Novel Precision Bias Control System For Class AB Power Amplifiers

Bruce E. Gillingham

July 07, 2022

1. INTRODUCTION

Stabilizing the output stage (OS) bias current in a solid state class AB audio power amplifier has been a subject of discussion for decades. This task is usually accomplished by designing a "bias generator" to have a voltage temperature coefficient close to that of the combined transistors in the OS, and mounting certain parts of the bias generator in close thermal proximity to the OS's transistors. In theory, this is not too difficult but in practice it is complicated by potentially large temperature differences among the transistors and difficulty in achieving tight thermal coupling of the bias generator and output transistors. The standard approach is a temperature compensation system and does not actually measure the parameter that we wish to control. Ideally, the bias current would be measured directly and stabilized by a closed-loop electrical feedback system. There have been numerous designs taking this approach with varying degrees of success, side effects, and complexity. I chose to refine the traditional approach with a combination of temperature compensation and closed loop techniques.

I designed, simulated, implemented, and tested a novel bias control system ("Bias Controller"). It provides highly effective bias stabilization, without undue complexity, using only linear analog processing. The controller features a simple temperature compensation calibration procedure (see Appendix A) and consumes little power (< 125 mW). Measured performance shows idle state (no amplifier load current) bias is held within $\pm$1% over a 40 °C heatsink temperature range. Performance with the amplifier driving a load is limited only by the thermal characteristics of the ThermalTrak™ transistors utilized in the design [1, 2, 3].

Although the Bias Controller was designed for use with ThermalTrak™ transistors, the basic technique will work with standard power transistors and external temperature sensing devices as well.

2. ENVIRONMENT

The Bias Controller was designed for a class AB Locanthi Triple amplifier Output Stage (OS) as shown in simplified form in Figure 1 [4, 5]. For clarity, only one pair of output power transistors (Q5, Q6) is shown.

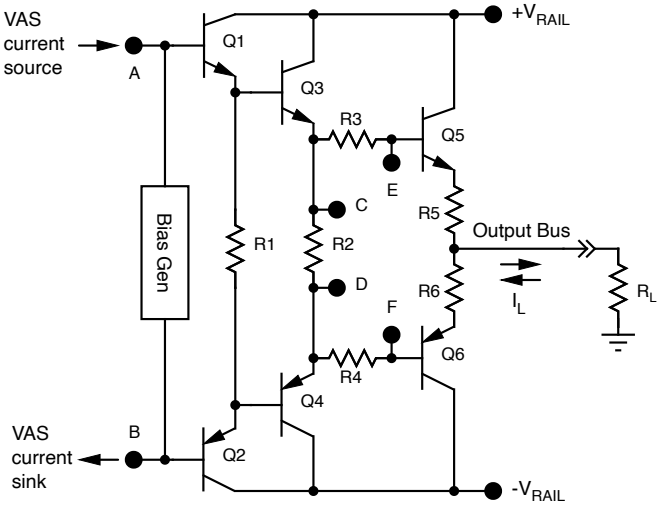

Figure 1: Output Stage

The OS receives current drive from the preceding Voltage Amplification Stage (VAS). The quiescent currents in the pre-driver transistors (Q1, Q2) and driver transistors (Q3, Q4) are set by the voltage across the bias generator (V_{A-B}) and resistors R1 and R2, respectively. The bias generator here is a voltage source that is adjustable via an optocoupler. Q1-Q4 operate in class A (always conducting) under all normal conditions. The final output transistor pair (Q5, Q6) operates in class AB. Both output transistors conduct under quiescent ($I_L = 0$) and low load current conditions. When the magnitude of the load current is large enough, one of the transistors will turn off. The quiescent (bias) current in Q5 & Q6 is critical to optimizing amplifier distortion performance. The proper value of bias current smooths the transistor turn-on-turn-off transitions and minimizes gm doubling. It is set by V_{A-B} and equal value resistors R5 and R6. Before being imposed across the R5-R6 series pair, V_{A-B} is reduced by the base-to-emitter junction voltages (V_{BE}) of all six transistors and the voltage drops across R3 and R4. These resistors are "base stopper" resistors used to dampen possible parasitic oscillations. As discussed below, they are also used here to measure the driver transistor load currents.

The goal of the Bias Controller circuitry is to maintain a specific voltage between the emitters of Q5 and Q6 when the output current (I_L) passes through zero. Per Barney M. Oliver, minimum distortion in the output pair is obtained when the emitter-to-emitter bias voltage is approximately 52 mV [6]. This bias voltage is actually the critical parameter, not the bias current in the output transistors, which may be set by the value of R5 and R6. Stabilizing the bias voltage is challenging as the desired voltage is quite small and V_{BE} of the transistors is strongly a function of temperature.

3. BIAS CONTROLLER APPROACH

The approach presented here is to stabilize the voltage between points C and D (V_{C-D}) in Figure 1 with a slow (relative to music frequencies) closed control loop. With transistors Q1 through Q4 within the loop, their temperature variations do not impact the output bias. The control loop leaves us with only the temperature variations of Q5 and Q6 to compensate for. This is aided by the use of ThermalTrak™ power transistors that include an internal temperature sense diode. The temperature of the diode tracks the the temperature of the transistor die fairly closely. These diodes are used to adjust the setting of the control loop around the pre-driver and driver stages. Thus, V_{C-D} changes to compensate for V_{BE} changes in Q5 and Q6. Both simulation and test data show that the quiescent current in the final output transistors can be maintained very precisely over temperature using this technique.

As in the standard approach, changing V_{C-D} also changes the bias currents in the pre-driver and driver stages, but amplifier performance is usually fairly insensitive to this as long as the currents remain high enough.

One complication with this approach is that V_{C-D} changes somewhat when the amplifier is driving a load. Although driver transistors Q3 and Q4 operate in class A, their emitter currents are not constant. Q3 must source the base current of Q5 and Q4 must sink the base current of Q6. These currents may be substantial. When a driver transistor's emitter current increases, its V_{BE} increases as does the voltage drop across its base resistor. This is also true, to a lesser extent, for the pre-drivers. These increased voltage drops lower V_{C-D}. With a sinusoidal load current, V_{C-D} is modulated with a full wave rectified version of the load current and its average value decreases (see Figure 2). With the traditional approach, this is not an issue because there is no load current at the time of crossover. With this approach, the control loop unnecessarily "corrects" the load-induced decrease in the average value of V_{C-D}, resulting in a higher than desired bias at the time of crossover.

