Thermal transient variation of power amp quiescent current
Instrumentation and findings

Jan Didden

Introduction
Every class-AB power amplifier needs a bias circuit that ensures a small but important quiescent or bias current through the output stage, even in the absence of a signal. There are two main (conflicting) requirements for the value of the bias current. The first one is that the bias current must be sufficient to ensure a smooth handover from one output device to the other when the signal current traverses through zero. Insufficient quiescent current would aggravate the dreaded cross-over distortion. The other requirement is that the bias current should be as low as possible to increase the amplifier efficiency and minimize dissipation losses. **Figure 1** shows a generic topology.

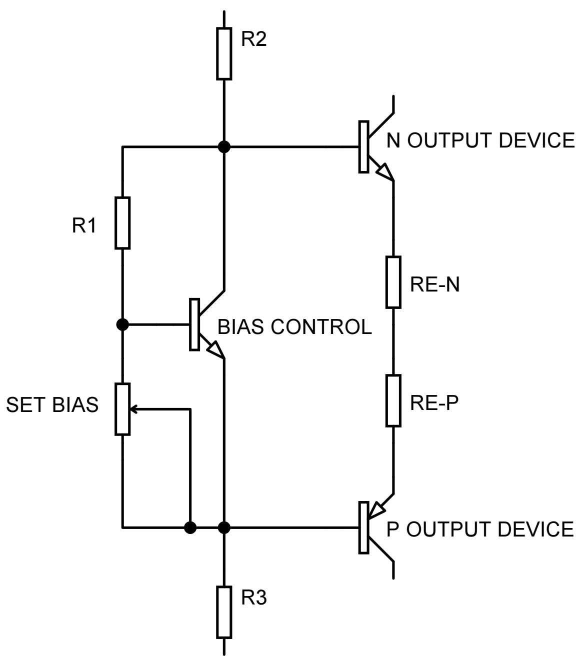

The bias set voltage must typically be set to a value of several Vbe thresholds in the output stage, plus the voltage across any output stage emitter resistors. Because of the spread in the Vbe values between devices, the bias current is normally adjusted when the amplifier is first switched on. However, when the devices in the output stage heat up with dissipation, the Vbe threshold required for a set bias current will be decreasing, and bias current will increase unless measures are taken. The inclusion of output device emitter resistors will help to stabilize the bias current with temperature as they provide a measure of DC local feedback. The value of these resistors must be kept low for output efficiency so their ability to sta-

Figure 1 Generic output stage biasing

Photo 1 – a representative example of thermal bias feedback measurements on build

bilize the bias current is limited.
Another stabilization measure generally taken includes some form of temperature feedback, for instance by physically mounting the bias generator device or a temperature sensitive device that is part of the bias voltage generator, isothermal to the output devices. For instance, a small TO126 transistor may be mounted on the output stage heatsink, close to the power device(s), or even bolted directly to an output device. **Photo 1** shows another representative example where an NTC resistor that is part of the bias circuit is bolted onto the case of one of the output devices with the same bolt that mounts the device on the heatsink. By suitable design of the bias control circuit, excellent long-term static bias current stability may be obtained.
There is, however, always a certain thermal lag between changes in the output device die temperature and the reaction of the bias control circuit. In fact, the thermal compensation loop is nothing more than a negative feedback loop, with a transfer function that may or may not be stable, exhibit lag, overshoot, etcetera. It is generally not easy to quantify the factors of this feedback loop and its behavior is generally not known with enough detail to make statements on the transient behavior of the loop. There have been several reports in the literature on thermal transient behaviour of bias control stages ([1], [2], [3]), and some simulation studies ([4], [5]) have also been performed. But actual output stage measurements have been sorely lacking.

In the first part of this article I will introduce a purpose-build piece of test equipment ('bTest'), designed to measure the actual bias current flowing through a class-AB output stage, under dynamic conditions[1]; dynamic conditions being defined as reproducing varying signal levels in a dummy load to reasonably mimic music reproduction conditions. I will discuss the design concept, the considerations for the selected topology, implementation and some representative measurements.
In part 2 I will present some measurements regarding thermal transients in so-called ThermalTrak transistors, and how these transients propagate to the on-die thermal sensing diodes.
In part 3 I will look at the distortion behaviour of a representative amplifier, under modulation of the bias spreader voltage, to emulate the bias variations as a result of the thermal transients as shown in part 1. I will finish with some conclusions.

[1] This article does not concern itself with the long-term stability of the bias current, where long-term is defined as a period exceeding several 10's of seconds. This issue has been addressed extensively by others, for instance by Cordell [4] and Self [5].

Part 1 – Bias current thermal transient variation

Measuring bias current

To judge the transient thermal behaviour of the bias control loop one could compare the actual bias current value with the actual temperature of the output device junction. The junction temperature is dependent on the power dissipation, and thus on the power output of the stage, which in turn is dependent on the output signal driving the speaker. We are not in itself interested in junction temperature; what we are interested in is how the set bias current value varies (or not) under varying output levels. The aim is a constant bias current whatever the output stage dissipation. Therefore, comparing the instantaneous bias current with the instantaneous output device dissipation on a suitable time scale should show no deviation from the set bias current value. The bTest equipment does not measure the output device dissipation; rather, the variations, if any, of the bias current are shown with respect to the output signal level in a resistive load which is of course a measure of device dissipation.

A suitable way to measure bias current is to measure the voltage across the output stage emitter resistors, and scaling the voltage for indication in mA. However, since the *output current* also flows through one of the emitter resistors, depending on output signal polarity, some way has to be devised to separate the output current from the bias current measurement. The problem is aggravated by the dynamic range of the currents involved; while bias currents might be as low as 20mA, output currents can exceed 10A. One option would be to measure the voltage across the emitter resistors and subtract a voltage representing the instantaneous output current. Such a compensating voltage could be measured across a small resistor in series with the load. The requirements in terms of accurate gain- and subtraction circuits, both for amplitude as well as phase matching are not trivial, however. Therefore, another concept was selected.

