
J. Gordon Holt

T H E C A R V E R A M P L I F I E R C H A L L E N G E

Is it possible to make a \$700 “mainstream-audio” power amplifier sound exactly like a high-priced perfectionist amplifier? Bob Carver, of Carver Corporation, seemed to think he could, so we challenged him to prove it.

The question posed above seems laughable. If it *were* possible to make an average, modestly priced amplifier sound just like state-of-the-art, wouldn't it already have been done? Of course it would. State-of-the-art sound would thereby become much more affordable, and high priced power amplifiers

[amplifiers] would become as extinct as *Diplodocus*.¹ That is the conventional wisdom. Bob Carver, founder and personification of Carver Corporation, has never been noted for his conventionality.

Ever since he introduced the first high-powered solid state amplifier in 1971, Carver has been laying waste to conventional wisdom with one brilliant design innovation after another – the “magnetic amplifier,” the “peak unlimiter,” the “sonic hologram generator,” the “auto-correlator,” the “asymmetrical charge-coupled FM Detector,” and the “digital time lens.”²

But everyone has his limits of capability, and pride goeth before a fall; when Bob claimed, some time ago, in conversation with Publisher LA., that he could make his \$700 Model 1.0 amplifier sound “indistinguishable from” any amplifier of our choice, we were confident that he was finally out of his depth. Carver Corporation is, after all, a “mainstream” manufacturer, not a “high-ender”. Bob’s designs are unabashedly aimed at the mass market, notorious for its lack of aural perspicacity. What, then, could he possibly know about the design subtleties that make a Stasis 500 sound different from an Eagle 7A? Bob’s claim was something we just couldn’t pass up unchallenged.

Our first task was to come up with a “reference” amplifier that would represent a genuine challenge – one as different from, and as superior to, his solid-state Model 1.0 amp as possible. One obvious contender was a large tubed amplifier we had on hand, but we soon realized that our choice would not be all that simple. There were, it seemed, some peripheral considerations.

We knew that Carver couldn’t possibly pull this off, at least not to the point where none of us would be able to distinguish between his modified 1.0 and our reference amp. After all, some of the most highly trained audio ears in the world would be listening for the differences. What worried

us was the possibility that Carver might come *so* close to matching the sound of our reference amp that its designer/manufacturer would be embarrassed, chagrined, and outraged. And, while not normally concerned about offending a manufacturer in a product, we *are* concerned about fairness.

In order to select a reference amp for this experiment, we would be obliged to “single out” one model of one manufacturer’s line. If Carver then managed to even approximate the sound of that amplifier, its manufacturer would quite naturally ask “Why us? Why did you single *us* out for ridicule?” And we would be hard put to answer without appearing unfair.

So, we decided to make an exception to our usual policy of forthrightness. We decided not to reveal the “reference” amp’s identity, saying only that the reference unit is a high-powered, very expensive stereo unit with a strong and unique sonic “personality”, and a penchant for being very finicky about the loudspeakers it works with. It was, we were gleefully confident, likely to be *very* dissimilar in sound from Carver’s own designs, and probably much more unpredictable in terms of its behavior with a given loudspeaker.

We then turned to the matter of loudspeakers. Again, we wished (with no implied malevolence) to make things as difficult as possible for Carver, and were fortunate this time in that two speakers which seemed to meet that criterion were among the six then in-house for routine testing. We’re not going to identify them, either. Suffice it to say that both are exceedingly revealing of subtle details in the sound, are in different ways “difficult” loads for an amplifier, and between them, excel in every aspect of loudspeaker reproduction. We were confident that we had effectively stacked the deck against Carver’s success.

Getting Started

Although both LA and JGH had met Bob Carver several times before, this was to be our first one-on-one association. We didn’t know what to expect. It turned out that

¹ *Diplodocus* was a dinosaur who hasn’t been around for about 80 million years

² Brilliant innovations some of them may be, but their names are notable more for catchiness in the marketplace than for descriptiveness of engineering innovation. –L.A

Carver, too had misgivings about us, based on past experiences with the “underground press and a normal anxiety about whether his success at meeting our challenge (about which he had no doubt) would be fairly reported.