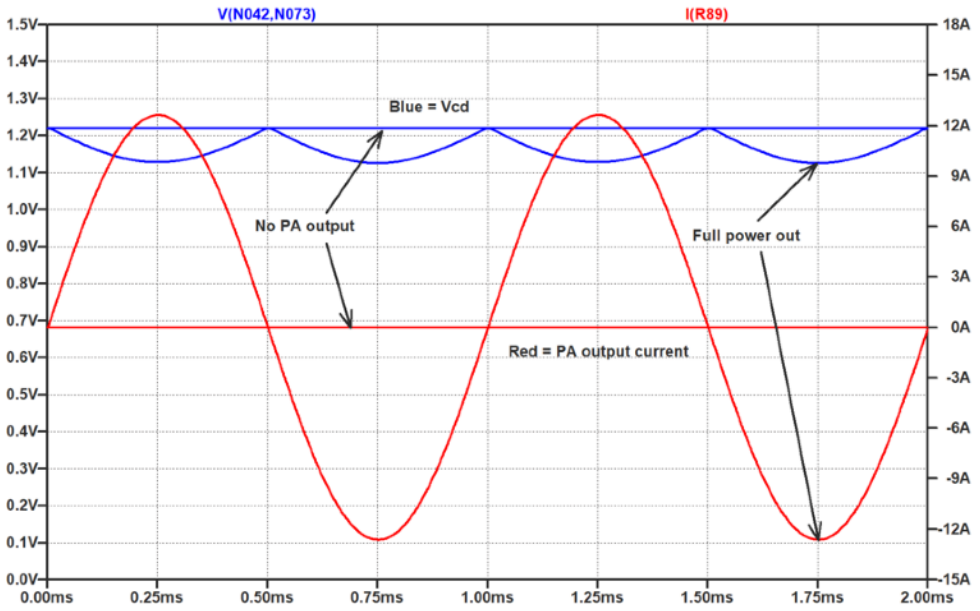

Figure 2: $V_{C\text{-}D}$ With and Without Load Current (Simulation)

Compensation for load current variation is performed by measuring the driver output currents and adjusting the control loop setting to compensate. $V_{C\text{-}D}$ - $V_{E\text{-}F}$ provides a voltage proportional to driver output current. Although this voltage is purely linear with load current and the the $V_{C\text{-}D}$ change is not, the differences are not significant for the range of currents encountered.

The Bias Controller was designed to optimize temperature coefficient matching and thermal coupling of the bias generator to the output transistor. Given this goal, the design is very successful. However, as it does not measure and control the bias current directly, it has some limitations. Its ability to track dynamic changes in output transistor die temperature is limited by the design of the ThermalTrak™ parts. As the thermal mass of the transistor die is not large, its temperature can change very rapidly in response to changes in transistor power dissipation. There is thermal resistance between the transistor die and the much larger thermal mass copper header. There is another thermal resistance between the copper header and the temperature sensing diode. This structure produces temperature differences (gradient) and a thermal delay (lag) between the two semiconductors when the amplifier is driving a load. Fortunately, the concomitant sensing errors result in a bias current higher than at no-load. This is usually less of a problem than a lower bias. The fact that bias increases when driving a load may contribute to the optimum idle bias setting obtained by measurement being considerably lower than the Oliver point. Note that the conventional approach to biasing also has this same issue.

4 IMPLEMENTATION

The Bias Controller circuitry (simplified) is shown in Figure 3. Each potentiometer has a series resistor (not shown) to limit its control range. The circuitry's local ground is referenced to the amplifier's Output Bus and it is powered by floating power supplies, $\pm V_F$, (not shown). This minimizes common mode voltages that must be rejected in order to measure the necessary voltages.

The basic control loop function is performed by differential amplifier A4 and loop filter A1. A4 measures -$V_{C\text{-}D}$ and provides a proportional current through R8 to the loop filter, an "ideal" integrator. The integrator compares this current with a reference current through R1 derived from a precision voltage reference, V_{REF}, that is highly stable over temperature. If the currents are not equal, the integrator's output will move toward one floating supply rail or the other. This changes the current in the LED of optocoupler Opto 1. The LED, in turn, changes the current flow in the associated phototransistor that is a part of the Bias Generator. The phototransistor

adjusts the Bias Generator voltage (V_{A-B}) to achieve the desired setting for V_{C-D}. The high DC loop gain provided by the operational amplifier integrator provides tight regulation of the average value of V_{C-D}.

