A 250 W Broadband Solid-State Linear Amplifier

By Dick Frey, K4XU

Editor's note: Section and figure references in this article refer to the 2013 edition of the ARRL Handbook.

Fig 17.43 — This 250-W amplifier for 160 through 6 meters provides a detailed design example as well as a practical project. Additional photos and information about the interior layout may be found elsewere on the CD-ROM.

Additional files, photos and other information mentioned in this article are available elsewhere on the CD-ROM. The amplifier described here and shown in **Figure 17.43** is neither revolutionary nor daring. It uses commonly available parts — no special parts and no "flea market specials." It is based on well-proven commercial designs and "best design practices" acquired over the past 30 years as solid-state technology has matured. This design project was undertaken for the 2010 Handbook as a detailed design example as well as a practical solid-state amplifier you can build. It is a continuing project.

A block diagram for the project is given in **Figure 17.44**. The amplifier is built on three PC boards — a PA module, a low-pass filter assembly, and a board for control, protection and metering circuitry. The *Handbook* CD includes *ExpressPCB* files for these boards, and the artwork can be used to have boards made in small quantities (see **www.expresspcb.com** for details).

The basic PA configuration has been in the *Handbook* since the 2010 edition. It is intended to be a "ham-proof" external amplifier for QRP transceivers that put out 15 W or less. It is designed for a gain of 30×, or 15 dB. Drive power of less than 10 W will provide 250 W output from 1.8 through 51 MHz. The amplifier provides exceptionally linear performance, necessary for high quality SSB and PSK modes, and is rugged enough to withstand the most rigorous contest environment.

Amplifier design tends to focus on the RF section, but a successful stand-alone solid-state amplifier is equally dependent on its control system. The control requirements for a tube amplifier are well known, while those for solid-state amplifiers are not. This is mostly because the functions of a solid-state amplifier's control system are generally transparent to the user. Parameters are monitored and protection is applied without any operator intervention. This must be. While tubes are fairly forgiving of abuse, semiconductors can heat so quickly that intervention *must* be automatic or they can be destroyed.

Transistors are sensitive to heat, so cooling and temperature compensation are critical to a successful design. Transistors require a heat sink. Power amplifier tubes have large surface areas and are cooled by air blown on or through them. Transistors are small. Mounting them on a heat sink increases their thermal mass and provides a much larger surface area so the heat dissipated in the devices can be removed either by convection or forced air. The thermal design of an amplifier is just as important as the electrical design. More information on thermal design may be found in the **Analog Basics** chapter.

Silicon's thermal coefficient causes the bias current to increase as the device heats up if the bias source is fixed. The increased current causes even more heating and can lead to thermal runaway. For stable Class AB linear operation, the gate bias for a MOSFET or bipolar transistor must track the temperature of the device. The control circuit typically uses another silicon device such as a diode thermally coupled to the amplifier heat sink near the transistor to sense the temperature and adjust the bias to maintain a constant bias current.

Transistor power amplifiers are designed to operate into 50 Ω. Operation into a VSWR other than 1:1 will cause an increase in device dissipation and other stress. The success of the solid-state transceiver is due to its integrated PA protection system. The temperature of the heat sink, the load VSWR, the output power and the supply current are all monitored by the control system. If any of these exceed their threshold limits, the RF drive is reduced by the transceiver's ALC system.

An external solid-state PA protection system must perform the same functions, but the driver's ALC circuit is not always available so other means must be used to protect the PA. This is usually accomplished simply by taking the amplifier out of the circuit. An indicator then tells the operator which condition caused the fault so appropriate action can be taken. Access to the driver's ALC system would make this protection task more automatic, smoother and less troublesome, but no two transceiver models have the same ALC characteristic. This makes the design of a universal ALC interface more difficult.

17.11.1 The 1.8 to 55 MHz PA — Detailed Description

Figure 17.45 shows the power amplifier (PA) schematic. Two Microsemi VRF151 MOSFETs are used in this amplifier. The circuit topology is a 4:1 transmission line transformer type, rather than a "tube and sleeve" type common in many PA designs and discussed earlier. This style offers more bandwidth, necessary to provide performance on 6 meters. Typical gain is 15 dB; 10 W drive will easily provide 250 W output with a 48 V supply. There is a lot of latitude in this design. It can even be operated on an unregulated supply. As long as the maximum unloaded voltage does not exceed 65 V, the transistors will not be overstressed. Other devices such as the MRF151, SD2931 or BLF177 would probably also work but have not been tested. They will require a regulated power supply, however.

FEEDBACK — TWO KINDS

The amplifier's gain is controlled by two kinds of feedback. Shunt feedback (from drain to gate) is provided by the link on T2 through resistors R5 and R6. It tends to lower the input impedance but it also helps to keep the gain constant over frequency and improve the linearity. Series feedback is provided by the 0.05 Ω of resistance in each source. This increases the input impedance, cuts the gain by 3 dB, and most importantly, it has a huge effect on the linearity.

Without any feedback at all, the amplifier would have more than about 30 dB gain (×1000) at some frequencies, tending to make it unstable — prone to parasitic oscillation. And the linearity would be terrible, –25 dBc or so IMD products. It would also be very sensitive to load changes. The input SWR is 1.2:1 on 160 meters and rises to 1.5:1 on 6 meters. The amplifier's gain is 15 dB ±0.5 dB over the same frequency range.

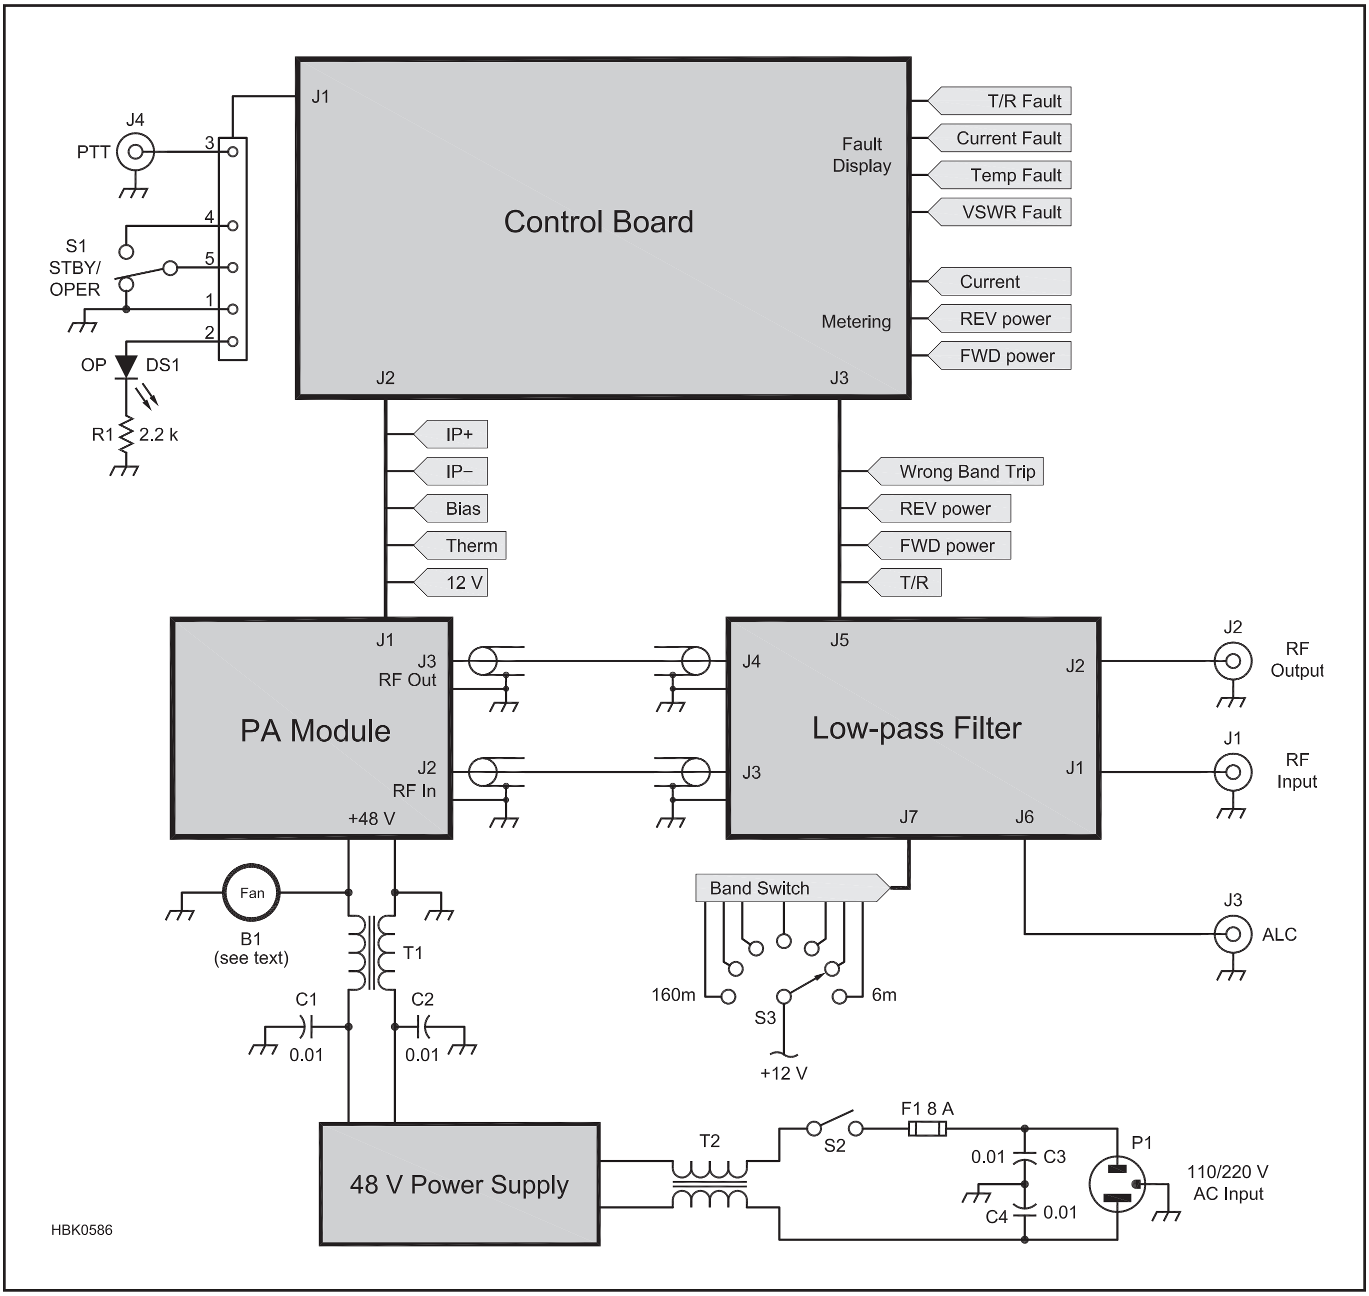

Fig 17.44 — Block diagram of the 250-W solid-state amplifier. It is built on three PC boards, which are described in the text and accompanying diagrams. T1 and T2 are 2 turns through a Fair-Rite 2643540002 ferrite bead if needed to suppress RFI from the switching power supply.

