Design Notes for "A Luxury Linear" Amplifier

Here are some of the design considerations that went into a recent amplifier project, along with an additional experimental circuit and some interface suggestions.

By Mark Mandelkern, K5AM

This article is a companion to the article "A Luxury Linear," which included a general description of the circuits, performance specifications and all the schematics for a 1500-W 2-m linear amplifier.[21] But it gave few particulars explaining the functioning of the circuits or the special roles of the various components. This article will fill in most of these details. One circuit in the amp, of a more experimental nature, was omitted from the previous article and will be described here. This is the heater idle circuit, which drops the heater voltage 12% during very long standby periods—aiming to dramatically increase tube life. Finally, some problems of interface with other station components will be discussed, including driver IMD, coax relay sequencing and ALC.

[1]Notes appear on page 20.

5259 Singer Road
Las Cruces, NM 88005
email: k5am@lascruces.com

RF Circuits

Most of the details for the RF circuits are given in the caption to Fig 4.[22] Here are a few additional suggestions. There can be quite a bit of RF voltage in the plate tank. The swinging link with its Teflon sleeving should stay, throughout its travel, at least 6 mm away from the tank coil. For avoiding excessive RF voltage and arcing, it is most important to keep the loading sufficiently heavy, especially on initial tune-up. The dc-grounded link is a major safety feature. There is no sliding connection at any point of the output circuit; this is important for efficiency and to prevent erratic tuning.

The only difficult item in the output circuit is the link tuning capacitor. There was arcing in the first capacitor I tried, which had less spacing. If you find a capacitor as specified, except with wider spacing, so much the better. But you don't want less capacitance—that would mean more RF voltage. Too much capacitance would limit the effectiveness of the link tuning capacitor as a loading control. To minimize the RF voltage across the plates, I try to keep the link tuning capacitor near maximum capacitance. I use the swinging link for coarse-loading adjustment and the link tuning capacitor for fine-loading adjustment. This works well. Use the link tuning capacitor for loading; if you approach maximum, nudge the link in a bit, and if you find you are down towards half capacitance, ease the link out a bit. Sounds complicated. but I haven't touched the swinging link in the past six months. Another advantage of this loading procedure is that adjusting the link tuning capacitor has very little detuning effect on the tank coil, whereas swinging the link has a noticeable effect.

The voltage divider for HV metering provides a 10,000 to 1 sampling ratio at the HV2 point. The 20-MΩ resistor in Fig 4 consists of two 10-MΩ special glass HV types, with neat spiral resistive traces wound inside the glass tubes. These types are found now and then in the surplus catalogs. A string of twenty 1-MΩ, ¼-W, 1% metal-oxide resistors would be suitable (Digi-Key #1.00MXBK).[20] The resistors are specified as 1% types not because accuracy is required here, since there is a calibrating adjustment in the metering circuit (Fig 10), but because the precision types will be more stable, thermally and over time.

Don't miss the note on C5 in Fig 4, about not connecting any additional bypass capacitor at the screen terminal. The 100-Ω resistor has an important decoupling function.

It may be possible to reduce input drive requirements below 30 W. Owners of 25-W transceivers might be especially interested. A heavy, wide silver-plated strap for the T-match coil would be the first thing to try. I did try a heavier wire in a hairpin loop, with no improvement. That's when I concluded that most of the loss was inside the tube. My driver is capable of 200-W output. So, except for trying to hit the magic 25-W input spec, I had little motivation to work further on the input circuit. The T-match tuning capacitors might be one place to try for improvement. Glass piston trimmers might be better. Trying to read RF voltages at each of the three terminals was wild! One high, one middling (the one fed), one almost nothing. Adding a heavy wire connecting the three grid terminals didn't help. Feeding just one grid terminal doesn't seem to make the other two jealous. The geometry of the socket and the tube base indicates that it may be pointless to worry about this asymmetry; the solid grid ring built into the tube base has less inductance than anything that could be built around the socket. This amp shows a noticeable improvement over the old 1000-W amp; I heard my EME echoes with a horizontal single Yagi at moonset—this is not a first, but it might be for an all-homebrew station.

Power-Supply Circuit

Most of the power supply is straightforward and routine. The heater regulator uses the ubiquitous 723 IC. A minor complication arises from having both sides of the heater above ground, due to the cathode current shunt and the construction of the tube, with the cathode internally connected to one side of the heater. The heater voltage is applied to the A+/A– points, while the S+/S– points provide heater voltage sensing at the socket. Thus the entire regulator circuit is referenced to the S– point. Current for the 723 does flow through the cathode metering shunt, but it hardly moves the plate meter. The heater power supply is returned to the A– point. As noted in "A Luxury Linear" (Note 21), the heater supply did not reach the 90- to 130-V ac operating goal, dropping out at 98 V. I was content with that and did not try to improve it, but it should be easy to do so. I used #18 wire for the A+/A– leads. Since the regulator senses voltage at the socket, drop in these leads is of no concern; but heavier wire could be used to improve low-line-voltage performance. I think, though, that the greatest drop may be in the heater fuse holder, an ordinary 3AG type—it gets pretty hot! An automotive type blade fuse and socket will probably have less voltage drop.