The requirement to compensate for the load current can be waived if the bias current is not measured continuously but only at the instants where the load current is zero. At those instants, the voltage across the emitter resistors accurately represents the bias current value. Since music and test frequencies always have zero crossings, the task then is to detect a zero crossing, sample the emitter resistor voltage and scale it suitably for readout. In this fashion, a cycle-by-cycle bias current value is obtained which should be sufficient to judge the transient thermal behaviour of the output stage. The main difficulty here is the consistent detection of the zero crossing of the signal current in view of large variations in signal level, frequency and slew rate, and the non-zero sampling time required to sample the emitter resistor voltage. Even if a zero crossing can be reliably and consistently detected, and a sampling event started, the output current may no longer be zero at the end of the sampling interval.

Bias current measurement test set

The solution I chose is to wait for a zero crossing and then forcefully maintain the output current at zero until the end of the sampling interval. This can be done with a high-speed solid state power relay, which disconnects the load from the amplifier at the start of the sampling. The relay is released again when sampling is terminated. Conceptually this is shown in **Figure 2**.

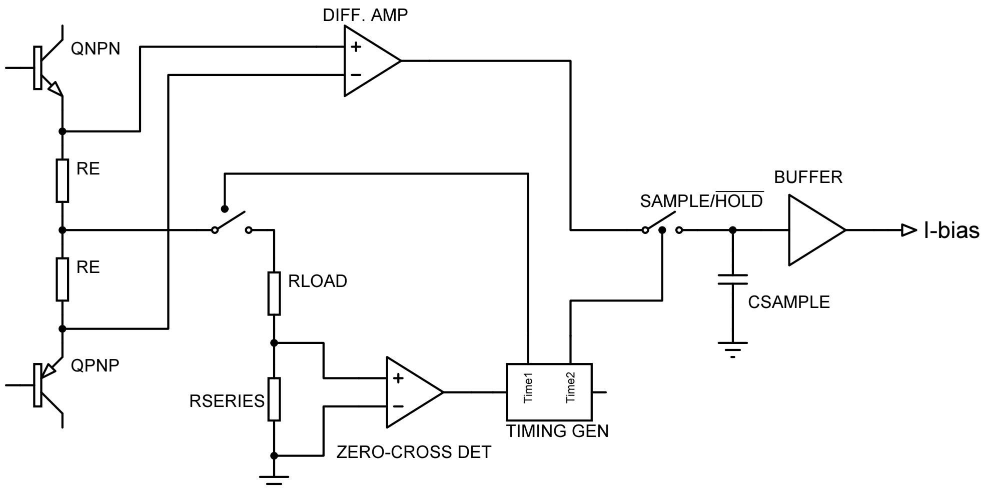

Figure 2 – Conceptual diagram of bTest dynamic bias test set.

Differential sense amplifier

The amplifier sensing the voltage across the emitter resistors is formed by several low-noise differential pre-amplification stages in series, referenced to the amplifier output voltage (the mid-point between the two emitter resistors). The preamp stages are shown in **figure 3**. J2 connects to the amplifier output, while J1 connects to the top and bottom of the Re string. The preamp stages also have gain switching provisions; one switch allows selecting for the emitter resistor value, providing normalization to 1 ohm. The other gain switching provides for switching the sensitivity in mA per Volt test output signal.

This is followed by a high-common-mode-rejection instrumentation amplifier to refer the sensed voltage to system ground, as shown in **figure 4**.

Sample/hold and buffer

The sample command for the sample & hold amplifier is provided by the zero crossing and timing circuit. The S&H is based on the LM198 integrated sample and hold IC. For accurate sampling, the sample period should be several microseconds, depending on the sample capacitor C3. The latter is a low-leakage 1000pF mica capacitor, which is a good compromise between fast charging during sam-

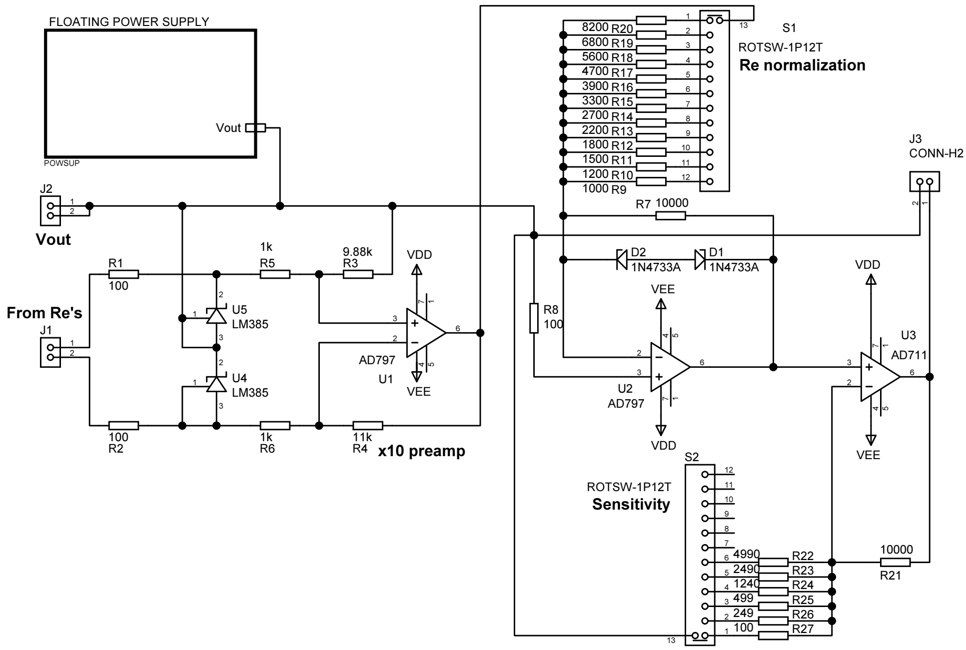

Fig 3 bTest preamp, normalizing and scaling circuitry.