PA INPUT

On the input side, T1 is a sleeve-and-tube RF transformer. Its small size and tight coupling make it suitable for very wide bandwidth operation. When properly compensated, this type is able to provide a match to a 12.5-Ω load over a very wide frequency range from 1 to 100 MHz. The problem is that the input impedance of the two transistors is not a flat resistive load.

The gate of a MOSFET is essentially a high-Q capacitor. A voltage greater than its threshold (V_{th}) applied between the gate and source will control the conductivity of the drain-to-source path. No power can be dissipated in a purely reactive element, so we cannot match to a capacitor. Fortunately, all real reactive elements have losses and it is this loss that we would match to in a MOSFET gate for single-band operation. But that is no good here because we want a broadband match.

There are several ways to match a MOSFET over a broad frequency range. The most common is to swamp the gate capacitors with resistors. If the resistor value is lower than the impedance of the transistor gate capacitance, it dominates what the transformer sees as a load at the lower end of the bandwidth. The gate capacitor impedance decreases at higher frequency so some series resistance is added, R3 and R4, so that there is always a minimum real part to the load on the secondary of T1.

For a 1:1 input SWR, T1 wants to be terminated by a 12.5 Ω load. More than half of this is provided by R3 and R4. The rest is through shunt loads R7 and R8, plus the impedance of the output passed through the

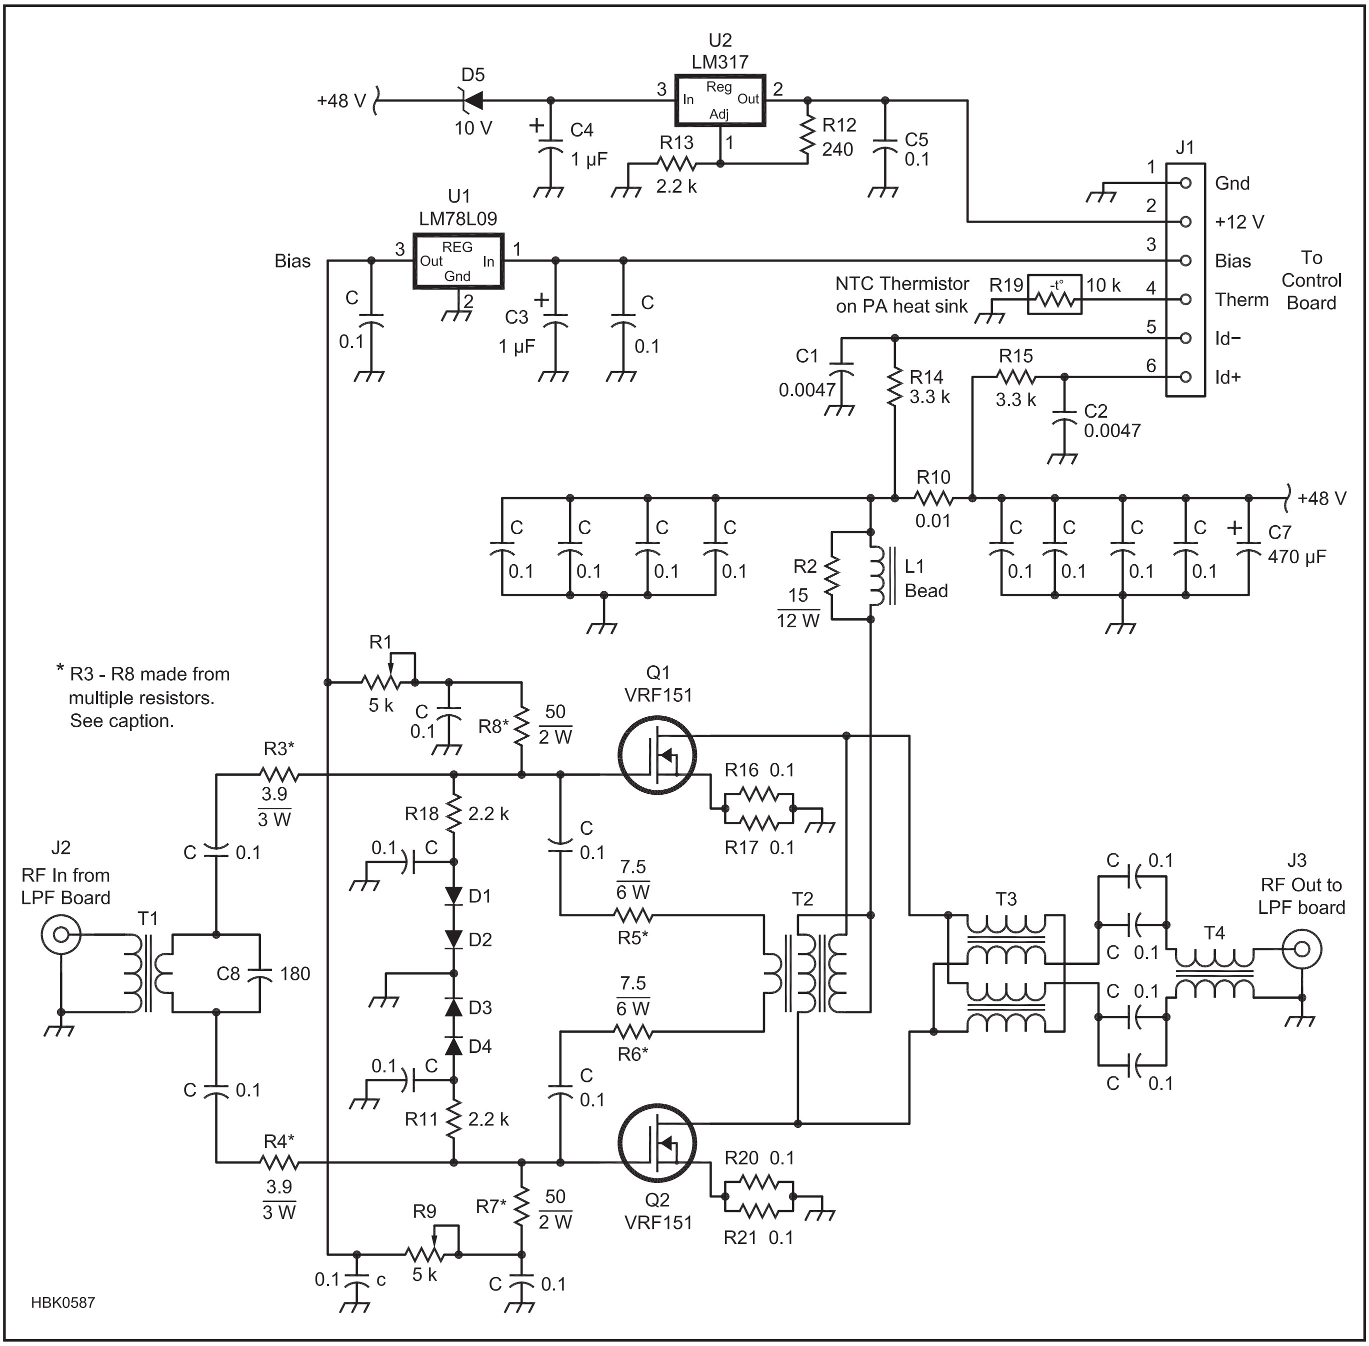

Fig 17.45 — Schematic diagram of the 250-W amplifier PA module. A complete parts list may be found elsewhere on the CD.

C — 0.1 µF, 100 V X7R 1206 SMT.
D1-D4 — SMT silicon PN-junction diodes such as 1N4148 or equivalent.
L1 — 2t on Fair-Rite 2643480102 ferrite bead.
R2 — 15 Ω, 12 W.
R3, R4 — Three 12 Ω, 2 W SMT resistors in parallel.
R5, R6 — Two 15 Ω, 3 W resistors in parallel.
R7, R8 — Two 100 Ω, 1 W SMT resistors in parallel.
R19 — 10 kΩ, 5% NTC thermistor, SMT (DigiKey 541-1150-1-ND).
T1 — 2t #22 wire on CCI RF400-0 core (or two Fair-Rite 2643006302 cores)
T2 — Primary, 1t #22 wire; secondary 8t #22 wire bifilar wound, on Fair-Rite 5961004901 core.
T3 — 2 × 3t 25 Ω coax on Fair-Rite 2861010002 core (see text).
T4 — 3t RG-188 coax on Fair-Rite 2643665802 core (on cable to LPF board, not on PC board).
Q1, Q2 — Microsemi VRF151 MOSFET.

feedback network. T1 has no center tap so a balanced load is forced by the action of resistors R7 and R8. Having "soft" center taps on both the input and output absorbs any differences between transistors and greatly improves the network's RF balance. This in turn improves the cancellation of even harmonics at the output. The gate impedance is raised by the effect of the source resistors, further improving the match.

