A Tribute to Bob Pease

Articles and Book Chapters

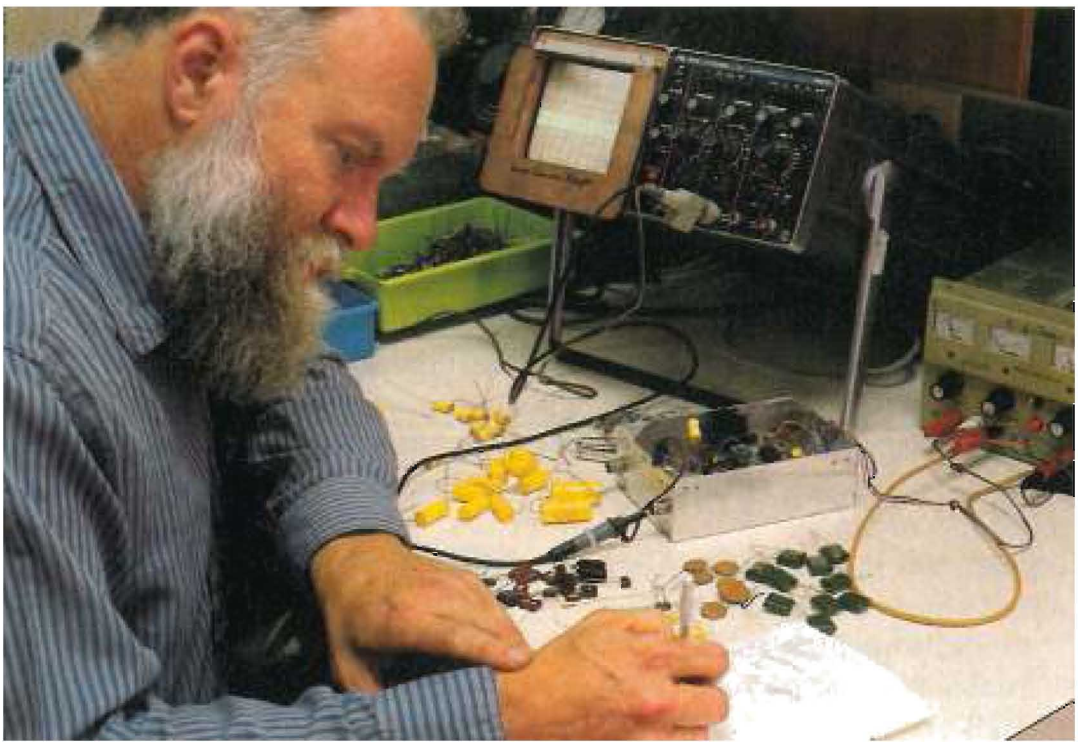

Readers' Notes:

A great many readers took the trouble to pass on a few kind words of appreciation for Dialogue. *These all said about the same thing, in many ways, and though they made fascinating reading (for us), it would serve no useful purpose to include them here. On the other hand, there were a couple that were unusually interesting, in that they discussed a point that was glossed over (and was further distinguished by a slide rule error).*

Your recent issue of *Analog Dialogue* concerning noise in operational amplifier circuits[1] was of great interest. It was an excellent discussion of noise in general and its measurement. I appreciated the differentiation between the various types of noise.

Since Hewlett-Packard's voltmeters are often used for noise measurements, I was particularly interested in your comments on page 7 regarding the characterization of noise. You mentioned that the indication on an average-responding RMS-calibrated meter should be multiplied by 1.11 to obtain the RMS value of the noise being measured.

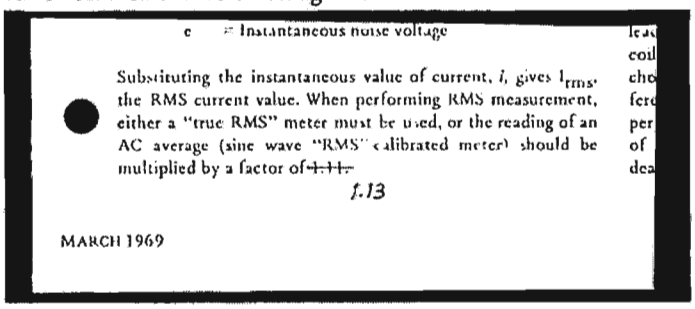

e = Instantaneous noise voltage

Substituting the instantaneous value of current, *i*, gives I_{rms}, the RMS current value. When performing RMS measurement, either a "true RMS" meter must be used, or the reading of an AC average (sine wave "RMS" calibrated meter) should be multiplied by a factor of 1.13

MARCH 1969

(from *Analog Dialogue*, March, 1969)

The average-responding meter *does* indicate a voltage that is 1.11 (or $\sqrt{2}\ \pi/4$) times the average or rectified value. This is derived from the fact that the average value of a sine wave is 0.637 times the peak and the RMS value is 0.707 times the peak. Thus the meter measures 0.637 and displays 0.707, a ratio of 1.11. This number is, therefore, a calibration constant that must be applied to calculate the effects of waveforms other than sine waves.

For white noise we can show that the average value is $\sqrt{2/\pi}$ times the RMS value. This is derived by integrating the probability distribution of random noise. In the actual measuring instrument, the integration is performed by the ballistics of the meter movement. Therefore, the indication on an average-responding RMS-calibrated meter will be $\sqrt{2/\pi} \times \sqrt{2}\ \pi/4$, which calculates to be 0.886, or −1.05dB. Thus the average-responding meter will indicate 0.886 of the true RMS noise voltage, and its reading must be multiplied by 1.13 to obtain the correct value.

You may wish to refer to the Hewlett-Packard *Journal*, Volume 6, Numbers 8, 9, 10 (*April, May, June, 1955*) for a further discussion of waveform errors in average-responding voltmeters.

Thank you for publishing *Analog Dialogue*. We enjoy reading it very much.

Don A. Wick
Customer Service Manager
Hewlett-Packard
Loveland, Colorado

Our thanks and appreciation to Mr. Wick for his sharp eyes and enlightening comments. If you would like copies of the H-P JOURNAL *issues referred to, we suggest you get in touch with your local H-P field office.*

A professional cartographer, when drawing up a new map, will have it copyrighted. Then he will add into his map a small but distinct error, an artifact.

Any person who steals this map is unlikely to be bright enough to find or correct the error. Wholesale copy of the map including the artifact is acceptable proof of infringement of copyright. The thief is caught with marked bills in his hot hand.

In my NEREM paper on noise[2], I deverly *(sic)* stated that the Fudge Factor to use (when trying to measure white noise on an average-rectified-responding AC voltmeter, with its calibration specified in terms of RMS value of sines) was +11%. *(Ed. The actual figure published in the paper to which Mr. Pease refers was 11.3%)* Actually, the correct value, as stated in our up-to-date applications article P/N-10 is +12.8%, as a thoughtful interpretation of Bennett's *Electrical Noise*[3], page 44, will prove out.

Now, in reference to your recent *Analog Dialogue* article (foot of first column, page 7), tell me there isn't any jam on your face!

Yours for Better Analogs
Robert A. Pease
Philbrick/Nexus Research
Dedham, Massachusetts

Mr. Smith, in his original manuscript, stated: "Most meters will read 12% low when measuring noise, but since this error is constant, a true RMS meter is not required." In converting this figure to a multiplying factor, which is more useful, our slipstick slipped. The correct figure is indeed 1.13, and we confess to having committed a 2% error. But, Mr. Pease, that redness on our face is a blush, not jam!

A further comment: If Mr. Pease is claiming that he has intentionally "rigged" a figure in a paper he has written for presentation at a technical society meeting, for the sole purpose of exposing copyright(?) violations, most readers would wish that it had been an obscure figure, rather than one which is pretty useful. Somehow, this sort of thing is unfair to the engineer who simply wants information, presented as accurately as possible. It also tends to impeach a source whom we know ordinarily to be of high technical integrity, a former colleague for whom we have great admiration. Ours was an honest mistake.

ERRATUM. In *Analog Dialogue*, Vol. 3, No. 1, page 4, Figure 1b, associated with "An op amp and a chopper give precision ramp generator," as reprinted from *Electronic Design*, there is an obvious error:
The 2N1131 transistor should be a PNP type.

[1] *Analog Dialogue*, Volume 3, No. 1, March 1969, "Noise and Operational Amplifier Circuits"
[2] "How to Characterize and Measure Noise in Operational Amplifiers" by R. A. Pease, *NEREM* 1968
[3] *Electrical Noise*, W. R. Bennett, McGraw-Hill

Voltage-to-frequency module serves diverse applications

by Robert Allen Pease
Teledyne Philbrick, Dedham, Mass.

With the addition of several external components, a versatile new voltage-to-frequency converter can be made to yield a broad range of communications and instrumentation functions at reasonable cost. The 1-cubic-inch model 4701 adopts the charge-dispensing approach to v-f conversion, and has a unit price of $59 in small quantities.

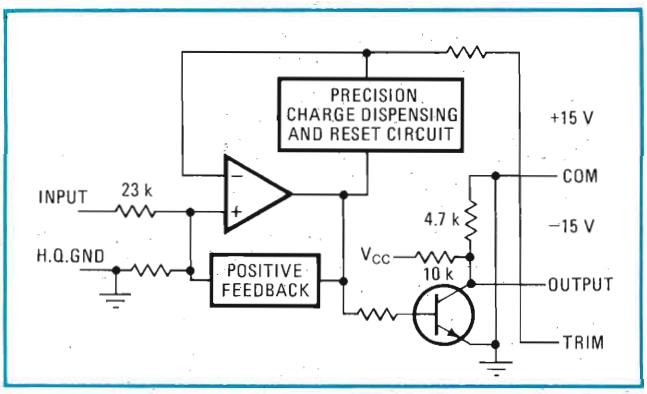

1. Converter. Frequency-to-voltage module follows charge-dispensing design approach. Output is buffered and TTL-compatible.

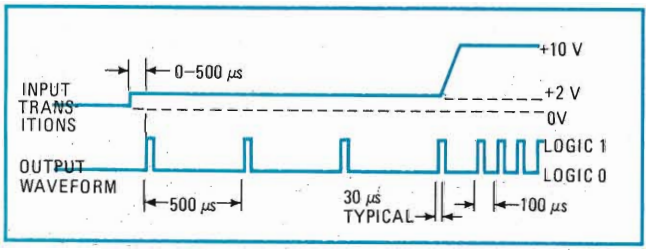

2. Proper timing. Typical waveforms show timing relationships between input and output for input levels of 2 and 10 volts.

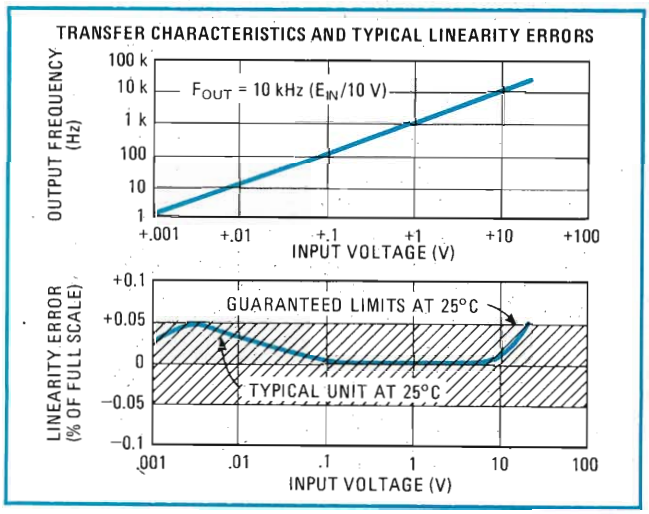

The module provides output pulses of about 30 microseconds width at a repetition rate that's directly proportional to the analog voltage level. As Fig. 1 shows, the amplifier functions as a zero crossing detector with a +13-volt quiescent output. When an input signal passes through the input resistor, the charge dispensing and reset circuit senses zero crossing and connects a large-value discharge capacitor. This capacitor quickly drives a precision timing capacitor below zero, cutting off the amplifier and thus resetting it rapidly to a +13-v output.

The negative-going pulse at the amplifier output is normally inverted in a TTL-compatible circuit. Typical input and output waveforms are shown in Fig. 2.

Three of the more interesting of circuit functions realizable with the v-f converter are: inexpensive digital voltmeters, ramp generators, and analog-to-digital converters. An inexpensive DVM can be constructed simply by adding a counter/display to the output of the v-f converter, taking advantage of the 4701's linearity, which is within 0.015%.

Programable, very linear ramp generators based on v-f techniques (Fig. 3) overcome a major difficulty of conventional designs—they do not suffer from charge leakage of the timing capacitor under temperature extremes. By driving a ripple counter with an R-2R ladder connected to its output, the 4701 converter generates a highly linear 10-v ramp that operates at frequencies to 80 hertz. Output impedance is 51 kΩ ±5%.