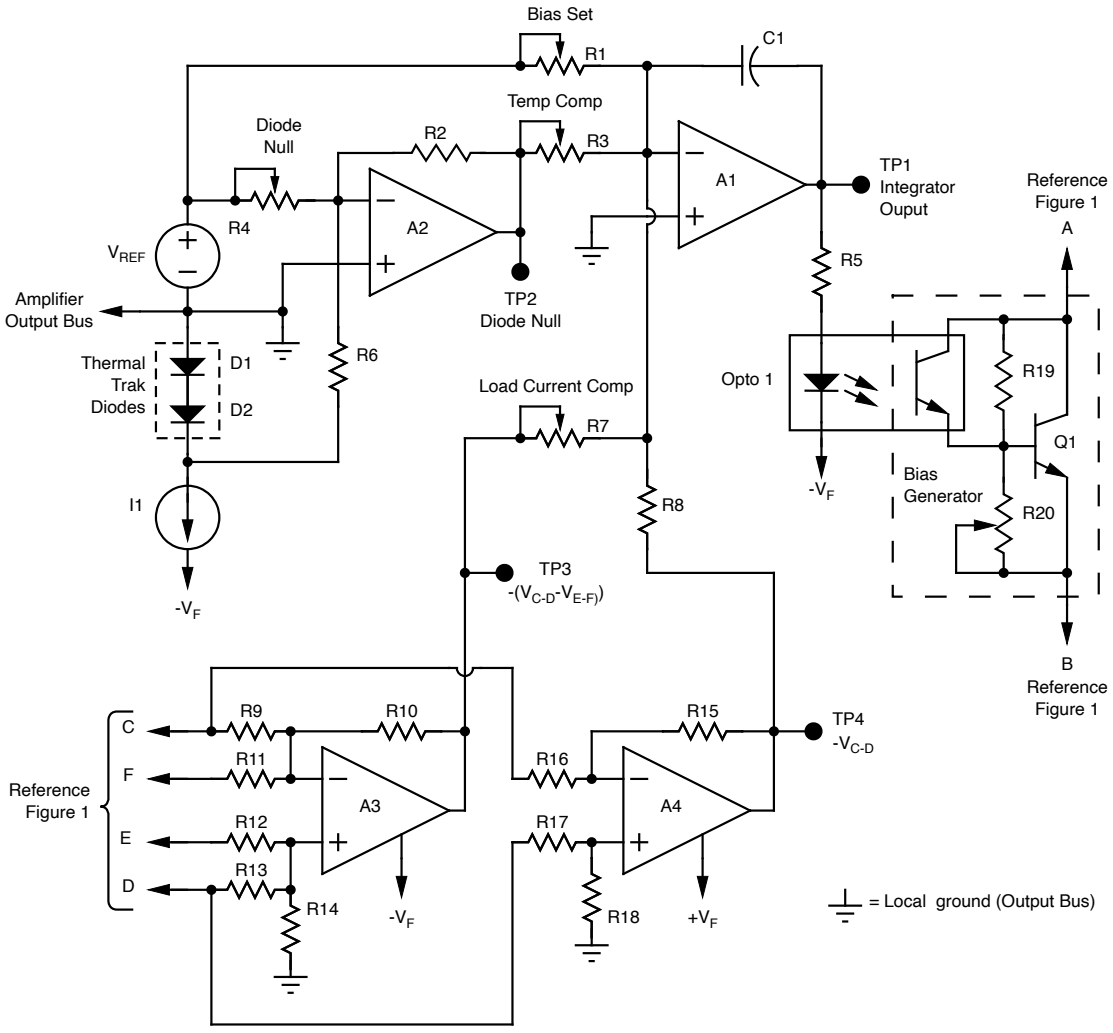

Figure 3: Bias Controller (Simplified)

To provide temperature compensation for Q5 and Q6, a precision, temperature-stabilized current source, I1, forward biases their two internal tracking diodes (D1, D2) which are connected in series. A2 measures the voltage across the diode pair and provides a proportional current to the integrator through R3. Adjusting R3 changes the gain from the diodes to the integrator, which effectively changes the temperature coefficient of the Bias Generator. This is a key feature of this approach. The compensation temperature coefficient can easily be adjusted to whatever is required. Temperature coefficient mismatches between power transistor and sensing diode can be corrected. R4 is adjusted to zero the output of A2 with the amplifier in a cool state, making the value of R3 irrelevant. The output transistor bias is initially set using R1 with the amplifier at that same “cool” temperature. This establishes one point on the linear compensation “curve.” As the output transistors and their tracking diodes heat up, the voltage across D1 and D2 decreases, driving A2’s output negative. This draws current from the integrator through R3, lowering the set point for V_{C-D} , and, hence, the output pair bias. This compensates for the decrease in output transistor V_{BE}. If R3 is then adjusted at a high temperature to restore the bias to the initial set point, a second point is established for

the compensation line. Two points define the straight line required to match the changes in V_{BE}.

Summing amplifier A3 monitors points C, D, E, and F and outputs -(V_{C-D} - V_{E-F}). This is equal to the inverted sum of the voltages across R3 and R4 [-(V_{R3} + V_{R4})], and is proportional to the sum of the load currents in the driver transistors. Under quiescent conditions, these current are small and approximately equal. With a sinusoidal load, the emitter currents in Q3 and Q4 increase on alternate half cycles. As both current increases contribute to a decrease in V_{C-D}, summing V_{R3} and V_{R4} catches them both. The output of A3 is always negative and draws current from the loop filter integrator through R7. This current increases as load current increases, compensating for the reduction in current drawn by A4. Appendix B addresses the calibration of this function.

Additional implementation detail is provided in Appendix C.

5 PERFORMANCE

The Bias Controller was first simulated in LTspice® and then implemented in two (identical) monoblock power amplifiers that I designed [7].

5.1 Simulations

The Bias Controller, integrated with an output stage, was simulated in LTspice®. Voltage dependent voltage sources were used as op amps to avoid model issues with floating power supplies and ideal voltage and current sources were used to simplify the simulation. Hence, the results are optimistic. Also, no attempt was made to incorporate electrothermal models in the simulations so they did not capture the dynamic thermal effects discussed in Paragraph 3.

5.1.1 Bias Stability in Idle State

In the idle state (no output load current) and all transistor junctions at the same temperature, simulated performance is excellent. A brief optimization of resistor values achieved a bias current deviation of approximately 0.2% over a 60 °C range of junction temperature (see Figure 4).

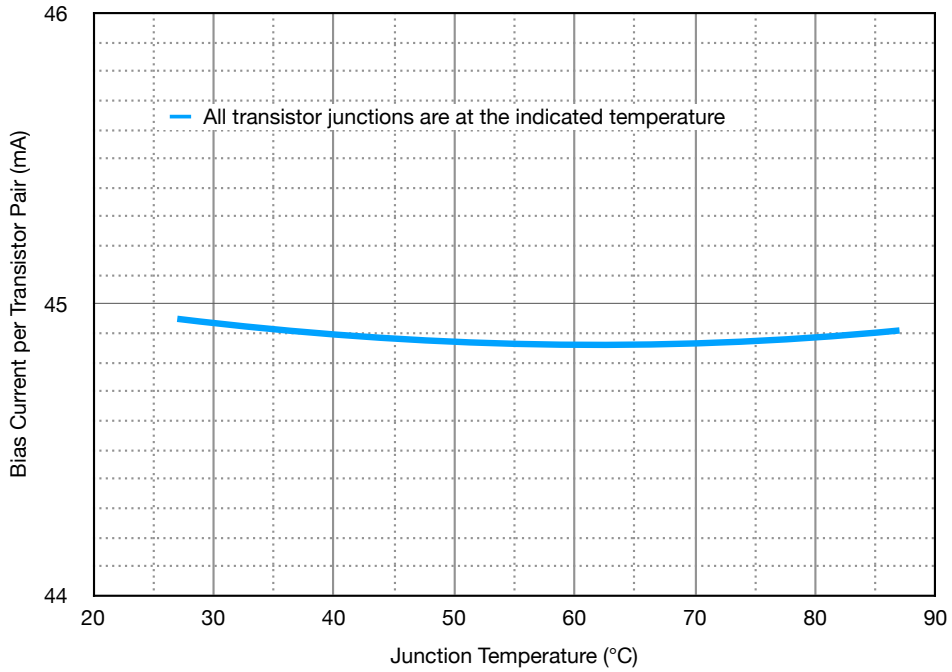

Figure 4: Bias Current vs. Temperature (Simulated)

Figure 5 shows the simulation results for a different thermal scenario. It assumes the amplifier has initially stabilized at 45 °C. Realistic junction temperatures are used for the bias generator, pre-driver, and driver transistors. These are fixed and the output stage junction temperatures are varied from 45 °C to 105 °C. Resistor values are the same as in Figure 4. Bias deviation from the initial value is less than 0.3%, another excellent result.