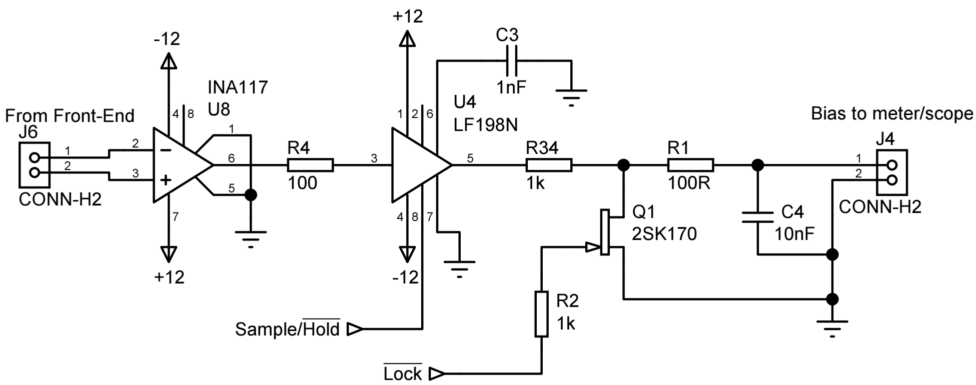

Fig 4 bTest ground-referencing differential amplifier, sample & hold and output circuitry.

pling, and low droop during hold. The charge on the sample capacitor is buffered by an internal buffer stage, and low-pass filtered by R34+R1 and C4. The output is connected to a BNC connector for metering and/or viewing. The transient variation of the bias current can be visualized with an oscilloscope. FET switch Q1 is switched on when there is no zero crossing detected for some time. This will reset the measurement output to zero to prevent ambiguous readings when no sampling is occurring.

Zero-crossing detector

The zero-crossing detector is based on a special very-low-level comparator, designed to switch fast and clean on low level signals, the ALD2331, and the circuit is shown in **figure 5**. One section is used for the zero-crossing detection. J6 connects to the Iout sense resistor in series with the amplifier load resistor. This input signal is amplitude-limited to prevent damaging the comparator. The comparator output signal is clamped for negative values and is send to the timing generator microprocessor connected to J5.

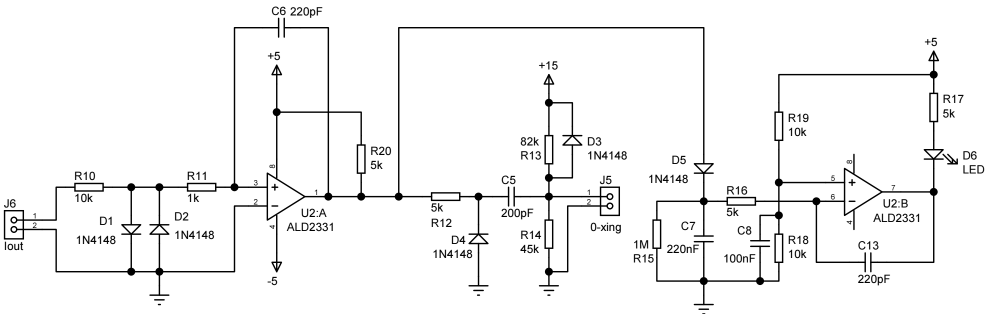

Figure 5 Zero-crossing detector and indicator.

The second half of the dual comparator is used to drive an indicator. The indicator will be ON during normal operation. However, when there are no zero crossings during a time set by the time constant of C7 and R15, C7 will discharge and the LED will extinguish, signaling that sampling the bias value is not taking place. This could for instance be the case when the output current is not sufficiently high to reliably detect zero-crossings. As mentioned before, the microcontroller will also reset the bias output value through Q1 in figure 4 to avoid ambiguous readings[2].

Timing generator

The timing generator is based on a small microcontroller with a simple firmware. The program waits until a zero-crossing is detected, and then opens the solid-state relay. It then sends out a sampling command of about 5uS. After sampling is terminated, the solid-state relay is closed again and a lock-out delay of about 5mS is started. After the lockout delay has expired, the controller waits again for a zero-crossing and the process starts again. An internal timing loop will switch on Q1 (figure 4) if no zero crossing has been detected for approximately 350 milliseconds.

Solid State relay

This is an adaption of a circuit published at diyaudio.com by user 'Chocoholic' [6]. The relay uses both a floating power supply as well as a floating MOSFET driver. It is powered by a miniature floating 5V switching supply and switched by a ground-referred control signal, yet is fully floating. In this application it is connected directly to the output of the amplifier under test. **Figure 6** shows the circuit.

[2] Of course, the microcontroller could also be used to signal the 'sampling' LED just as it resets JFET Q1 – but the use of the 2nd half of the comparator for this is a remnant from an early prototype and I have left it as it was.

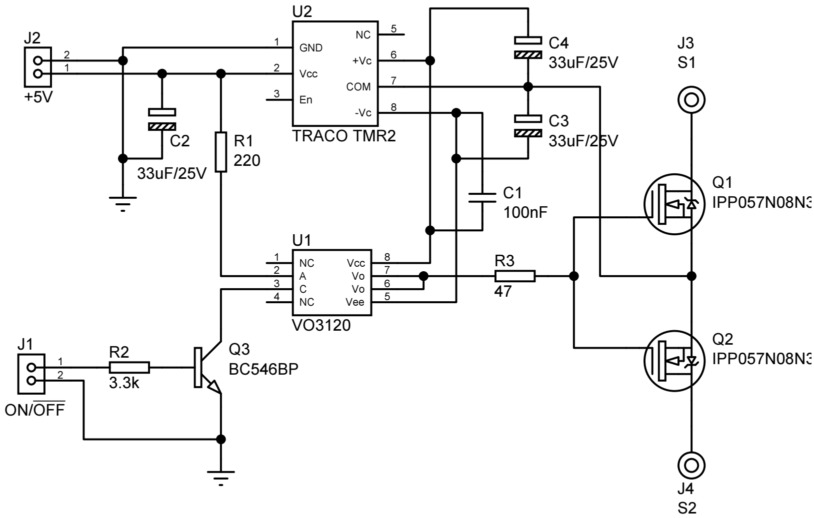

Figure 6 Floating high speed solid state relay with ground-referenced control.

Power supplies

The preamp stages that measure the voltage across the emitter resistors referred to the amp output voltage will swing many 10's of volts while their input signals are in the mV range. There are opamps available that can be fed from more than +/-50V but these fall short in terms of noise and speed. Therefore, the power supply voltage for those preamp stages is provided by an isolated +/-15V DC-DC converter, with its center tap connected to the amplifier output terminal. A small resistor (10 ohms) is used to isolate the (small) parasitic capacitance of the DC-DC converter from the amp output.
Although not necessarily required, the rest of the equipment is also supplied from a similar DC-DC converter referenced to ground, so the total test set can be fed from a small 5VDC wall wart.