We found Bob to be a friendly and personable gentleman, powerfully built, outgoing in manner, and just as serious about the reproduction of sound as are we. It took only an hour or so of relaxed banter before he confessed that he, too, was pleasantly surprised – to find that we didn’t have horns or cloven hooves.

Before Bob started work in earnest, it was necessary for us to all agree on certain ground rules, so that we could ultimately agree as to whether or not he had succeeded in accomplishing his goal. After some amicable discussion, we agreed on the following:

- The objective was to make the two amplifiers sound absolutely identical, or at least similar enough in sound that none of us could tell one from the other with better than 50% (pure chance) consistency.
- The reference amplifier should not have a higher power-output capability than the Carver. If it had, Bob would be obliged to beef up his own power supply, which would take additional time and prove nothing. (If cost-effectiveness is no consideration, any knowledgeable designer can put together a “perfect” power supply, given time and the necessary parts.)
- If we felt there was *any* audible difference between the amplifiers, Bob would be allowed 48 hours to eliminate that difference. If he proved unable to do so within that time, we would declare the game over and him the loser.
- If Bob felt that he had duplicated the reference amplifier, and we still heard differences, we would be subjected to a blind A?B test in which the only criterion would be whether we identified the reference amplifier correctly more than 50% of the time.

Because none of us figured that this project would be rapidly concluded, we had reserved a room for Bob in Santa Fe’s La

Posada hotel. After Bob and his fifteen numbered cardboard cartons of equipment were settled in, we unboxed one of his 15 amplifiers and headed to my place for some preliminary listening.

We were pleasantly surprised. The Carver amp had none of the usual “mass-fi” solid-state hardness, but was, in fact, very listenable, with good depth, quite good detail, and only a modicum of that high-end dryness and laid-back midrange which characterize medium-priced solid-state amps.

Not su[r]prisingly, the reference amplifier sounded *very* different and, in our opinion (shared, in most respects, by Bob), much better. We noted, with interest, that he immediately heard every difference that we observed between his amp and the reference.

The Approach

I had assumed that Bob would simply listen at length to our reference amplifier, make a measurement or two, then try various means to duplicate what he had heard and measured. His approach turned out to be much less scattershot than that. I don’t think we had listened for more than an hour when Bob suggested that he “get to work.” We transported him and the two amplifiers back to his room, leaving him to his own devices for the rest of the day.

Next morning, Bob called to say he had something for us to hear. How soon? As soon [as] we could get down to his room.

The hotel room was a shambles! Across one end was a long table buried in oscilloscopes, distortion analyzers, voltmeters, the two amplifiers, a soldering iron, a white noise generator, two unidentifiable chassis full of inductors, resistors, and capacitors, a large table fan (there was no air conditioning), a half-dozen partially-drained Diet Coke cans, and perhaps 50 feet of audio cables, test leads, and clip-lead interconnects. The adjacent sofa and table were covered with countless little plastic bags of resistors and capacitors, several schematic diagrams, and sheets of paper crammed with arcane numbers and calculations. On the floor under the table was a Rogers LS-3/5a

loudspeaker which appeared to be connected to both amplifiers at once.

Bob explained that this would be a different kind of listening test. We would not be listening to his modified 1.0 or our own reference amplifier. We would be listening to the *difference* between them. He explained that he had tacked two identical loads, each approximating a loudspeaker, to one channel each of his and our amplifiers. He had then connected the LS 3/5A and a sensitive voltmeter between the Tot terminals going into those dummy loads. This simple hookup would allow him to hear and measure the amplitude of any differences between the signals appearing at the amplifier outputs.