DC FEED TRANSFORMER

In addition to providing the link for the feedback, T2 also acts as the dc feed choke. It is wound with two parallel bifilar #22 wires and a single turn for the feedback. At dc, the current flows in opposite directions through each #22 wire so the net current is zero and the core does not saturate with dc. At RF, the choke with its ferrite core provides at least 50 Ω of inductive reactance making it essentially invisible to the RF signals across it. As the feedback transformer it provides a ¹⁄₁₆ sample of the drain-drain voltage to the gate feedback loop.

It is not often mentioned in the literature, but the dc feed transformer T2 acts as a 180° hybrid combiner (**Figure 17.46**). A hybrid combiner has four ports: two inputs, the sum port and the difference port. The sum of the two input signals appears on the sum port and the difference between them (differences due to voltage or phase) appears at the difference port. In this case T2 is terminated with a 1:4 balanced transformer that brings the output impedance up to 50 Ω. The sum port is across the whole secondary. The two input ports are between each end of the secondary and ground. The difference port is between the center tap and ground.

No two transistors or their circuit layouts are exactly equal. In a push-pull circuit these slight differences gives rise to imbalance between sides which gives rise to even harmonics at the output. By not placing the usual heavy RF bypassing at the center (the difference port) of the bifilar winding in T2

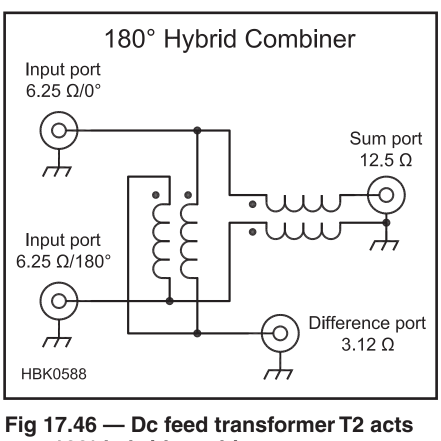

Fig 17.46 — Dc feed transformer T2 acts as a 180° hybrid combiner.

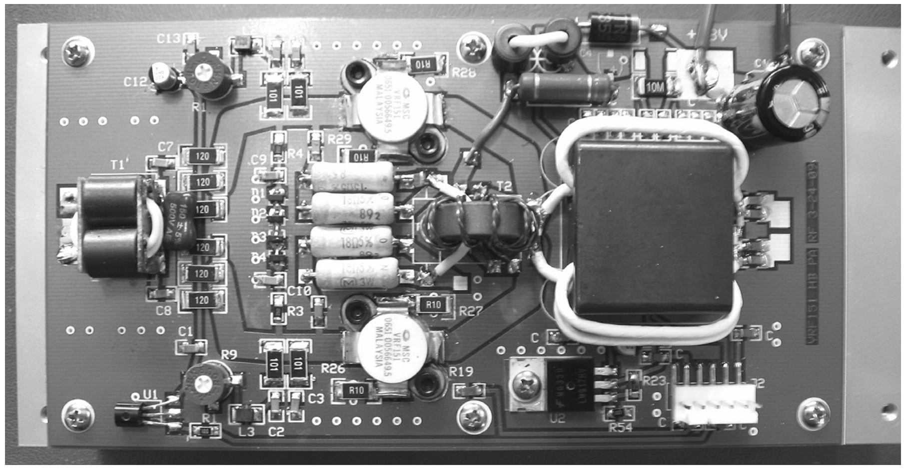

Fig 17.47 — The power amplifier module PC board assembly mounted on the heat sink.

the difference between the two drain voltages shows up across L1 and is dissipated in the parallel resistor R2. This enables the amplifier to achieve better than 40 dB suppression of even harmonics, helps the efficiency, and greatly improves its stability when driving mismatched loads.

THE PA OUTPUT TRANSFORMERS

The output circuit of the PA uses a 4:1 impedance transformer. We have two VRF151 transistors that will each put out 150 W if properly cooled. That means 300 W output power for the pair operating on 48 V. The classic formula relating drain-to-drain load impedance to the supply and output power is

$$R_L = \frac{2\ V_{dd}^2}{P_O} \qquad (10)$$

For 48 V this gives 15.36 Ω. But the drain cannot swing all the way to zero because of the finite on-resistance of the MOSFET. The actual available swing is typically 90% of the supply voltage so a more practical formula is

$$R_L = \frac{2\ (0.9\ V_{dd})^2}{P_O} \qquad (11)$$

A 4:1 transformer gives 12.5 Ω. This is a nice load for these devices.

The 4:1 impedance ratio is performed by a simple transmission line transformer. The impedance of the coax used should be the geometric mean between the two impedances to be matched:

$$Z_0 = \sqrt{50 \times 12.5} = 25\ \Omega$$

Coax with this characteristic impedance is not a common stock item but it is available as p/n D260-4118-0000 from Communications Concepts, Inc. (**www.communication-concepts.com**). Two feet are required. An acceptable alternative is #22 shielded 600 V Teflon-insulated wire such as Belden 83305-E whose physical dimensions result in approximately the same characteristic impedance.

These two coax lines are connected in parallel on the drain end and in series on the 50 Ω output end. Any voltage on the input is put in series at the output, giving a 2× voltage ratio or a 4× impedance ratio, exactly what we want. If coiled up in separate coils, the transformer will work without any ferrite. The two coils cannot be allowed to couple so they cannot be on the same form.

In order to get a wider frequency response, the inductance of these coax coils is increased by winding them on a ferrite core. The core used here has two holes, so independent coils can be wound on each side without any coupling between them and it makes a nice neat package. Separate cores would work just as well. The ferrite core is type 61 material with a permeability of $\mu_i = 125$. It is a binocular bead but could be replaced by two 0.5 dia × 1-inch long sleeves of the same material. The ferrite "load" on the coax makes its outside shield a high impedance from end to end, while inside the shield the coax maintains its 25 Ω characteristic impedance between center and shield. As long as the coax lines of T3 are wound as two non-coupled coils, the amplifier will operate from 21 to 80 MHz without any ferrite at all.

The advantage of the transmission line type of RF transformer is that it does not have the leakage reactance that plagues the tube-and-sleeve type of transformer used on many solid-state amplifiers. Simply stated, a parasitic leakage inductance is introduced in series with the primary due to incomplete coupling of the flux between the primary and secondary winding. This increases the apparent impedance of the low-Z side of the transformer as the frequency increases. The output impedance of the transistors decreases with frequency — a double hit of mismatch that causes the gain to drop off quickly.

There is one disadvantage to the transmission line transformer. It has a balanced input and output. Sometimes designers will ground one side of the output and rely on the ferrite loading to decouple the ground side of the output. This has a negative effect on the balance of the amplifier, and on those even harmonics we want to minimize. It also doubles the flux stress in the ferrite causing it to heat more. The solution is T4, a simple current balun — four passes of the 50 Ω output line through a toroid of type 61 ferrite. With T4 in place, a balanced load on the output of T3 is maintained, the even harmonics are suppressed, and the efficiency is 5 to 10% better on most bands.

PA LAYOUT

The PA board (**Figure 17.47**) was designed with all parts mounted on the top surface — no through-hole parts at all. The back side of the PC board is a continuous ground plane and is mounted directly to the heat sink without spacers.

17.11.2 Control and Protection

The control board appears far more complicated than the PA but in reality, it is just a few analog and logic ICs. The various control and protection circuits are shown in **Figure 17.48** and the LED displays and drivers (also on the control board) are shown in **Figure 17.49**.This circuitry monitors several parameters, displays them, and if necessary, puts the amplifier into standby if one of them goes out of range. This control system could be used on any amplifier. All solid-state amplifiers need similar protection. The amplifier is protected for:

1. Over temperature, by a thermistor on the heat sink and setting a limit.
2. Over current, by measuring the PA current and setting a maximum limit.
3. High SWR, by monitoring the reflected power and setting a maximum limit.
4. Selection of a low-pass filter lower than the frequency in use.

Each of these fault trips results in forcing the amplifier into the standby position and out of the RF path, and lighting an error LED. There is also an ALC level detector that generates a negative-going feedback voltage for the driver when the RF drive goes above the level corresponding to maximum power. If this PA were part of a transceiver, the several faults described above would generate inputs into the ALC system and turn back the drive rather than taking it off the air. We do not always have that luxury so the best course is to take it off line until the cause can be fixed.

THERMAL COMPENSATION

MOSFETs are sensitive to temperature. If a fixed bias is used on the gate to set the quiescent bias at 100 mA when the device is cold, the current will increase as the device heats up. In some devices, it will cause thermal runaway. The hotter it gets, the more current it draws, causing even more heating and so on. The solution is to sense the temperature of the device and reduce the gate bias as it heats up. The VRF151 is relatively insensitive compared to similar high power RF MOSFETS.

The compensation system is quite simple, effective and foolproof. It relies on the thermal characteristic of silicon diodes that as a diode heats up, the forward voltage across it goes down approximately 2.4 mV/°C. Two diodes in series are used at the bottom of each gate voltage divider (D1-D4, any silicon PN-junction diode in a suitable SMT package will work). Mounted on the PA board, they heat up along with the transistors and reduce the gate voltage by a proportional amount. The 100 mA of bias at 25 °C is less than 150 mA at 200 °C. This compensation system may not work as well if other MOSFET types are substituted because they have different thermal coefficients of V_{th} and may require more aggressive thermal compensation. Gate bias voltage is provided via a PTT-activated 9 V regulator.

A simple way to check or adjust compensation is to place the amplifier module, board and heat sink, on an electric frying pan. Set it to 212 °F (100 °C) and monitor the drain current. It should stay within 150% of the cold setting. Be patient. This takes a while because the response time of this arrangement is quite long. Under some operating conditions, it is possible for the transistor to get very hot before the heat travels to the diodes. For this reason, we have additional means to protect the transistors.

OVER-CURRENT PROTECTION

One simple protection method is to limit the power into the PA module by limiting the maximum supply current. The drain current is sensed across shunt resistor R10. The sense voltage is amplified by U2A and sent to the current meter. It also goes to comparator U2B where it is monitored and compared to the limit voltage set by R13. If it exceeds this limit, the comparator trips the fault latch U1, lights the OC LED, and opens the PTT.

