What's All This Ripple Rejection Suff, Anyhow?

Bob Pease
ED Online ID #1509 #1780 #1870

(Part 1 – February 4, 2002)

Recently, I helped a guy who needed low output ripple on a power supply. His 1000-V output required low, submillivolt ripple. I designed a couple of circuits for him. Details soon. To make sure that I wasn't missing any tricks, I looked up a voltage regulator, the Philbrick 5910, designed by my old colleague Bruce Seddon about 42 years ago. It was optimized to provide ±300 V dc at 100 mA (see the figure), or 300 mA in a larger R-300 power supply when additional output tubes are paralleled.

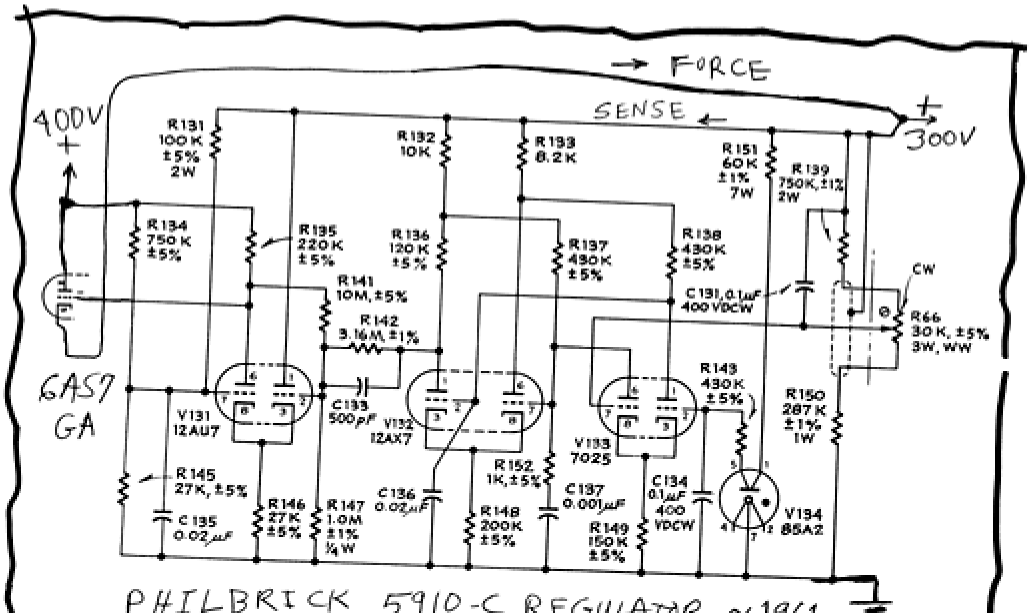

Anybody can design an operational amplifier (op amp) with a gain of 100,000 or 1,000,000 at dc. But this one needed to swing its output 80 V dc plus 15 V p-p at 120 Hz, with a summing-point error of less than 150 µV ac rms. That's 100 dB of gain at around 120 Hz! Forty years ago, most regulator amplifiers had two-stage amplifiers. But this one used three dual triodes for each supply, + and - 300 V. Two honest stages of 12AX7 provide a lot of gain (µ = 100), and the 12AU7 (µ = 20) gives a good healthy drive to the grid of the output tube (6AS7GA).

Additionally, the positive feedback of R132, R133, and R141 provides much more gain-even at 120 Hz, and at dc. Although the output pass tube (6AS7GA) has a µ of only 2, this amp can easily drive the grid to any necessary voltage, whether no-load or full-load, low-line or high-line. Also, it has submillivolt gain error, for line, load, or ripple.

There are good bypass capacitors, such as at C131 and C134, to filter and bypass the noises for frequencies above 4 Hz. That helps keep the output's noise below 250 µV rms. This amplifier was also optimized for fast bandwidth. The output bypass capacitors (not shown) were 150 µF at 525 VDCW, specified with good low RS.

The 5910 dual regulator amplifier was built in a little boxy subassembly. The tubes stuck up above the box, with the passive components mounted on turret terminals along the insides of the box.

