Stan’s Safari Part 3

(First published in Hi-Fi Critic in 2008)

After finishing my article for the last issue I thought I’d put to rest the ghosts of valve amplifiers past but it was not to be. In the days that followed I found myself digging out a number of old work-books where I was startled to discover just how many such amplifiers I had designed in my younger days. One design in particular stood out. It comprised of a huge radio transmitter valve mounted behind a modified “Vent-Axia” kitchen fan which attempted to counteract the heating effects of the several amperes of current flowing through this single ended Class A stage. The rest of this ludicrously simple amplifier comprised of no more than a substantial output transformer; a single driver stage valve and a handful of passive components. And that my friends was it. With an output comfortably over 100 watts it sounded simply magnificent to my ears and was the best valve amplifier I ever built. Maybe not very carbon friendly these days but I was certainly up for building another pair until I found I hadn’t retained the transformer design notes. And so it came to pass that once again I was able to avoid the dark side and return to the slightly more sane world of solid state amplifiers.

Over the next few pages and beyond I will try to take the reader through some of the problem areas of transistor amplifiers and the considerations that enter the designer’s mind. What I am not going to do is to give a step-by-step lesson in amplifier design as this will go over the heads of many readers and probably generate too many challenges from armchair theoreticians for my liking. What I will highlight will be some of the things that can screw up the “perfect” transmission of a music signal through an amplifier whilst asking the question “do they screw up the signal enough to be a concern?” Time and time again in the history of hi-fi a problem has been found which may have the capability to degrade the signal and that is all that is needed by some designers who then go off on a quest of outstanding single-mindedness to bring us a product completely free of that particular ill. Yet perhaps nobody else was actually finding that particular problem to be anything that interfered with their listening pleasure. For me the art of amplifier design is to look at all the areas of potential degradation; establish which are relevant to the product brief and then reduce them to a level where they cease to be noticed thereby avoiding the costs needed to produce a notionally “perfect” amplifier. Other than to satisfy engineering ego there is absolutely no point in reducing a distortion or colouration below the level at which it creates an audible effect if it is going to add cost or complexity to the finished product.

When I designed Rotel’s products I termed this philosophy the “Balanced Design Concept” and it is a concept I still adhere to. Some say it is all that remains of a 1960s hippy approach of doing no more work than is necessary but I prefer to see it as a Yorkshireman’s pride in spending no more brass than is necessary.

The majority of us own transistor amplifiers and some of us actually believe the stories of design originality spun by the manufacturers. Yet in many respects audio amplifier design has largely stood still over the last 40 years. Now I know this is an extraordinary statement, but consider the facts. The earliest transistor amplifiers used inter-stage transformers, had very limited output powers and were none too reliable. Then in 1956 Dr. Hung Chang Lin of RCA laboratories published a simple and elegant circuit of a directly coupled quasi-complementary amplifier; an amplifier which can justifiably claim to be the grandfather of almost every one made today. (Incidentally RCA continued to help drive the design of audio amplifiers, notably producing a seminal booklet in the early 1960s, the “RCA High Fidelity Amplifier Circuits Manual” as I recall and Julian Verecker used the contents as a basis for the first Naim amplifiers. He was not alone in finding inspiration in those pages.) For some years most transistor designs followed a similar topology to the Lin design; an early example in the UK being the singularly awful sounding Leak Stereo 30. Then at some stage in the mid 1960s a variant appeared with the single transistor input stage being replaced by a differential stage; the classic “long-tailed pair” of transistors.

I’m not sure which designer made the first such audio amplifier but I never saw one before the mid-sixties and I’m pretty sure that when they appeared they were based upon the circuit topology of Bob Widlar’s 1965 Fairchild μA709 integrated circuit. Wildar was a circuit design genius and his 709 product and the later ubiquitous 741 model actually incorporated most of the design structure (the “Operational Amplifier” topology) which you’ll find in today’s amplifier designs. For sure there have been detail improvements such as the many possible alternative output stages; fully complementary input and driver stages; alternative output devices (Mosfets etc.) but essentially the core design of most amplifiers is the same. Now obviously you can view this outcome in two ways. Either the following two or three generations of designers have conspicuously failed to come up with anything new and original or the early Wildar design got it so right that those following have simply needed to refine and develop it year on year to achieve better and better performance.