An analog-to-digital converter is constructed from a 4701 and several inexpensive components. In Fig. 4, the converter output feeds a counter which is reset by a 1-

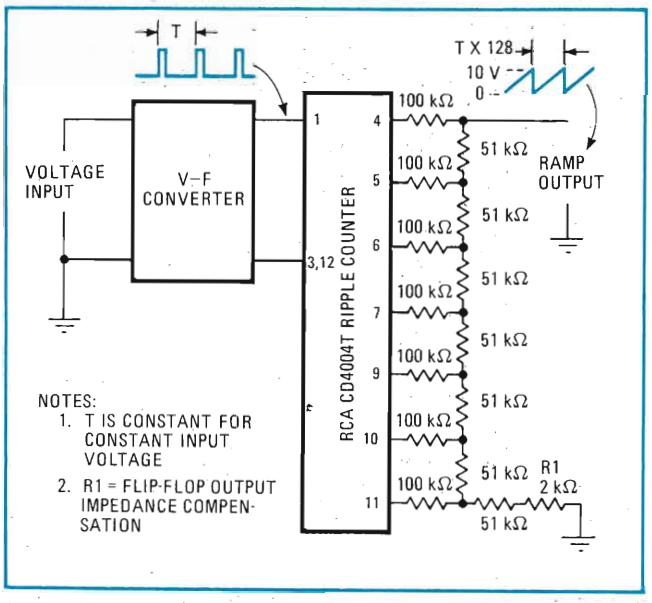

3. Better ramp. Voltage-controlled ramp generator doesn't share the charge-leakage problem of conventional generators.

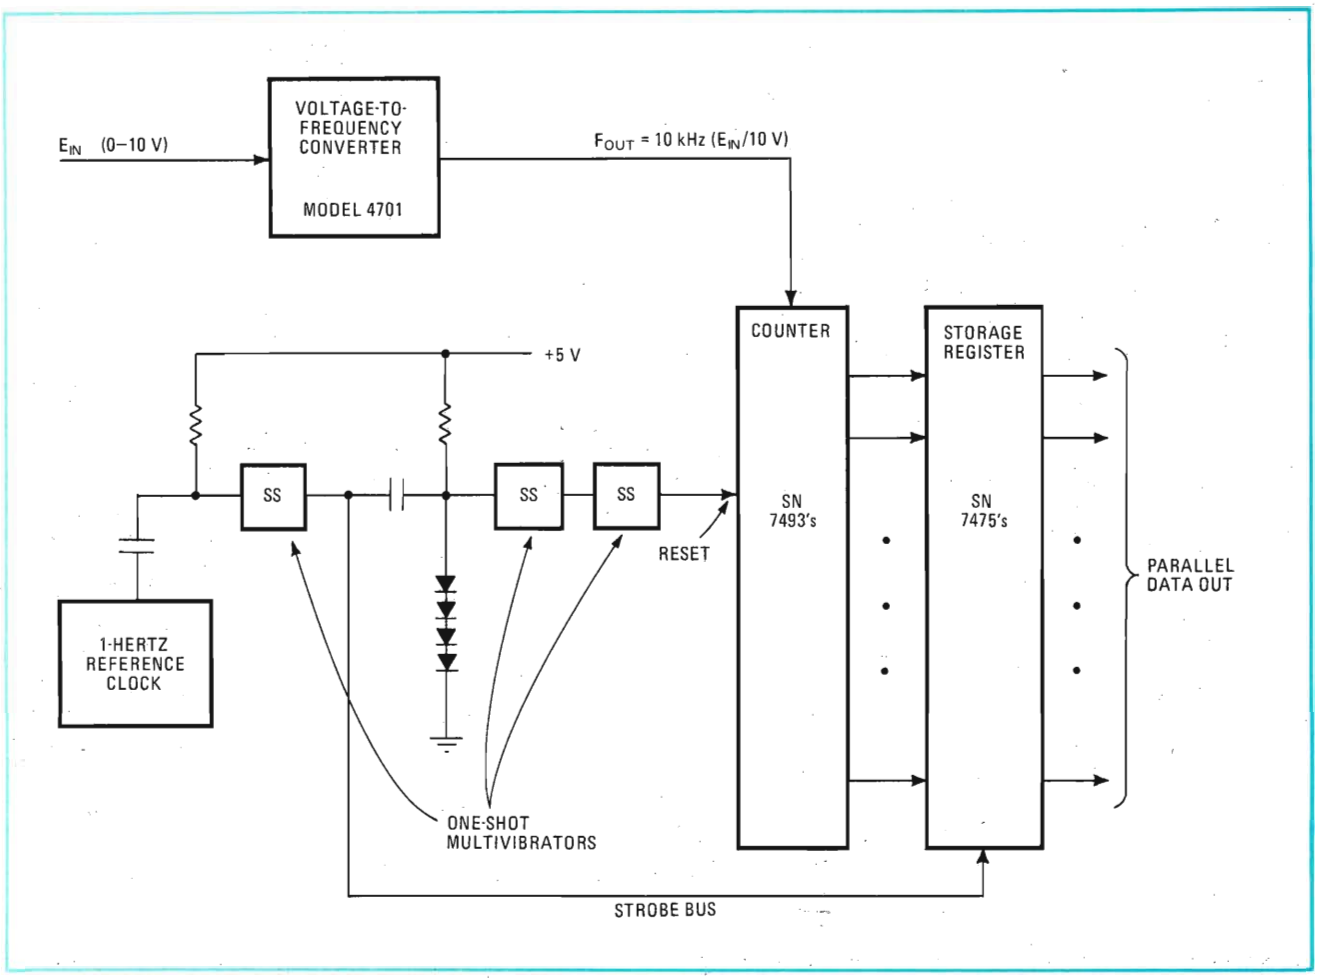

4. Low-cost conversion. Analog-to-digital conversion is inexpensive with the voltage-to-frequency converter and several other components. For fast computer processing, the output data is available in parallel from a single strobing.

Hz clock pulse. Just before reset, the peak counted value of the input frequency is strobed into the storage register. The data can be read out of the register (or a following parallel-to-serial shift register) in a few microseconds, so that processing may be performed on virtually any digital computer on a time-shared basis. □

Much ado about much accuracy

Leonard Accardi's Idea for Design, "Self-Stabilized Zener Insures Constant Current in Op-Amp Voltage Reference" (ED No. 26, Dec. 21, 1972, p. 66), suggests that an IN829 with a Sprague op amp yields 150 μV of drift vs 1000 μV with a 741 op amp. But neither the 741 nor the Sprague op amp is guaranteed to be that good. Those are just typical data. And use of the IN829 zener diode causes a guaranteed maximum drift of 7000 μV. So why quibble over typical values of 1000 μV or 150 μV.

Robert A. Pease

Teledyne Philbrick
Allied Drive at Route 128
Dedham, Mass. 02026

The author replies

Mr. Pease's comments repeat, for the most part, what was said in the article. The IN829 and IN-829A have the lowest guaranteed drift available, 5 mV (not 7 mV) at 7.5 mA over the wider temperature range of 155 C (−55 to 100 C). Since typical values are normally well below maximums, this justifies the use of the better op amp.

This circuit accomplishes what no other reference circuit has ever done, reducing the error to essentially that of the zener itself by use of a remarkably simple circuit. I invite Mr. Pease and anyone else to try to design a circuit that does the job either better or simpler. I don't think it's possible.

Leonard Accardi

66-30 54th Ave.
Maspeth, New York City 11378

Reminder circuit saves car batteries

Robert A Pease
National Semiconductor, Santa Clara, CA

Whenever you turn on your car's headlights in the daytime, whether for rain, fog or emergency conditions, there's a good chance you'll forget to turn them off, and the battery will run down. The circuit shown in the **figure** connects to your light switch and produces a loud buzz when you turn off the ignition, if you've failed to turn the lights off.

In this circuit, the ubiquitous LM555 forms a conventional multivibrator; the 1N4002 diode ensures that it gets power only when the lights are on and the ignition off. The 100Ω resistor and 10-μF capacitor protect the IC from transients on the 12V bus (although the transients are not very severe—the circuit has 12V across it only when the ignition is off).

The 47Ω resistor in series with the 5-μF capacitor permits the circuit to drive more than 50 mA into the car radio's speaker without significantly loading the radio's usual output. If you connect one speaker lead to the circuit output and the speaker remains silent, try the other lead. Tape and insulate all connections to avoid blowing fuses. **EDN**

To Vote For This Design, Circle No 457

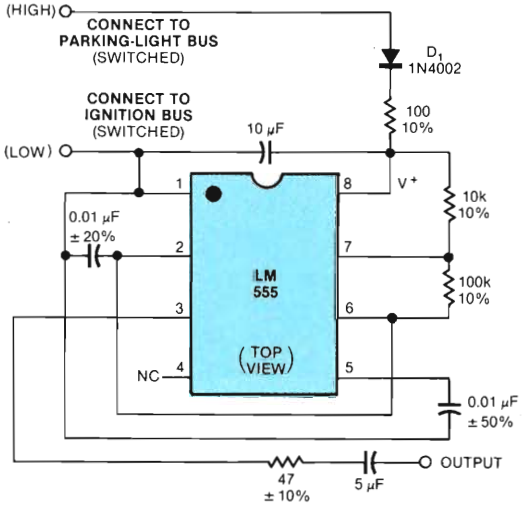

Save your car battery *by connecting a multivibrator to the light and ignition switch. A lights-on/ignition-off condition sends 50 mA to your car radio's speaker and produces a warning buzz.*

'Dithering' display expands bar graph's resolution

by Robert A. Pease
National Semiconductor Corp., Santa Clara, Calif.

Commercially available bar-graph chips such as National's LM3914 offer an inexpensive and generally attractive way of discerning 10 levels of signal. If 20, 30 or more steps of resolution are required, however, bar-graph displays must be stacked, and with that, the circuit's power drain, cost and complexity all rise. But the techniques used here for creating a scanning-type "dithering" or modulated display will expand the resolution to 20 levels with only one 3914 or, alternatively, make it possible to implement fine-tuning control so that performance approaching infinite resolution can be achieved.

The light-emitting-diode display arrangement for simply distinguishing 20 levels is achieved with a rudimentary square-wave oscillator, as shown in Fig. 1. Here, the LM324 oscillator, running at 1 kilohertz, drives a 60-millivolt peak-to-peak signal into pin 8 of the 3914.

Now, the internal reference circuitry of the 3914 acts to force pin 7 to be 1.26 V above pin 8, so that pins 4 and 8 are at an instantaneous potential of 4.0 mV plus a 60-mV p-p square wave, while pins 6 and 7 will be at 1.264 V plus a 60-mV p-p square wave. Normally, the first LED at pin 1 would turn on when V_{in} exceeded 130 mV, but because of the dither caused by the ac component of the oscillator's output, the first LED now turns on at half intensity when V_{in} rises above the aforementioned value. Full intensity is achieved when V_{in} = 190 mV.

When V_{in} rises another 70 mV or so, the first LED will fall off to half brightness and the second one will begin

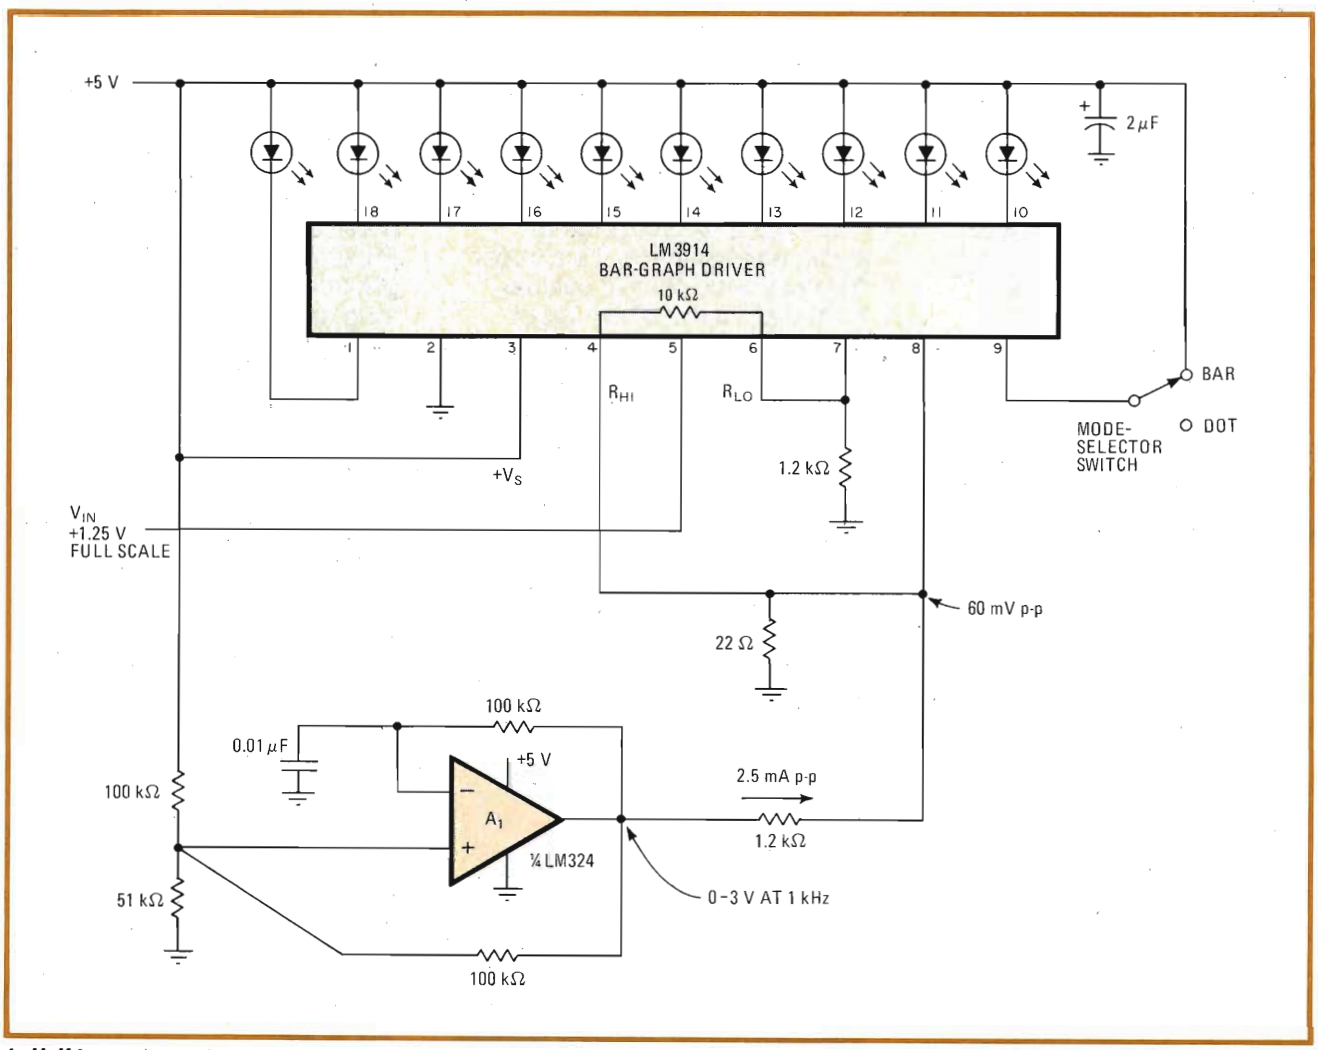

1. Half tones. Input-signal biasing on LM3914 bar-graph chip is set by the instantaneous output of a low-amplitude square-wave oscillator so that bar-graph resolution can be doubled. Each of 10 LEDs now has a fully-on and a partially-on mode, making 20 states discernible.