I bought a 30-year-old R-300 that was still in very good shape. I used a Variac to turn up the line voltage very slowly to "form up" the electrolytics. It would be harmful to apply full line power right away. I fired it up and it regulated nicely.

After I did a general check-out, it was time to do noise testing. I used a series stack of three 25-W light bulbs to draw ~190 mA of load. The ripple voltage on the main ("upstream") filter capacitors rose to 9 V p-p. The output ripple-plus-noise increased from 100 µV rms to perhaps 120 µV-truly negligible, barely 3 ppm (p-p) of the dc output, even at full load. The ripple was barely 40 µV p-p. Not bad! The load regulation was submillivolt, and the line regulation was very good too.

So it's encouraging that 40 years ago, a high-gain three-stage op amp could provide excellent ripple rejection, and that old machine still runs well today. Soon, I'll show how to get low ripple voltage by adding an add-on circuit to an existing high-voltage supply. Perhaps using a fast FET op amp running on ±6 V.

I'm looking forward to using that old R-300 to run a whole bunch of tests on various old vacuum-tube operational amplifiers. I've been waiting to do that for years! We rarely use vacuum tubes these days, but here's a good example of how tubes could do some very good work. I'll have more comments on the old art of designing with tubes, and the early days of operational amplifiers, 40 or more years ago.

(Part 2 – November 11, 2002)

As I said a few months ago, the study of ripple rejection will continue. Some guys asked me, "How can I convert a high-ripple voltage supply to a low-ripple one, or to a very low-ripple one?" I am reminded of the little girl who pondered over the Sunday School teacher's question: "What must we do to receive forgiveness of sin?" After some thought, she replied, "First, you have to sin...."

So first we have to make some ripple. Figure C, located on my Web site at www.national.com/rap/ripplerej.html, is a floating, battery-powered sawtooth oscillator, which I will call a NoiseMaker. Its task is to add ripple on top of the output of a quiet high-voltage supply and make it noisy enough that we can see if our ripple-rejection circuits are working-and exactly how well.

The circuit of Figure D (found on the above Web page) is not much more than an ac-coupling network found in any ac oscilloscope. But it includes safety factors and safety networks to protect

the scope. After all, most scopes won't stand 300 V dc on their ac inputs. The coupling capacitors aren't rated for that. It also helps to define the bandwidth.
The circuit of Figure E (same Web page) is an ordinary ac pre-amp, because the supply we studied in February had just 100 µV ac rms of noise. But when we get through, the noise will be so low that we'll need a pre-amp to evaluate it.

Now, how does the circuit of the figure on this page (Ripple Rejection Type A) work? Every op amp, if fed a noisy feedback path, can be made to try to amplify a counter-error signal to cancel out that noise. Even if the NoiseMaker produces 0.25 V rms of noise, the error voltage at the input of the op amp need not be more than a millivolt. That's enough to make the output move to cancel out that noise, as seen at its input.

This circuit's limitation is the op amp's gain-bandwidth product. So the amplifier must have a lot of GBW product, but its CMRR or PSRR are not very important. I think I'll try an LF411, with good bandwidth, or an LMV751, with 7 nV/^-Hz and low I_B. If we can keep the op amp's gain above 100 at 10 kHz, we can get a 100:1 improvement at that frequency, and even more at lower frequencies. Could we use a bipolar op amp with lower V_N and worse I_N? Not impossible, but unlikely. Hey, I'll try anything-once....

NOTE: In every application, the best choice of amplifier, and the best way to optimize the noise, depends on the impedances and on the specified bandwidth for which low noise is required. When things start changing, the optimization changes a lot. The rules for this example are: For a 3-dB bandwidth of 20 Hz to 10 kHz, get the lowest noise possible with the NoiseMaker specified-95 mV p-p at 26 kHz, plus the 100-µV rms broadband noise of the basic Philbrick R-300.

What are the results? What's the best noise, with the best amplifiers? I didn't get this optimized before the deadline for this column. Further, I want to try out a special version of the Sallen-Key filter. We'll call that Ripple Rejection Circuit Type B.

I'll have the results by next month. Then we'll see which circuit is good, which is better, and which performance is best, with what amplifier, to reject the noise of the NoiseMaker. We'll find out which circuit gives the lowest noise with the NoiseMaker turned off. Have you ever seen a 300-V power supply that had less than 1 µV rms of noise? See you in a month!