The choice of heater transformer is the main factor determining line regulation, however. I used a Signal #36-6. Rated 18 VCT at 12 A, it is much heavier than needed in the FWCT circuit used. But on the surplus market (searching 20 catalogs) it cost less than the next smaller size available, which was a bit too small.

Screen Regulator

Not so commonly discussed, and of special interest (judging from inquiries) is the screen voltage shunt regulator. Power tetrodes, and especially the 4CX1000A, commonly exhibit negative screen current flow due to secondary emission. Feisty electrons from the cathode hit the screen and knock off more electrons, even more than those arriving. Like, throw one ping-pong ball, forcefully, into a bucket of ping-pong balls, and see a dozen bounce out. To the amp builder (who, in my case at least, disavows knowing anything about what's going on inside the tube) it seems as if current is coming out of the tube. (Current in my shack flows from positive to negative.) This negative screen current flows into the screen supply and tends to increase the screen voltage, which if unchecked will destroy the tube. (The actual source of the negative screen current is the plate supply.)

A series regulator will not suffice to deal with this negative screen current problem; only a shunt regulator will do. VR tubes are traditional, and Zener diodes have been recently in favor. I wanted a fully adjustable supply to enable experimentation with different operating parameters. I also felt that power transistors are inherently more reliable than Zeners. And the transistors give you more watts per dollar. Also, the transistors used operate very

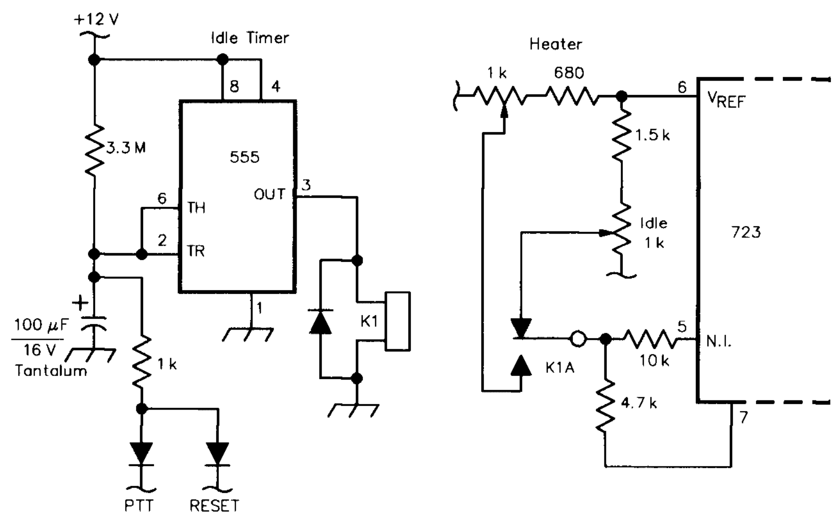

Fig 11—Schematic diagram of the heater-idle circuit. This is a modification of a portion of the heater-regulator circuit in Fig 5. Refer also to the caption for Fig 6.

far from their maximum ratings while Zeners, in this application, typically operate very near their ratings. The shunt transistor regulator was only a bit of extra trouble and has worked flawlessly.

Well, in fact, the first shunt regulator tried, using a 723, didn't work at all! It worked fine on the bench, in all positive and negative current tests, but the shunt transistor shorted when I put the amp in the rack and fired it up. It's best to write about what works, rather than what didn't work, but in this case the story includes a few warnings and explains the reason for some components in the final circuit. I made a half dozen changes, and the second regulator, as shown in Fig 5, performed perfectly from the start. It's hard to say exactly what went wrong with the first circuit, but here are the guesses: oscillation in the 723, possibly triggered by the very jumpy screen current variations during CW, SSB or pulse tuning operation. RF getting into the sensitive 723? (That would explain erratic behavior, but would not be enough by itself to cause the transistor to short out.) Excessive current in the shunt transistor. (In a 5-A transistor?—Note the energy stored in the 22-µF electrolytic capacitor before the 3-kΩ resistor was added.) Puncturing of the mica under the shunt transistor—even two micas. (At 300 V?—What's going on here?—I wish I understood transients better.) Dynatron oscillation? (See page 54 in the reference in Note 3.) I never took a course in electronics, so I can only guess what true failure analysis in an industrial environment would involve: digital storage scopes, chart recorders and a bucket full of transistors (paid for by the company). So, after losing three of the four transistors I had on hand, all I could do was go on hunches and try something a bit different.

The screen regulator that works doesn't use a 723, but merely a single, fairly high-gain transistor as driver for the shunt transistor. A certain amount of gain is necessary, but more than is needed might bring in stability problems. The Zener to the left of the 2N2222A in Fig 5 is merely to provide collector voltage within the transistor ratings; using a high-voltage transistor here would have meant lower gain. The other 15-V Zener is the regulator reference, along with two V_{be} drops in the transistors. Say the reference is then 16.2 V. The output voltage will be this reference multiplied by the ratio set up by the voltage divider at the output, including the screen voltage adjustment pot on the front panel. This ratio having the range 12 to 22, the nominal output voltage is about 200 to 360 V.