If both amplifiers had exactly the same gain (amplification), and were fed exactly the same signals to their outputs at exactly the same instant, the signals appearing at one amp's Plus terminal would be exactly the same as those appearing at the other amp's Plus terminal. That is, there would be no voltage difference between those terminals, and no signal would appear across the monitoring loudspeaker and voltmeter. No sound would be heard and no voltage would be read on the meter. Any sound, or voltage reading, would thus reflect a difference between the signals at the amplifier outputs – a difference which is was now Bob's stated objective to eliminate.

In essence, this is a test of the ability of one amplifier (the Carver) to cancel the output signal of the other (the Reference). Or, as Bob expressed it, to compare the *transfer functions* of the two.

A transfer function is nothing more than a statement of the relationship between the signal fed into a device and the signal that comes out of it. For example, a frequency response specification is a description of the transfer function telling us how much an input signal of fixed amplitude and varying frequency will vary in amplitude at the output.

Bob's test hookup would show much more than frequency response differences. In fact, one of his most interesting statements, for those of the "every amplifier

is the same except for frequency response" school, was that varying frequency response between the 1.0 and the reference amp made up only about 25% of the significant differences. Relative phase shift, source impedances (damping factors) – in short, every electrical difference between the amplifiers – would produce a signal at that test point between the Plus output terminals. When the amplifier outputs were identical, in all respects, there would be total cancellation – a null – of the difference signal. Bob's goal was a 70 dB null, or .03% difference between the two amps. Just to indicate how ambitious a goal this is, Bob quoted a figure of 48 dB as the null you might normally hope to produce between two channels of the same amplifier! The meter would measure the voltage difference between the two hot terminals, and thereby the degree of cancellation in decibels; the speaker would reproduce this signal to give an idea of how audibly significant the differences were. (To check how loudly the music produced the difference signal, you had but to disconnect one of the hot leads; in that situation there was a 0 dB null.)

This output nulling technique is not a Carver innovation. It has been known for years to be a possible way of comparing amplifiers, at least in theory. But it could never be made to work with amplifiers having slightly different group delay and phase-shift characteristics, because any loss of signal synchronism impairs the effectiveness of the cancellation. In other words, it wasn't used *because* all amplifiers are very different – the test was too sensitive! But phase shift happened to be only one of the many parameters for which Bob planned to compensate. Hearing of this level of sophistication made LA and I begin to suspect that Bob just might be able to pull this off after all.³

³ Actually, I was impressed – but I still doubted the relevance of the null test to the actually driving of loudspeakers. Bob's imitation loudspeaker might not stress an amp or store energy and feed it back to an amp to nearly the same degree that our real reference loudspeakers would do. Plus, I had once upon a time picked up the differences between ½" of steel lead from a (continued)

We were still pretty confident that he couldn't, though. After all, 66 years of amplifier design have *still* not resulted in any way of pinning down the subjective effects of every measured imperfection -- even if we had measurements for them all, which we don't. The beauty of Bob's approach, however, was that he didn't need to know what all those objective imperfections were doing; all he had to do was eliminate them.

Neither LA nor I had any idea what "adjustments" would be involved, but I, for one, was convinced that the area that would ultimately stymie Bob was that of harmonic distortion content. I have long believed that some of the major sonic differences between amplifiers were related to the relative and absolute amplitudes of their harmonic distortion components. (It is known, for example, that the amplitude of the high-order harmonics -- the third, fourth, fifth, and sixth harmonics -- become progressively weakened in the signal from a tubed component, and remain relatively constant from a solid state device.) I was a little shaken when I learned that a half-dozen small potentiometers that Bob had wired into his amplifier were "distortion pots," which enabled him to change the amplitude of any "spurious" harmonic as desired, independently of the other harmonics!