At first glance the requirements made of an audio power amplifier are not too demanding. It must amplify in a linear fashion keeping distortion to low enough levels; it must deliver enough volts and amps to drive the loudspeaker load to the required level; it must be stable under all normal conditions; it must have a wide enough bandwidth to deliver audio music signals in their entirety and it must be reliable. With the possible exception of the latter it was the case that the designers of the 1970s felt they were fulfilling all the above requirements yet it was apparent to many listeners that different amplifier models produced different sounds; often massively different! And so began an acrimonious conflict between the “subjectivist – how does it sound” engineers and the “Ohm’s Law – how does it measure” engineers; a conflict which still rumbles on today to judge from the often offensive letters which appear in “Electronics World” the successor to the much beloved “Wireless World” of my youth.

If we are to start anywhere why not the output stage which control the flow of current from the power supply to the loudspeaker. If we look at the behaviour of low level signals we soon come across crossover or switching distortion. With a conventional “Class B” push-pull output stage one transistor is conducting and the other is not so when the signal passes through the “zero” point there is a dead spot until the signal rises to a level where the other transistor starts to conduct current. In this region transistors are very non-linear so a high level of distortion is produced and it is rich in the high order harmonics which are quite audible and unpleasantly so. Furthermore because there is no current flow there is no effective negative feedback; feedback that would normally reduce the distortion to an acceptable level. There is, of course, a very simple cure to this problem and that is to operate the output stage in Class A where a high standing current maintains all the output transistors in their linear operating regions. Unfortunately such amplifiers typically need to dissipate twice the power of their rated output making them expensive to run and impractical in the summer months. (The figures stack up even worse for single ended Class A designs such as the valve amplifier mentioned earlier. The standing dissipation rises to four times the rated output so a pair of my glorious amplifiers become a reasonably effective 1 Kw fan heater.)

Over the years endless candles have been burned by engineers trying to solve the problem of how to get near to Class A performance whilst keeping the standing current acceptably low. These have ranged from Nelson Pass’s dynamically biased output stage in the original Threshold amplifiers to the Quad “Current Dumper” with numerous other ideas in between. Yet most designers have simply increased the standing current in the output transistors (so-called Class AB) in the belief that there is now a continuous current flow and hence no “dead zone”. Maybe so but what they do have is two transistors now operating in a really non-linear part of their operating curve and the distortion is worse than it was. Such distortion is rich in high harmonics but it can be reduced by negative feedback except that in the great majority of amplifiers the amount of feedback available decreases with frequency. In fact what usually happens is that the signal distortion (and the “nastiness” of the distortion) rises as the signal increases in amplitude until it reaches a level where the output stage is linear when the distortion falls away.

Tinkering with the quiescent or standing current moves this zone around a bit but doesn't actually remove it. So here we have our first common assumption in amplifier design whose validity is certainly open for question and there are plenty of others to follow.

Incidentally you don't need a bank of test gear to evaluate the performance of different amplifiers operating at low signal levels. Simply connect up the amplifier to the loudspeakers; set the volume to almost inaudible levels then also connect the output wires (via a wideband isolating transformer) to the inputs of another amplifier whose volume is wound up to a decent listening level. Then you can play around with level setting on the amplifier under test and hear, in many cases, just how bad the music sounds. In some cases it sounds very bad but then some readers may say "I only listen to house music where the playback level is 120 dB from the off. Why should I need to worry about very low level distortion?" Well here we come to value judgements. Should the designer reduce the distortion to zero; reduce it to about the same level as the residual noise or not worry too much about it? It is the designer's call and ultimately you will judge if he made the right call.