There are two details that may be crucial. I used no micas, but mounted the TO220 shunt transistor directly onto a 2×3×1/8-inch aluminum plate, which in turn is mounted on the side wall using ceramic standoff insulators. Second, I added the 3-kΩ, 20-W resistor in the shunt transistor collector lead. This limits the peak (transient?) shunt-transistor current to about 100 mA (at 300 V) and prevents a possible momentary saturated crowbar short to ground. The resistor could be lower valued, and sink more current, but the amp includes an overload circuit set to 30-mA negative screen current, so a regulating range of 100 mA is adequate. Normal operation results in about 20-mA negative screen current peaks during SSB operation. The 0.1-µF capacitor on the collector and the 22-µF capacitor at the output are intended to soften transient pulses which were thought to make the transistor unhappy.

The regulation obtained is only about 1%. This is intentional, as I was most concerned with stability, which is inversely related to regulation. The degree of regulation depends on the gain of the transistors and the resistance of the voltage divider at the output since base current for the first transistor flows through the 120-kΩ resistor. Lowering the divider resistance would improve the regulation, but 1% regulation is more than adequate. It's sometimes difficult, in these cyberdays, to remember that the super-precision available is not always needed—or desired.

From the E = 400-V supply, the R = 5-kΩ (10-W) series resistor limits the forward screen power to $P = E^2/4R$ = 8 W, well within the 12-W rating. Thus no forward screen current overload circuit is needed.

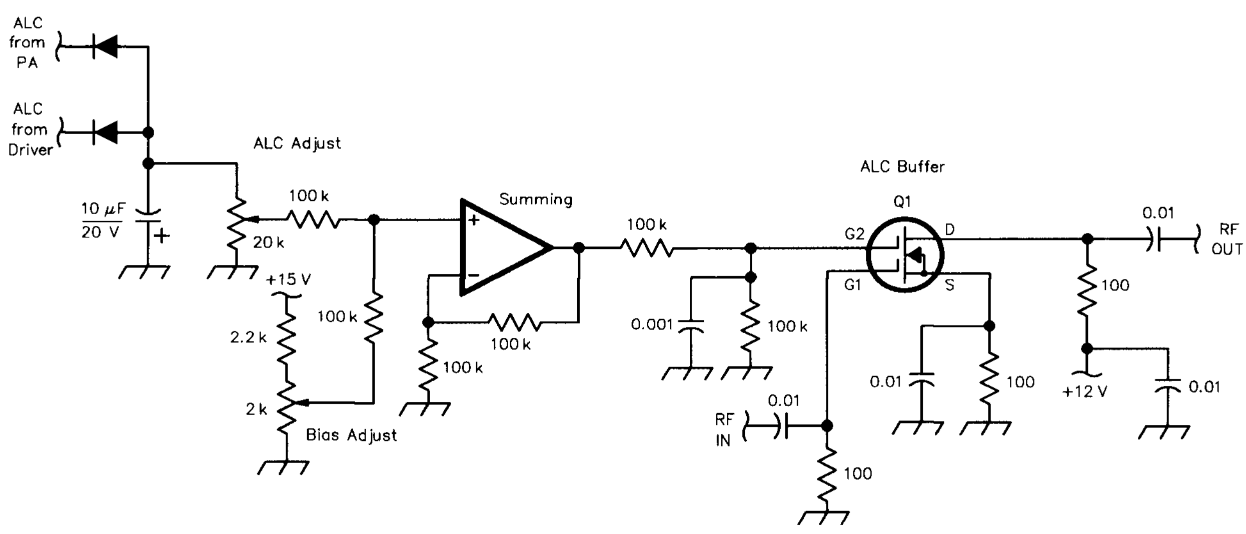

Fig 12—Schematic diagram of an ALC circuit that may be added to a transceiver or transverter. The MOSFET Q1 may be any one of dozens of small-signal RF types (40673, 3N211, ...) often found on the surplus market. The NTE 454 replacement type is available from Hosfelt.[19]

Control Circuits

The warm-up timer in Fig 6 is powered by the +12H rail, only after the air-flow switch, at terminal AIR, is actuated. Several seconds are required for the blower to get up to speed after the line switch is thrown; then the blue light comes on and the timer starts. But only after voltage appears at the heater. This is sensed by the op amp connected to the S+ terminal. So, for example, if the fuse is blown, the timer will not start until after you replace the fuse.

The cool-down latch in Fig 6 is a bit tricky. The basis of operation is the fact that when the line switch is thrown off, the +12-V rail drops more quickly than the +24-V rail (which has minimal load except while transmitting). This throws the latch high and starts the cool-down timer. Since the timer starts in the low state, the discharge pin at the COOL terminal is low, and the cool-down relay K3 in Fig 5 is energized. After about three minutes the timing capacitor charges enough to bring the trigger terminal low and set the 555 high, opening the discharge terminal and allowing the cool-down relay to drop out.

An unexpected anomaly showed up when this simple timer was first tried. Not a very serious defect, but perhaps an amusing incident in the quest for perfection. Here is the situation. After working a few choice DX stations, the band seems dead and you throw the line switch off. Before the cool-down circuit times-out, you hear another DX station, change your mind, and throw the line switch ON again. After allowing the tube to warm-up (see page 8 in "A Luxury Linear"), you work the station. Congratulations! Now you throw the line switch off and are surprised to hear the blower immediately stop! What happened? The cool-down circuit was in COOL mode and timed-out while you were working the new one. That's the reason for the connection to the +24RLY line. What does the cool-down circuit have to do with the coax relays? The +24RLY line resets the cool-down timer (if latched high) the next time you transmit; this was the simplest band-aid solution.