That first listening test in Bob's room was an ear opener. He had already achieved a surprisingly effective null -- a 50 dB reduction below the level measured at each amplifier's output. But there was still a substantial amount of sound coming from the Rogers speaker, and that sound was some of the filthiest, dirtiest crud I have ever heard!! Bob explained that he had nulled out most of the things that both amplifiers were doing right, leaving only such things as distortion and frequency-response deviations. Yes, I thought, and those are going to be the hardest things of all to null out.

Bob explained that a 50 dB null meant that the difference between the two ampli-

fiers amounted to about 0.3% of the total output of each. The dramatic audibility of that 0.3% came about because he was driving the amplifiers to rather high output levels, and because of the ugly nature of what was left in the sound.

At this point we ran into a problem. The AC line voltage at La Posada was quite low, meandering around 106 volts much of the time. This would quite obviously throw off both amplifiers, enough so that they would probably not perform the same way with a more normal line voltage. I loaned Bob my variac.

The next day he had managed to boost that 50 dB figure to 70 dB, and felt ready to try some listening. By this time the difference signal between the amplifiers was audible only with an ear glued to the Rogers LS 3/5A, even with the output of the amps cranked up. There was no doubt that Bob had achieved something impressive, but we questioned whether it would translate into true duplication when driving real-world (and difficult-load) loudspeakers. We moved the project to my listening room.

The Listening Comparisons

The signal sources for our listening tests were to be both CDs and LPs. The CD player used was a Sony 520-ES, the analog player a SOTA Sapphire turntable with Well-Tempered Arm and Ortofon MC-2000 cartridge, with Ortofon's T-2000 step-up transformer. The preamp was a Conrad Johnson Premier Three.

Program sources were as follows, for the following specific sonic attributes: "The Portrait" and "Peter the Hermit," from *Growing Up in Hollywood Town* (Sheffield CD-13 and Lab 13) for depth and perspective, HF naturalness, bass heft and tightness; Respighi's *Church Windows* (Reference Recordings RR-15) for breadth, depth, bass range and control, and massed string tone; Beethoven & *Enesco Violin & Piano Sonatas* (Wilson Audio Specialties W-8315) for tonal accuracy, depth, and imaging specificity and stability; "Improvisations" by Jim Keltner, from *The Drum Record* (Sheffield CD-14/20) for high-end openness

capacitor to a crossover as opposed to 1/2" of copper lead -- and these two amplifiers had much bigger differences than 1/2" of wire.

& timbre and low-end attack, control and range; and McBride's "Mexican Rhapsody," from a badly worn copy of *Fiesta In Hi-Fi* (Mercury Living Presence SR90134) for treatment of HF stridency and mistracking.

We made no effort to do A/B testing, since we feel it does not replicate normal listening conditions, and there is still insubstantial evidence that A/B testing reveals small differences as well as does prolonged listening to each unit under test. In our tests one amplifier would be wired into the system and auditioned as long as we wanted, using a wide variety of program material that always, however, included the material listed above. Notes were made of anything we heard that we thought different from the other amplifier, and those specific points were checked again when we went back to the other amplifier,

A Good Beginning

We were not too surprised to find that there was no longer a dramatic difference between the 1.0 and the reference amp. In fact, what surprised us was just how similar they sounded. They were almost a perfect match, except for a slight difference in perceived depth and perspective, a marked difference in low-frequency range and control, and a noticeable difference in high end smoothness. We were pretty taken aback by the similarity, but, because the differences were reliably audible, we were still confident of our abilities to hear differences between the two amps. And, because the differences were important in type, though small in degree, the expensive reference amp was unthreatened. In spite of the really amazing feat he had pulled off so far, Bob was disappointed.

With 70 dB of null, he assured us, they should sound identical. They didn't; it was back to the test bench and soldering iron for Mr. Carver.