P.S. When I built this circuit, I carefully put a 330-‡ resistor in the input path, in series with each of the 400-V, 0.1-µF mylar caps. But I neglected to draw them into the published schematic. Please add those in. Otherwise, the high charge during turn-on would fry the 1N914s!

(Part 3 – December 9, 2002)

Okay, I wish I could tell you guys that I have improved the old R-300's noise down to 1 µV rms, on top of its 300-V dc output. I tried. I applied both the Type A servo amplifier from last month (electronic design, Nov. 11, p. 84) and the new Sallen-Key filter, Type B, shown in the figure. It's battery powered, and the LMV751 runs well on three AA cells. (If you want long and cheap battery life, use C or D cells instead.) I also used the level-shifter, and the pre-amp and NoiseMaker shown at www.national.com/rap/ripplerej.html.

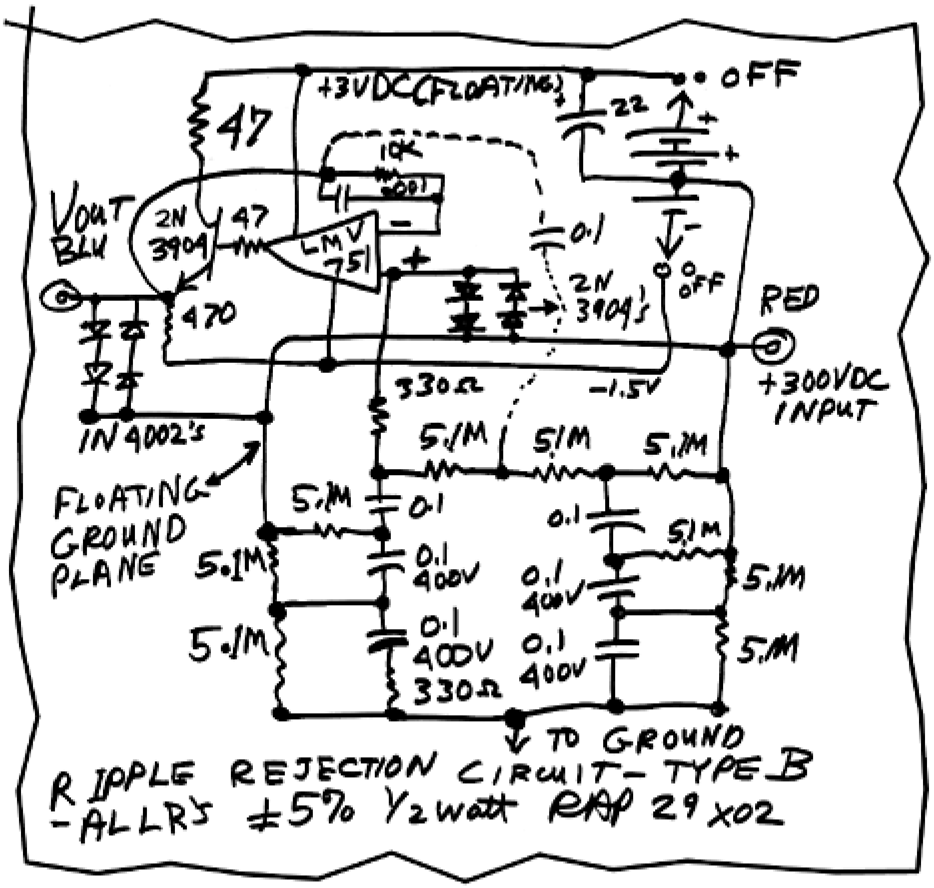

Actually, I was able to work the noise down from 75 to about 9 µV rms, using each of the Type A and B circuits. I tried to get it lower, but the spatial orientation of the magnetic flux in the neighborhood would have made it very hard to get anything below 9 µV. I'll never say never, but it's not easy. Magnetic fields are nasty and hard to shield or screen out. Layout is critical and not easy. Maybe I'll just put longer wires on the pre-amps and move them far from the R-300, the transformers of all the power supplies, and the voltmeter. Might work!