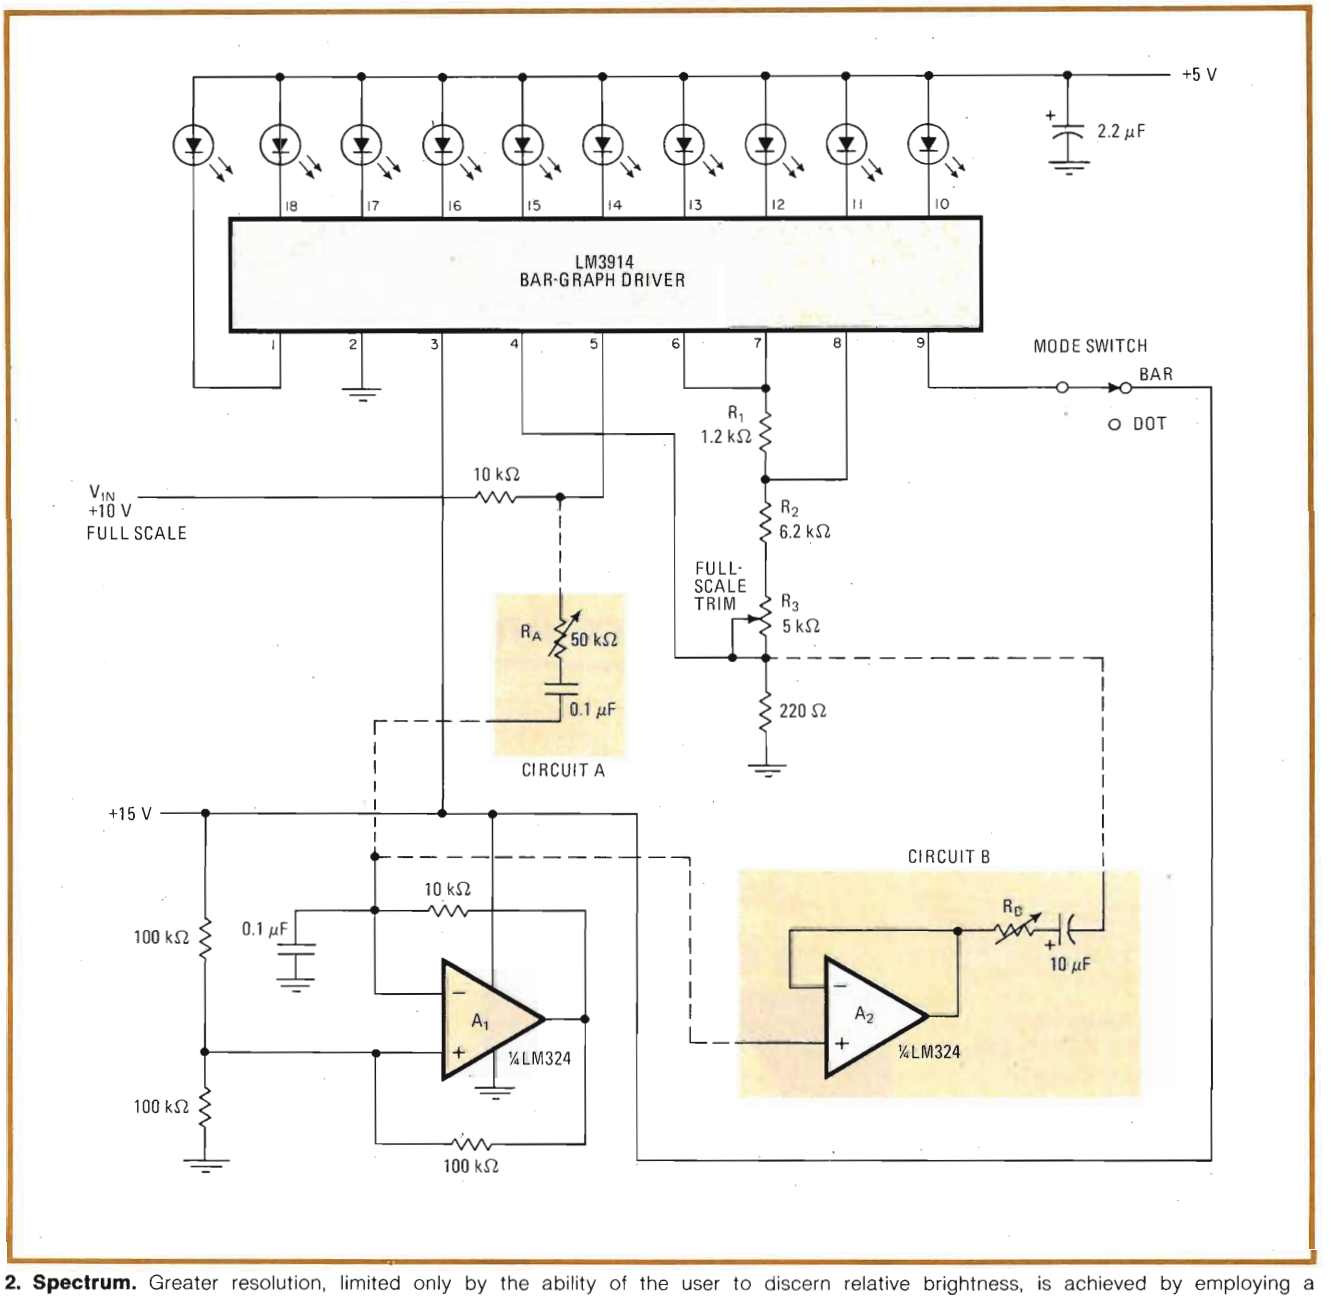

2. Spectrum. Greater resolution, limited only by the ability of the user to discern relative brightness, is achieved by employing a triangular-wave oscillator and more sensitive control circuitry to set the voltage levels and thus light levels of corresponding LEDs. Two RC networks, circuits A and B, provide required oscillator coupling and attenuation. B replaces A if oscillator cannot suffer heavy loading.

to glow. When V_{in} reaches 320 mV, the first LED will go off, and the second will turn on fully, and so on. Thus 20 levels of brightness are easily obtained.

Similarly, greater resolution can be achieved by employing a triangular-wave oscillator and two simple RC networks as seen in Fig. 2. Here, by means of circuit A, this voltage is capacitively coupled, attenuated, and superimposed on the input voltage at pin 5 of the LM3914. With appropriate setting of the 50-kilohm potentiometer, each incremental change in V_{in} can be detected because the glow from each LED can be made to spread gradually from one device to the next.

Of course, if the signal-source impedance is not low or linear, the ac signals coupled into the input circuit can cause false readings at the output. In this case, the circuit in block B should be used to buffer the output of the triangular-wave oscillator.

The display is most effective in the dot mode, where supply voltages can be brought up to 15 V. If the circuit's bar mode is used, the potentials applied to the LEDs

signetics FULLY ENCODED, 9046 x N, RANDOM ACCESS WRITE-ONLY-MEMORY | 25120

FINAL SPECIFICATION[10]

DESCRIPTION

The Signetics 25000 Series 9046XN Random Access Write-Only-Memory employs both enhancement and depletion mode P-Channel, N-Channel, and neu[1] channel MOS devices. Although a static device, a single TTL level clock phase is required to drive the on-board multi-port clock generator. Data refresh is accomplished during CB and LH periods[11]. Quadri-state outputs (when applicable) allow expansion in many directions, depending on organization.

The static memory cells are operated dynamically to yield extremely low power dissipation. All inputs and outputs are directly TTL compatible when proper interfacing circuitry is employed.

Device construction is more or less S.O.S.[2].

FEATURES

- FULLY ENCODED MULTI-PORT ADDRESSING
- WRITE CYCLE TIME 80nS (MAX. TYPICAL)
- WRITE ACCESS TIME [3]
- POWER DISSIPATION 10uW/BIT TYPICAL
- CELL REFRESH TIME 2mS (MIN. TYPICAL)
- TTL/DTL COMPATIBLE INPUTS[4]
- AVAILABLE OUTPUTS "n"
- CLOCK LINE CAPACITANCE 2pF MAX.[5]
- $V^{CC} = +10V$
- $V^{DD} = 0V \pm 2\%$
- $V^{FF} = 6.3V^{ac}$ [6]

APPLICATIONS

DON'T CARE BUFFER STORES
LEAST SIGNIFICANT CONTROL MEMORIES
POST MORTEM MEMORIES (WEAPON SYSTEMS)
ARTIFICIAL MEMORY SYSTEMS
NON-INTELLIGENT MICRO CONTROLLERS
FIRST-IN NEVER-OUT (FINO) ASYNCHRONOUS BUFFERS
OVERFLOW REGISTER (BIT BUCKET)

PROCESS TECHNOLOGY

The use of Signetics unique SEX[7] process yields Vth (var.) and allows the design[8] and production[9] of higher performance MOS circuits than can be obtained by competitor's techniques.

BIPOLAR COMPATIBILITY

All data and clock inputs plus applicable outputs will interface directly or nearly directly with bipolar circuits of suitable characteristics. In any event use 1 amp fuses in all power supply and data lines.

INPUT PROTECTION

All terminals are provided with slip-on latex protectors for the prevention of Voltage Destruction. (PILL packaged devices do not require protection.)

SILICON PACKAGING

Low cost silicon DIP packaging is implemented and reliability is assured by the use of a non-hermetic sealing technique which prevents the entrapment of harmful ions, but which allows the free exchange of friendly ions.

SPECIAL FEATURES

Because of the employment of the Signetics' proprietary Sanderson-Rabbet Channel the 25120 will provide 50% higher speed than you will obtain.

COOLING

The 25120 is easily cooled by employment of a six-foot fan, ½" from the package. If the device fails, you have exceeded the ratings. In such cases, more air is recommended.

BLOCK DIAGRAM

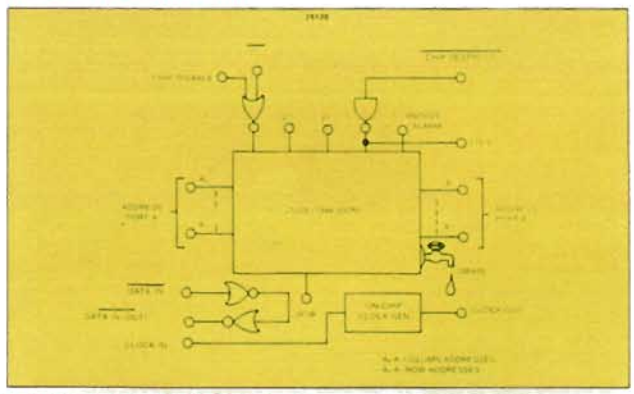

PART IDENTIFICATION

TYPE	"n"	TEMP. RANGE	PACKAGE
25120	0	0 to −70°C	Whatever's Right

1. "Neu" channel devices enhance or deplete regardless of gate polarity, either simultaneously or randomly. Sometimes not at all.
2. "S.O.S." copyrighted U.S. Army Commissary, 1940.
3. Not applicable.
4. You can somehow drive these inputs from TTL, the method is obvious.
5. Measure at 1MHz, 25mVac, 1.9pF in series.
6. For the filaments, what else!
7. You have a dirty mind. S.E.X. is Signetics EXtra Secret process. "One Shovel Full to One Shovel Full", patented by Yagura, Kashkooli, Converse and Al. Circa 1921.
8. J. Kane calls it design (we humor him).
9. See "Modern Production Techniques" by T. Arrieta (not yet written).
10. Final until we get a look at some actual parts.
11. Coffee breaks and lunch hours.
12. Due credit to ELMA for inspiration.

TYPICAL CHARACTERISTIC CURVES

BIT CAPACITY VS. TEMP.