The complete unit was built over several iterations in a 1U 19 inch case, the contents of which itself is in some sort of transient state as well, see **photo 2**.

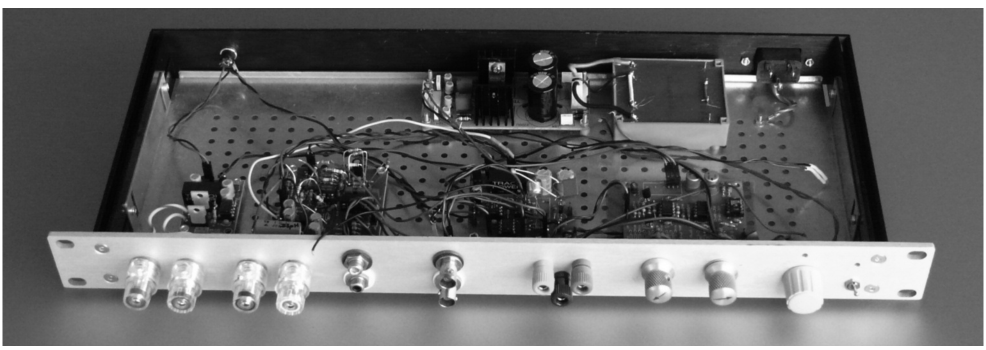

Photo 2 Prototype measurement unit 'bTest'.

Measurement results – TO-3

Most measurements were done on a vintage SONY TA-N86 power amplifier. This amp has one pair of BJT output devices in a metal TO-3 can, which are bolted to a heatsink, one for each channel at each side of the enclosure, as shown before in photo 1. bTest was set up to provide an output of 100mV/mA; the bias current was adjusted, with no signal after a suitable warm-up period, to 60mA. The bTest nominal output voltage thus was 6VDC.

Looking at the bias current output DC-coupled doesn't show anything of interest as the small transient variations are a few percent of the nominal value only. We need to look at the AC-coupled output, which shows the deviation from the nominal 60mA/6VDC value.

Figure 7 shows the variation in bias current with an output signal that is switched between two different levels. It is clear that the different signal levels, causing different dissipation levels in the out-

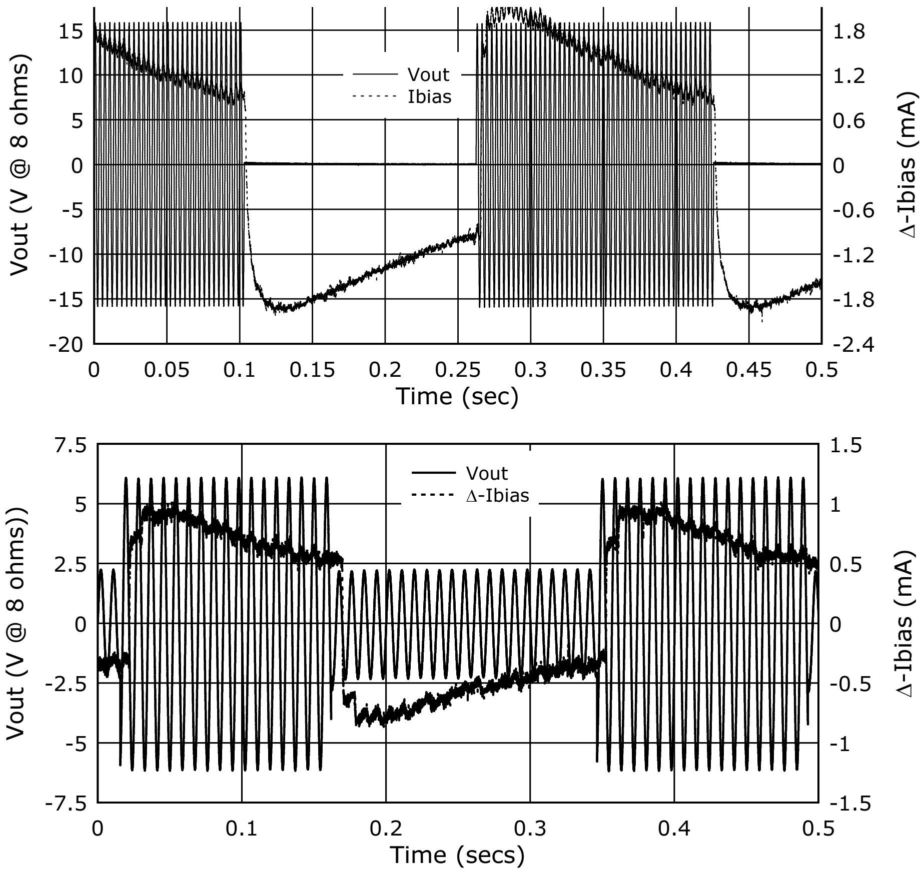

Figure 7 a, b – Bias current variation with output level for different amplitude shift keyed output signal levels, 8 ohm resistive load.

put stage, also cause a deviation of the bias current. Since the measured variation of the emitter resistor voltage is only a fraction of a mV, the test output signal is somewhat noisy, but the variation can clearly be seen and interpreted. What is interesting is that the output device die seems to heat up very quickly as a result of a dissipation change. It should be noted that sampling takes place at a falling signal zero crossing. This makes it look as if the bias variation is delayed by one cycle. Although the average bias current changes slowly over time, small-scale transient changes of up to 2mA peak (2 to 3%) are seen.

You can see that the initial transient slowly returns to an average level, as the rising temperature propagates through the device to the NTC sensor (photo 1) and to the bias control circuit over the course of several 100 milliseconds.

Figure 8 shows a similar result but now with an output signal level that is modulated by a low-frequency sine wave. Again, the instantaneous changes in bias current are clearly seen to follow the dissipation envelope.