OVER-TEMPERATURE PROTECTION

The heat sink temperature is monitored directly by a negative temperature coefficient thermistor, R19, mounted on the PA assembly. It is the lower half of a voltage divider sensed by U3A and fed to comparator U3B. When the thermal sense voltage exceeds the limit set by R14, the U1 fault latch is tripped and the PTT line opened. Any time the fault latch is tripped it lights an LED to indicate the cause of the fault so the operator can address the problem. Cycling the OPERATE - STANDBY switch resets the fault latch and restores normal operation.

VSWR PROTECTION

The amplifier is protected for high VSWR. The amplifier is designed to operate into 50 Ω. It acts somewhat like a constant voltage source. It will dutifully try to put the same voltage across whatever load it is given, even a short. The SWR bridge produces a voltage proportional to both the output power and the relative mismatch between the load and 50 Ω.

T1 on the LPF board is the heart of a dual directional coupler that provides both the VSWR detection and the forward power monitoring. It is really two transformers on one two-hole core, Fair-Rite # 2843010402. Each has 17 turns of #28 AWG as the secondary and a single “turn” of #22 AWG Teflon hookup wire as the primary.

T2 is a single directional coupler that looks at the power reflected from the low-pass filter. Its secondary has 20 turns of #28 AWG wire, and the primary is a single pass of #22 AWG wire. In order to minimize the impedance bump it places in the RF signal path between the PA and the LPF, the core is mounted in a small window cut in the board. This allows the toroid to be mounted so the primary wire can go straight through the center of the toroid. Detailed photos of all the transformers are contained elsewhere on the CD.

Several load conditions can produce the same value of indicated VSWR. At one mismatch load condition the amp might be trying to put out way too much power. At a different reflection coefficient, it might see a very high reflected voltage. This could raise the peak voltage on the drains past the voltage breakdown limit of the MOSFET. The voltage from the detector on the reflected power port of the SWR bridge is brought to the meter and to a comparator that will trip the fault latch when the voltage is past a limit. Again the PTT is opened. The operator can either lower the SWR by improving load match or reduce the output power to limit the reflected power from the bad match. The amplifier is happy either way.

Transformer T1 on the LPF board is the heart of the directional coupler used for VSWR protection and power metering. It consists of identical transformers wound on each side of a two-hole ferrite core. The secondary is formed from 17 turns of #28 AWG enameled wire. The primary is a single pass of #22 AWG insulated hookup wire though the hole. A picture of the transformer is contained elsewhere on the CD.

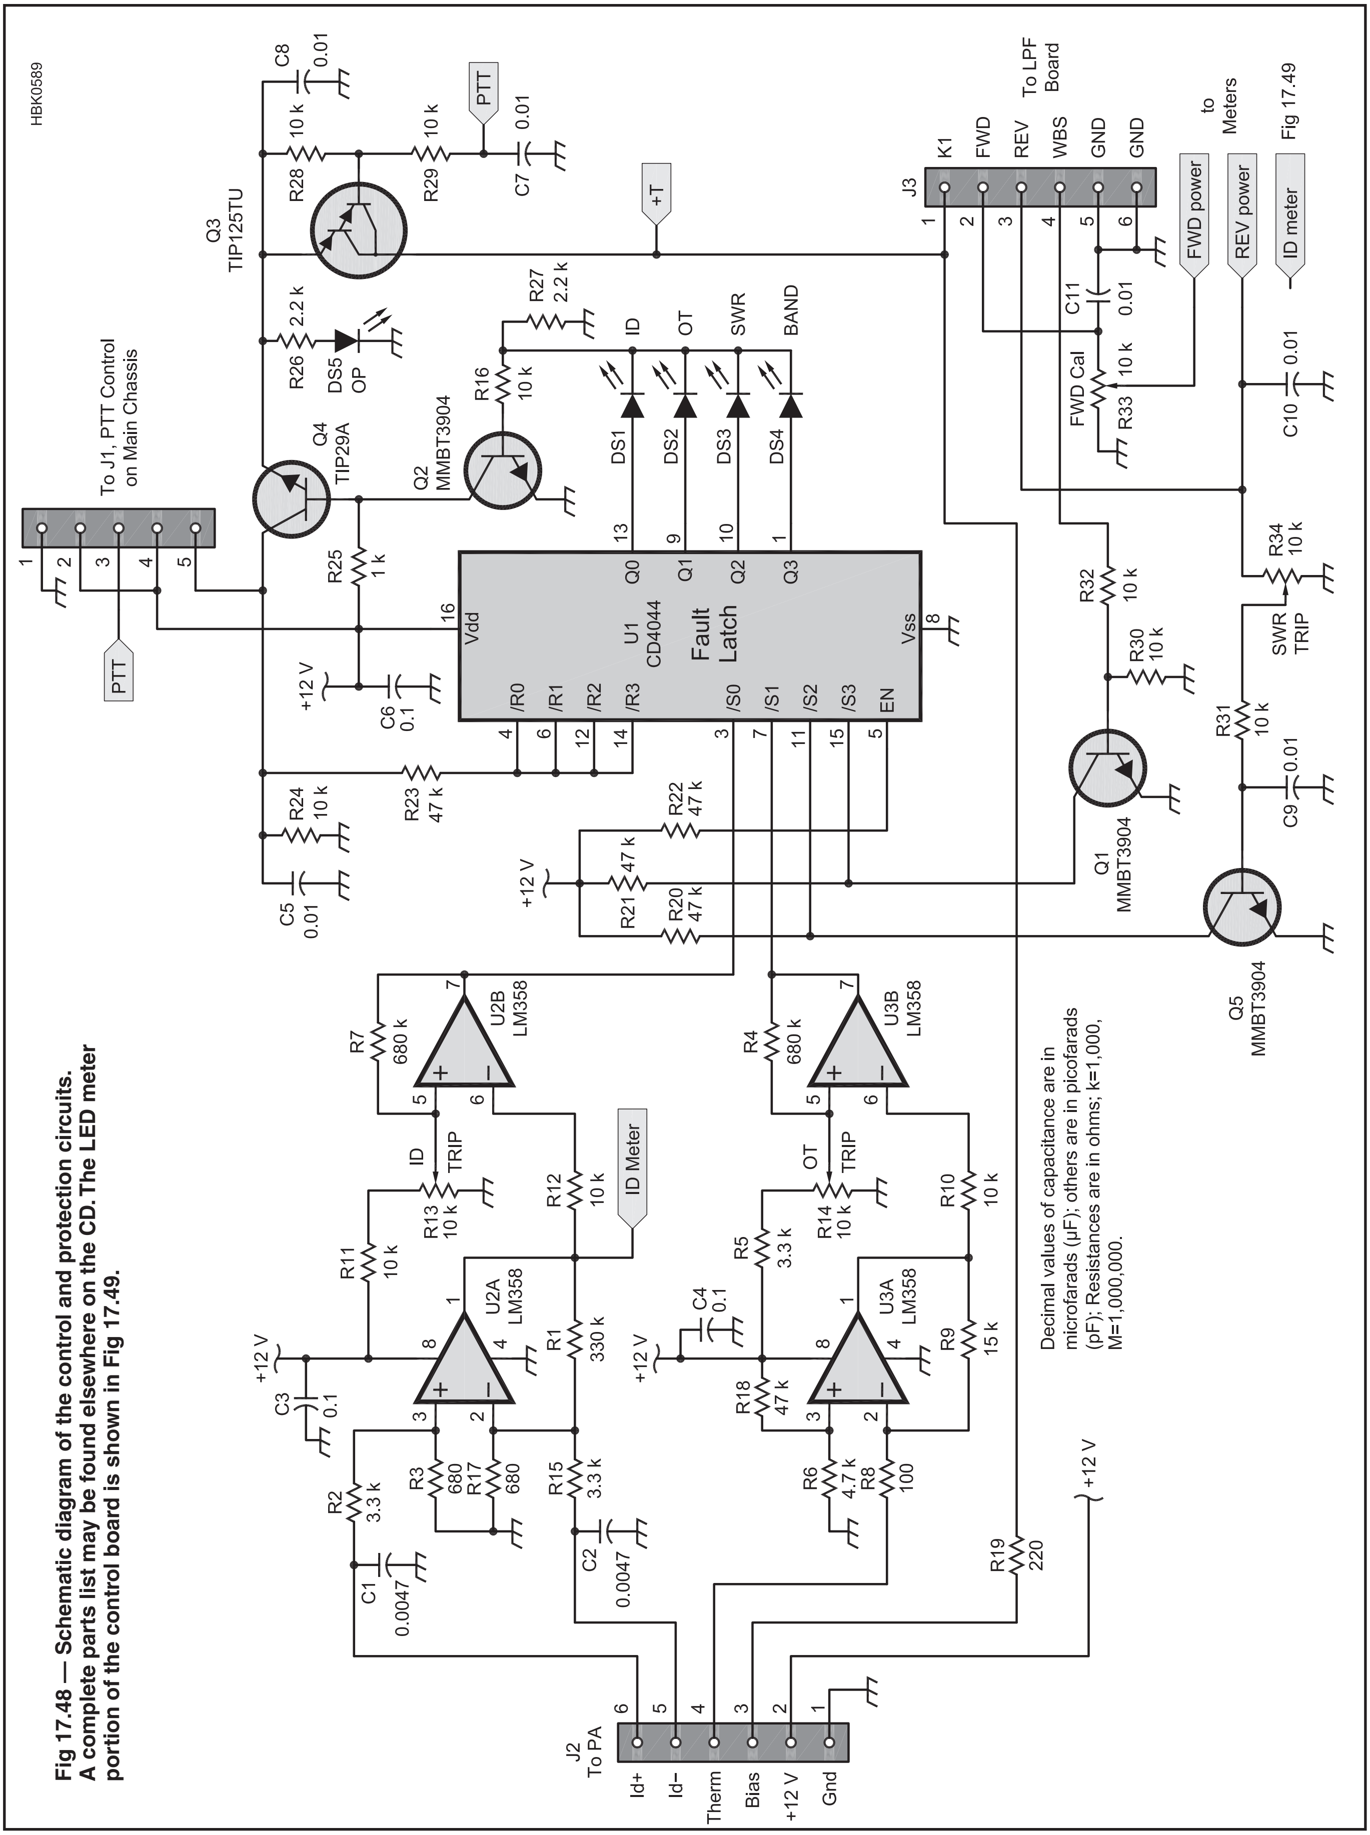

Fig 17.48 — Schematic diagram of the control and protection circuits. A complete parts list may be found elsewhere on the CD. The LED meter portion of the control board is shown in Fig 17.49.