The regulated heater supply, with 10-V at the unregulated point, made over voltage protection (OVP) absolutely essential, to protect against a short-mode failure in the pass transistor Q1 in Fig 5. After hearing of troubles with fold-back and SCR methods of OVP, and considering the various failure modes, an op-amp circuit and a relay seemed to afford the highest reliability. At the bottom of Fig 6, the first op amp acts as a comparator. When the heater voltage exceeds the preset limit, the op-amp goes low and pulls the OVP latch low, removing the ground return for the heater relay coil at point HR in Fig 5. The 10-µF capacitor in the OVP circuit ensures that the OVP latch goes high on start-up when the line switch is thrown ON.

The 2.2-µF capacitor in the overload circuit (Fig 9) forces the overload latch low (normal) on power-up. When low, the latch holds the flash oscillator high and the red light off. I suppose the inverter could have been avoided if, when I got to that point, surprised to find myself standing on my head, I had gone back and redesigned the entire control board. But a low parts count was not a major design goal.

The air-flow switch protects the tube in three ways. Not only does it control the supply voltage (+12H) to the 723 in the heater regulator, and by the same line enable the operate circuit in Fig 7, but it also controls the coil voltage to the heater relay in Fig 5.

Coax Relay Sequencing

VHF operators are especially interested in antenna-relay sequencing devices. In my station (all homebrew), sequencing is included internally in the transceiver, each transverter, each driver amp and each kW amp. This provides the highest level of protection. An MGF-1402 GaAsFET preamp is still working perfectly after ten years with the original transistor. The sequencing circuits used here may be incorporated in other transceivers and transverters for fool-proof coax relay and preamp protection. The sequencing circuits are very simple and much more reliable than many excessively complex systems.

The two op amps shown in the sequencing circuit in Fig 7 use very simple R-C networks to obtain switching delays. Notice that the diodes are reversed. The PTT op amp shifts to an output of about +11 V in transmit. At the coax relay op amp, the diode allows this to charge the capacitor almost instantly, so the coax relay can begin to close. But at the screen relay op amp, the positive control signal must pass through the 68-kΩ resistor, so the screen relay is delayed. At the end of the transmission, the PTT op amp shifts back to an output of about −11 V, and the procedure is reversed.

Keying the amplifier on CW requires special care to avoid any modification of the keying waveform produced by the transceiver. Assuming that the transceiver has the desired keying waveform (my homebrew transceiver has separate internal adjustments for make and break), we want the amplifier to pass this along to the antenna without alteration. The keying circuit at the bottom of Fig 7

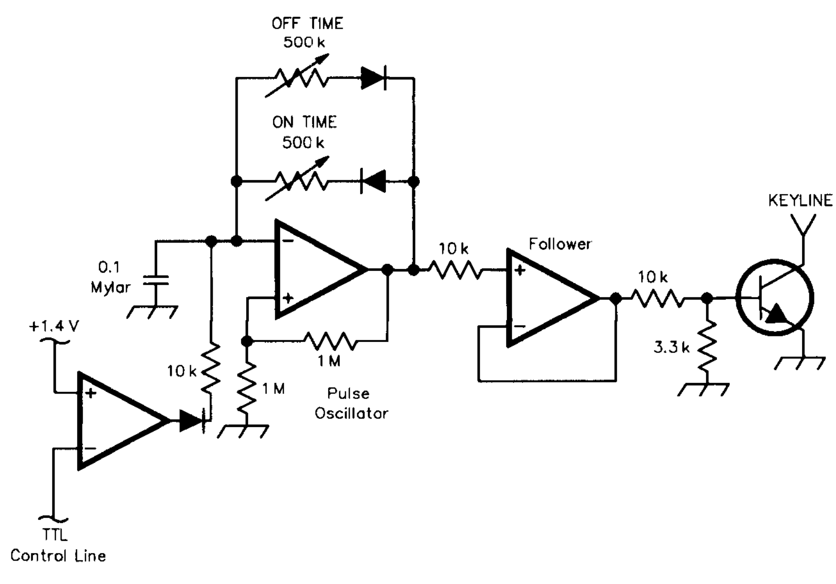

Fig 13—Schematic diagram of the pulse-tuning circuit used for low-duty-cycle tune-up and adjustment. Refer to the caption for Fig 6.

provides for this. The diode allows the key closure signal, a negative output from the first op amp, to almost immediately key the bias circuit (at the KEY2 point), so the amp is up and ready for the dit when the transceiver finally gets around to sending it (about 2 ms later). When the key opens, however, the positive output from the op amp must force its way through the 6.8-kΩ resistor to charge the cap. This allows sufficient time to keep the amp keyed while the dit is decaying. The second op amp, with hysteresis provided by the 1-MΩ feedback resistor, switches sharply when the capacitor crosses the zero potential point. This is a fast-switching delay circuit; it does not produce a keying waveform. It has a square-wave output which encloses the transceiver keying waveform, allowing the transceiver alone to shape the transmitted CW element.