Further: Look at the careful filter structure of the schematic diagram of the figure, Type B. There should not be much 1/f noise. But I had plenty of jitter, wobble, and 1/f noise-much more than I expected. There was maybe 10 to 20 µV of jitter. I considered the layout and the characteristics of the mylar capacitors (400-V capacitors running on less than 150 V dc), yet I couldn't find much of a clue as to what was causing the noise.

Then I thought about the resistors. Can an AB 1/2-W, 5.1-M‡ resistor running at 150 V generate enough 1/f noise-current noise-to make a poor reading? Maybe so. I will study this later, just not this week. There are some trick tests that I'll want to run.... Hey, I never looked this closely before!

I did get some good data on the ac ripple-rejection, using my NoiseMaker (schematic is on my Web site) to try to cut the ripple of an 85-mV p-p (26-mV rms) noise at a 26-kHz nominal switching frequency, with lots of harmonics at 52 and 78 kHz. Although the Type A circuit improved the 26 mV to 0.88 mV, Type B did a bit better at 0.83 mV. But that was NOT unexpected. So it's not too hard to get a 30-dB noise improvement by employing a cheap, low-noise op amp, such as LMV751.

Both schemes worked pretty well. But Type A didn't like to tolerate a lot of capacitive load from the op-amp's output to ground-that is, the capacitance from the power-supply low terminal to ground. It was able to drive the R-300's capacitance. But adding a coax cable over to the voltmeter made it grouchy at 4 MHz. So I just had to bring the voltmeter (HP3400A) over by the amplifier.

Yet Type B would surely be grouchy if there was a lot of capacitance from $+V_{OUT}$ to ground. So any fast, nimble, low-noise amplifier may get grouchy about capacitive loads. In some cases, a series R-C damper can help. There are no easy answers. If I got a "hotter" op amp with a lot more than 4 MHz of GBW product, that would certainly help; and some 26- to 52-kHz noise can be filtered with extra L and C. So this analysis is encouraging.

In all of these studies, I never got a shock, never caused a BANG, and never blew up any components or op amps (except when my thumb absent-mindedly nudged a ±2-V supply up to ±12 V). So I don't feel bad, figuratively and literally.

These are the Schematic Diagrams of the 3 circuits I employed, NoiseMaker and Level-Shifter and Preamp Circuits, to use with "Ripple-Rejection Stuff" columns, to evaluate the low-noise, low-ripple circuits in Electronic Design.