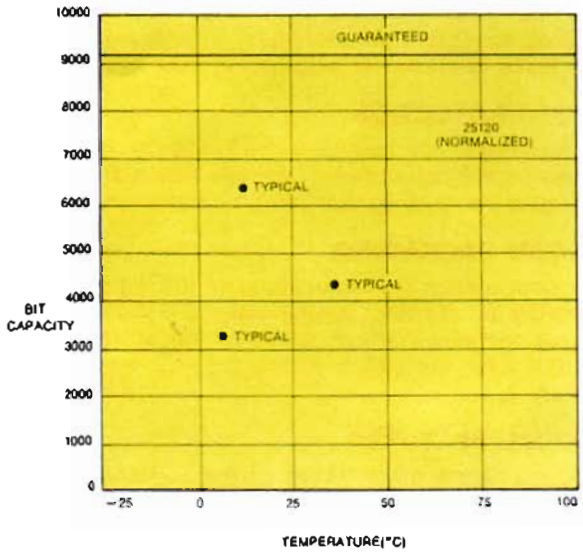

I_{FF} VS. V_{FF}

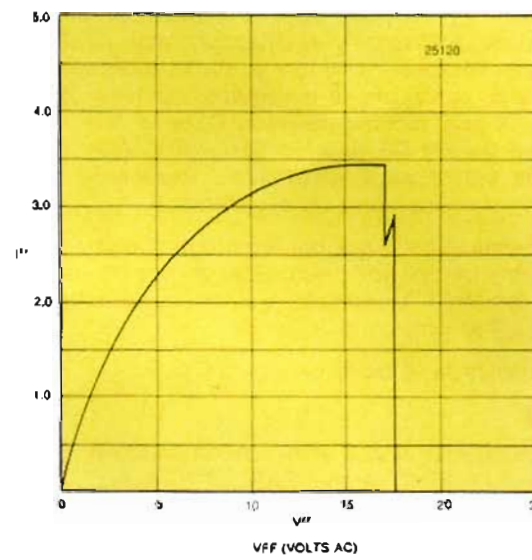

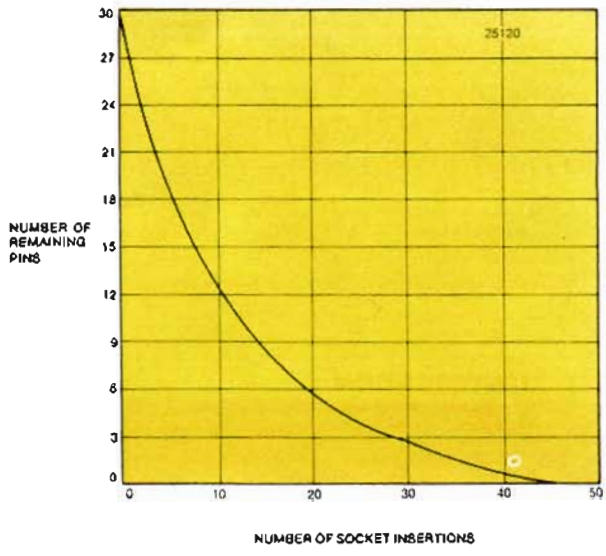

AOL VS SELLING PRICES

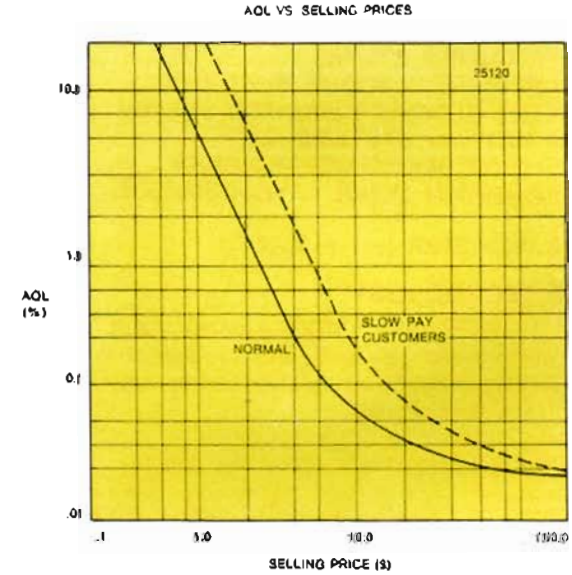

A test circuit of Fig. 1 was constructed using standard 741 type OA's, monolithic transistor arrays and 0.1 percent resistors. Initially, with $i = 0$, the mirror gains were equalized by removing R_1 and R_4 (thus assuring $I^+ = I^-$) and adjusting R_5 to obtain $i_O = 0$. For VCCS operation, G_v was measured leaving port B-B' open ($i = 0$), while for CCCS operation, port A-A' was shorted to ground making $v_1 = v_2 = 0$. To minimize scaling error due to finite gain A of the OA's, the ratio $(R_2 + R_3)/R_1$ was restricted to a maximum value of 100. For $A \sim 10^5$, this assures less than 0.1 percent error at dc due to finite OA gain. Using various resistance combinations, values of G_v up to 0.5 mho and A_i as large as 400 were obtained with a scaling error of the order of 1 percent, which was largely due to nonunity gain of the mirrors. Individual trimming of the mirrors to precisely unity gain (±0.1 percent) resulted in scaling errors of the order of resistance tolerances. The small signal performance of the convertor is shown in Fig. 2(a) for the case where $R_1 = \infty$ and $R_2 = R_3 = 0$. Here the circuit is a VCCS with $G_v = (1/R_4)$. For large R_4 the bandwidth is equal to the GBW of the OA's (~1 MHz). The observed decrease in bandwidth for smaller R_4 is due to the increasing output resistance of the OA's at high frequencies.

In the previous analysis, perfect matching of the input offset voltages of the OA's ($V_{OS1} = V_{OS2}$) was assumed. In practice, a finite mismatch $\Delta V_{OS} = V_{OS_2} - V_{OS_1}$ is amplified by the circuit transconductance resulting in an output offset current i_{OS}. The contribution due to finite and unequal input bias currents, however, can still be made negligible by proper choice of resistances (see (6)). Fig. 2(b) shows the measured output offset as a function of the mismatch in input offset voltage ΔV_{OS}. Here, ΔV_{OS} was varied by adjusting the offset control of one of the OA's in the circuit. For $G_v \leq 10$ mmho, an output offset less than 10 μA was achieved without any special effort at V_{OS} matching.

The CMRR of the circuit was measured at 100 Hz for various resistance combinations and a 20 percent mismatch in mirror gains. For values of G_v from 10^{-3} mho to 0.2 mho, the CMRR was virtually constant at 85 dB (±2 dB), which is comparable to the 90-dB specification of the 741. Larger CMRR would be expected with the use of higher quality OA's.

The proposed convertor circuit may also be used with single-ended input sources, by appropriately terminating the unused input terminals. Terminating the output terminal in a resistance R_O gives an output voltage $v_O = i_O R_O$ and extends the application of the converter to include VCVS and CCVS operations. In this case, however, an additional buffer OA may be required in order to provide low enough output resistance.

References

[1] B. L. Hart, and R. W. J. Barker, "A precision bilateral voltage-current convertor," *IEEE J. Solid-State Circuits*, vol. SC-10, pp. 501-503, Dec. 1975.

[2] D. J. Hamilton, and K. B. Finch, "A single-ended current gain cell with agc, low offset voltage, and large dynamic range," *IEEE J. Solid-State Circuits*, vol. SC-12, pp. 321-322, June, 1977.

[3] S. Pookaiyaudom and W. Surakampontorn, "An integrable precision voltage-to-current convertor with bilateral capability," *IEEE J. Solid-State Circuits*, vol. SC-13, pp. 411-413, June 1978.

[4] B. L. Hart and R. W. J. Barker, "Universal operational-amplifier convertor technique using supply current sensing," *Electron. Lett.*, vol. 15, pp. 496-497, Aug. 1979.

[5] ——, "DC matching errors in the Wilson current source," *Electron. Lett.*, vol. 12, pp. 389-390, 1976.

[6] K. G. Schlotzhauer and T. R. Viswanathan, "Improved integrated unity gain current controlled sources," *Int. J. Electron.*, vol. 36, pp. 97-100, 1974.

[7] J. G. Graeme, G. E. Tobey, and L. P. Huelsman, *Operational Amplifiers Design and Applications*. New York: McGraw-Hill, 1971, pp. 206-207.

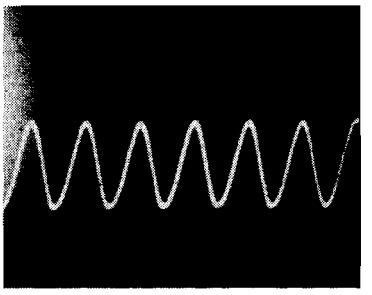

Fig. 1. Oscillator network in the above letter tuned for an oscillation frequency of 10 Hz. (x-axis; 50 ms/cm; y-axis: 2 V/cm)

the input of the next "G" block, will cause phase errors, causing the sine amplitude to be damped out. (These losses do, in practice, always occur.) Thus the output will not have stable gain nor amplitude unless positive feedback is added. It may, in fact, never even start oscillating, or if it does start, it may gradually die out. This is bad enough in an audio oscillator. For a subaudio (amHz? 1 μHz) oscillator it would be disastrous to have to wait for, say, half an hour for a cycle before finding that there was not one.

I am not sure of exactly what form of DVCCS is intended, but if distortion occurs at large levels, and noise predominates with low levels, then I fear that this circuit will not give satisfactory waveforms at any frequency, slow or half-fast.

In other words, despite all the author's discussion of realization, I question whether the author has actually built and tested the circuit.

Reply[2] *by R. Nandi*[3]

The oscillator circuit described in the above letter[1] was built by realizing DVCCS's using μA741 type OA's and a few resistors. Good quality sine-wave oscillations were obtained and the frequency was observed to be smoothly variable in a range of about 1:6 from a typical design value of 10 Hz in the lower side. An experimentally observed waveform is given in Fig. 1 here.

Now, with the availability of high input impedance ($\approx 10^{10}$ ohms obtainable at present with integrated voltage-follower modules [1]) active DVCCS/DVCVS device in IC form [2], the low-frequency oscillator realization proposed in the above letter is probably one of the simplest existing configurations with a grounded resistor control (often required for the electronic control in VCO's).

The comments of Pease on the workability of the circuit are appreciated since they give an opportunity for clarification of its performance through this discussion.

References

[1] "A fast integrated voltage-follower with low input current," in *Linear Applications*, vol. 1, *National Semiconductor*, p. AN5-1.

[2] A. M. Soliman, "A grounded inductance simulation using DVCCS/DVCVS," *Proc. IEEE*, vol. 66, pp. 1089-1091, Sept. 1978.

[2] Manuscript received February 26, 1980.
[3] The author is with the Department of Electronics & Telecommunication Engineering, Jadavpur University, Calcutta 700032, India.

Comments on "Integrable Insensitive Subaudio-Frequency Sine-Wave Generator Using DVCVS/DVCCS"

ROBERT A. PEASE

There is nothing wrong with what the author has said about the circuit in the subject letter.[1] It is what he has not said that bothers me. Any lossiness in the capacitors, or any resistive loading of the capacitor by

Manuscript received November 30, 1979.
The author is at 682 Miramar Ave., San Francisco, CA 94112.
[1] R. Nandi, *Proc. IEEE*, vol. 67, pp. 1568-69, Nov. 1979.

Digitally Programmable Gain Amplifiers with Arbitrary Range of Integer Values

BARRY B. WOO

***Abstract*—A new method for designing programmable gain amplifiers which provide a range of selectable integer gain values is presented. The design uses a single noninverting op-amp with a switched resistor**

Manuscript received January 21, 1980; revised March 13, 1980.
The author is with Fluke Automated Systems, Inc., 630 Clyde Ave., Mt. View, CA 94043.

0018-9219/80/0700-0934$00.75

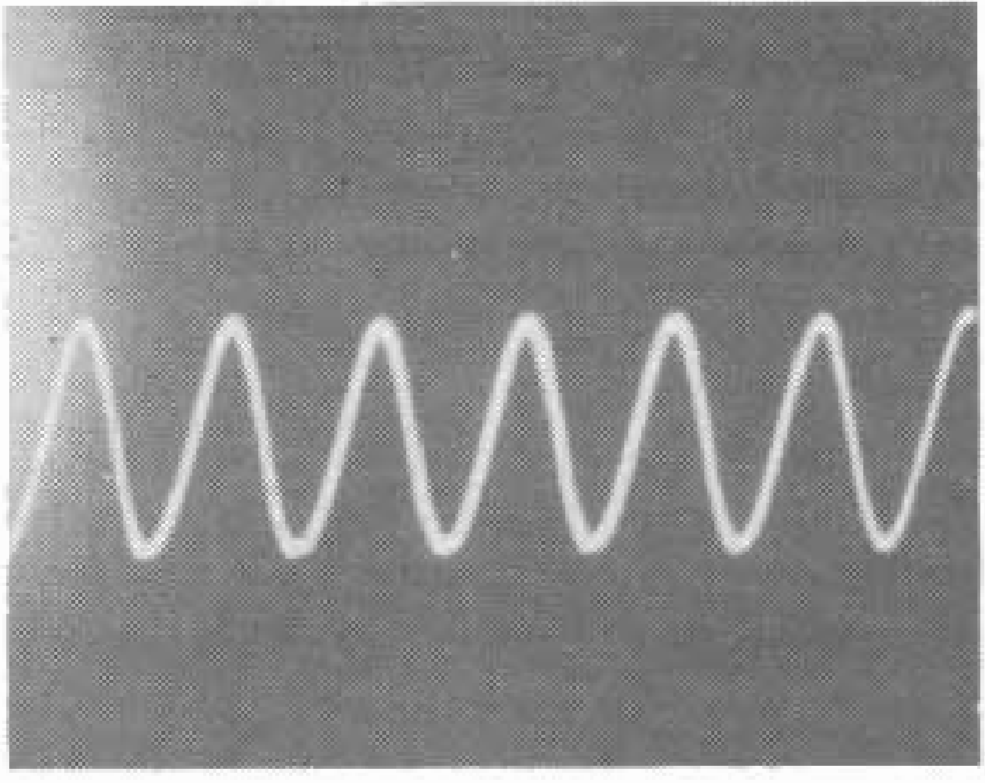

Fig. 1. Oscillator network in the above letter tuned for an oscillation frequency of 10 Hz. (x-axis; 50 ms/cm; y-axis: 2 V/cm)

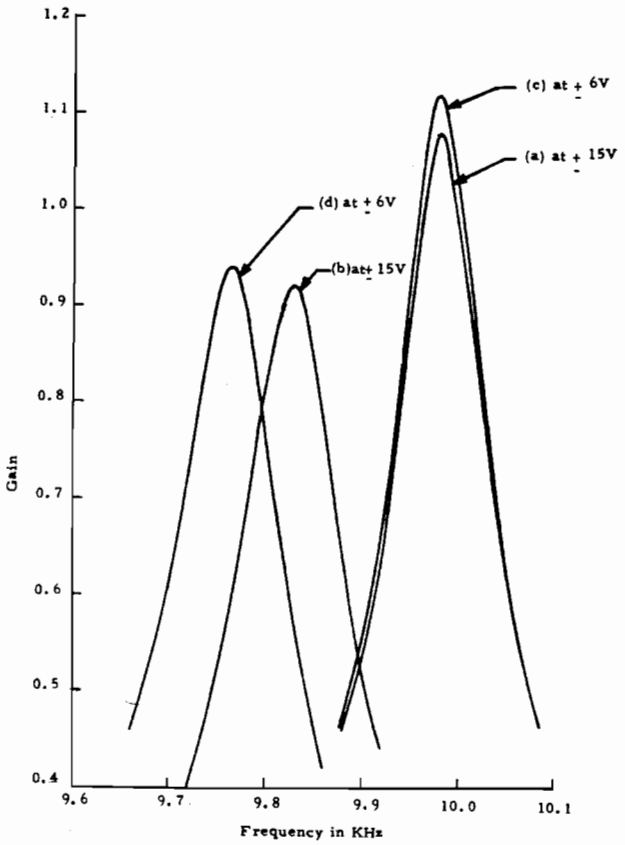

Fig. 5. Magnitude characteristics of double integrator filters. (a) and (c) for the new filter and (b) and (d) for the earlier circuit.