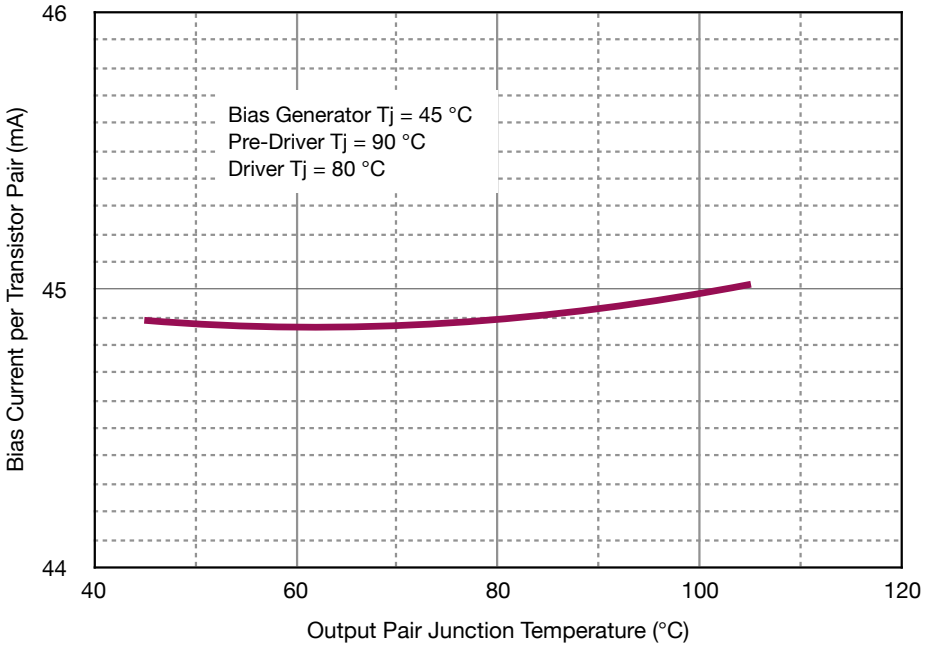

Figure 5: Bias Current vs. Temperature (Simulated)

5.1.2 Bias Stability With Sustained Power Out

Figure 6 shows the results of a simulation with a continuous sine wave output into a 4 ohm load as the peak current is varied. As the simulation does not include electrothermal models, the only errors are those from the Bias Controller. With the load compensation circuitry disconnected (blue line), the error is substantial. With it connected (green trace), the bias variation is negligible.

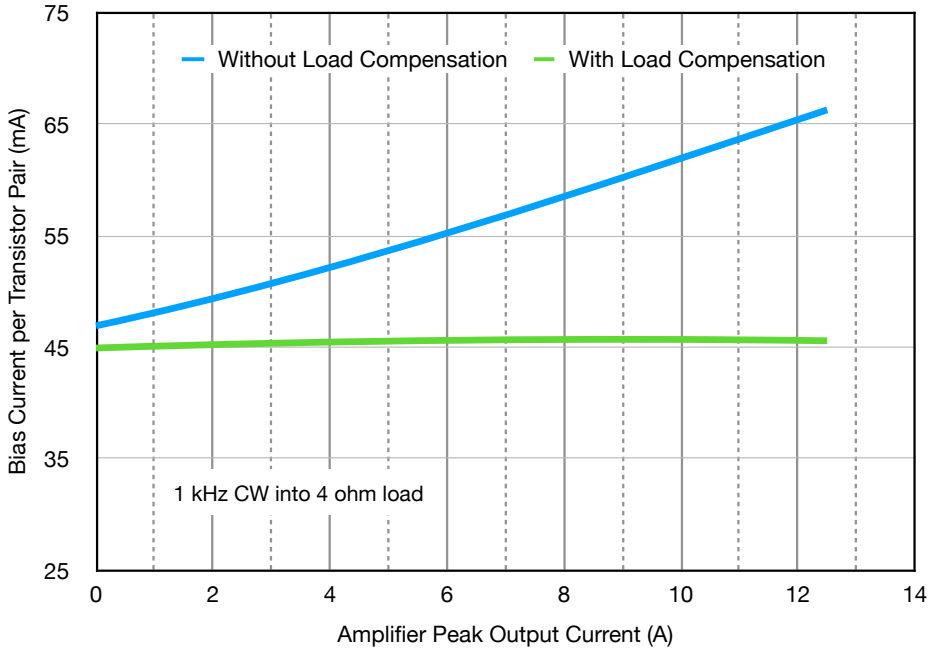

Figure 6: Bias Current vs. Amplifier Output Current (Simulated)

5.2 Hardware Measurements

The Bias controller was implemented in two 300 W (4 ohms) monoblock “test” amplifiers. There were no significant performance differences between the two amplifiers. Specialized test equipment was employed to enable measuring bias during power output conditions (see Appendix B).

5.2.1 Bias Stability in Idle State

As in the simulations, measured performance in the idle state is excellent. The test amplifier was warmed up from a room ambient start by driving a load. The main heat sink temperature was monitored by a LM35 sensor mounted on it [8]. The amplifier drive was removed briefly prior to recording each data point. This allowed the bias to stabilize following the output stage thermal transient. The measured bias variation was approximately ±1% over a 40 °C range in heat sink temperature. Figure 7 shows measured data from one test amplifier. It is probable that stabilizing the amplifier at each measurement temperature in an oven would result in even less variation.

