abosulute performance, but they are not fit to drive on the street. But car companies learn a lot of things from race cars that eventually shows up on production cars.

My newest preamp is extremely exotic. I don’t use any negative feedback at all. This is really an experiment. It has local feedback, but it has only a resistor that sets the gain. It'd like to see if that makes that much difference. This project isn't practical, it's way too expensive, but it helps me to learn how far I can go toward creating the ideal preamplifier.