The key sequencing also has a connection to the screen-relay sequencing circuit. This makes the keying dependent on the operate, PTT, sequencing and overload circuits, for an extra bit of protection.

In connecting the amplifier key line to the station, the easiest way to avoid unwanted interaction would be to build an interface circuit that allows the station keyer to separately, but simultaneously, key the transceiver and amplifier. However, for many transceivers it may be possible to simply tie the key lines together. The transceiver schematic must be checked; can it handle 12 V on the keyline? Can an isolating diode be added?

Metering Circuits

This was the really fun part! Imagine: training op amps to do arithmetic! The op amps here are used in three ways: the noninverting amplifier, the inverting summing amplifier with virtual ground and the differential amplifier. These amplifier circuits are described fully in Chapter 8 of the *Handbook*.[7] In the following, refer to Figs 8.44 and 8.46 to 8.49 in the *Handbook*, along with the associated formulas and notation.

The simplest noninverting amplifier used here is the one for HV metering. The meter is to read 5000-V full-scale, and the divider inside the plate compartment (Fig 4) yields a 10,000 to 1 ratio, so 5000 V would produce 0.5 V at the HV2 point in Fig 10. The 10 kΩ resistor in the noninverting (+) input, and the 10 kΩ resistor inside the plate compartment (which is for filtering, in conjunction with feed-through and bypass capacitors not shown on the diagram) have a negligible effect on the circuit, as the op amp input is of very high impedance. Our op amps on the +12/–12 rails have maximum outputs of about +11/–11 V, so we ask for +10 V output at full-scale, using a 10-kΩ meter multiplier resistor and a 1-mA meter. We'll need a gain of 10/0.5 = 20. Arbitrarily choosing 10 kΩ as input resistor R_i, we calculate 190 kΩ for the feedback resistor R_f. I usually choose R_i = 10 kΩ because I have a whole box of 10-kΩ resistors. For R_f, I choose a trimmer at about half the nominal value, and then a fixed resistor, chosen from standard values, to obtain the best range. In this case, a 100-kΩ trimmer and a 150-kΩ fixed resistor yield a range of 150 kΩ to 250 kΩ. That's a fairly good bracketing of the 190-kΩ goal; it doesn't usually work out so well. A 200-kΩ trimpot would not permit sufficient adjustment range, allowing for the tolerances of the circuit components. So omitting the fixed resistor would mean using a 500-kΩ trimpot and having a very touchy adjustment situation.

Plate-current metering is similar. The cathode shunt at the socket (Fig 4), with 1 A of plate current, will produce 0.33 V at point S– in Fig 10.[23] A gain of 10/0.33 = 30 will yield the 10 V we want for the meter with its multiplier. With R_i = 10 kΩ again, we'll need R_f = 290 kΩ. "Shunt" is the traditional term, as if we were to connect a meter directly, but here the shunt merely converts current to voltage, which the op amp amplifies, and the meter multiplier then converts back to current. The 1-µF capacitor provides some meter damping, giving a less wild indication on SSB.

Screen-voltage metering is also similar to the above. So are forward- and reverse-power metering, although the amplification needed for these may depend on the particular construction of the directional coupler. The forward-power meter is peak-indicating. Although very simple, it is extremely useful in conjunction with pulse tuning (see pp 11-12 in "A Luxury Linear" and the "Operation" section below). The diode enables fast attack with a positive signal from the forward power coupler, while the 1-MΩ/1-µF network provides the hold function. The time constant is chosen so that the meter holds long enough for full-power indications under pulse tuning and CW operation, but reacts fast enough to follow tuning adjustments.

Zero-Center Screen Metering

Perhaps more interesting are the zero-center and expanded-scale metering circuits. Zero-center screen-current metering is essential with the 4CX1000A. I could find no zero-center meter on the surplus market to match the others on hand. Finding five similar-looking meters on the surplus market was the hardest part of the whole project. The alert reader will have already noticed that there are actually three different Triplett types (with identical cases) in the cover photo for "A Luxury Linear".

The inverting-summing amplifier with virtual ground is one of the most useful of the op-amp-arithmetic connections; we use it for zero-center screen current metering. The secret is the virtual ground. The noninverting input (+) is really ground, while the op amp with inverse feedback tends to keep the inverting input (–) at the same voltage level as the noninverting, namely zero! So the inverting input is always at 0 V; it acts like a ground. The main advantage of this circuit, compared to the noninverting summing amplifier, is that there is no interaction between the several inputs.

We want a 50-0-50 mA screen meter. The 100-Ω screen-current shunt in Fig 5 will produce –5 V at the –E2 input for 50-mA of positive screen current. Thus we need a gain of –1 to obtain the desired +5 V for half-scale deflection in this inverting circuit. Again with R_i = 10 kΩ, R_f will have a nominal value of 10 kΩ; we use a 5-kΩ trimpot and an 8.2-kΩ fixed resistor.

Now we center the meter. We need only fool the circuit into thinking there is 50 mA flowing. We want an op-amp output of +5 V with no screen current. Using the –12-V rail as an input for this purpose, we need a gain on this branch of +5/–12 = –0.417. With R_f already chosen as nominally 10 kΩ, we calculate R_i = 24 kΩ in this branch; the components shown give a range of 18 kΩ to 28 kΩ. The zero setting hasn't moved a hairline in the two years since the initial adjustment.