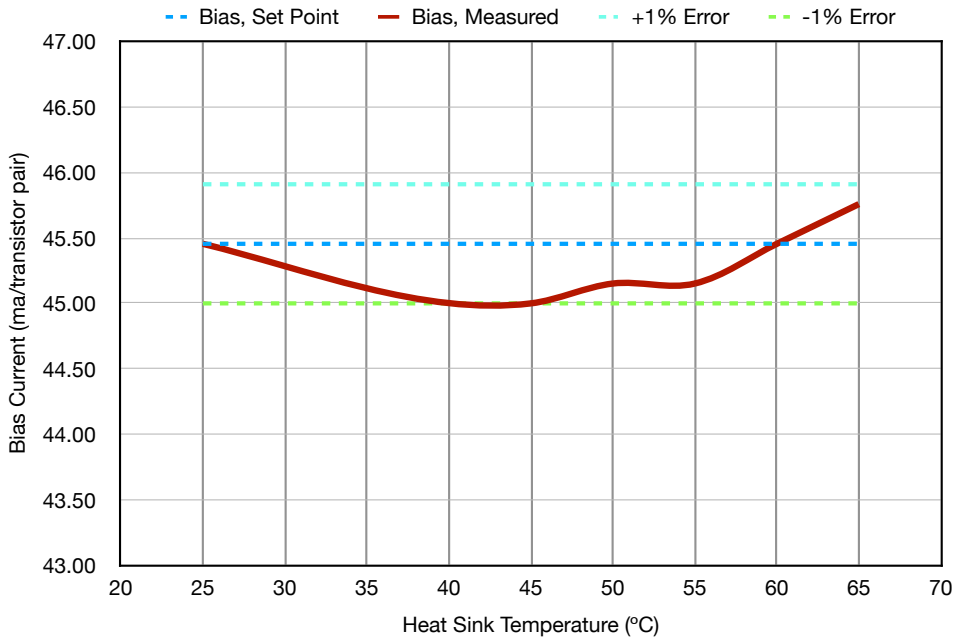

Figure 7: Idle Bias Current vs. Temperature

5.2.2 Bias Stability During Sustained Power Output

The transistor die in the ThermalTrak™ transistors can heat up very rapidly in response to a power transient. The time constant is on the order of 20 ms. This is an order of magnitude, or more, faster than the sensor diode and Bias Controller loop response times. Thus, at the start of a tone burst into a load, the bias will initially rise very quickly, and then decrease as the thermal feedback takes effect. However, it will not decrease to the idle state value as long as the output transistors are dissipating significant power. This is due to the temperature gradient between transistor and sense diode. This behavior can be seen in Figure 8. Drive to the amplifier was applied at T=0. At output powers between 2 W and 300 W, the bias current jumps sharply at turn-on. Bias rises 35-39% above idle level in the range of 100-200 W output power. Maximum output transistor power dissipation occurs within this range. Over bias is significantly reduced at 1 minute and approaches a new plateau in 2- 3 minutes. This is much faster than what can be achieved with a temperature sensor external to the power transistors. The worst case (100 W) stabilizes at 29% above the idle level. This behavior is not addressed by the Bias Controller and is shared by the conventional approach to bias regulation.

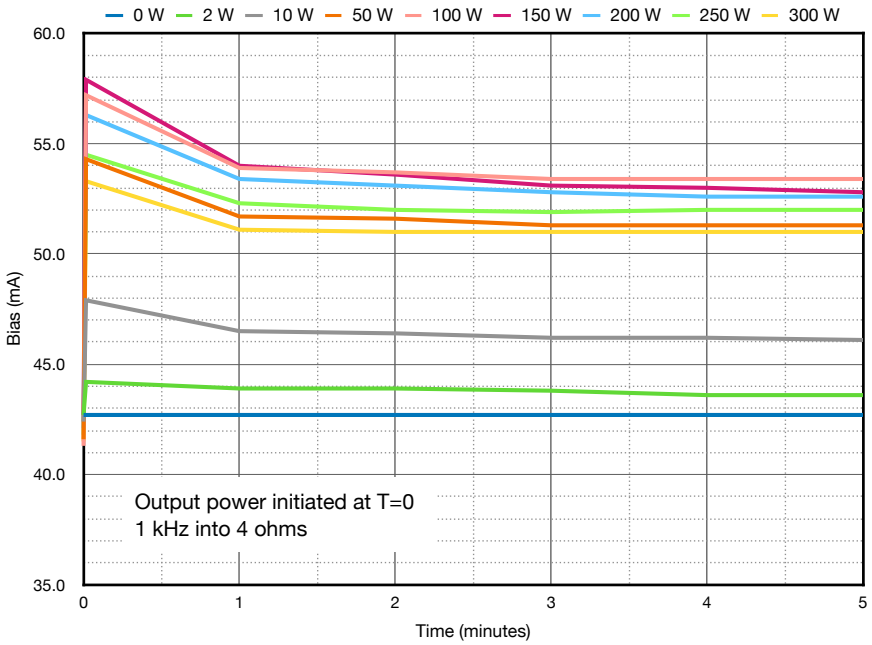

Figure 8: Bias Current vs. Time and Output Power Level

5.2.3 Bias Current Recovery Following a Tone Burst

Figure 9 shows the recovery of the test amplifier from an over bias condition caused by a tone burst consisting of 2000 cycles of 1 kHz in a 4 second burst period (note: burst waveform is continuous). The output power level during the burst was 150 W into 4 ohms (near worst case power dissipation). Bias during the burst (not shown) rose from 42.8 mA at the start to 59.6 mA at the end. Figure 9 shows a very steep initial decline in bias current at the end of the burst (T=0). This is the transistor die rapidly cooling to the header temperature. The slower decline that follows is a combination of the transistor and header cooling between bursts and the Bias Controller responding to the header's temperature. Bias recovers quickly (~1.5 seconds).