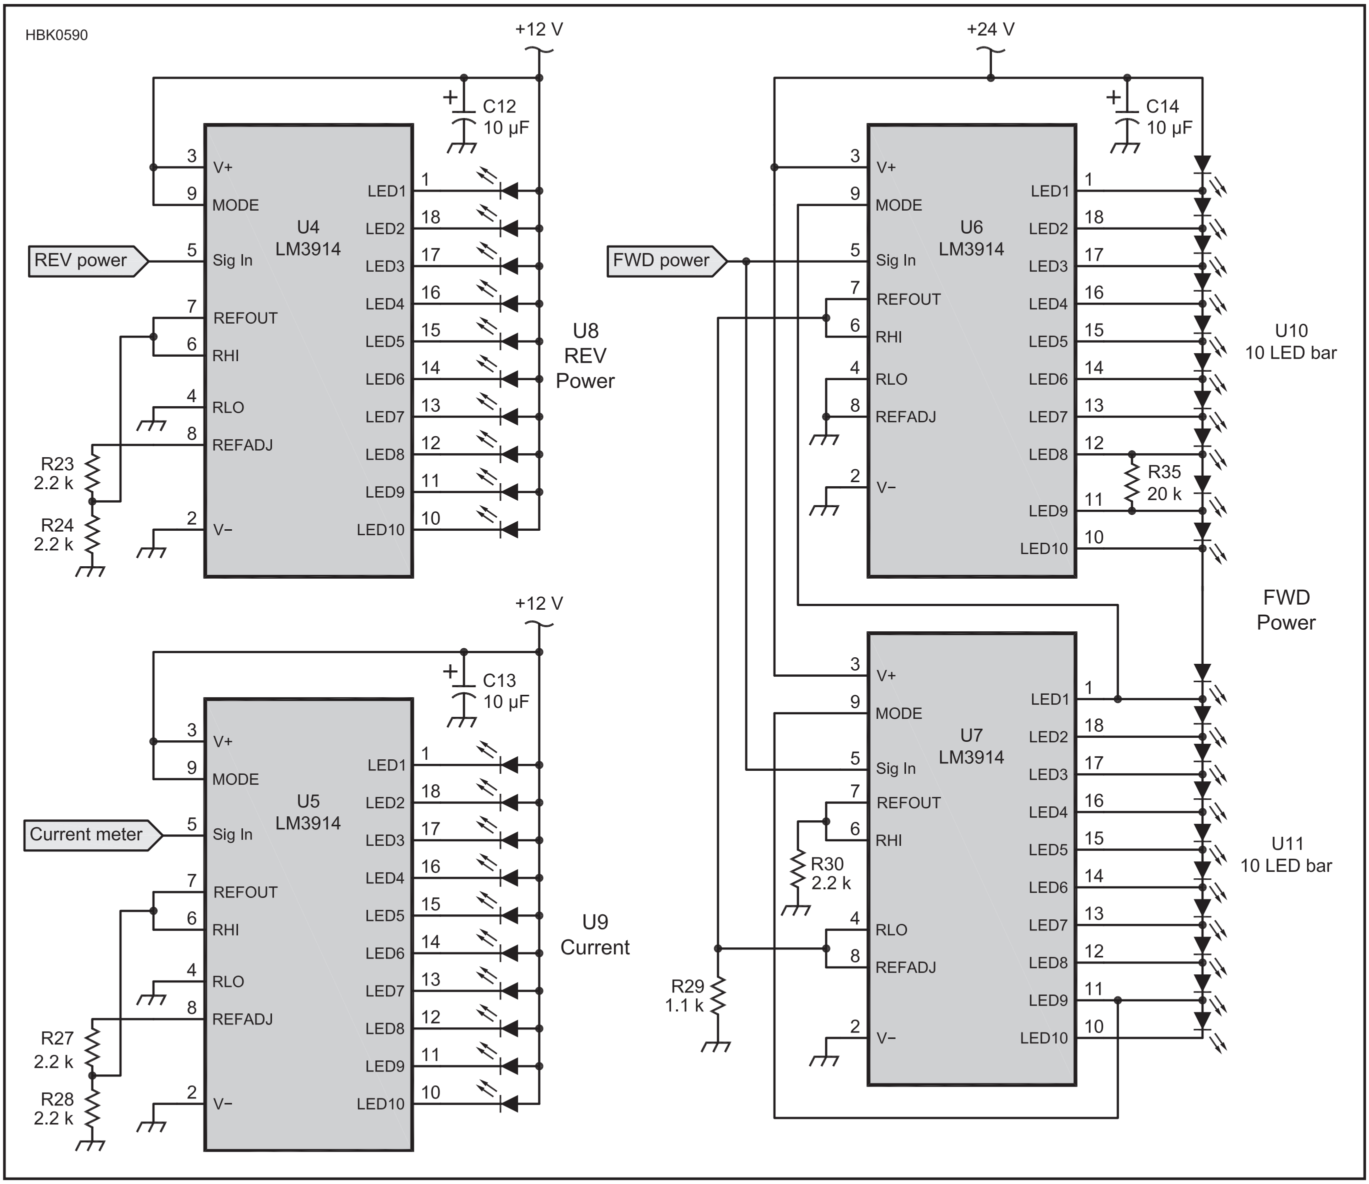

Fig 17.49 — Schematic for the LED meters, which are also located on the control board with the circuitry shown in Fig 17.48. A complete parts list may be found elsewhere on the CD.

BAND FAULT PROTECTION

One problem unique to external broadband amplifiers is that they generally do not care what frequency they amplify, but their low-pass filters on the output certainly do. The problem arises when the operator forgets to change the band switch on the amplifier when moving to a higher band. All the power from the amp is reflected back from the LPF and the amplifier is distressed — all that power with no load. The solution is to place a reflected power sensor between the PA and LPF. A monitor will see this condition and trip the fault latch for Wrong Band Selection.

T2 is single directional coupler used for reflected power only. The transformer is mounted in a small window cut in the board that allows the primary wire to go straight through the center of the toroid. This coupler normally sees all the reflected harmonic power from the LPF. The maximum harmonic power from the PA output is down 13 dB from the fundamental. See Figure. 17.50. When the wrong filter is selected, all of the power is reflected by the filter so the threshold is set to detect this 20× (13 dB) difference.

17.11.3 Low-Pass Filter

For amateur use, FCC §97.307(d) requires that harmonics be suppressed at least 43 dB for operation below 30 MHz, and at least 60 dB on 6 meter and higher frequency bands. The output signal from a broadband ampli--fier itself contains harmonics and needs to have a separate filter to meet the FCC's requirements for harmonic suppression. **Figure 17.50** shows the harmonic output of this amplifier without any low-pass filtering.

In a well designed solid-state push-pull amplifier, the second harmonic is 30 to 40 dB below the fundamental if the balance is good, but the third harmonic is only down 13 dB. This means the low-pass filter needs to supply about 10-13 dB of attenuation at the 2nd harmonic and 30 dB or more to the rest. On 6 meters the filter must provide at least 35 dB attenuation at the 2nd harmonic and 50 dB to the rest. None of this is difficult for a properly designed low-pass filter. In a commercial application, all harmonics must be at least 60 dB down. This usually

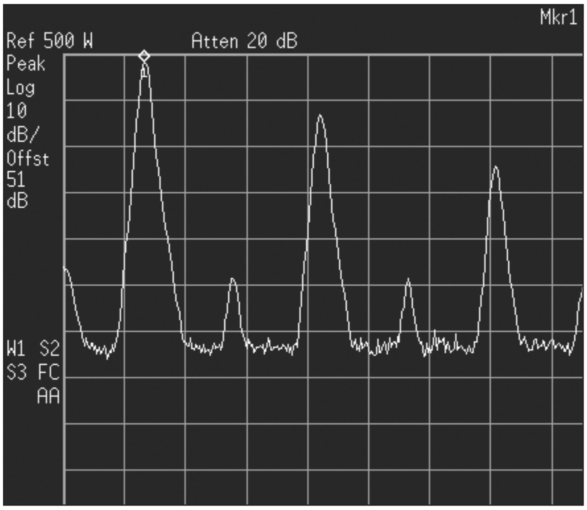

Fig 17.50 — The output of the 250 W amplifier without a low-pass filter. The even harmonics are suppressed by more than 40 dB. The odd harmonics are not suppressed by the push-pull balance and must be attenuated in the LPF. Harmonics after the LPF are all greater than 60 dB down.

requires a more complicated filter, especially if continuous coverage is desired.

The largest board in the amplifier, shown schematically in **Figure 17.51**, contains the low-pass filters, TR antenna relay, SWR bridge and the ALC detector. Amateur bands are harmonically related, so a filter for 40 meters will not do anything to reduce the second harmonic on 80 meters. If an amplifier is going to cover 160 to 10 meters, it will need at least five filters. With 30, 17 and 12 meters, the number of filters is usually increased by one to provide better suppression of harmonics. When 6 meters is added another filter is required, and it must be able to bring all 6 meter harmonics to –60 dB. This requires a more complicated filter.

FILTER DESIGN

The HF band filters can easily meet their requirements with simple five-element, 0.044-dB ripple Cauer filters. These filters use an elliptic topology. The nulls can be arranged to provide specific treatment of the third harmonic. Also, the insertion loss is the lowest of the several common filter types. The *SVC Filter Designer* software provided on the CD was used to design the LPFs used here. (See the **RF and AF Filters** chapter for more information on Cauer and other filter types.).