It should also be noted that the noninverting integrator requires only a grounded capacitor, which is an additional advantage from the IC point of view. Finally it must be noted that any change in power supply voltages and/or changes in environmental conditions will affect both ω_p' and Q_p' only to the extent of second-order terms. Thus the active sensitivity of both ω_p' and Q_p' w.r.t. the GB products will be less by an order of magnitude in the proposed filter circuit. In the case of the recently reported filter circuit [1], which also does not require matched amplifiers, the error in ω_p, though it might be small, causes the problems if one compares in terms of the magnitude of the transfer function, which is of final interest in filter applications.

II. Experimental Results

In order to compare the performance the newly proposed circuit with that of the earlier circuit, both the circuits were constructed and tested with randomly selected OA's. The OA's were of LM307N (National Semiconductor) type and their GB products were only in the order ranging from 500 to 800 kHz. Both the filters were realized by choosing the nominal values for Q_p and f_p equal to 100 and 10 kHz, respectively. The f_pQ_p product was twice as much as the lowest value of the GB products, thus creating a worst case situation. The nominal value of gain was chosen to be 1. The decade capacitances were set as close as possible to the calculated values and they were of 1 percent capacitances. However, this author did not have the facility of using precision resistors and the resistances of ±10 percent accuracy were measured for their values. No further adjustments were done in either of the circuits. It should also be mentioned that the same OA's were used in the same order in both the circuits. The supply voltages to the OA's were first maintained at ±15-V dc and the magnitude characteristics of the bandpass output were obtained. These characteristics are given in Fig. 5. In the case of the newly proposed circuit, the realized values of f_p and Q_p were found to be 9.98 and 106.2 kHz, respectively. The Q enhancement is due to the second-order effects of the GB products. In the case of the earlier circuit, the same values were found to be 9.83 kHz and 94.5, respectively, indicating a frequency shift close to 2 bandwidths. There is also a reduction in gain at f_p' in the earlier circuit and this is due to the effect of the nondominant poles which was also expected because of the low GB products of the OA's.

Next in order to find the low sensitivity nature of the proposed circuit, the supply voltages to the OA's were reduced to ±6V dc and the magnitude characteristics were obtained in both the circuits. These characteristics are also given in the same Fig. 5 for purposes of comparison. The maximum change in magnitudes were obtained in the passband of each set of characteristics. They were found to be 3.7 percent in the new circuit and 41.5 percent in the case of the earlier circuit. The superior performance of the newly proposed circuit is clearly indicated by the comparison of these two values and also by comparing the set of characteristics in Fig. 5. The active sensitivity is reduced by an order of magnitude. Though overdrive does not affect the stability of the new circuit, there is an occasional oscillation during switching on the dc supplies. This problem can, however, be avoided by connecting silicon diodes across the resistance "Q_pR" as in other circuits of this category [3].

III. Conclusions

In this letter, a new three-port high gain amplifier using two OA's has been developed. Actively compensated inverting and noninverting integrators have been developed using such a three-port and both of them do not require matched amplifiers. The double integrator-loop filter employing these integrators is found to possess superior performance and this has been brought out through both theoretical as well as experimental results.

References

[1] K. R. Rao *et al.*, "An actively compensated double integrator filter without matched operational amplifiers," *Proc. IEEE* (Lett.), pp. 534–535, Apr. 1980.
[2] S. Natarajan and B. B. Bhattacharyya, "Design and some applications of extended bandwidth finite gain amplifiers," *J. Franklin Inst.*, vol. 305, no. 6, pp. 321–341, June 1978.
[3] D. Akerberg and K. Mossberg, "A versatile active-*RC* building block with inherent compensation for the finite bandwidth of the amplifier," *IEEE Trans. Circuits Syst.*, vol. CAS-21, pp. 75–78, Jan. 1974.

Comments on "High-Input Impedance Inverting/Noninverting Active Gain Block"

ROBERT A. PEASE

In the above letter,[1] while it is not unreasonable to *wish* that a unity-gain inverter could be made as easily and accurately (without resistors) as a unity-gain follower, it is absurd to *pretend* that it can be so made. Messrs. Nandi and Bandyopadhyay are merely running the A_2 amplifier open loop. To postulate that the open-loop gains of A_1 and A_2 are matched, does no harm at all. To reason that therefore the outputs will track, however, is preposterous. An engineer would maintain this argument only so long as he has not tried to make it work. While there are some circuits so complex that it is not reasonable to expect the author to have built and evaluated them, a circuit comprised of just two operational amplifiers (OA's) does not fit that description. Yet it is obvious that the authors have not built this circuit, because it will not work. Even if carefully trimmed, the output will drift all over the place, no matter whether the OA's are good high-gain or good low-gain amplifiers. The forte of the OA is *feedback*. But A_2 has *no feedback* around it. As every engineer who has tried to run an OA with no feedback knows, you cannot expect a stable operating point for linear operation, as it will drift badly.

Allow me to point out the following specific criticisms: I will agree that $A1$ and $A2$ are well matched, so that $Av_{A1} = Av_{A2}$. Let's pick a reasonable value for a typical OA, 100 dB or 100 000. Then it is reasonable to assume that the outputs will track (will be equal but

Manuscript received November 20, 1979.
R. A. Pease is at 682 Miramar Avenue, San Francisco, CA 94112.
[1] R. Nandi and A. K. Bandyopadhyay, *Proc. IEEE*, vol. 67, p. 690, Apr. 1979.

0018-9219/80/1200-1548$00.75 © 1980 IEEE

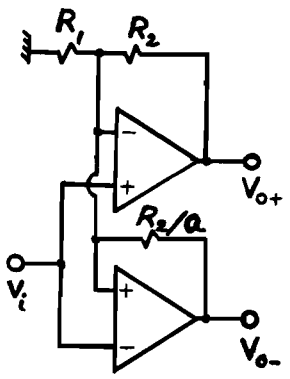

Fig. 1.

opposite) when

$$-\frac{V_{\text{out}}^-}{Av_{A_2}} \cong \frac{V_{\text{out}}^+}{Av_{A_1}} \cong V_{\text{error}}.$$

The outputs would theoretically track if the input voltage of each amplifier is amplified by equal (but opposite) gain. But is it reasonable to expect the voltage *noise* of A_1 to match and cancel out the *noise* of A_2? If A_1 and A_2 each have extremely low-voltage noise, only 0.3 μV in a 10-Hz bandwidth, there would be 240 mV peak-to-peak noise at A_2's output. This would not be usable.

What if A_1 drifts 5 μV (not an unreasonable typical short-term stability for a fairly good OA) and A_2 does not? The V_{out}^+ will only drift 5 μV, but the V_{out}^- will drift 0.5 V. This is not exactly what you expect of a good counter-phase (push-pull) amplifier. It is this kind of problem which has led to the popularity of OA's with resistive feedback around them, and which explains the unpopularity of open-loop amplifiers with no negative feedback.

For Messrs. Nandi and Bandyopadhyay to submit this circuit is an example of unscientific reasoning, and/or wishful thinking. For PROCEEDINGS OF THE IEEE to publish it, shows a noncritical editorial policy. The combination of these two characteristics is unfortunate, but remarkable.

Reply[2] *by R. Nandi and A. K. Bandyopadhyay*[3]

The comments made by R. Pease on the above article[1] appear to evolve from some confusion about the differences between an active-gain element (or gain building block) such as an OA in the open-loop and an active VCVS (or finite gain voltage amplifier) such as an operational amplifier with a closed-feedback-loop so as to bring down the very high voltage-gain into useful and controllable ranges [2].

An inverting/noninverting VCVS can be realized in the building block proposed by the authors[1] by mere insertion of a suitable resistor (R_2/a) as shown here in Fig. 1. The voltage gain assuming matched amplifiers with identical open-loop gains (A) is

$$\frac{V_{0+}}{V_i} = -\frac{V_{0-}}{V_i} = \frac{A(1+a+x)}{A(1-a)+(1+a+x)} \tag{1}$$

where

$$x = R_2/R_1. \tag{2}$$

If $A >> 1$, equation (1) reduces to

$$\frac{V_{0+}}{V_i} = -\frac{V_{0-}}{V_i} = (1+a+x)/(1-a). \tag{3}$$

The circuit in Fig. 1 was built using μA741 amplifiers and selecting $x = 1$, $a = \frac{1}{2}$ satisfactory stable amplification at around 1 kHz was verified. Note that it still has one resistor less than the usual configuration [2] and R_1 (grounded) controls the transmission.

In this connection, the authors like to state that in the above letter, an opening idea was given on the supposition that application alternatives follow the existing design techniques, and hence typical basic controlled source designs were not explicitly reported. However, details on these networks, both in the above letter and the one proposed here, along with their applications are now under investigation/communication [3] stage.

Further Comments[4] *by Robert A. Pease*

Messrs. Nandi and Bandyopadhyay claim that I appear to be confused about the use of an OA used open-loop versus a VCVS or finite-gain voltage amplifier. On the contrary, Messrs. Nandi and Bandyopadhyay were very explicit in their original letter about the requirement for an OA (otherwise their unity-gain amplifier would not have worked). By introducing VCVS, they are trying to cloud and confuse the issue.

The original letter argued that their OA circuit would avoid any need for resistors and the matching thereof. Now they have decided to recommend a circuit with resistors in it. It is they who are confusing the issue.

I constructed the circuit illustrated in Fig. 1, using $x = 1$, $a = 0.5$, $A_1 = A_2 =$ LM 741 and/or LF356. The gain was nonlinear and distorted and the outputs had huge offsets which varied with the signal level. I would not call this satisfactory at 1 kHz or at any other frequency.

(If the input offsets of the two amplifiers were meticulously zeroed out by adding trim pots, which were not shown in the drawing, then the distortion was decreased, but the outputs had a few volts of randomly varying offset, plus 1-V p-p of jitter, wobble, and noise. That is what happens when an OA is run open loop. I do not consider this satisfactory either.)

Further Reply[5] *by R. Nandi and A. K. Bandyopadhyay*

We were carrying out investigations to derive a universal high input impedance inverting/noninverting voltage amplifier/follower building block in our proposed circuit and at first hand obtained satisfactory performance as reported earlier. Fuller descriptions on practical results along with theoretical analysis on the stability and the effect of the amplifier pole are expected to be forwarded for publication at some later time.

REFERENCES

[1] G. E. Tobey, J. G. Graeme, and L. P. Huelsman, *Operational Amplifiers—Design & Applications*. N.Y: McGraw-Hill, 1971 (See Appendix A).
[2] J. G. Graeme, *Applications of Operational Amplifiers*. New York: McGraw-Hill, p. 51, 1973.
[3] A. K. Bandyopadhyay, R. Nandi, and D. Chakrabarty, "A low sensitivity low-frequency single grounded resistance controlled *RC*-oscillator"—Communicated.

[4]Manuscript received May 12, 1980; revised June 6, 1980.
[5]Manuscript received July 28, 1980.

Further Comments on "On Divergence and Probability of of Error in Pattern Recognition"

STEPHEN D. STEARNS

In [1] Chitti Babu presented a relationship between the divergence of a pattern and its nearest neighbor classification risk. Babu's proof of this relationship was later refuted in [2] by Bhattacharya and Toussaint, who supplied an alternate proof in [3]. In this letter, we show that Bhattacharya and Toussaint's refutation of Chitti Babu's proof is logically incorrect. In addition we give a proof of the relationship that

[2]Manuscript received February 25, 1980.
[3]R. Nandi and A. K. Bandyopadhyay are with the Department of Electronics & Telecommunication Engineering, Jadavpur University, Calcutta 700032, India.

Manuscript received June 5, 1980.
S. D. Stearns is with GTE Products Corporation, Sylvania Systems Group, Western Division, P.O. Box 188, Mountain View, CA 94042.

Further Comments on "Integrable Insensitive Subaudio-Sine-Wave Generator Using DVCVS/DVCCS"

ROBERT A. PEASE

I now see that in the above letter[1] Nandi has a circuit for a sine-wave oscillator which can be tuned from 60 Hz all the way down to 10 Hz, although I would hesitate to claim that as "subaudio." But Nandi has not chosen to tell us what is the material or the lossiness of the capacitor he has chosen to operate at 10 Hz. He has not indicated that he considers this capacitor "integrable." He has not shown how or whether the circuit which gave the waveform in Fig. 1. does start oscillating by itself, nor has he indicated what he would do if the oscillator's amplitude failed to start promptly. Again, it is not what he says, but what he does not say, that causes concern.