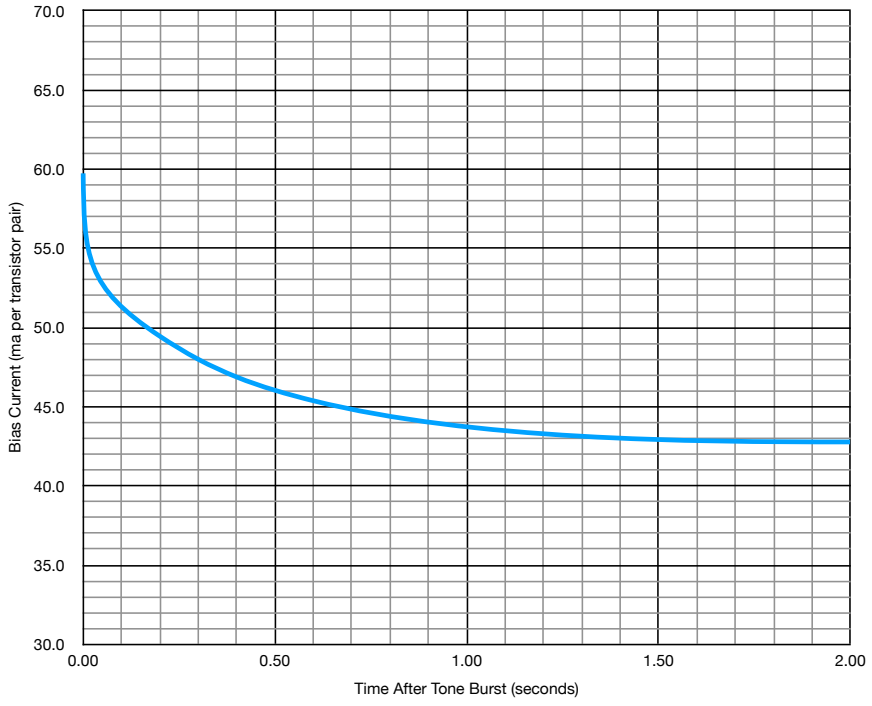

Figure 9: Bias Current vs. Time Following a Tone Burst

5.2.4 Amplifier Distortion

The purpose of bias stabilization is to prevent thermal runaway and to minimize amplifier distortion. Sensitivity to crossover distortion is a function of output level. At very low output levels, the amplifier operates in the class A region where there is no crossover distortion. At high output levels, crossover distortion is a small percentage of the output and may be masked by other distortion mechanisms. In his book *Designing Audio Power Amplifiers*, Bob Cordell states that 4.5 W into 8 ohms is in a range sensitive to crossover distortion [9]. One of the test amplifiers was measured at that level and also at ±3 dB and ±6 dB relative to that level. The test frequency was varied from 20 Hz to 6 kHz. The results are shown in Figure 10. Distortion is around 0.001% and is nearly independent of both level and frequency. This is a very good result and shows no indication of crossover distortion issues. Figure 11 shows test results with the same frequencies and voltage levels into 4 ohms, which doubles the output power. The distortion is approximately doubled into 4 ohms, but the curves are even more tightly grouped. Figure 12 shows 1 kHz THD for both test amplifiers over a wider range of power levels (0.2 W to 355 W) into 4 ohms. The distortion is virtually constant over more than a 30 dB (1000:1) range of output power. Again, there are no indications of crossover distortion.

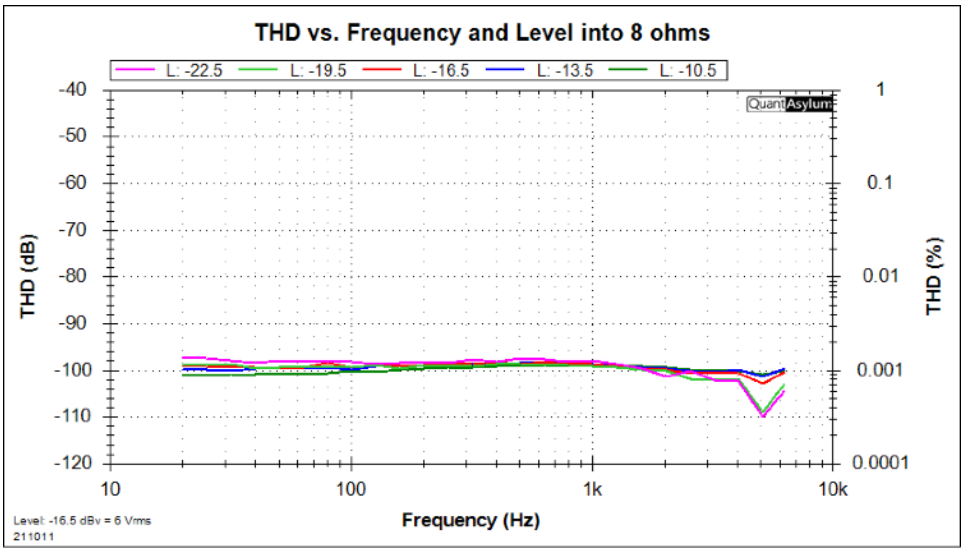

Figure 10: Test Amplifier Distortion 1.1 W to 18 W (8 ohms)

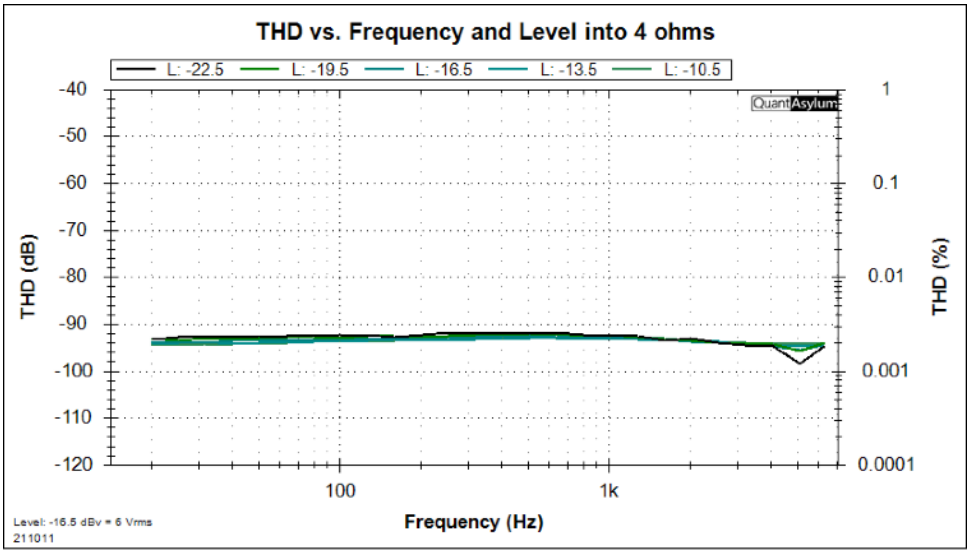

Figure 11: Test Amplifier Distortion 2.2 W to 36 W (4 Ohms)

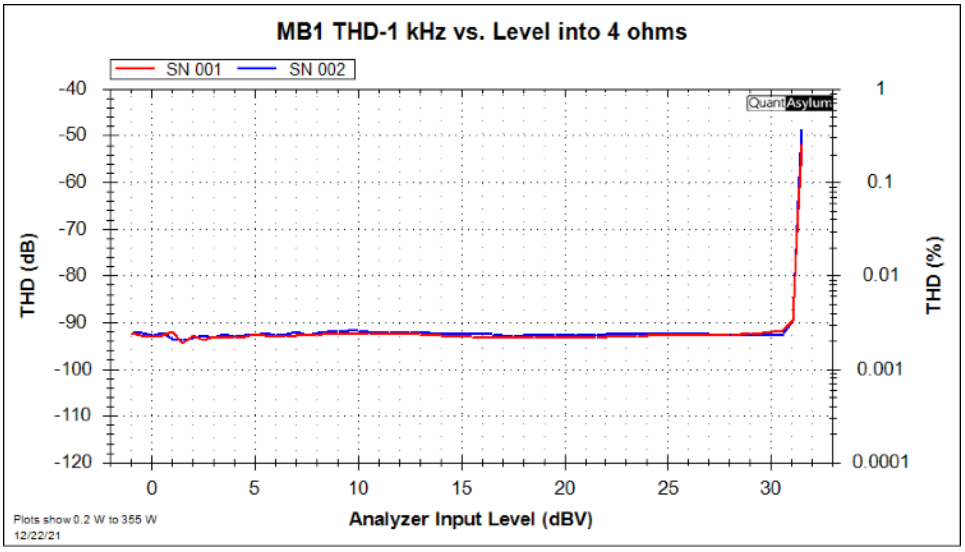

Figure 12: Test Amplifier Distortion 0.2 W to 355 W (4 Ohms)

6.0 CONCLUSIONS

This new approach to class AB bias stabilization has excellent performance and no known significant negative effects. Its only limitations are those associated with the thermal lag and temperature gradient inherent in the physical design of the transistor/sensor system.