It took another day to find the source of the trouble and work on correcting it. The trouble, it seemed, came from my variac, which could not deliver enough current to meet the brief, but very high demands of the reference amplifier when playing [when

playing] music into demanding loudspeakers rather than mockups. Back at the hotel room, Bob had been trying to match his amp to one that was working with one hand tied behind its back.. The matching which had produced a 70 dB null in the hotel collapsed to 35 dB in my home, so it was necessary to produce a new model of the reference amplifier as it performed with adequate current availability. Fortunately, my line voltage was normal (115 V), so the variac could be dispensed with. Bob was discouraged at having to do his entire analysis and modeling over again, but glad of a problem concrete enough to be addressed.

A Second Stab

After another day, Bob seemed convinced that he had done it. We gathered for another listening session, and, indeed, it sounded as if he had. The high end stridency we had noticed in the 1.0 was gone (or, as it turned out, was just as present in the reference amp). Depth presentation, midrange solidity and 3-dimensionality, imaging, high end sweetness – in short, all the characteristics one normally finds important in amplifier evaluation – were identical. But, as we relaxed with a Sheffield jazz record, we thought we picked up a difference between the reference and the 1.0. With the reference, the low range of the guitar was a bit ill-defined; with the 1.0, you could "count the cycles." Granted, in this respect the 1.0 was *better* than the reference, but that was beside the point! We were looking for duplication. Bob reached into his bag of tricks and dropped the output of the 1.0 from 500 watts below 30 Hz, to a mere 65-100. Believe it or not, even though we were listening at subdued levels, that did the trick: the 1.0 was now a bit muddy and ill-defined through its lower range, just like the reference.

More interesting, though, and disturbing, was that the soundstaging had now changed, and the two amps were no longer the same. It turned out that Bob had to go back and diddle some more, exhausting his 48-hour limit.

The Final Achievement

After this last bit of tweaking, where Bob was able to reinstate his 70-dB null while driving a very difficult load, we now had what sounded like two absolutely identical amplifiers. No matter what speakers we used, every “difference” we thought we had isolated turned out to be there, in equal quantity, when we swapped amplifiers.

This time, the listening went on through the whole afternoon and much of the evening, until all of us were listened out. More leisurely listening, refreshed by a good night’s sleep, failed to turn up anything. As far as we could determine, through careful comparisons and nit-picking criticisms, the two amplifiers were, in fact, sonically identical. It is a gross understatement to say that we were flabbergasted!

The next morning, I told Dick Olsher over the phone what we had found. “Bullshit!” was his reasoned response. “That just can’t be.” But it was, it was then that we started to realize that, in reporting the outcome of this Challenge, we were going to have more to contend with than outrage and wonderment. We were going to have to contend with incredulity.

On the face of it, what Bob Carver pulled off should be impossible. You can’t make a silk purse from a sow’s ear. What about the audible differences between transistors, capacitors, internal wiring – all the things that we *know* contribute to the superiority of no-holds barred amplifiers? What about all the things that amplifier designers have learned during the past twenty years, which enable them to build better amplifiers (at whatever price) than have ever been built before? How could all of these things have been factored into the relatively quick and painless transformation of an average amplifier into a world-beater? But, of course, the “factoring-in” was the key to all this.

You see, Bob didn’t have to concern himself about quality capacitors, minimal internal wiring, gold connectors, or any of those things; all he needed to do was duplicate, at the output of his amplifier, the sum of their effects at the output of the

reference amp. Once he had obtained the necessarily deep null between those amplifiers, it was his belief that ears were not going to pick up on what was left. To do this, he needed only (!) to know how to change practically any parameter of his amplifier’s performance – a knowledge which we must now acknowledge is his.

After the second day of listening to his final design, we threw in the towel and conceded Bob the bout. He packed up his equipment and limped triumphantly back to his Lynnwood, WA home base. (He had single-handedly hoisted the hefty reference amp onto a table at one point during the proceedings and injured his back.) The question remains whether or not we might have eventually picked up some miniscule but repeatedly audible difference between the amplifiers, had we been able to listen longer.