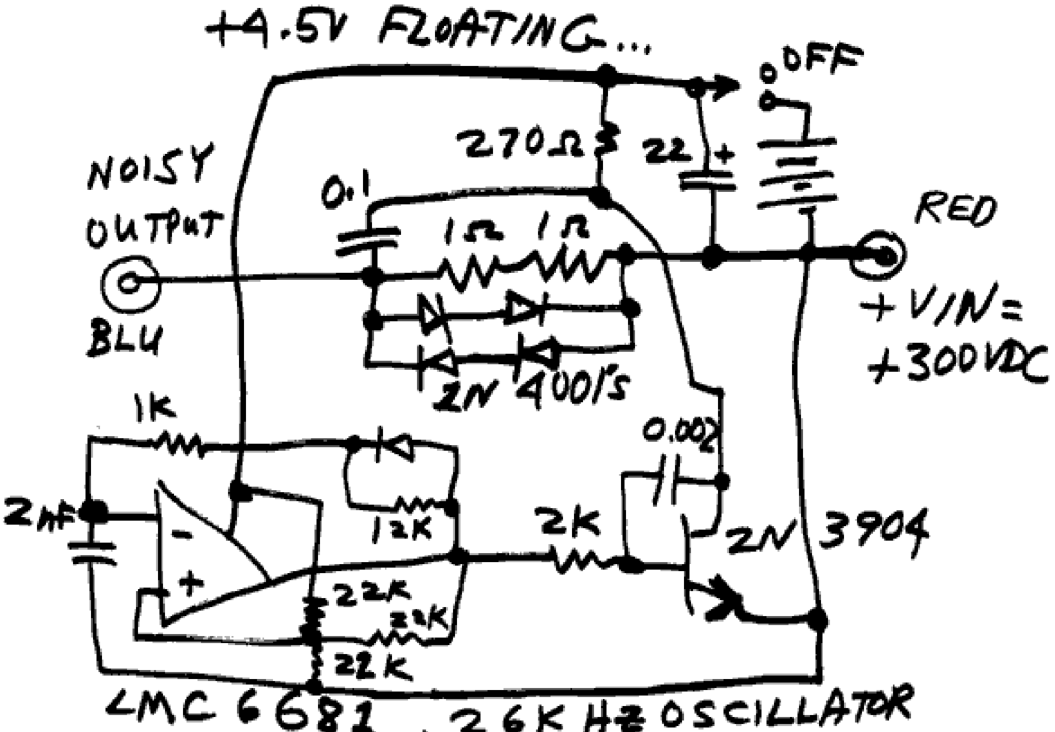

FIGURE C., NOISE-MAKER

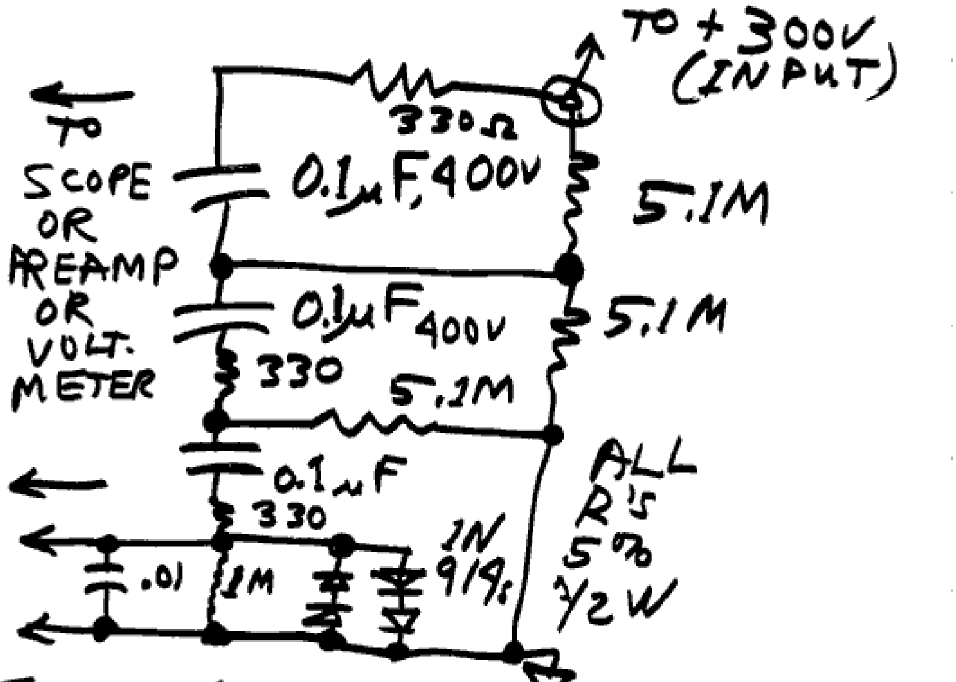

FIGURE D - AC COUPLING NETWORK

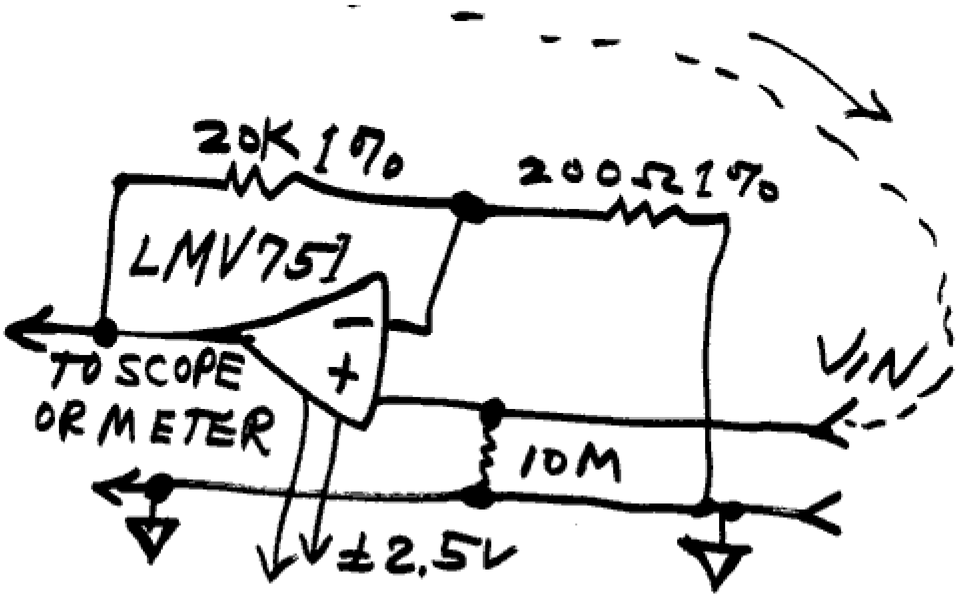