Now let's look at the other extreme, the handling of high level signals. Early amplifiers were notoriously poor at driving low impedance loudspeaker loads, a classic example being the original Quad 405 whose protection circuits would spring into action whenever the volume level rose above string quartet levels. The root difficulty was that the early power transistors could not carry high currents; a problem given that the output current doubles with a halving of the load impedance. The solution was to use loads of output transistors in parallel but the manufacturing cost of the product then rose rapidly; more transistors; bigger heat-sinks; more labour. There was a legitimate problem which was easily demonstrated by measuring the maximum current output of an amplifier usually in the form of a pulse into a very low value resistor. I was still able to measure amplifiers in the 1990s which could drive no more than 3 or 4 amps into a load and that really isn't enough. The main audible effect is that of compression; the loud notes go so loud and no more whatever you do with the volume control. There are also a number of more subtle effects but the compression was the one that everyone could hear. Fortunately the demands of the pro-audio market for extremely high power amplifiers led to the availability of lots of low-cost high power transistors. And so began another of the hi-fi industry's specification wars with manufacturers claiming maximum current outputs of 100 amps or more. Sometimes the manufacturers are missing the point. How clean is that current delivery? How linear is the amplifier at such high output levels? Is the current delivered "instantly" or does it trundle down the wires with all the urgency of a watched kettle?

Amplifier distortion rises as the output current rises. Some designers tackle this problem by adding extra power transistors on the basis of a theory that the current sharing will result in a lower current through each transistor and thus the linearity will improve. Except all those output devices have to be driven and the driver stage is still carrying the same load. The old, old problem; you change one thing and you change everything. Personally I'm of the opinion that the high current demands of the load have to be delivery virtually instantly and in a linear fashion. So rather than focus simply on maximum current output I do put effort into ensuring the current has a clean and linear path from the power-supply to the loudspeaker by, for example, placing the reservoir capacitors next to the output transistors and by having a very low resistance path from the transistor pins to the output terminal. We'll hopefully come back to this interesting topic in a later issue.

The final consideration for this issue is the thermal behaviour of the amplifier. At the initial theoretical stage of the design the matter of temperature hardly merits consideration. The heat-sink is designed to do its job then made slightly smaller to save space and money and the design caravan moves on into the night. But temperature rises do things to transistors. Too much heat and they suffer thermal runaway and blow up. But at lower extremes their specifications such as current gain; speed; and junction voltage all change with temperature. These changes are particularly problematical in the output stage because the transistor chip is usually small and inside a casing which is then bolted to a heat sink which is designed to dissipate the heat. But in practice a short period of high current flow will lift the chip to a very high temperature whilst the heatsink and the transistor case are still only luke-warm to the touch. In the meantime the performance of the amplifier can have changed. The bias points may have changed and with it the quiescent current settings. The open-loop bandwidth may have changed and with it the amount of negative feedback available; and so on and so forth. Has it changed enough for the effect to be significant? Again there seems to be two prevalent attitudes. Most designers ignore the effect whilst one or two others go on about the insidious effects of "thermal distortion" where even the changes of signal level though a humble resistor are said to change its temperature and so its resistance value and hence the performance of the circuit. To such designers I'm tempted to say, "Get a life" because these temperature effects, where they matter, can be minimised and where they don't matter they can be safely ignored.

I designed a high-end power amplifier not so many years ago and painstakingly dismantled a number of output transistor types to look at the chip themselves. I chose the device with the largest substrate and then went to some lengths to get the heat from the substrate to the heat-sink as efficiently as possible. I was then able to measure a clear improvement in the amplifier's thermal stability. Did it make a difference? Well I like to think so but only because I'd previously expended a lot of effort in reducing all other known forms of "distortion" to negligible levels.

We've hardly touched the surface of amplifier design yet already we've got plenty of things to consider and you know what makes it really interesting? You improve one aspect of an amplifier's performance and there's a good chance that another aspect will get worse. Back to that balanced design again. You'd be forgiven for wondering how most designers get their amplifier designs to work. Actually it's not too difficult because the bog-standard amplifier circuit is amazingly tolerant and flexible and will usually give an acceptable performance even if the designer gets a few things wrong. In fact I'd go so far as to say that there are plenty of designers out there who don't really know what they are doing but who have discovered that a little bit of tinkering can go a long way.

Of course so many of the problems including that of thermal behaviour don't need to be considered with a properly designed Class A amplifier because, for example, the output stage current is a constant regardless of what the signal is doing and after an initial warm-up period the temperature of all the components reaches some sort of equilibrium from which it never shifts. This is a refrain you are likely to hear time and time again throughout this series of articles; it's not a problem with Class A. But then we all know that such amplifiers just aren't practical. The're big; can run excessively hot; are inefficient; and environmentally unsound. They deserve to go the way of coal fired generators. But try as I might I do love coal fires and I've got quite a soft spot for Class A amplifiers.