Expanded-Scale Heater Metering

The differential amplifier is the circuit of choice for measuring a voltage between two points, neither of which is at ground. Alternatively, the virtual ground in the summing amplifier part of this circuit could have been made virtual S–. But while that would have worked for measuring the voltage, the expanded-scale part of the circuit uses the +12 rail, which is ground-referenced. (The regulation of these rails is essential to the stability of these circuits.) That means another regulator would be needed, referenced to S–. So the differential amplifier is really the simplest solution. The circuit used here in Fig 10 is particularly simple. The heater voltage sensing is at terminals S– and S+; let me call the voltages at these terminals E_1 and E_2, respectively. With all four resistors equal (the value doesn't matter), the output of the differential amplifier is $E_1–E_2$. The differential amplifier therefore simply measures the heater voltage and converts it to a ground-referenced voltage.

The summing amplifier is used to obtain the expanded-scale heater metering, 5.0 to 6.0 V. We want a 1-V change in heater voltage to produce a 10-V change at the op-amp output, for full-scale indication. Thus a gain of –10 will be needed. The differential amplifier has been arranged to invert the heater voltage once, so once more and we're on our feet again. With R_i = 10 kΩ in the measuring branch, we want R_f = 100 kΩ. For an expanded scale indication, we use the expanded scale branch to give the circuit a tendency to indicate –5 V (of course the left meter pin gets in the way); then it will take +5-V input at the measuring branch to indicate a composite zero. The meter has a 1-V full-scale range, so this means a hypothetical –50 V at the op-amp output. Using the +12-V rail for this, we want a branch gain of –50/+12 = –4.17, and we need R_i = 24 kΩ. The individual calculated outputs due to the two input branches are superimposed—they add. Thus the output is $10(E_2–E_1)–50$. The output is zero when the heater voltage $E_2–E_1$ is 5.0, and the output is 10 V when the heater voltage is 6.0.

Adjustment of these circuits is easier than it might seem at first. All adjustments are made with the amp on the bench using small test power supplies to simulate the parameters. Not with high voltage while transmitting! For most tests and adjustments the tube is even cold; a switch turns off the blower, so there is no noise on the bench and the heater is off. In each circuit, first adjust the CAL trimpots (for correct deflection, no matter what part of the scale), and then the CENTER or LEFT trimpots. (I should have put the CAL trimpots at the input—then there would have been no interaction between the CAL and ZERO or LEFT adjustments.)

All the trimpots are at the front edge of the control board. For touch-up adjustments, the amplifier may be slid a few inches out of the rack, leaving all cables attached. The trimpots are then accessible through the top cover vent holes.

Construction

Blower mounting was one of the toughest jobs. The low-noise requirement demands mounting on rubber. The suggestions here are not exactly the way I installed the blower—they are hopefully easier. A good source for material is your local auto parts supplier. Vacuum hose is very thick-walled, with a small ID. Slicing lengthwise through one wall yields a heavy piece of rubber that can be slipped over the edge of the blower outlet. Fig 3 shows that I supported the blower by three screws, with rubber grommets, to the rear panel, but my suggestion is to fabricate a bracket inside the amp and support it from above. The weight of the blower, and a bit of pressure from the bracket, will keep the rubber gasket snug and air-tight.

Inside the grid compartment, across the air inlet, is an RF shield made of copper screening. The air-flow switch is mounted inside the grid compartment; the steel actuating wire on the switch passes through a small hole in the copper screen to the air vane, which is fully inside the blower outlet tube. Alternatively, and more easily, the air flow switch may be mounted on the side of the outlet tube.

After a few clumsy attempts, a very simple method for assembling the input circuit T-match (with insulated above-ground rotors) was found. The key idea is that the two rotors connect together and to the coil (Fig 4 shows the rotor of C2 incorrectly). The two capacitors mount on a piece of one-sided copper circuit board (Fig 2), which is mounted on ceramic insulators. This results in a very low-inductance connection between the two capacitor rotors. The coil connects between the circuit board and a ground lug. The lug is copper; use a small magnet to reject the steel solder lugs in your junk box. The silver-mica blocking capacitor, seemingly redundant, is to protect the tube from loss of bias in the event of a short in C2. The shafts in both grid and plate boxes are of Delrin rod from Small Parts.[16] There might be a better material; there was a recent discussion on the rec.radio.homebrew Usenet Newsgroup about the best materials in RF environments, but I don't remember any conclusive consensus. I do know that there is a glob of melted nylon sitting on my desk as a reminder of the first attempt to connect to the input tuning capacitor shafts.

To wind the ¼-inch tubing for the plate tank without kinking (I made a new coil after the photo!), seal one end in a vise, fill with sand, seal the other end, wind, cut off the ends, and then return the stolen sand to your kid's sandbox.

Construction of the control board entails a certain dilemma. Many projects are essentially experimental; the builder wants to try a few new methods. If one goes through all the trouble of etching a board based on the first draft of the circuit, the many subsequent modifications will shortly turn it into a nightmare. If a different construction method is used, then, when all the final changes have been made, the amplifier is finished and there is no longer a need for an etched board.