Low-pass filters are precision-tuned circuits. If the values are not right, the filter will have high loss and/or high VSWR in the passband. The calculated capacitor values are rounded to the closest 5% standard values. This is done by the *SVC Filter Designer* program. In a 5th order Cauer filter, there are two parallel resonant circuits that set the nulls in the response. If the exact calculated C value is not used, its paired L must be adjusted so the desired null still hits at the proper frequency. **Table 17.6** gives all the LPF capacitor and inductor values, the corner frequency (F_c), and the frequencies of the nulls (F1 and F2). L1-C2 resonate at F1 and L2-C4 resonate at F2.

LPF Inductors

Inductor winding details are given in **Table 17.7**. The low-frequency coils are wound on Micrometals T80-2 powdered iron toroid cores. For 20/30 and 15/17 meters the mix is changed to T80-6. Use of toroid cores keeps the Q of the coils high, makes physically smaller coils and provides magnetic shielding. This construction also helps to prevent the various sections of the filter from "talking to each other" and causing "suck-outs" in the passband or "lumps" in the stop band.

The 10/12 and 6 meter coils are self-supporting air-wound types with no cores which gives the highest possible Q. Coil

Table 17.6
Low Pass Filters
5th order Cauer
Ap = 0.044, As = 40 dB

Band (m)	*Fc (MHz)*	*C1 (pF)*	*L1 (µH)*	*C2 (pF)*	*F1 (MHz)*	*C3 (pF)*	*L2 (µH)*	*C4 (pF)*	*F2 (MHz)*	*C5 (pF)*
6	57.4	36	0.1571	6.8	154	75	0.1245	20	100.85	27
10/12	30.9	68	0.307	12	82.9	150	0.2387	36	54.3	51
15/17	22.265	120	0.424	22	52.09	220	0.338	62	34.771	91
20/30	15.053	180	0.619	33	35.217	330	0.458	100	23.508	130
40/60	8.837	300	1.058	56	20.674	560	0.831	160	13.801	240
80	4.778	560	2.027	100	11.177	1000	1.517	300	7.461	430
160	2.243	1200	4.181	220	5.248	2200	3.329	620	3.503	910

Table 17.7
Low Pass Filter Inductor Winding Details

Band (m)	*L1 (nH)*	*Core Type*	*Wire Size (AWG)*	*No. of Turns*	*Inside Dia. (in.)*	*L2 (nH)*	*Core Type*	*Wire Size (AWG)*	*No. of Turns*	*Inside Dia. (in.)*
6	157	—	16	5	0.312	124	—	16	4	0.33
10/12	307	—	16	7	0.35	238	—	16	3	0.33
15/17	424	T80-6	18	10	—	338	T80-6	18	9	—
20/30	619	T80-6	18	12	—	458	T80-6	18	10	—
40/60	1058	T80-2	18	14	—	831	T80-2	18	12	—
80	2027	T80-2	20	19	—	1516	T80-2	20	17	—
160	4180	T80-2	20	28	—	3329	T80-2	20	24	—

All of the cores are wound with enameled copper wire. The size is as large as will fit to maximize coil Q.

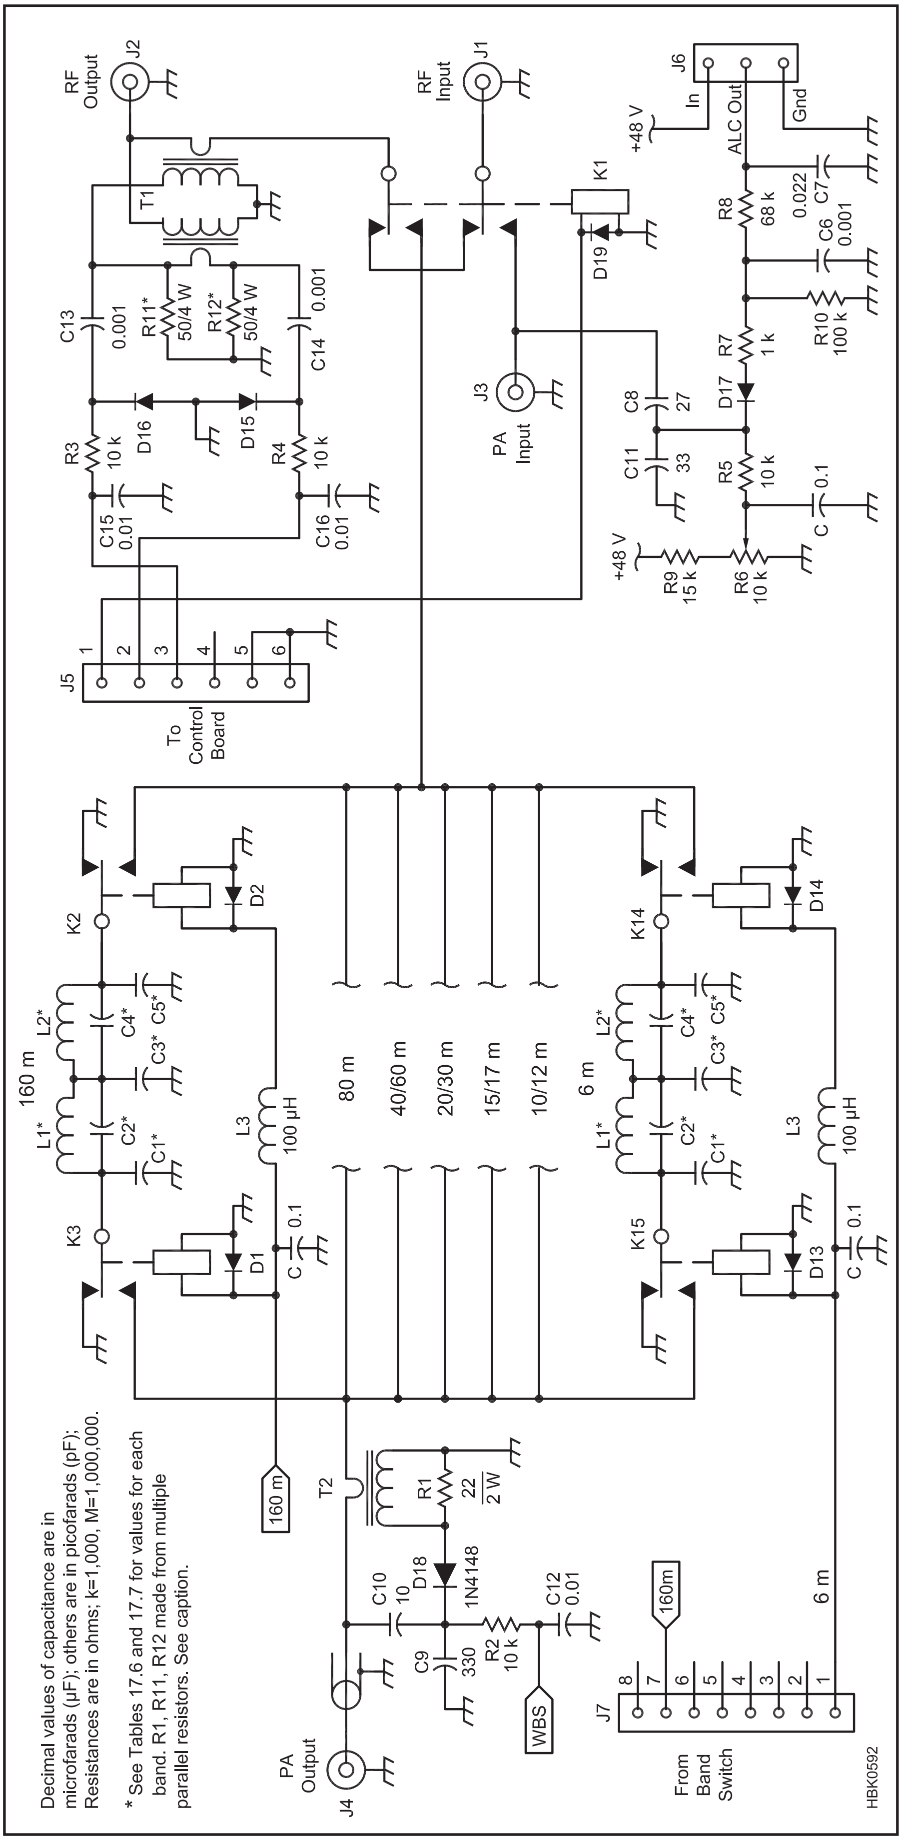

Fig 17.51 — Schematic for the low-pass filter board. The filters are described in detail in the text, and values are shown in Tables 17.6 and 17.7. This board also contains the TR relays, directional couplers and ALC circuitry. A complete parts list may be found elsewhere on the CD.

K1 — DPDT, 12 V coil, 8 A contacts (Potter & Brumfield RTE24012F).
K2-K15 — SPDT, 12 V coil, 10 A contacts (Omron G5LA-14-12DC).
R1 — Two 47 Ω, 2 W resistors in parallel.
R11, R12 — Two 100 Ω, 2 W resistors in parallel.
T1 — Dual transformer (see text).
T2 — Secondary: 20t #28 AWG on Fair-Rite 5961000201 toroid. Primary: 1t #22 AWG (see text).

adjustment is done by compressing or spreading turns on the cores. On the high bands, the coils are 5% high when wound tight. Spreading the coils slightly brings them to the proper value.

As mentioned before, it is important that the nulls in the Cauer filter response occur at the right frequency. Since the capacitor values have been rounded to the closest 5% values, the value of each parallel inductor has to be tweaked to set the null on the proper frequency. This is easy to do with a network analyzer but rather difficult for the home builder because the nulls are at various frequencies up to 154 MHz.

LPF Capacitors

Selecting capacitors for the LPF is a bit trickier. At 300 W of RF, the requirement is 123 V RMS or 174 V peak. A 500 V capacitor will easily handle this. We are also looking for RF current handling capability. At 300 W of RF, the requirement is 2.5 A RMS into 50 Ω. If a capacitor carrying this current has an equivalent series resistance (ESR) of 0.4 Ω it will dissipate 2.5 W. Harmonic currents can increase the heating dramatically.