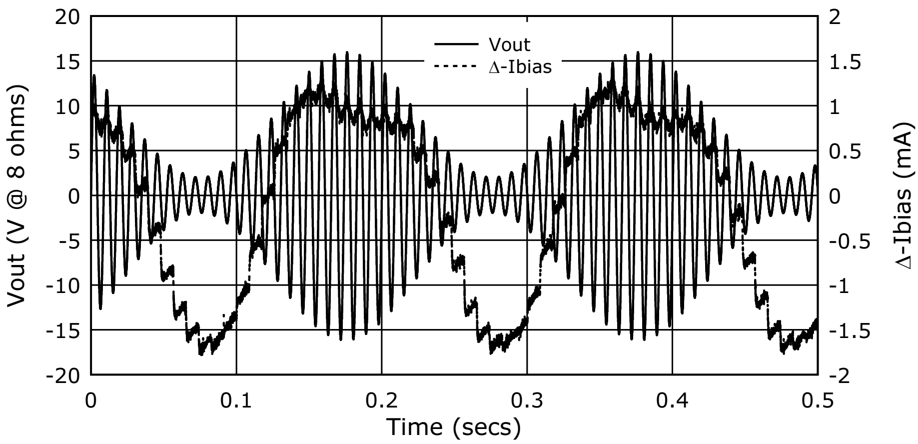

Figure 8 Bias current variation with output level, AM modulation, 8 ohm resistive load.

Measurement results – TO-247

To get a feel for the influence of the output device heat sinking on the transient bias variation, the metal TO-3 devices in photo 1 were replaced with a plastic TO247 device (ML3280) as shown in **photo 3**. Note that there is no heat sinking at all.

The resulting graph is shown in **figure 9**.

The thermal modulation is similar in shape to the case of the fully heat sunk TO-3 from figure 8 and the peak-to-peak bias variation is even somewhat smaller, which appears counter-intuitive. However, if one accepts that the transient thermal modulation is the result of locally heating of the die and its immediate thermal mass, it could make sense. The thermal mass of a die covered by the volume

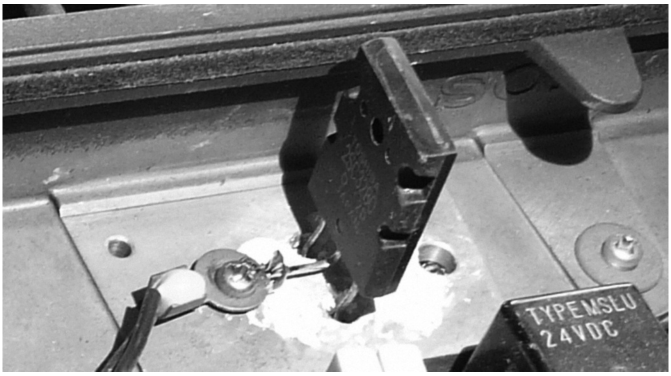

Photo 3 Replacing the metal TO-3 output devices with plastic TO-247 with no heat sinking.

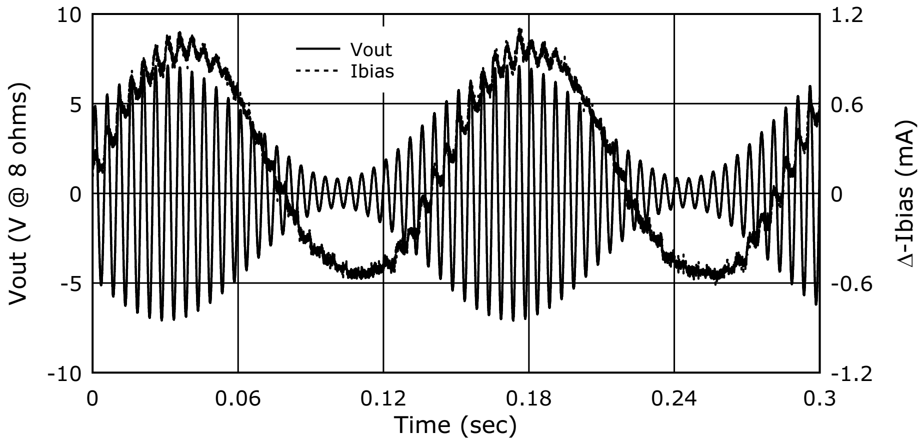

Figure 9 Thermal modulation with TO-247 plastic output devices, no heat sink.

of a plastic TO-247 case might be larger than the thermal mass of a die in the otherwise empty space of a metal TO-3 case. Of course, the measurements on the non-heat-sinked TO247 had to be performed quickly as the lack of heat sinking caused an appreciable and potentially dangerous runaway of the nominal bias current.

Measurement results – single-ended class-A

The next test was done on a small single-ended MOSFET class A output stage running with 1.2A bias, **figure 10**.

This is a SE amplifier with a unipolar power supply; the output load was 8 ohms resistive in series with a 3300uF DC blocking capacitor. Hence there is a phase shift between the output voltage and the output current, especially at the low modulation frequencies used in these measurements. Therefore the bias modulation, which results from the dissipation modulation, leads the output voltage. Output devices are TO-247 MOSFETs on a heatsink.

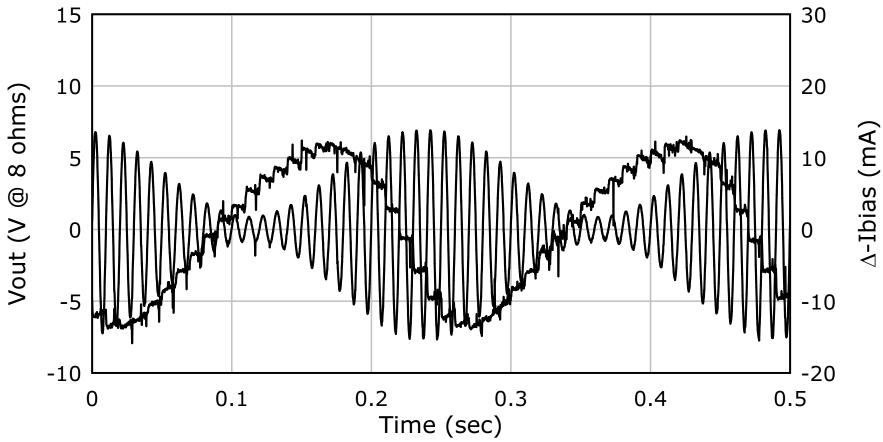

Figure 10 Transient bias modulation in a class A output stage, capacitively coupled with 3300uF to an 8 ohms resistive load.