Somehow I doubt it. We had thrown some of the most revealing tests that we know of at both amps, and they came through identically. Even on the subliminal level – the level at which you gradually get the feeling that one amplifier is more “comfortable” than another – we failed to sense a difference between the two amps.

It is true that there were no “controls” here – no double-blind precautions against prejudices of various kinds. But the lack of these controls should have, if anything, influenced the outcome in the other direction. We *wanted* Bob to fail. We *wanted* to hear a difference. Among other things, it would have reassured us that our ears really *are* among the best in the business, despite “70-dB nulls.”

There were times when we were sure that we *had* heard such a difference. But, I repeat, each time we’d put the other amplifier in, listen to the same musical passage again, and hear exactly the same thing. According to the rules of the game, Bob had won.

Disquieting Implications

The implications of all this are disquieting, to say the least. If, after only four days of work, it is possible for someone – design genius or not – to make a \$700

amplifier sound exactly like a state-of-the-art amplifier costing many times as much, what does that say for the cost-effectiveness of the latter?

Carver claims that the original, unmodified 1.0 amplifier had been designed to sound “the way he wanted it to.” If, in fact, he could make it sound any way he wished, as seemed to be proven with his success in this experiment, why then did he elect to go with a typical mid-fi “solid-state sound” instead of emulating the sound [of] one of the best-sounding solid-state or tubed amplifiers on the market? There were, it turns out, some good reasons.

Bob admits that he is not sure himself about the audible effects of some of the parameters he juggled to match the transfer functions of his amp to that of our reference. Had he been using this trimming technique to produce a certain desired combination of sonic qualities, using only his ears to evaluate what was going on, the task would have been quite a bit more difficult and time-consuming, the results far less predictable. This, in fact, is what he did with the 1.0 amplifier, which in his opinion still sounds excellent on the loudspeakers with which it will most likely be used (if not on the loudspeakers we used).

Secondly, Bob had never before had a chance to listen critically to a “world-class” amplifier like the one we chose as our reference, and ended up admitting that there were things about its sound that he preferred to his own amp. He might, he averred, “do some things differently in future designs.”

Does that mean that Carver Corporation might consider producing, commercially the modified 1.0 whose “sound” Bob had, quite literally, pirated from that state-of-the-art amplifier? Maybe yes, maybe no.

Is It Theft?

The possibility of Carver’s manufacturing his modified amplifier raises some very knotty questions concerning morality and legality. Does an amplifier manufacturer who designs something from scratch, coming up with a sound unique to that product,

have the exclusive right to that sound? In other words, is it dishonest or even illegal for someone to use a technique such as Carver’s transfer function analysis to duplicate that “unique” sound, without having done all the usual homework involved in designing an amplifier from scratch? There has never been a legal decision about this, but an analogy from computer software may shed some light.

Some years ago, a firm called Micro Pro started marketing the first automated spreadsheet for microcomputers. Called Visicalc, this program allowed a vast number of rows and columns of figures to be set up, by the user, to perform in mere seconds spreadsheet calculations that would have taken an accountant hundreds of hours to do with pencil and paper.

When Visicalc came out, there was nothing else like it. Within months, however, it was followed by the first of what soon became a flood of imitations, each capable of the same functions as Visicalc, but each using somewhat different ways of accomplishing the same end. Those “copycat” programs are still around, because the law deemed the functions which could be performed by Visicalc to be not copyrightable; only the specific program for accomplishing that function could be copyrighted. Thus, it is likely that Carver, or anyone else with his technical smarts, would be legally free to duplicate the sound of any amplifier, as long as different circuitry was used to do it.

But whatever Bob, and others who can match his technical virtuosity, choose to do with the results of this project, I think that the field of high-end audio amplifier manufacture will never be quite the same again. High price and high status will continue to be handmaidens in audio, but the knowledge that high performance and high price need no longer be inseparable cannot help but impair the glamo[u]r of cost-no-object power amps.

We’re still a little bewildered around here about how all this turned out. Not the way we expected. But that’s the way it was.