Manuscript received May 12, 1980.
R. A. Pease is at 682 Miramar Avenue, San Francisco, CA 94112.
[1] R. Nandi, *Proc. IEEE*, vol. 67, pp. 1568–1569, Nov. 1979.

0018-9219/81/0200-0285$00.75

A New Fahrenheit Temperature Sensor

ROBERT A. PEASE

Abstract—**A new monolithic integrated circuit provides an output voltage scaled proportional to the Fahrenheit temperature without the necessity for the user to subtract a large constant offset (as with most Kelvin-scaled temperature sensors). An advantage of this circuit is that the gain is inherently calibrated (e.g., to +10.00 mV/°F) by simply trimming the offset error at room temperature. This behavior is shown to be a predictable function of the inherent bandgap voltage of silicon, +1.220 V.**

The circuit has been designed to permit trimming of accuracy (gain and offset) by blowing metal links and by the familiar Zener-zap technique. The sensor achieves a typical accuracy of ±0.5°F at room temperature, and ±1.0°F over a rated temperature range from −50°F to +300°F, without any requirement for laser trimming. The practical limitation on accuracy is the production testing of devices under conditions of precisely known temperature. Applications circuits are also presented to show where this new sensor has advantages in accuracy and convenience over present sensors.

I. Background

THE voltage of a forward-biased semiconductor junction diode V_F has long been known to be a stable and fairly linear temperature sensor [1]. However, the wide production spread of this V_F has prevented the diode from becoming a popular temperature sensor because it is not easily specified or calibrated tightly. Other drawbacks of the silicon diode are its low sensitivity (only 2 mV/°C) and its nonlinearities on the order of 1 to 3 percent over a 200°C temperature range.

When two silicon junctions are operated at different current densities ($J1$, $J2$), the differential voltage ($V_{F1}-V_{F2}$) is a predictable, accurate, and linear function of temperature:

$$V_{F1}-V_{F2}=\frac{kT}{q}\ln\frac{J1}{J2},$$

where k is Boltzmann's constant, T is the absolute temperature (°Kelvin), and q is the electron charge.

This relationship was first exploited in the LX5600 transducer [2], which used conventional monolithic linear integrated circuit technology to generate an output voltage:

$$V=K_2(V_{F1}-V_{F2})=\left(K_2\cdot\frac{k}{q}\cdot\ln\frac{J1}{J2}\right)T_{\text{Kelvin}}.$$

In this circuit, K_2 was chosen to provide a scale factor of 10 mV/K [3] so that the output voltage would be 2.73 V plus 10 mV/°C. The Kelvin thermometer was further popularized with the introduction in 1978 of the LM135, which has the characteristics of a two-terminal (temperature-sensitive) Zener diode [4] (see Fig. 1). But a specific drawback of any Kelvin thermometer is the way that its output consists of a large constant voltage (e.g., 2.732 V at 0°C) plus a relatively small temperature coefficient (e.g., 10 mV/°C). So if the output voltage changes from 2.98 to 2.99 V, we cannot be sure whether the temperature rose 1°C, or if there was a shift in the offset or gain stability. Even if we trust the gain stability of the sensor, an ordinary Celsius thermometer function requires us to subtract 2.732 V from the output of a typical thermometer circuit, with an offset stability better than 0.3 percent, or else 1°C of error will result. Analog circuit engineers will recognize that this is a nontrivial stability to attain, especially at low cost.

Manuscript received April 30, 1984; revised July 19, 1984.

The author is with National Semiconductor Corporation, Santa Clara, CA 95051.

II. Circuit Approaches and Principles

To avoid the need for the user to subtract this large offset, G. C. M. Meijer of the Delft University of Technology, The Netherlands, devised a Celsius (self-offsetting) sensor in 1979 [5] whereby a current proportional to absolute temperature (I_{PTAT}) is generated, and a current proportional to the V_F of a diode is subtracted from it. Meijer showed that the resulting current is inherently calibrated, when it is properly trimmed at any one temperature. This calibration will be valid despite all normal production variations in transistor and IC characteristics. Refer to Figs. 2 and 3. It can be shown that the total output current is scaled proportional to the Celsius temperature, at every temperature.

This circuit has several weaknesses. It is necessary to trim accurately both the gain and the offset, for at any given temperature one cannot tell which is contributing more to the output current. Also, a current-mode output can easily be contaminated by capacitive coupling from any source of ac, or by leakage of dc current. Worst of all, when the resistors have a linear temperature coefficient, they can cause nonlinear (quadratic) current errors over a wide temperature range, in addition to slope errors.

III. Circuit Description

A new circuit takes advantage of that basic principle, avoids those drawbacks, and provides other benefits also [6]. The block diagram of Fig. 4 shows that a large V_{PTAT} of about 1.59 V at room temperature is generated, and two

0018-9200/84/1200-0971$01.00

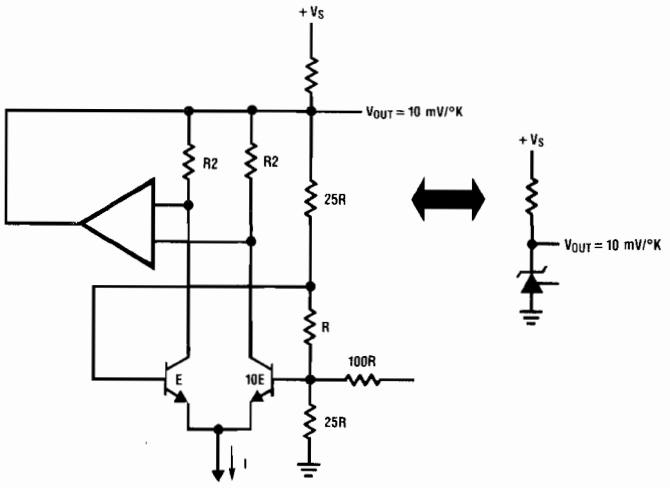

Fig. 1. An integrated-circuit Kelvin-scaled temperature sensor.

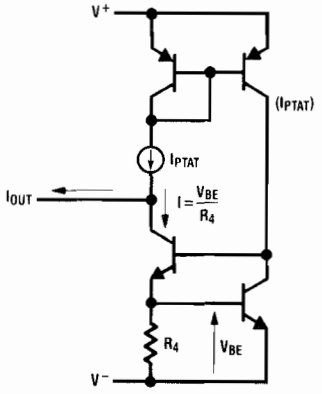

Fig. 2. Principle of a new temperature sensor per Meijer.

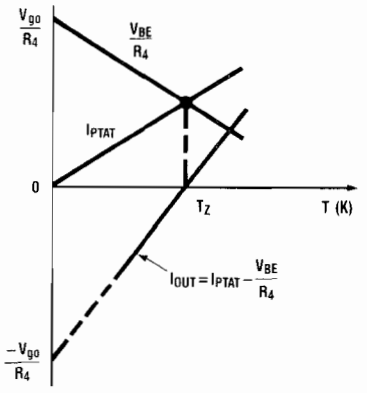

Fig. 3. Graph of currents versus temperature for the circuit in Fig. 2.

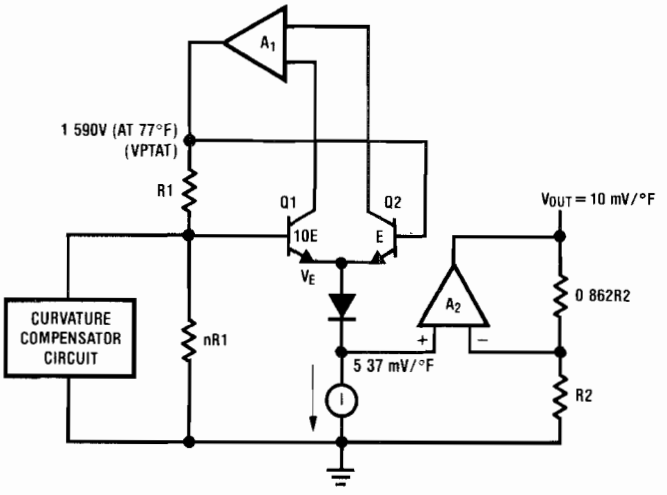

Fig. 4. Block diagram of Fahrenheit temperature sensor.

V_F's are subtracted from it. The resulting voltage is amplified by $A2$ to provide an output voltage that is linearly proportional to Fahrenheit temperature. There is also a curvature-compensation circuit added to compensate for the inherent nonlinearities of V_F versus temperature (patent pending).

In Fig. 4, at +77°F the two transistors require a 60 mV (V_{PTAT}) offset to be imposed across R_1. The amplifier $A1$ will enforce this condition and will servo the base of $Q1$ to a level of $n \cdot 60$ mV, also V_{PTAT}. We choose $n = 26.5$, so at +77°F, $V_b = +1590$ mV + 2.963 mV/°F. Then from this voltage two V_{be}'s (V_F's) are subtracted. Their +77°F value of (588.2 mV − 1.2032 mV/°F) each provides a 77°F result of +413.5 mV plus 5.37 mV/°F at the positive input of $A2$. When this voltage is amplified by a gain of 1.862, the output voltage at 77°F will be +770 mV plus a temperature coefficient of +10.0 mV/°F, as required.

If there is an error in the output voltage at any particular temperature, this error can be fixed by adjusting the ratio ***n*** of Fig. 4. This is easy to do by opening or shorting the links of a quasi-binary trim network in series with one end of the resistor, as shown in Fig. 5. These links are opened by blowing an aluminum fuse ($F1$, $F2$), or by shorting out a resistor with a well-known "Zener zap" ($Z1$–$Z4$) [7]. The output voltage of each die is trimmed while the wafer is being probed at wafer-sort, mounted on

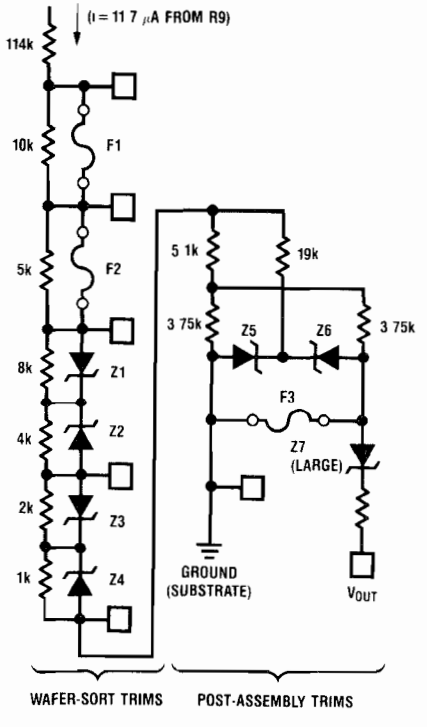

*PATENT PENDING

Fig. 5. Details of trimmable resistor "$nR1$."

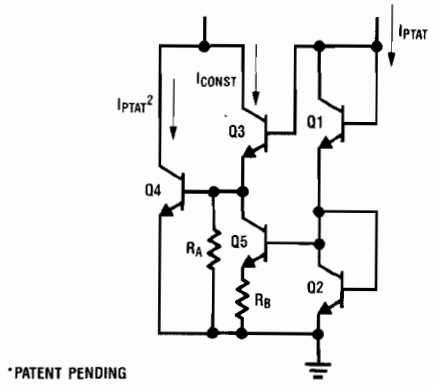

Fig. 6. Details of curvature correction circuit.

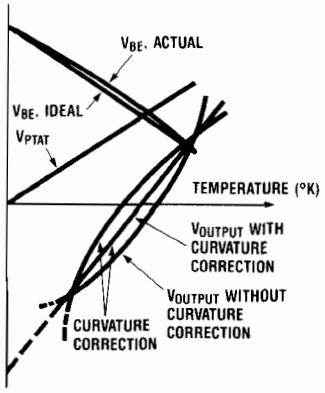

Fig. 7. Graph of voltages versus temperature showing curvature and curvature correction (exaggerated).

a chuck whose temperature is known with an accuracy of $\pm 1/4°F$. When this trimming is completed with minimum error at the ambient temperature of this test, the analysis (in the Appendix) shows that the thermal gain (mV/°F) is optimized to the desired temperature coefficient—a singularly convenient situation.

However, this offset trim procedure does require the electrical gain of the output amplifier to be quite close to the theoretical value of 1.862, so the resistors "$R2$" and "0.862 $R2$" are trimmed using metal links to attain the necessary gain within ± 0.006 before the offset trimming is performed.

The output voltage of the circuit will now be a linear function except for the nonlinear (quadratic) error due to the forward voltage of the transistors (V_F's). In a bandgap reference, where the V_F's are *added* to the V_{PTAT}, the curvature error is *negative* at hot and cold temperatures, referred to room temperature. Here the V_F's are *subtracted*, so the curvature errors would be *positive* (see Fig. 7, bottom curve).

IV. Special Curvature Correction Circuits

In order to correct for this inherent nonlinearity of the V_{be}-versus-temperature curve, the circuit of Fig. 6 is connected into Fig. 4 (patent pending). In this circuit, a current I_{PTAT} is fed through two transistors $Q1$ and $Q2$. This bias establishes the current through the other three transistors:

1) The current through R_B and Q_5 will also be PTAT.

2) The current through R_A will be proportional to V_{be}, as a function of temperature.

3) When these two currents are properly summed into $Q3$'s emitter, $Q3$'s collector current will be constant versus temperature.