REFERENCES

1. U.S. Patent 7,279,983, "Output Amplifier Structure with Bias Compensation," October 9, 2007.
2. ON Semiconductor datasheets for NJL3281D and NJL1302D, June 2006. Available at www.onsemi.com.
3. Mark Busier "ThermalTrak™ Audio Output Transistors," ON Semiconductor, AND8196/D, February 2005. Available at www.onsemi.com.
4. Bart Locanthi, "Operational Amplifier Circuit for Hi-Fi," *Electronics World*, pp. 39-41, January 1967.
5. Bart Locanthi,"An Ultra-low Distortion Direct-current Amplifier," *Journal of the Audio Engineering Society*, vol. 15, no. 3, pp. 290-294, July 1967.
6. Barney M. Oliver, "Distortion in Complementary Pair Class B Amplifiers," *Hewlett Packard Journal*, pp. 11-16, February 1971.
7. LTspice®, developed and distributed by Linear Technology Corporation/Analog Devices, Inc. Available at www.analog.com.
8. "LM35 Precision Centigrade Temperature Sensor," Datasheet, National Semiconductor/TI, December 2017. Available at www.ti.com.
9. Bob Cordell, *Designing Audio Power Amplifiers (2nd ed.)*, Routledge, New York, 2019.
10. Jan Didden, "Thermal transient variation of power amp quiescent current instrumentation and findings," *Linear Audio*, vol. 9, pp. 171-188, April 2015.
11. Bob Cordell, *Designing Audio Power Amplifiers (2nd ed.)*, p. 414, Routledge, New York, 2019.
12. On Semiconductor datasheets for KSA3503 and KSA1381 available at www.onsemi.com.
13. Texas Instruments datasheet for OPA4197 available at www.ti.com.

APPENDIX A: TEMPERATURE COMPENSATION CALIBRATION

Calibration of the temperature compensation function is straightforward, and requires minimal test equipment, a major feature of this approach. The procedure first sets the bias at one temperature with the voltage from the temperature sense diodes nulled out and then resets the bias to the same level at a different temperature using the voltage from the sense diodes. Although any two temperature points may be used for calibration, it is recommended that the initial point be room ambient to minimize drift during the process. Note that the steps below describe the use of two voltmeters. This is for ease of procedure description and process execution, but a single meter can be moved between the two measurement points.

1. Turn the power amplifier off.
2. Remove any signal input to the power amplifier.
3. Allow the power amplifier to stabilize thermally at ambient room temperature.
4. Connect DC Voltmeter 1 between Bias Controller local ground and TP2 (Diode Null).
5. Connect DC Voltmeter 2 from Q5 emitter to Q6 emitter (bias voltage).
6. Initially, adjust R3 to the approximate center of its adjustment range.
7. Initially, adjust R4 to the approximate center of its adjustment range.
8. Initially, adjust R20 to the minimum bias setting (max resistance). [Note: The bias generator must be designed such that the manual bias adjustment (R20) dominates the optocoupler loop control and can completely reduce the bias current to zero.]
9. Initially, adjust R1 to its maximum bias setting (minimum resistance).
10. Turn the amplifier on. The output bias current should be near zero.
11. Using R4 and Meter 1, set the voltage at TP2 as close to zero as possible.
12. Using R20 and Meter 2, set the bias voltage to approximately twice the desired value.
13. Using R1 and Meter 2, quickly lower the bias voltage to the desired value.
14. Allow the amplifier to warm up and thermally stabilize. The warm-up may be accelerated or extended by driving a load, but the drive must be disabled several minutes for stabilization before progressing to the next step.
15. Using R3 and Meter 2, adjust temperature compensation by resetting the bias to the value obtained in step 13.

APPENDIX B: LOAD COMPENSATION CALIBRATION

Calibration for driver loading effects is not as straightforward as it is for temperature. I spent a great deal of time working this issue. A reasonable value for R7 may be obtained from simulation. This may be the best approach for many. Measuring the relationship between driver output current and $V_{C\text{-}D}$ on the actual hardware and calculating the necessary compensation is problematic. The modulation on $V_{C\text{-}D}$ is not large and it rides on the much larger amplifier output signal. Also, as a driver transistor's output current varies, so does its power dissipation and, hence, die temperature. The temperature-induced variations in V_{BE} may obscure part or all of the load induced variations.

There are many ways to approach the calibration and at different levels of amplifier construction/integration. I chose to calibrate with a fully integrated amplifier using a dynamic test with carefully chosen parameters to help separate the non-thermal, load-current-induced variations in $V_{C\text{-}D}$ from dynamic thermal variations.

To aid in the measurements, I constructed a specialized piece of test gear modeled on that described by Jan Didden in his *Linear Audio* article "Thermal transient variation of power amp quiescent current" [10]. The basic concept is a high-gain differential amplifier with its ground referenced to the amplifier's output bus. The differential amplifier measures the voltage across the two series-connected emitter resistors in the output stage. The amplified signal is then translated from output bus reference to amplifier power ground reference with a unity gain, high-voltage-compliant differential amplifier. This signal can be observed directly on a DC coupled oscilloscope. The Didden design also looks for output signal zero crossings and samples the voltage across the resistors at that point in time. This is the theoretically ideal way

to measure bias under dynamic conditions. To insure there is time for an accurate sampling, the output signal is briefly turned off after the zero crossing detection.

My test gear is less complex. It forgoes Mr. Didden's selectable gain and internal timing generator features. Instead, the gain is fixed and timing for the sample and hold is derived from the arbitrary waveform generator (AWG) used to generate the test signals. A measurement gap of any width can be inserted at the desired point in the test waveform. This works very well, but sacrifices the ability to do measurements with actual music signals.