Interestingly, the absolute value of the transient bias modulation is much larger than in the previous example at some 28mA pk-pk. However, the *relative* variation is quite similar to the previous cases at between 2 and 3%.

Measurement results – Sanken ThermalTrak devices

Finally I measured my paX amplifier [7] which uses Sanken 'ThermalTrak' output devices. These are Darlington output devices which have several diodes in the same TO-247-type case as the output devices; these diodes are a part of the bias current control loop. The general idea is that this concept gives much tighter control of the bias current because of the intimate contact of the diodes with the output device die. As **figure 11** shows, the thermal transient is smaller at 1mA pk-pk than for instance

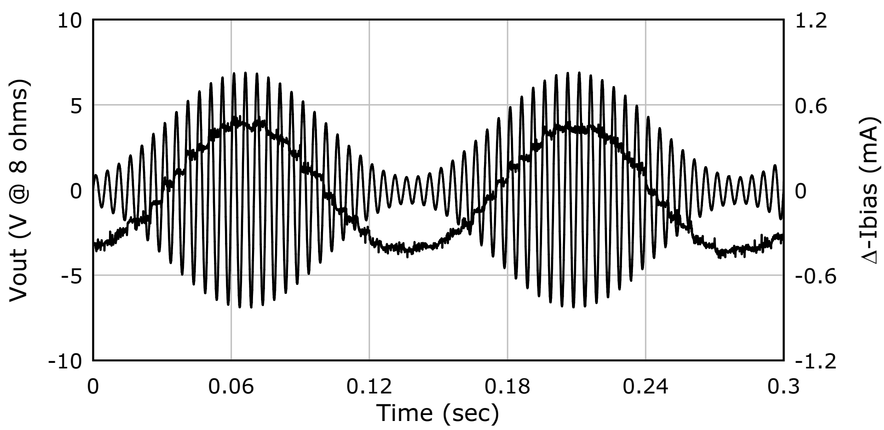

Figure 11 – transient bias modulation in 'ThermalTrak' output devices (Sanken STD03N and P).

as shown in figure 9 at 1.6mA pk-pk, but is still appreciable. Note also that because the ThermalTrak output stage has a much lower nominal bias current of 35mA versus the 60mA of the TA-86N (figure 9), the *percentage* variation is still very similar. The integrated diodes do make a difference for long-term stability but apparently cannot prevent the phenomenon of transient thermal bias modulation.

Conclusions from part 1

A test set has been presented that allows measurement of power amplifier bias current variations under transient conditions. It has been shown that the bias current in class AB and class A amplifiers reacts rather quickly to transient changes in output stage dissipation, within 10's of milliseconds. Although over a longer period, static bias current will remain reasonably constant in a well-designed thermal control loop, transient changes of up to 3% of the set value were observed.
The changes appear independent of the actual heat-sinking of the output devices, but appear to depend on the thermal mass of the device itself. So-called 'ThermalTrak' output devices show largely similar percentage variations in bias current, and do not seem to be effective to prevent thermal transient variations in bias current.

Part 2 – ThermalTrak diode thermal transient behaviour

The results of part 1 of this article, in particular with the ThermalTrak transistors, made me look for a way to find out how fast the diodes on the die of a ThermalTrak transistor react to changes in the die temperature. For this purpose I build up a small test jig shown in **figure 12**
Conceptually, the difference in diode voltage between two ThermalTrak power devices (MJL1302D) is compared with a gain-of-12 amplifier formed by an NE5534 opamp (U2). The difference is low-pass filtered and can be viewed on a 'scope. Both TT-transistors are mounted isothermal on a large

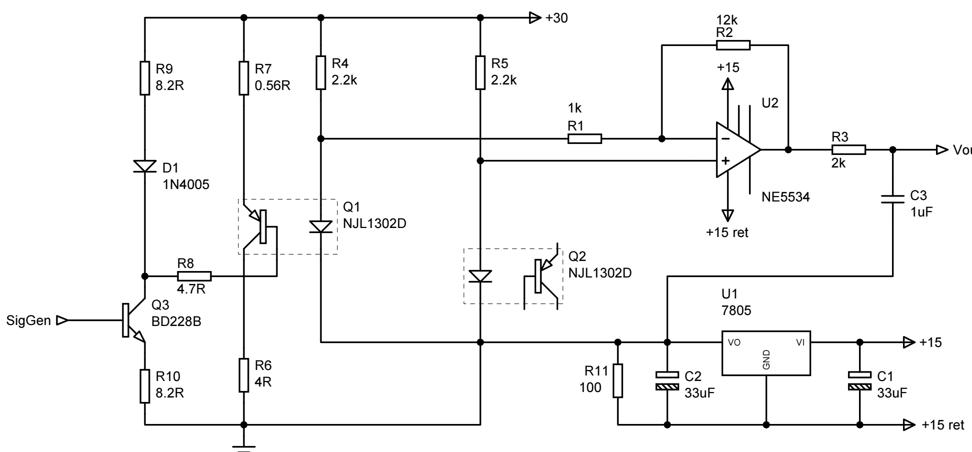

Figure 12 Test jig for measuring ThermalTrak diode thermal transient behaviour.

heatsink, and both TT-diodes carry the same nominal current; therefor, they will have substantially the same diode voltage.
For the measurement TT-device Q1 is repeatedly switched on and off into a load of 4 ohms (R6). This modulates the temperature of Q1's die and the junction temperature of the on-die TT-diode. This should result in differences in Vdiode between the dissipation-modulated Q1 diode and the fixed-temperature diode of Q2.

Measurement results

Figure 13 shows the results from modulating Q1's dissipation at a 5 second on-off rate, between zero and about 33W watts dissipation.