4) Because of the transistor's logarithmic relationship, since $V_{be_4} = V_{be_1} + V_{be_2} - V_{be_3}$ then

$$I_4 = \frac{I_1 \cdot I_2}{I_3} = \frac{(I_{\text{PTAT}})^2}{I \text{ constant}}$$

and a highly desirable square-law characteristic is obtained from $Q4$. When this current is added into the node V_b of Fig. 4, the square-law errors inherent in the V_F's of the transistors are cancelled.

Previous designs by Dobkin and Palmer [8] and by Song and Gray [9] attempted to cancel the quadratic errors of a bandgap reference by adding a (linear versus temperature) imbalance current to the collector currents of the transistors which generate the ΔV_{be}. These circuits were able to cancel out the quadratic errors, but left residual errors with a shape proportional to $(T - 25°C)^3$, and a magnitude of 0.1 percent of voltage reference, peak-to-peak, over a temperature range $-55°$ to $+125°C$. Such a 0.1 percent error in a bandgap reference is equivalent to a 0.43°F p–p curvature error in this Fahrenheit temperature sensor. The actual nonlinearity of the Fahrenheit sensors (from data such as shown in Fig. 12) normally shows a factor of 2 or 3 improvement over that figure, thus indicating superiority over older circuits.

V. Circuit Layout Topics

The complete schematic diagram of the temperature sensor is shown in Fig. 8, and a photomicrograph of the integrated circuit is shown in Fig. 9. The die size is 44×72 milliinches. Various details of the circuit which involve both the schematic and the die layout will be discussed here.

Normally, $Q10$ is a considerably larger device than $Q9$, and its collector-to-base leakage and collector-to-substrate leakage would be larger than $Q9$'s leakages. To avoid a resulting imbalance at high temperatures, $Q9$ includes a large transistor (the same size as $Q10$) and $Q10$ includes a small transistor (the same size as $Q9$) with open-circuited emitters, so that all leakages will tend to be equal and cancel.

It may be instructive to observe the actual operating points and current biases which are controlled by a rather devious servo loop. Assume that $Q9$'s base is held at $V_{\text{PTAT}} = 1.59$ V at $+77°C$. Then $Q4$ must provide 11 μA for $R9$ and $R10$, and about 5 μA for each of $Q11$ and $Q12$. When this 21 μA is fed to $Q1$, the current sources $Q2$ and $Q22$ will run at 21 and 10 μA, respectively. Of the 21 μA that flows through $Q2$ and $Q3$, how much will be split

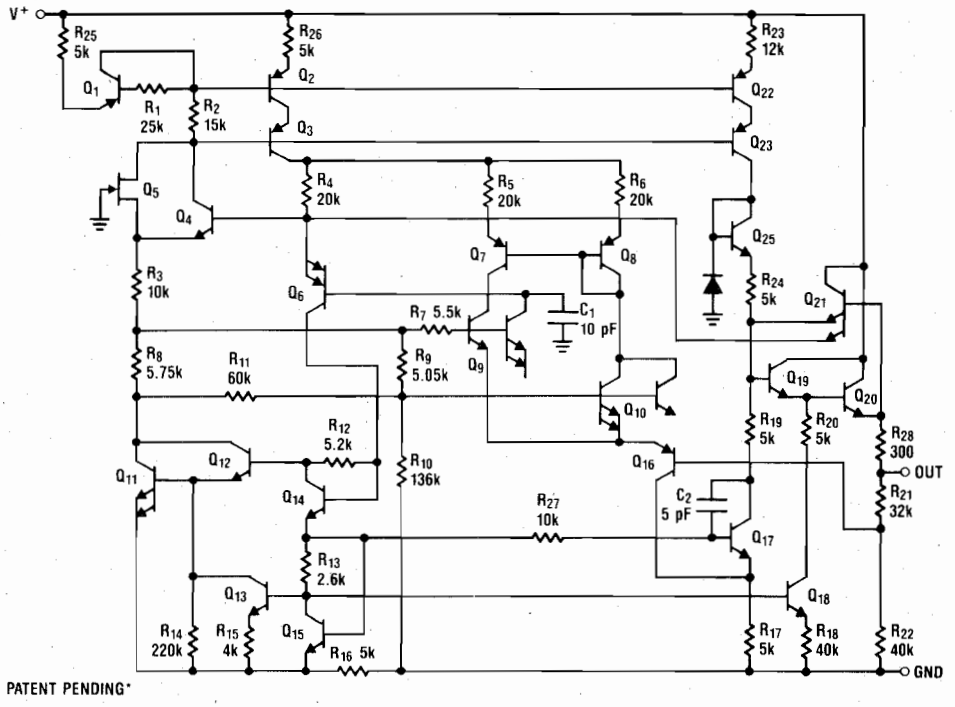

Fig. 8. Schematic diagram of Fahrenheit temperature sensor.

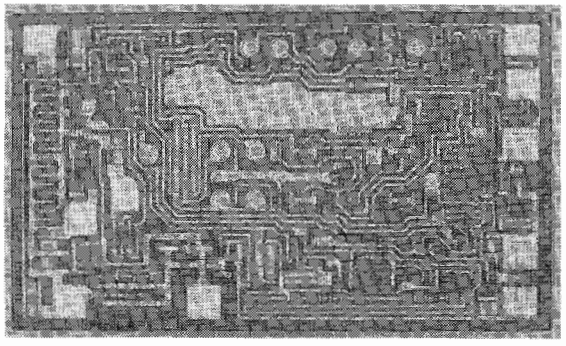

Fig. 9. Photomicrograph of Fahrenheit temperature sensor.

through $R4$, $R5$, and $R6$? The solution is that $Q9$, $Q10$, and $Q16$ will have to run on the current that is left over when I_{CQ23} is subtracted from the current through $R17$. Then $Q6$ will run only on the current that remains when I_{EQ16} is subtracted from I_{CQ3}. It can be readily seen that, under worst-case conditions, I_{CQ16} may cover a wide range. Consequently, $R12$ and $R13$ have been added to the square-law circuit, as "gm compensators," to prevent V_{CE} of $Q15$ from rising as I_{CQ6} increases, and thus to make I_{CQ11} invariant of I_{CQ6}. Also, $Q2$ and $Q22$ have been cascoded by $Q3$ and $Q23$, to maximize the output impedance of those current sources. After it is conceded that this circuit actually does appear to work and to give stable dc operating points, it is not difficult to analyze the two servo loops based on the premise that they are orthogonal and thus largely noninteracting:

1) The differential-mode loop, whereby any differential imbalance of I_{CQ9} and I_{CQ10} acts to drive the voltage on the integrating capacitor $C1$ in the correct direction so that (with the help of some buffering by $Q6$ and $Q4$) the base of $Q9$ is servoed to the desired voltage: $V_b = (n+1)(kT/q)\ln J1/J2$.

2) The common-mode loop, whereby the total current shared by $Q9$ and $Q10$ can increase or decrease to supply the current required by the node at $Q17E$. The collector of $Q17$ is buffered by $Q19$, $Q20$, and $Q16$ to force $Q16$'s collector to provide this current. While it is uncommon to see two interleaved and orthogonal (noninteracting) servo loops, this is an excellent example which actually does perform well. In practice, $Q6$, $Q16$, and $Q23$ each run at around 11 μA.

Also, the curvature-compensation current ($I_{CQ11} + I_{CQ12}$) does not flow directly to the base of $Q10$, but only about 1/10 of this current flows there. The rest is attenuated through $R8$, $R11$ and is dumped unceremoniously into the emitter of $Q4$. This 10:1 attenuation is helpful at keeping a reasonable current level in $Q12$, as $R14$ is already quite large (220 kΩ) and 2.2 M would be ridiculous.

Although the block diagram shows a "diode" connected to the emitters of the n-p-n temperature sensors ($Q9$ and $Q10$ in Fig. 8), it is obvious that the p-n-p transistor ($Q16$) is appropriate because its base draws very little current and thus can be driven below ground even though $R21$ and $R22$ provide a fairly high impedance level. For example, at -50°F, $Q16$'s base voltage will be at -268 mV, but most of its 8 μA emitter current will flow into its collector which is at $+80$ mV dc. Thus, $Q16$ will still have a valid V_{CE} operating bias and adequate beta, even though its base voltage is biased considerably below the substrate (ground).

In Fig. 5, note that access to each pair of Zener diodes is by means of just two probe pads rather than the customary three probe pads. While it is natural to wish that two probe pads can do the job normally done by three, it is fair to observe that this can only be done if a special structure is

used for the Zeners. In practice, if a minimum-geometry device is used for each Zener, when current is fed through a Zener diode pair, the forward-biased diode is almost as likely to short out as the Zener, due to uneven current sharing. To prevent this, the n^+ double-cathode of the Zeners is connected to a wraparound ring which encloses the anodes. When current flows through a Zener, it is localized near the tip of the n^+ which protrudes into the base material, but when it flows through a *forward*-biased diode, it spreads out over the entire periphery of the diode, and is unlikely to blow it out (patent pending).

The multiple probe pads at the right end of the die are the probe pads for the trimming of $R10$. Conversely, the large transistor structures at the left end of the die are *not* the output transistors, but part of an abortive plot to trim the devices after assembly. Fortunately, it has not been necessary to trim the devices after assembly because in most cases the Zeners did not zap, nor did the fuses blow as planned, due to sneak paths through the substrate, etc.

Specifically, it has been found that a double Zener zap such as $Z5$, $Z6$ in Fig. 5 will not function if one end is connected to ground. For example, when a large positive current is applied through $Z7$ to $Z6$ and $Z5$ in an attempt to cause $Z5$ to zap shorted, the $Z6$ junction will be forward-biased and will inject heavily in the mode of a vertical p-n-p transistor, so most of the current will thus be diverted to the substrate, and $Z5$ will not get enough current to zap. Conversely, if a large negative current is applied through $Z7$ to $Z6$ and $Z5$ (to try to zap $Z6$) the tub of $Z7$ will forward-bias versus the substrate, and no significant current can flow through $Z5$ or $Z6$. Even if a large negative current is applied directly to $Z6$, the tub of $Z6$ will forward-bias versus the substrate, which will inject holes upward and cause this p-n-p current to be collected by $Z6$'s anode (n^+). Again, most of the current will avoid $Z6$, which will refuse to zap. But these limitations do not apply to a case such as $Z3$ and $Z4$, where neither terminal is connected to the substrate. Thus, this trim scheme may be used for the post-assembly trim of other circuits (patent pending) when none of the trim pins is tied to ground.

VI. Experimental Results

The circuit has been constructed using conventional bipolar linear IC processing, with Sichrome resistors added on top of the oxide. The characteristics of the integrated sensor [10] are shown in the table of Fig. 10. The low quiescent current of the circuit will cause less than 0.2°F of junction temperature rise above the ambient. The accuracy of the sensor is better than ±1°F over a wide temperature range for a significant fraction of the production distribution (see Fig. 11). Error versus operating temperature is shown for a few samples in Fig. 12. This graph was made by comparing the output of the sensor to the output of a Rosemount platinum sensor/bridge whose accuracy was known within ±0.05°F over the temperature range −50 to +300°F.

V_{OUT}	0V + 10.0 mV/°F (-50°F to +300°F with 0.1 mA pull-down)
Accuracy at 77°F:	±0.5°F max (0.3 typ)
Accuracy for Rated Temperature Range:	±1.0°F max (0.6 typ)
Line Regulation:	±0.1°F ($3.5V \leq V_S \leq 10V$)
Load Regulation:	±0.1°F ($0 \leq I_L \leq 0.5$ mA)
Quiescent Current:	90 μA max (60 μA typ)
Packages:	TO-46 hermetic (-50°F to +300°F) TO-92 plastic (-40°F to +225°F)

Fig. 10. Characteristics of Fahrenheit temperature sensor.

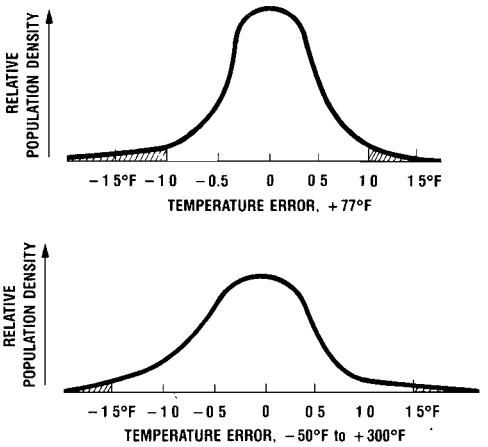

Fig. 11. Distribution of temperature accuracy.

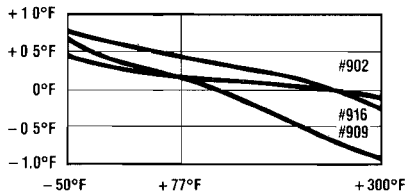

Fig. 12. Typical accuracy versus temperature.

This circuit by itself can, of course, put out only a positive voltage which corresponds to a positive temperature in the range +5 to +300°F. But when a pull-down resistor is connected as shown in Fig. 13 or 14, the circuit can provide negative as well as positive output voltages, referred to its ground terminal. A range of +3.00 to −0.50 V corresponds to +300 to −50°F.

VII. Conclusions

The sensor is easy to apply in ordinary analog or ADC systems. It can provide a positive or negative voltage, and is also usable in a current mode (Fig. 15). Its output can be fed directly to an analog-to-digital converter, to a DVM, or to a voltage-to-frequency converter (VFC), as in Fig. 16. The demands on accuracy for a VFC are much less stringent, even at +100°C, than for a Kelvin thermometer because the VFC does not have to convert the 2.73 V base

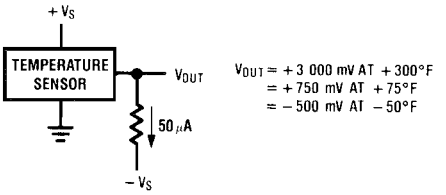

Fig. 13. Pull-down for operation at negative temperatures.