For the past seven years I've used wire-wrapping for control boards in all my gear. There are quite a few advantages. Quite dense packing is possible, much denser than ordinary perf-board construction. I even install all the resistors and other small parts in sockets, along with the ICs. This makes it a trivial matter to change a component value, facilitating circuit development. Although it does add a bit to the cost, it can save hours and hours of time. Making a wiring change cannot be said to be easy, working in the maze of wires under the board, but it is possible, clean and neat. It does require some concentration, a small price to pay for having no unsoldering to do. The board is mounted on hinges for easy access to the wire side. I've never seen a wire-wrap connection fail.

All the wire-wrap supplies are available from Digi-Key.[20] Every socket and pin is numbered using a simple matrix scheme, and the numbers are noted on the schematics in the notebooks. Point 237 is pin 7 of U23, the third IC in the second row. Resistor R456 is plugged into the fifth socket in the fourth row, with one lead at pin 6 (by convention, this will be the left or upper end in the schematic). Trimpots can also be fitted onto sockets, if the sockets are of the machined-pin type. Depending on the trimpot type, a touch of solder at each lead may be warranted. The dip relays naturally plug into sockets. A few larger components, such as the coax relay driver and its 5-W base resistor, are soldered to separate wire-wrap pins at the edge of the board (at the bottom in Fig 1). Each (stranded wire) lead to the board is filtered as noted in the caption to Fig 6; the rows of bypass capacitors and RF chokes consume a surprisingly large portion of the board. Each lead to the board passes through two holes before soldering to its wire-wrap pin. This is for strain-relief, so a wire cannot be bent at the point where it is soldered. At certain parts of the amplifier, tiedowns are used for this purpose; Hosfelt has handy stick-on types.[19] DX chasing and contest work require the utmost level of reliability; one broken wire can spoil your whole run.

Operation

Providing suitable drive power is important. Problems can arise when a solid-state driver is used. These "bricks" often exhibit greater IMD (splatter) when operated at reduced output. This will be the case no matter how linear the final amplifier is; it amplifies whatever you feed it. It's just like the old computer cliché, "garbage in, garbage out." It sounds contradictory—a driver operating with less power should produce less splatter. But IMD is relative, and we are to amplify whatever comes out of the driver. One solution would be an attenuator, perhaps built around a dummy load, so the brick may be run at a more linear level. A better solution would be adjusting the bias in the driver. The bricks are usually rated for 100% duty cycle on FM, a steady carrier. SSB and CW operation is much gentler, so the driver idling power may be considerably increased, which should improve IMD performance.

My driver is capable of 200-W output; it's a homebrew conduction-cooled tetrode amplifier with neutralization and 26 dB of gain. As is typical for a class AB_1 tetrode, it is linear at any output level. The transverter maximum output is 2 W, running class A at the output transistor. About 14 dB of gain reduction is needed. This is done in the transverter at the milliwatt level by a resistive panel control in the RF path. The common practice of using high ALC levels to reduce gain often produces IMD (splatter).

Heater Idle Circuit

There are several modes of station operation that call for very long periods—hours and hours—when an amplifier must be ready for near-instant operation, but is rarely used. In my case, with this 2-m amp, the main situation is VHF contest operation (in the sparsely settled Southwest). The contest starts at noon Saturday—hopefully with a big blast on a wide-open 6-m band. There is not much doing on 2 m until around sunset when tropo improves and operators in neighboring states start swinging antennas in all directions. But all afternoon we must be ready on 2; there may be sporadic-E at any moment. And then the same all day Sunday, with the 2-m tubes mostly just sitting and cooking away. A similar situation arises on any summer day when 6 m is open; if the skip is shortening we want to be ready on 2. A more common situation, which would apply to an HF amplifier, involves the DXer who spends almost all the time just listening (as the very best operators do), but must be ready for that new one.

The answer to these problems is the heater idle circuit. This is a simple timer that drops the heater voltage to a specified low level after no transmissions have been made for a specified time. I chose a 12% drop after 5 minutes. Because the heater voltage is regulated, this was easy to do with an IC timer, a DIP relay and a trimpot. In other amplifiers, an IC timer, a relay and a resistor in the filament transformer primary circuit would serve the same purpose.

The idle circuit schematic is shown in Fig 11. If there is no transmission for 5 minutes, the 555 timer switches the 723 regulator in Fig 5 from the panel pot, which sets operating heater voltage, to the trimpot inside that sets the idle voltage. The circuit resets to operating voltage automatically whenever the PTT line (mike button, foot switch or semi-QSK circuit) is keyed, or whenever the reset button on the amplifier front panel is pushed.

How long does it take for the heater to reach operating temperature after reset? I made a number of tests, aiming to determine whether I could forget about the idle circuit, and just grab the mike or key whenever I heard a new grid square, leaving the automatic recovery feature to restore operating heater voltage. The idea (unsubstantiated by any manufacturer or authority) is that the heater loses heat mainly through cathode current. Under this hypothesis, during receive periods the heater gets too hot, hotter than in continuous transmit operation. Thus, it should be possible to maintain operating temperature with less voltage when not transmitting, and full heater voltage would be needed only at the instant a transmission begins.

Does it work? Well, without asking the factory to saw open the tube to investigate, I can only conclude that since the tube has survived all the testing, it works fine. Still, I usually don't wait for the automatic reset. If I hear someone I want to call, I touch the foot switch for half a second to reset the heater voltage. But often I forget to reset manually and the automatic reset circuit works fine.