Figure E, LOW-NOISE PREAMP
G = 100

(ALL POWER SUPPLIES WELL-BYPASSED)

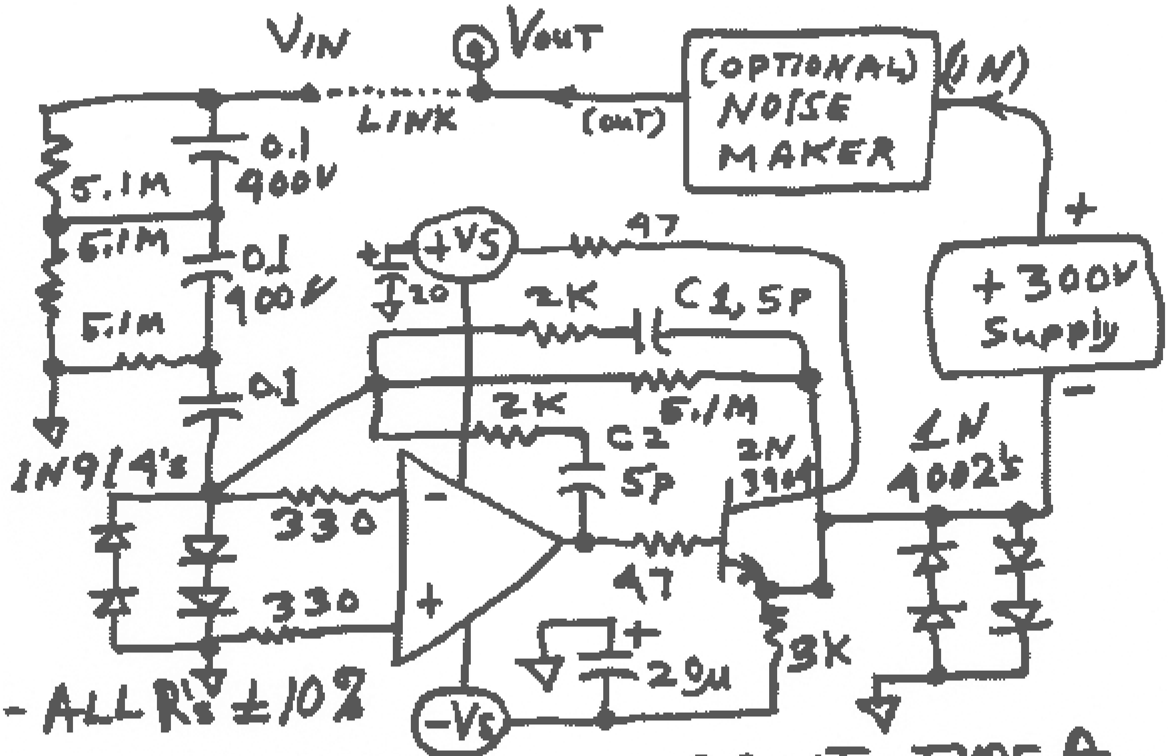

- RIPPLE REJECTION CIRCUIT - TYPE A
RAP 28ix02