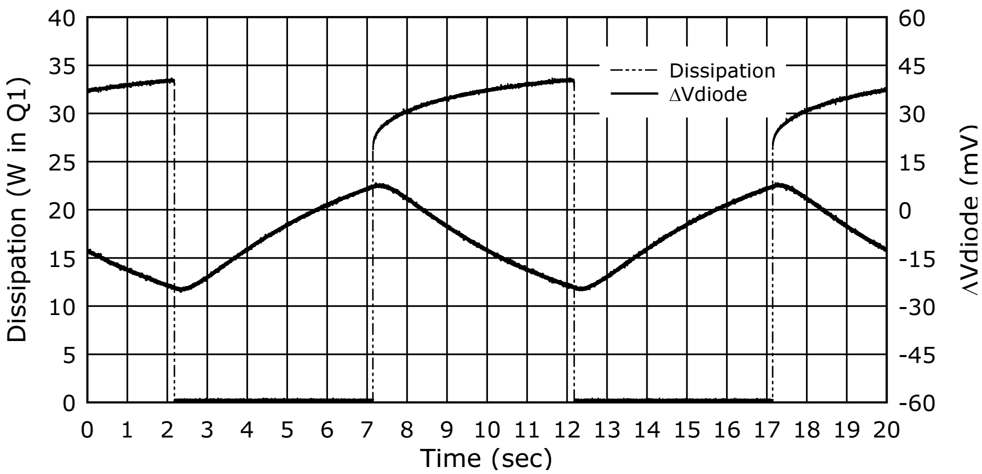

Figure 13 Change in on-die diode threshold voltage resulting from output device thermal transients at 50% duty cycle.

Clearly, Vdiode follows the dissipation envelope, Vdiode dropping steadily under dissipation, while rising again after the dissipation is switched off. The variation is relatively small, about 6% on a nominal 550mV with the diode current shown in figure 12. What I find surprising is that the diode voltage continues to change over a 5-second period. That means that the changes in bias current shown in part 1, which are manifest after a fraction of a second, are not instantly compensated for.

To look more into that time constant, the test was repeated with a 10% duty cycle modulation, as shown in **figure 14**. Although now the dissipation is zero for a period of nine seconds, Vdiode is still not fully at equilibrium at the end of that period and continues to rise. This shows that while the on-die thermal track diodes do help with keeping the bias current constant over longer periods of operation, they do not help with countering the shorter thermal-transient induced bias current changes as shown in part 1.

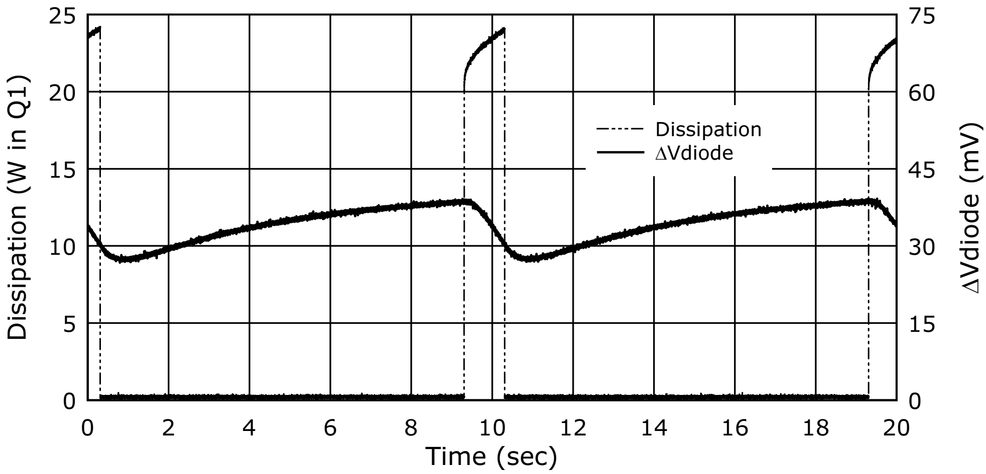

Figure 14 Change in on-die diode threshold voltage resulting from output device thermal transients at 10% duty cycle.

Looking at figure 14 more closely, it also appears that there is a delay between the change in dissipation and the reaction of the diode voltage. This can be easily seen when we expand the center part of figure 14 as shown in **figure 15**. The response of the diode to the changes in dissipation appears to be delayed by approximately 150 to 200 milliseconds. Note that this is not a phase shift between the dissipation and the diode voltage; if that would be the case, the diode voltage change would change in slope at the same moment that the dissipation changed, but we do not see any change in the rate of change of Vdiode at the moment of dissipation change. I do not know the mechanism for this delay, but it again shows that the thermal diode cannot be counted on to cancel thermal transient bias shifts in the output devices.

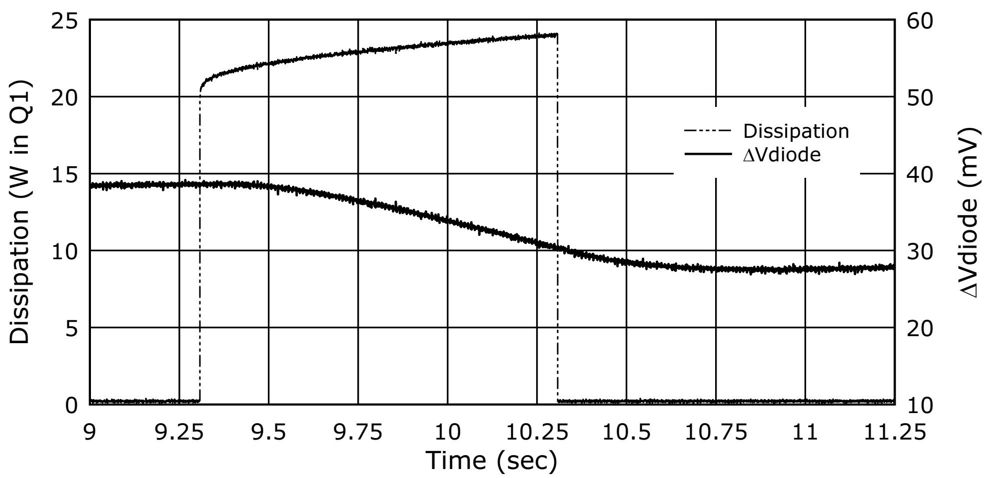

Figure 15 Expanded portion of figure 14 showing 150-200msec delayed diode response.

Conclusions from part 2

So called ThermalTrak devices, with an on-die diode that can be included in the bias control loop, offer no advantage regarding compensation for thermal transient bias current changes due to dynamically varying dissipation in the output devices. These diodes react to dissipation changes with a delay of 150 to 200 milliseconds. Furthermore, their reaction appears heavily low-pass filtered at fractional Hz frequencies, and is much slower than the transient thermal changes in the output devices themselves.