The procedure I used is as follows:

1. Disable the control loop by not installing/removing R5.
2. Disable temperature compensation by not installing/removing R3.
3. Connect a resistor the same value as R8 in parallel with C1. This changes the integrator into a unity gain low pass filter.
4. Monitor A1's output with specialized test equipment similar to that discussed above.
5. Turn the amplifier on.
6. Manually adjust the bias with R20 to the desired setting. This may have to be reset periodically as the amplifier warms up.
7. Adjust R1 to zero the output of A1. This may have to be reset periodically as the amplifier warms up.
8. Connect the amplifier input to a tone burst generator and the output to an appropriate resistive load and an oscilloscope.
9. Drive the amplifier to full peak power with a tone burst of 25 cycles of 1 kHz in a 500 ms period. These parameters allow the output of A1 to settle to within 10% of its target value at the end of the burst [in my implementation where R8*C1 = 10 ms], while minimizing the impact of dynamic thermal effects.
10. Adjust R7 to minimize the peak-to-peak deviation of A1's output during the burst interval without incurring any initial rise at the start of the burst. It may be necessary to reset the bias and adjust R1 to maintain the trace on the oscilloscope screen. Do not use AC coupling on the scope to accomplish this.
11. Remove the resistor in parallel with C1.
12. Install R3 and R5.
13. Calibrate temperature compensation (see Appendix A)

APPENDIX C: ADDITIONAL DESIGN CONSIDERATIONS

C1.0 Bias Controller Loop Bandwidth

There are at least three considerations for selecting Bias Controller's loop bandwidth:

1. The control loop bandwidth should be high enough as to not slow correction. Bob Cordell provides a thermal model for the ThermalTrak™ transistors in his book, *Designing Audio Power Amplifiers* [11]. Although the time constant for the tracking diode is 350 ms, its response to transistor die temperature changes is actually much slower (several seconds) due to other time constants in the model. Thus, it appears that the control loop need not be super fast.
2. The control loop bandwidth should be low relative to audio frequencies to avoid interactions and possible distortion.
3. The control loop bandwidth should be low to minimize coupling of the output signal into the bias loop.

After much experimentation, I arrived at a loop time constant of approximately 240 ms (loop bandwidth ~0.66 Hz). This is a little faster than the tracking diodes and provides approximately 30 dB of attenuation to the lowest audio frequency (20 Hz).

C2.0 Bias Controller Ground Reference

The discussion above assumes that the Bias Controller "rides" on the amplifier's output bus. This may be the most esthetically pleasing connection point as it is very low impedance and can source and sink large currents. It is also relatively insensitive to stray capacitance. However, I chose to connect the controller to the midpoint of the resistors between the driver transistor emitters. This connection was used to minimize the amplitude of the common mode signals that must be rejected by the Bias Controller. It also provided some advantages during amplifier build and checkout given the test amplifier's partitioning.

C3.0 Floating Power Supplies

My choice of connection at the driver output gave greater impetus to finding a low-noise, low-capacitance-to-ground solution for the floating power supplies. The approach used was linear current sources referenced to the front end boosted power supplies ($\pm$75 V) and zener diode shunt regulators (6.2 V) in the Bias Controller. The current sources (15 mA) were constructed with low-output-capacitance transistors (KSA3503, KSA1381) to minimize capacitive loading on the driver stage [12]. A trimmer was added to balance the current sources, avoiding any significant DC current load on the driver.

The choice of linear current sources and shunt regulators is obviously inefficient and raises total power consumption of the bias controller to 2.25 W. This was not an issue in the test amplifier. Modern low noise switching regulators may be a viable option.

C4.0 Bias Generator Optocoupler

Special attention should be given to including an optocoupler in the bias generator. These devices often specify a very small coupling capacitance between the phototransistor and LED. Also, these two internal devices largely move together as the amplifier's output voltage changes, aiding isolation. However, it takes very little capacitance from the transistor base to ground to couple the output signal into the bias control loop. The voltage swing is very large and the transistor's base impedance is very high. The coupling rises with frequency. Even though the control loop's rejection of this signal also rises with frequency, there can be significant signal-related modulation of the bias. If there were no phase shift between the modulation and the output signal, this would not be a concern, but that is not the case at higher frequencies. Limiting the optocoupler's bandwidth by capacitively loading the base connection was investigated and found to be ineffective. This point is extremely sensitive and variable, preventing a stable cut-off frequency. Instead, a capacitor should be added across the transistor's collector and emitter pins. The resulting pole formed with R19 should be > 10X the Bias Controller's loop bandwidth to avoid stability issues. In addition, there should be no ground plane under the optocoupler package and consideration should be given to selecting a device with no external base connection.

C5.0 PWB Impact

Although the Bias Controller may seem complex, it can be quite made quite small. My implementation was built around a single quad precision op amp in a 14-pin SOIC package (OPA4197) [13]. The circuitry, less bias generator and current sources for the floating power supplies occupies less than 4 square inches of board space (see Figure C). It could be considerably smaller if all surface mount components were used. Power consumption, without floating power supplies is less than 125 mW.

If the proper amount of load compensation is to be determined by measurement/adjustment rather than by analysis/simulation, it is suggested that the PWB design include features to aid in this process, such as test points, jumpers, and potentiometers.

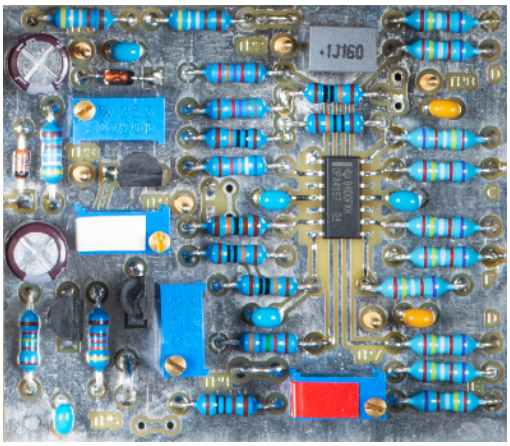

Figure C: Bias Controller Layout

C6.0 Further Thoughts on Load Compensation Calibration

The load compensation calibration approach that I used (see Appendix B) involves a considerable investment in time and some money. There is a promising alternative approach that I looked into but did not pursue to conclusion. This would be to close the feedback loop around the driver stage with the output stage disconnected. The driver stage can then be resistively loaded to ground. With the proper choice of test driving waveform and driver stage load resistance, the relationship between driver output current and $V_{C\text{-}D}$ might be determined directly without specialized test equipment. The required driver load current can be estimated by measuring the DC current gain of the the output stage. Making the driver load resistor and test signal small avoids the need to reject high-level common mode voltages. A DC test might work, but this is unlikely due to heating in the driver transistors. A pulse waveform, however, should allow one to measure the $V_{C\text{-}D}$ drop on the pulse edges with a standard two channel oscilloscope in differential mode. Once the relationship between driver output current and $V_{C\text{-}D}$ is known, the value of R7 (see Figure 3) can be calculated.