Ceramic capacitors come in several grades of dielectric quality — X7R, Z5U and NP0 (or C0G). The NP0 is called a “Class 1” dielectric (see the **Component Data and References** chapter for more on capacitor characteristics). NP0 (C0G) capacitors are more expensive and harder to find, but they are the only type suitable for use in RF power filters. The best way to make capacitors for a PA low-pass filter is to use several in parallel. This spreads the RF current across several units and permits obtaining odd values by combining standard value parts.

RF capacitors are very difficult to find, and all the required values are never in the dealer’s stock. Several capacitor manufacturers were tried and most quoted 6 to 14 week lead times and had minimum order quantities of $100 to $500 of each value! The best ceramic RF capacitors are made by ATC. Fortunately they also had most of the values in stock and in reasonable minimum quantities. Several values were paralleled when a particular value was not available. They are quite expensive, about \$3 each for 500 V units below 200 pF and \$8 each for the larger units. The filter set uses 29 different values. A commercial amplifier manufacturer would have to make a large capital investment to stock an LPF production line, another barrier for a solid-state legal-limit amplifier.

An alternative capacitor solution would be to use 500 V silver-mica capacitors. These are suitable for all but the 10 meter and 6 meter filters where a single unit's RF current rating might be exceeded and the lead parasitics changes the net value considerably. While SMT silver-micas are available, they are also hard to find in all the values needed. The leaded versions are far more common and are made by several companies. The downside is that the LPF board was laid out for surface-mount capacitors, all mounted on the bottom side of the board. Using leaded through-hole parts will require laying out a new LPF board and it will have to be larger.

With the capacitors used, the filter is suitable for a full kW output in a system with an SWR protection circuit as implemented in this amplifier. The coils might need to be implemented on the next larger size core.

LPF Construction

The LPF board, shown in **Figure 17.52**, has a continuous ground plane on the top side. All the coils are mounted through the board. As noted, all the capacitors are leadless SMT types mounted on the bottom side. Plated-thru vias complete the circuit to the top side ground. This arrangement with the coils on the top and capacitors on the bottom provides the best space efficiency and gives repeatable filter performance. Repeatability is not so much a concern for the home builder, but critical for commercial equipment.

To prevent coupling between bands, the filters are not in order across the board. On the two high band filters, the axis of each coil is perpendicular to the next to prevent coupling between them. Finally, each filter is shorted out by the selection relays when not in use, which helps prevent "sneak paths" around the selected filter.

The filter selection relays are used in pairs. Each is rated for 10 A, overkill for this application. They are quite inexpensive, around \$1 each in quantities of 25, and have very low loss and parasitic inductance. The 10/12 meter and 6 meter filter values were tweaked for lowest loss and best efficiency after the whole filter was built. Computer modeling of the stray capacitances and inductances proved tricky and in the end it was better to temporarily replace the end capacitors with trimmers and tune for best performance to find the final values and then replace them with fixed units.

ALC DETECTOR

There is also an automatic level control (ALC) detector on the LPF board. ALC provides a feedback signal to the exciter to prevent overdriving the amplifier, and it was once a very common connection between the exciter and amplifier. Now that most solid-state transceivers have an output power control, it is used less often but still useful.

ALC interfacing is difficult because there is no standard ALC system specification among different transceiver brands and models. The detector here is a generic circuit with median values that generates a negative-going ALC signal. D17 rectifies the drive signal after a detection threshold, set by R6, has been exceeded. The attack time and filtering time constants are set by the RC network following the diode. Some adjustment of the R/C values may be required to work with a particular model of transceiver because of loop dynamic stability issues.

17.11.4 Power Supply

The PA requires 48 V at 10 A peak. It does not have to be regulated provided the following conditions are met. The maximum no-load voltage of 55 V is preferred, and it cannot exceed 60 V under any condition to provide a suitable safety margin. It must have at least 10,000 µF of filtering to prevent noticeable hum modulation. A simple power transformer feeding a 20 A bridge rectifier and a filter capacitor is sufficient. It may require a step-start circuit to ease the high current surge on the line power switch. See the **Power Supplies** chapter for more information.

The alternative is to use a switching supply. Supplies rated for 48 V and 500-600 W supplies regularly go for less than \$100 at online auction. They are small and light. You may have to deal with RFI issues on some of the cheaper models, though. This usually just requires a ferrite bead and additional bypass

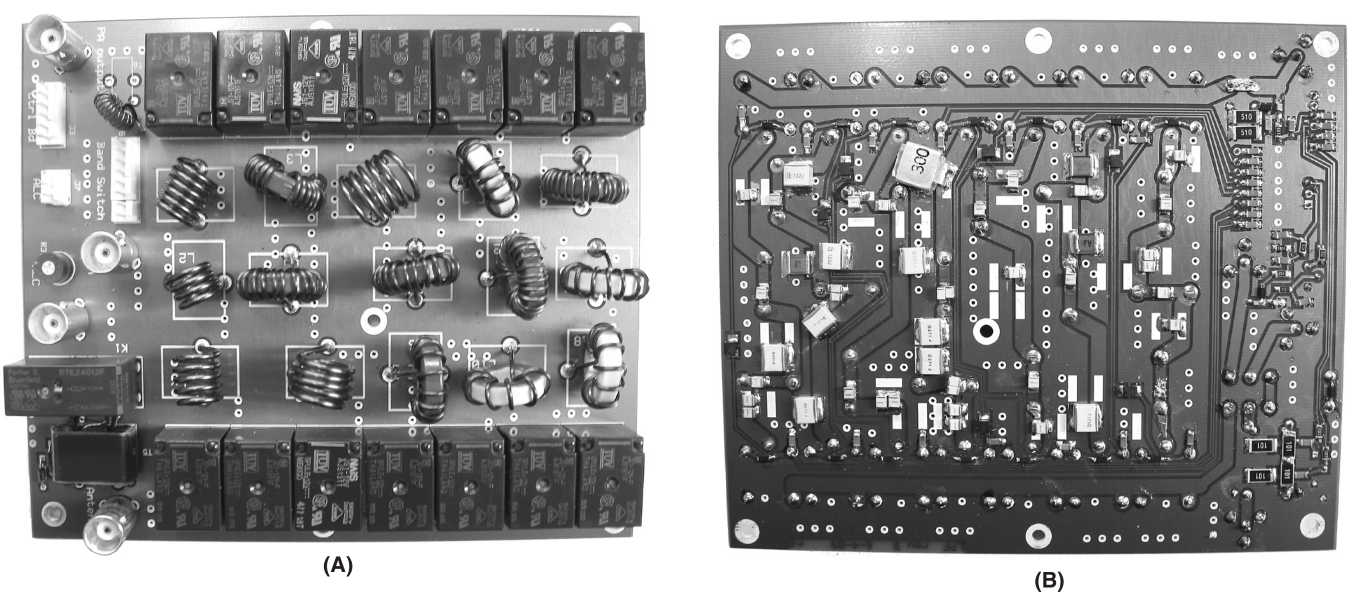

(A) (B)

Fig 17.52 — The assembled low-pass filter board with the coils, relays and other components mounted on the top side (A) and the LPF capacitors mounted on the bottom (B). The prototype used available capacitors — two sizes of ATC ceramic RF chip capacitors, a metal clad mica and three unencapsulated SMT mica units. All are serviceable, but the ATC ceramics work best.

capacitors on both the ac input lines and the dc output lines as well.

Power output varies with the square of supply voltage. The amplifier will put out 75 W if run from 25 V. The gain changes very little, just the output power. This squared relationship makes it more sensitive to ripple modulation, so make sure that the power supply, no matter how it is made, has reasonable filtering. It does not need to be current-limited since the PA control has this built-in. It does need a suitable fuse, however. A 15 A low-voltage fuse will do nicely.

The power supply used in the test amplifier was purchased on eBay. It was called "500 W 48 V 10.4 A Switching Power Supply for Radio" and was less than $100 including shipping. It is 8.5 inches long, 2 inches high and 5.54 inches wide. It was the best value of all power supply options investigated. However, as in all compromise decisions, it was not without a downside as will be described later.

17.11.5 Amplifier Cooling

Cooling is the single most important requirement for a transistor's reliability. Tubes can get hot and angry, but they are made from high-melting-point materials so they can stand the heat. Transistors cannot. They will melt. The active area of a VRF151 die is 0.034 square inch, about the area inside the letter P on a computer keyboard. It must dissipate up to 150 W when key down. The heat dissipated in the transistor die is conducted through to the transistor's package base. Here it is conducted to the heat sink which, because of its large surface area, can effectively couple the dissipated heat into the surrounding air.

Figure 17.53 shows the VRF151s on the PC board and heat sink. There is no substitute for ensuring a good mechanical fit before mounting the transistors. The heat sink surface must be as flat as possible and the transistors must also be flat. A sheet of 600 grit sandpaper on a glass plate makes a good way to check the flatness of high power metal-backed transistors. A couple of light strokes across the paper will reveal any high spots on the bottom of the device. Keep rubbing until the back is all the same color. The plating on the back is not necessary once it is mounted — bare copper is fine.

Similarly, the heat sink must also be flat. Extruded aluminum heat sinks are notoriously "lumpy." A careful swipe across the width of the sink with a large flat file, perpendicular to the direction of the fins, will reveal any ridges and valleys. File it until they disappear. Be careful to keep the file clean or you can do more harm than good. In commercial practice, extruded heat sinks are usually milled flat before use.

The interface between the transistor and the heat sink is always greased with thermal heat sink compound. Thermal grease is a suspension of zinc oxide powder in silicone or mineral oil. It is very similar to the white sun screen paint used to protect your nose at the beach. Here it is the oil that does the work. The zinc oxide is simply a filler to keep the oil from running away. Oil is not a good conductor of heat but is much better than air. The thermal grease is used to fill any microscopic gaps and scratches between the sink and the transistor. Use only as much grease as needed to fill the gaps. Spread an even thin coat of grease on the bottom of the transistor with a knife blade then put the part on the sink. Before putting in the screws, wiggle the part around on the sink to force out any trapped air and to make the layer of grease is as thin as possible. You should be able to feel the sink grab the part as you move it. It does not want to "float" on the grease!