The tests involved going from the idle condition instantly to full-peak power (pulsed), or full-power CW (dits), to see if full power was instantly achieved. Yes, it was, as quickly as I could read the meter. This power-output test for adequate heater temperature is consistent with tube manufacturers' suggestions for heater voltage adjustment in VHF operation: reduce the heater voltage gradually until the power output drops slightly, then bring it back up a bit. In other words, if you're getting full output, the heater is hot enough. Any more heater voltage and you're just cooking the life out of the poor tube.

Using the heater voltage and current meters, the heater resistance may be calculated. From this the heater temperature may be inferred, relatively, although I don't know the exact relationship. The 4CX1000A is rated nominally at 6.0 V, with advice for lower voltage if adequate, especially on VHF. I operate at 5.8 V (where it draws 9.2 A) and idle at 5.1 V (where it draws 8.4 A). This is a 20% drop in heater power during idle periods, enough to expect a vastly increased tube life. The calculated heater resistance drops from 630 to 607 mΩ.

Since the heater has lower resistance during idle mode, the current will be slightly higher than normal immediately after reset. The time taken for the heater current to drop to normal after reset is another test of the idle-circuit idea. Although difficult to clock, because it happens so quickly, it takes less than 2 seconds.

Note that the heater drops to idle only after 5 minutes without a transmission, not between each transmission during normal operation. And, if you are still worried, you can always reset manually with the mike button, the foot switch or the reset button a few seconds before resuming operation.

ALC for Transverters

For fighting splatter on the ham bands, ALC is the most powerful weapon. The ALC circuit in Fig 8 will produce a negative voltage for driving the ALC-controlled stages in a transceiver or transverter. The R-C network at the follower input prevents keying transients from generating ALC voltage. The pot on the front panel sets the ALC threshold at 0.1 mA of grid current. The panel adjustment may be used for experimentation. The 4CX1000A is rated for zero control-grid dissipation; that is, zero current. However, the spec sheet says a few milliamperes on peaks is okay. I wanted to see if on CW, where IMD is not a consideration, a few milliamps would result in higher efficiency. (I used a 5-mA grid meter initially.) The results: no, there is no improvement in efficiency with higher grid current.

Some rigs (eg, the FT-1000MP) do not have an ALC input line which controls the RF level at the transverter output jack. ALC may be added to a transverter with the circuit shown in Fig 12. The MOSFET buffer stage is designed for unity gain, but this may depend on the individual device. Increasing the drain resistor will increase the output, if necessary. The RF input level should be not more than –16 dBm. The gate 2 bias adjustment is set so that Q1 operates at the knee of its characteristic curve.[24] If the bias is set for maximum gain, the ALC voltage will be forced to rise to a high value before any significant gain reduction is obtained. My transverters have two or three ALC-controlled stages; this means that there is less gain reduction at each stage, and therefore less possibility of ALC-induced IMD. The ALC adjustment in Fig 12 is set so that –5 V applied at either ALC input will produce 20 dB of ALC compression. Excessive ALC may result in IMD. In operation, the transverter gain is adjusted for a maximum of 3 dB of ALC compression, metered on the transceiver panel.

Pulsed Tune-Up

The dit tune-up procedure, with amplifier keying, was described in "A Luxury Linear." In fact, the method used in my shack goes even further. My home-brew transceiver has a built-in pulser, set for a 33% duty cycle. The TUNE switch on the panel automatically shifts the radio into CW, silences the sidetone, hits the PTT line and starts the pulser. So testing at 1500-W PEP output involves only 400 W of average plate dissipation.

The schematic for the pulser, which could be built into a transceiver or as a separate device, is shown in Fig 13. A scope with calibrated time-base is useful for precisely setting the desired pulse rate and duty cycle, but a simple RF monitor scope is sufficient. Merely setting the OFF trimpot at maximum and the ON trimpot at midrange will be adequate.

Improvements

The main improvement I would suggest is in the swinging link mechanism. The simple shaft and insulator method described in the caption to Fig 4 works well enough. However, the link needs to swing only about 30°, and the heavy coax-braid connections make it a bit stiff. In operation, I use the swinging link only as a coarse loading adjustment, so there is no practical problem. Since I haven't touched the swinging link in the past six months, this is one of those situations where improvement is not worth the trouble. But for someone trying this circuit anew, I would suggest some sort of gear drive, for smoother control and finer adjustment.

I have only one other suggestion: build very carefully! Don't lay yourself open to a remark like I got from one of the locals here. All he could say when he saw the new amplifier was, "You got one of the meter labels crooked."

Notes

[21]"A Luxury Linear," *QEX*, May, 1996, p 3-12. (Photos also in *QST*, July 1996, p 19.)

[22]References to Figs 1-10 and Notes 1-20 are to the previous article (see Note 21 above). This article begins with Fig 11 and Note 21.

[23]Correction for "A Luxury Linear," Fig 10: Reverse the ± markings at the inputs of the plate current-metering op amp.

[24]A typical MOSFET characteristic curve is shown as Fig 10 in "A High-Performance AGC System for Home-Brew Transceivers," *QEX*, October 1995, p 12-22.