Part 3 – Impact of bias modulation on amplifier linearity

Having established that there exists transient thermal bias modulation in an output stage, with or without integrated ThermalTrak diodes, the important question is: does it impact the linearity and performance of our amplifiers?
To answer this question, another test setup was made, to modulate the bias voltage of an output stage while monitoring the distortion residual at the output. By choosing a high enough signal frequency with constant level, actual thermal bias modulation due to dynamically varying dissipation is avoided, so any possible changes in the distortion level must be caused by the externally-induced bias modulation.

My trusty Sony TA-N68 was set up with a simple modification enabling to change the bias level with the flick of a switch, see **figure 16**.

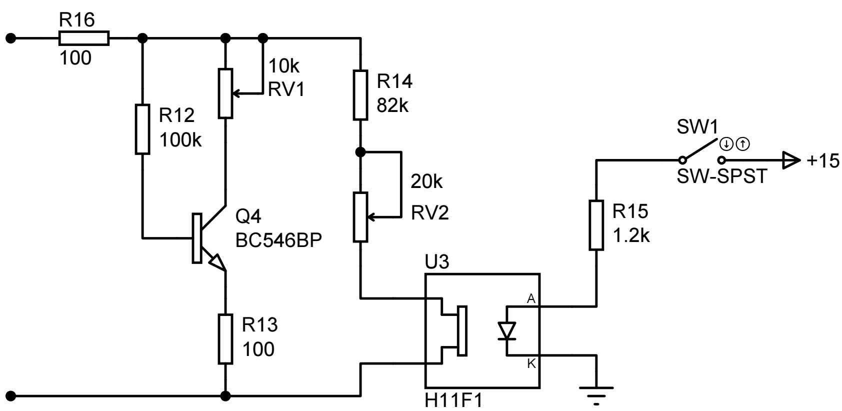

Figure 16 Bias spreader with modification for two discrete settings.

The circuitry around Q4 is part of the original bias spreader; added are U3 and its activation. Closing SW1 lowers the bias current in the output stage from 60mA to 52mA. An analog distortion analyzer (modified Boonton 1130) was used to look at the distortion level as well as the residual.
The analyzer indication did not change with the changes in bias current; in both settings the 1 kHz

THD was -91.3dB. The residual in both cases is shown in **figure 17**. Although in some places the high-bias THD seems higher than the low bias THD, it would have been easy to capture another segment where the situation would be reversed. The differences in the residual are quite small and indistinguishable from the stochastic variation normally observed.

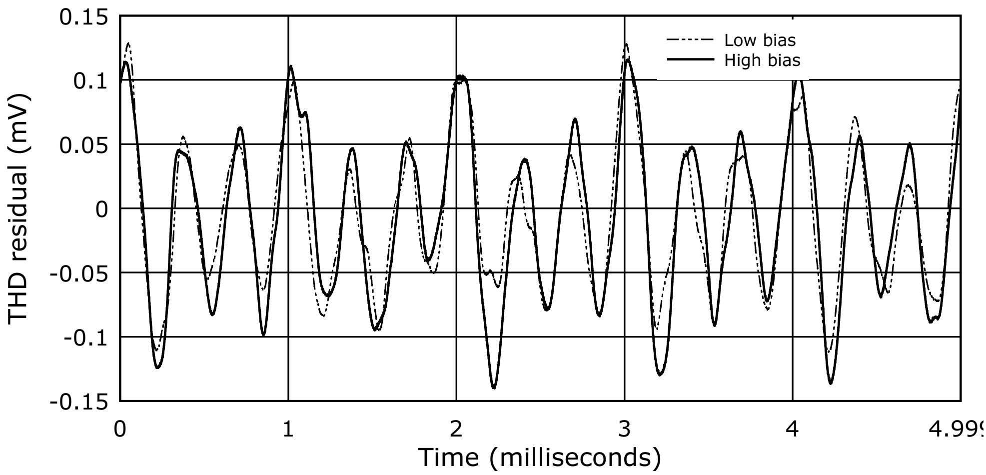

Figure 17 Distortion residuals for the two bias settings.

Conclusions from part 3

The small bias current variations resulting from thermal transients in the output stage were emulated by switching the bias current between two discrete levels and looking at changes in distortion level and shape. No changes could be identified beyond the stochastic variations. It can be concluded that bias changes due to output stage thermal transients do not lead to measureable changes in amplifier linearity.

Summary and Conclusions

The changes in bias currents resulting from thermal transients in audio output stages have been measured for output stages with bipolar output devices, in low to medium output power amplifiers. Thermal transients of up to 100 msec duration resulting from output level changes in amplifier output stages cause corresponding transient changes in bias current level. These changes, with output level variations corresponding to transient dissipation changes of some 10-20W are limited to 2 or 3% of the nominal bias current.

So-called ThermalTrak transistors (ON semi NLJ3102D and Sanken STD03) with on-die diodes which can be incorporated in the bias spreader circuit cannot compensate for these thermal transients. There is a delay between the occurrence of the thermal transient and the reaction of the diode voltage, which is too long to react on transients of less than 100 to 150 msec.

Transient changes in bias current of up to 10% do not cause any measureable changes in amplifier linearity.

Based on the above results, it can be concluded that transient thermal bias changes in power amplifier output stages do occur but do not lead to measureable effects on the signal, and are therefore unlikely to be audible.

References

[1] *Amplifier transient crossover distortion resulting from temperature change of output power transistors*, T.Sato *et al*, AES preprint 1896, October 1982.

[2] *La distortion thermique*, Hephaïstos, L'Audiophile N° 32, May 1984.

[3] *La distortion thermique, tube contre transistor*, Hephaïstos, L'Audiophile N° 33, September 1984.

[4] *Designing Audio Power Amplifiers*, 1st edition, Bob Cordell, McGraw Hill, 2011, ISBN 978-0-07-164024-4, Chapter 14.

[5] *Audio Power Amplifier Design Handbook*, 6th edition, Douglas Self, Focal Press, 2009, ISBN 978-0-240-52613-3, Chapter 22.

[6] http://www.diyaudio.com/forums/solid-state/212150-output-relay-power-amp-2.html.

[7] *paX – A Power Amplifier with Error Correction*, Jan Didden, Elektor 4/2008 (part 1) and 5/2008 (part 2); also in AudioXpress 10/2008.