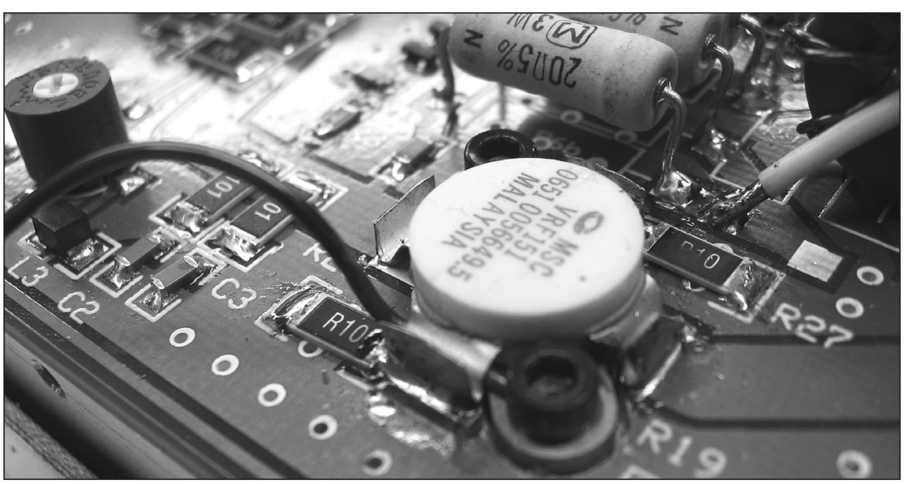

Fig 17.53 — It's important to make sure that the VRF151 power transistors mount flush to the heat sink, without gaps or air pockets that would impede heat transfer (see text for details). Once the transistors are mounted, a piece of wire is used as a lead forming device to bend the transistor leads down to the PC board pads for soldering.

THERMAL DESIGN

The VRF151's maximum junction temperature rating is 200 °C, but device lifetime is seriously degraded at this temperature. Industry design standards typically aim for 150 °C absolute maximum, and typical operating temperature about 130 °C. The heat sink chosen for this amplifier is a 7-inch length of aluminum extrusion 3.25 inches wide and 1 inch high with nine longitudinal fins. This provides a lot of surface area, but in order to remove the 250 W of heat dissipated when the amplifier is running at maximum CW power, the heat sink must have air forced through it. It is impossible to dissipate this much heat by convection alone. The forced air cooling keeps the junction temperature below 150 °C at all times. The combination of a fairly large heat sink area and a relatively small fan allows the amplifier to "keep its cool" without making a lot of noise.

A single 3-inch 24 V dc fan (47 CFM) provides pressurized air across the sink. The closed chassis forms a pressurized plenum. The air comes in through the fan but can only leave the chassis after traveling down the length of the heat sink and out the 3 × 1 inch opening in the rear panel. Ducting the air this way makes very effective use of the fin area on the heat sink.

The original prototype of the amplifier had two smaller cooling fans mounted on the rear panel. When the amplifier was put in its place on the operating table, it was overheating because both the air intake and exhaust were on the same end of the chassis. It was sucking in its own hot air. A new chassis was bent up and this time the fans were mounted on one side of the cover. This eliminated the self-heating problem and allowed a much cleaner layout of the rear panel. It provided enough space to mount the fuse holder outside rather than inside. The final version uses a single high capacity 24 V fan. A 24 V 5 W Zener diode is placed in series to run it from the 48 V supply. A 48 V fan would be a better choice, but a suitable unit was not available from distributor stock, only by special order.

Note of caution: The amplifier cannot be run at power without the top cover in place because the chassis must be pressurized in order to force the cooling air thorough the fins of the heat sink.

The fan runs at full speed all the time. It would be more "operator friendly" to run the fan at a very low speed to start off with and use the temperature sensing circuit to increase the fan speed when the heat sink reaches a moderate temperature, say 100 °F. A second threshold point would run the fan at full speed when the sink temperature reaches 120 °F. This is a project for the next version. As it is, the fan does not require any controls. It will continue to run after the supply is switched off until it bleeds down the filter capacitors.

17.11.6 Metering and Calibration

The cost of quality moving-needle meters has gone beyond reasonable, in part because digital panel meters have become more popular. However, you need to be able to assess the health of the amplifier at a glance and digital meters are not good for this. LED bar graphs are very inexpensive and are used for metering in this amplifier. They are easy to read and because of their low cost we can have one for each parameter being measured. Output power is displayed on a 20 LED string (made from two 10 LED bars). There is a "hang" built in so the SSB voice peaks can be seen easily. The drain current and reflected power are each shown on 10 LED bars. The output of the thermistor on the heat sink is also available for display by changing a jumper or incorporating a selector switch.

Absolute meter calibration is not so important in an amplifier like this. As long as none of the displays overflow in normal operation, they will serve as a way to easily monitor the operation of the amplifier.

Calibration of the meters will require using external standards. An ammeter in series with the power supply will be enough to calibrate the current meter. It should be set to display 15 A at full scale. R1 is adjusted to light bar #10 at 15 A. It would work out exactly right if all the resistors were 1% values. Forward power similarly requires a calibrated power meter on the output. R33 should be adjusted for 250 W when 18 bars of the 20 are lit. The top two LEDs are red and the others green so the operator can readily tell when the drive is too high.

Calibrating the reflected power is a little more difficult. The simplest method requires a calibrated external wattmeter. Run 30 W of RF power from another source backward through the amplifier, from antenna to input, *with the operate switch in the standby position*. Set R34 so the SWR LED trips at this point. This corresponds to an SWR of 2:1 at 250 W output. The amplifier should be stable and totally reliable up to this mismatch. If your load has a higher VSWR than 2:1, you can back off on the drive and continue to operate. As long as the reflected power is less than 30 W, the amplifier is happy.

Protecting the amplifier from damage requires a combination of careful attention to the operating conditions and reliance on the automatic limits in the control system.

17.11.7 Chassis

As seen in **Figure 17.54**, the chassis is a double clamshell configuration made from aluminum sheet stock. The top cover is 0.031-inch thick and the chassis bottom is 0.062-inch thick. This provides both adequate strength and workability with simple hand tools. Aluminum 12 × 24 inch sheets are available from several suppliers. They were sheared to size and bent into U-shaped parts at a local sheet metal shop.

The power supply and the fans are mounted on the cover. It overhangs the chassis base by $\frac{1}{16}$ inch on all sides. The top is attached by 4-40 threaded L-brackets, DigiKey part 612K-ND, that are riveted to the main chassis. A drawing of the chassis layout and front and rear panels is included elsewhere on the CD.

17.11.8 Performance

The maximum for the PA design itself is 300 W. Increasing its output past 300 W to make up for filter loss quickly degrades the IMD performance. It needs some headroom. So, in very un-amateur fashion, this amplifier is conservatively rated at 250 W output. This provides a clean signal and plenty of margin for wrong antenna selection, disconnected feed lines, and all the other things that can kill amplifiers that are run too close to their limit.

The PA will provide 250 W PEP for sideband or PSK and 250 W CW. The design goal for this PA was to make it reliable and at least as good as any competitive transceiver. The harmonics are –60 dB on HF and –70 dB on 6 meters. Transmit IMD is >38 dB down from either tone as shown in **Figure 17.55**.

Parasitics are not usually a problem in broadband amplifiers because of the feedback used. The prototype was tested into a 3:1 SWR load at all phase angles without breaking into parasitic oscillation anywhere.

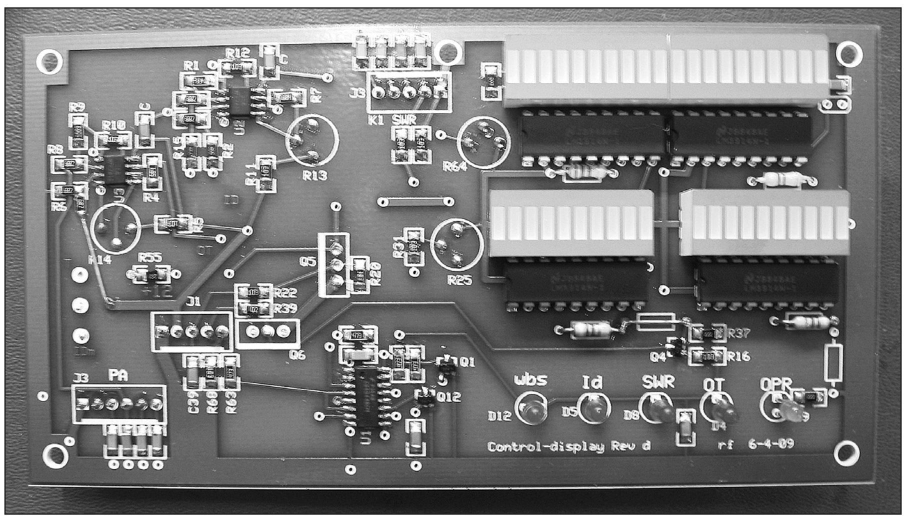

Fig 17.54 — The control and display circuitry mounts on a PC board that attaches to the front panel. LED bar graphs show forward and reflected power and current.

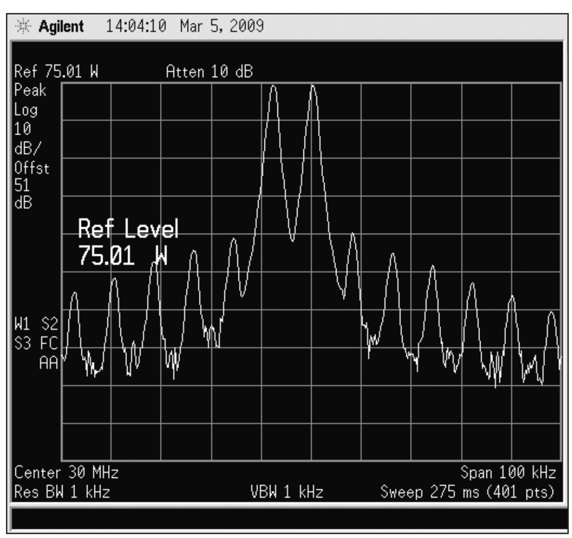

Fig 17.55 — Output of the 250 W amplifier during two-tone IMD testing. All IMD products are better than 38 dB down from either tone.