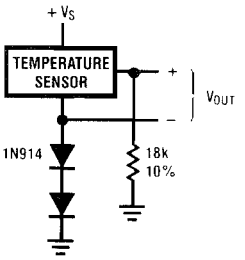

Fig. 14. Single-supply operation for negative temperatures.

line; it only has to convert a 1.00 V signal, which it can do much more accurately than a 3.73 V signal. These results all taken together show that an electronic Fahrenheit thermometer can be constructed and applied with accurate results, with no need for calibration, offsetting, or trimming by the user.

Appendix
Theoretical Derivation of Inherent Calibration of Fahrenheit and Celsius Temperature Sensors

It has been shown by Widlar [11] that a good approximate expression for the V_{be} of a transistor is

$$V_{be} = V_{GO}\cdot\left(1-\frac{T}{T_o}\right)+V_{beo}\left(\frac{T}{T_o}\right) + \frac{nkT}{q}\ln\left(\frac{T_o}{T}\right)+\frac{kT}{q}\ln\left(\frac{IC}{I_{co}}\right)$$

where T is the temperature in °Kelvin, T_o is an arbitrary reference temperature, V_{GO} is the bandgap of silicon, 1.22 V nominal, and V_{beo} is the transistor's base–emitter voltage at the reference temperature. The last two terms are relatively small and will be considered later. We will now define a voltage V_b, which is linearly proportional to absolute temperature and which can be represented by

$$V_b = C_1\cdot T.$$

A circuit to generate such voltage is quite familiar, as shown in Fig. 4. Now the voltage at the transistor's emitter will be

$$V_e = V_b - V_{be} = C_1\cdot T - V_{GO}\left(\frac{1-T}{T_o}\right)-V_{beo}\left(\frac{T}{T_o}\right).$$

We now have two terms which can be manipulated, V_e and

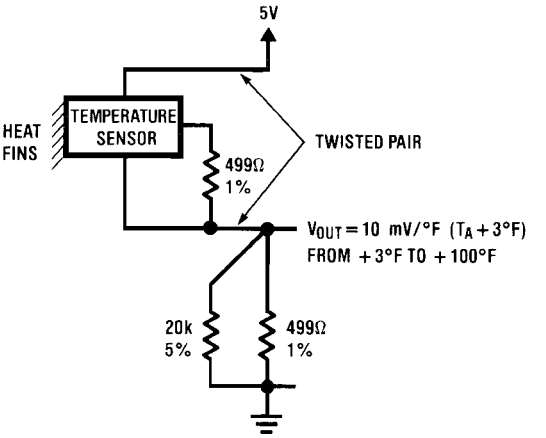

Fig. 15. Two-wire remote temperature sensor (current mode).

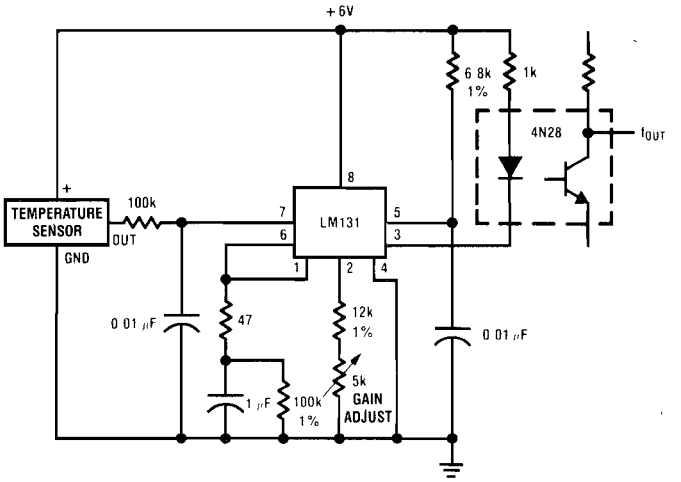

Fig. 16. Fahrenheit-to-frequency converter with isolated output (10 Hz/°F up to +300°F).

T. Let us define $V_e = C_2$ at $T = T_o$ and solve for C_1:

$$C_2 = C_1\cdot T_o - V_{GO}\left(\frac{1-T_o}{T_o}\right)-V_{be}\left(\frac{T_o}{T_o}\right)$$

$$C_2 = C_1\cdot T_o - O - V_{beo}$$

$$C_1 = \frac{V_{beo}}{T_o}+\frac{C_2}{T_o}.$$

We substitute this expression for C_1 into the equation for V_e, and get

$$V_e = \frac{V_{beo}\cdot T}{T_o}+\frac{C_2 T}{T_o}-V_{GO}\left(1-\frac{T}{T_o}\right)-V_{beo}\left(\frac{T}{T_o}\right).$$

The first and last terms cancel, giving $V_e = C_2(T/T_o)+V_{GO}(T/T_o-1)$. When we differentiate with respect to T, we obtain

$$\frac{dV_e}{dT}=\frac{C_2+V_{GO}}{T_o}.$$

This implies that if V_b is adjusted at T_o to give $V_e = C_2$, the rate of change of V_e with respect to temperature will be constant at $(C_2+V_{GO})/T_o$, independent of the value actu-

ally required for V_b. A conceptually useful form of this scheme would be to trim V_b for the condition of $V_e = C_2 = 0$ at $T = 0°C$. This would give

$$\frac{dV_e}{dT} = \frac{V_{GO}}{273.16°} = 4.47 \text{ mV/°C}.$$

This 4.47 mV/°C will be invariant of the processing, the transistor's beta or V_{be}, or the amount of trimming of V_b that may be required, as long as V_e is trimmed to 0 V at a temperature of 0°C.

Still, it would be quite inconvenient to have to hold a circuit to exactly 0°C before you can trim it. Fortunately, it can be seen from the equation $dV_e/dt = (C_2 + V_{GO})/T_o$ that we can accomplish this trimming at any temperature simply by trimming V_e to be equal to 4.47 mV per each degree above 0°C.

The circuit of Fig. 4 is a variation of this, where V_b is twice as big as in the circuit mentioned above, and where two forward-conducting diodes are connected in series to provide an output of 8.94 mV/°C, which happens to be convenient when a final output of 10.00 mV/°C is desired.

The two nonlinear terms from the first equation

$$V_{be} = (\text{linear terms}) + \frac{nkT}{q} \ln\left(\frac{T_o}{T}\right) + \frac{kT}{q} \ln\left(\frac{I_c}{I_{co}}\right)$$

are found in practice to be approximately quadratic for positive temperatures. Also, Tsividis [12] has shown that the assumption that the bandgap voltage V_G is constant with temperature is unfounded and causes nonlinear terms which are especially significant below 0°C. The total of all these error terms does have an approximately square-law characteristic, and has been successfully compensated out by the circuit of Fig. 6 (see text).

While the foregoing analysis is specifically written for the case of a Celsius-scaled device, it is generally applicable to a Fahrenheit-scaled device as well. In particular, the applicable statement would be

$$\frac{dV_e}{dT} = \frac{V_{GO}}{459.7°F} = 2.656 \text{ mV/°F}.$$

ACKNOWLEDGMENT

The author wishes to express his appreciation to C. T. Nelson, who designed the circuit; K. Cavanaugh, who designed the masks and layout; and D. O'Neill and E. Dietz, who planned and performed the testing of the breadboards and preproduction samples.

REFERENCES

[1] R. Pease, "Using semiconductor sensors for linear thermometers," *Instrum. Contr. Syst.*, June 1972.
[2] R. Dobkin, "Monolithic temperature transducer," in *Dig. Tech. Papers, Int. Solid-State Circuits Conf.*, 1974, pp. 126, 127, 239, 240.
[3] National Semiconductor, LX5600, LX5700, and LM3911 Product Data Sheets.
[4] National Semiconductor, LM135 Product Data Sheet.
[5] G. C. M. Meijer, "An IC temperature sensor with an intrinsic reference," *IEEE J. Solid-State Circuits*, vol. SC-15, pp. 370–373, June 1980.
[6] R. Pease, "A new Fahrenheit temperature sensor," in *Dig. Tech. Papers, Int. Solid-State Circuits Conf.*, 1984, pp. 292–293.
[7] B.-S. Song and P. R. Gray, "Precision curvature-compensated CMOS bandgap reference," *IEEE J. Solid-State Circuits*, pp. 634–643, Dec. 1983.
[8] R. C. Dobkin and C. R. Palmer, "A curvature corrected micropower voltage reference," in *Proc. Int. Solid-State Circuits Conf.*, Feb. 1981, pp. 58–59.
[9] D. Soderquist and G. Erdi, "*The OP-07 ultra-low offset voltage op amp*," Applications Note AN-13, Precision Monolithics Inc.
[10] National Semiconductor, LM34 Product Data Sheet.
[11] R. J. Widlar, "An exact expression for the thermal variation of the emitter–base voltage of bipolar transistors," *Proc. IEEE*, Jan. 1967.
[12] Y. P. Tsividis, "Accurate analysis of temperature effects in I_c-V_{BE} characteristics with application to bandgap reference sources," *IEEE J. Solid-State Circuits*, pp. 1076–1084, Dec. 1980.

Robert A. Pease was born in Rockville, CT, on August 22, 1940. He graduated from M.I.T., Cambridge, MA, with the B.S.E.E. degree in 1961.

He was employed by George A. Philbrick Researches, and designed many early solid-state operational amplifiers, analog computing modules, and voltage-to-frequency converters. In 1976, he joined the Advanced (Industrial) Linear IC Group, National Semiconductor Corporation, Santa Clara, CA, where he has designed several voltage regulators, references, and temperature sensors, and has written several technical articles on linear circuits and ferroequinology. He has also designed low-noise micropower op amps for the lunar seismometers, and temperature-to-frequency converters which were carried to the summit of Mt. Everest.

PROVOCATIVE TOPICS ON LINEAR INTEGRATED CIRCUITS

Robert A. Pease
Staff Scientist,
National Semiconductor Corp.
Mail Stop C2500A
2900 Semiconductor Drive
Santa Clara CA 95051

Abstract

A number of provocative topics will be discussed such as: relative priorities, Check Lists, Design Reviews, Design Rule Checks, digitally computerized simulation, analog computer simulations, breadboards, test strategies, and metal-mask options. These are demonstrated to be valuable for enhancing the probabilities of good results in high-volume linear Integrated Circuits (ICs).

Historical Considerations

It is well known that the design of linear integrated circuits has been, historically, an arcane art, accomplished by bizarre persons with strange personalities. While there may be an element of truth to this, it is also fair to say that reasonable, rational, and sane persons can also design successful ICs. But for best results, they should take up some of the more fanatical attitudes that the old-timers used to have. A mere enthusiasm is not necessarily sufficient for good results, but a meticulous, skeptical, suspicious approach sometimes is helpful to find and avoid possible problems that could wreck or delay the project. This is especially true for "large-volume" linear ICs where it has been observed that a few "minor" problems can turn the project into a "low-volume" one. I shall try to present examples and approaches which have been found useful to avoid or solve the problems.

Note that I am dwelling primarily on linear ICs. This is partly because the design of large digital ICs has often turned into a committee effort, with many man-years focused in a few months. Linear ICs are still usually designed by a small team of 1 or 2 engineers. And, unlike the reliance of digital designers on cook-books, standard cells, and libraries of basic building blocks, linear ICs are usually designed from scratch and customized because the performance depends critically on the details of the transistor and resistor design and layout. This must come under the purview of one person. The point I want to make is, that the techniques necessary for successful digital design and for successful linear design may have some areas of overlap, but many areas where there is nothing in common.

Planning

Whether we do it explicitly or implicitly, every good linear design starts with a data sheet. If a circuit is worth doing, that is because eventually an engineer will pick up the data sheet and say, "I really have to have this circuit, it's exactly what I need". The circuit must include features and specifications and applications circuits which are excellent and will exert a magnetic pull on the prospective user, and these features must make a good total package when put on a data sheet. You don't have to write up the entire data sheet at the start of the project; it may be enough to just envision it, and write down the key features, but the designer must be aware that the data sheet will not only be the birth of the product, and the life of the product, but, if the data sheet is deficient, it can be the death of the product. So, the designer (with some help from his friends in the marketing department) must be sure that the circuit can lead to a data sheet that will have all the right features, and not any significant drawbacks. A single specification is not enough. You have to have a reasonable package of specs and features, and you have to have applications notes so the user will not be able to fail to apply it successfully. At every stage, the designer (or the project manager) must be aware of the nascent data sheet. Any decision that would impact the data sheet must be considered. For example, if a test engineer begs to drop a test because it is too difficult, and this would impact the data sheet seriously, the designer must resist this. [1]

Of course, the "silicon" must be able to do its job, but after that is accomplished, the data sheet must do its job. For example, the old UA709 amplifier had some good features, and some characteristics that were not so good. But the UA709's low noise was never mentioned on its data sheet. So, to this day, people ask me, "Where can I find a low-noise op-amp with lots of gain-bandwidth product?". And I have to explain that the UA709 really does have good low noise voltage, lower than all the 741s or the BIFETs. So, in some applications, the 23-year-old UA709 is still the best IC for the job. But the data sheet kept that a secret.

Design Reviews

When the design of a circuit is nearly complete, the designer will normally hold a Design Review for an audience of his peers. The complete design is presented, and all the engineers are invited to marvel at the good design, and to criticize the bad design. On a good day, nothing is left but the good stuff. But, we all know that occasional design flaws slip past the jury of peers. To help minimize this, at our company, we have appointed several Czars who are intended to be experts in one particular narrow field of design. I am, for example, the Czar of Bandgaps, and I am responsible for keeping track of all our expertise in the design of Band-gap reference circuits. Every new circuit, every old circuit, every good

circuit, every bad circuit, and every "fix" of a bad circuit is supposed to be filed with me. That way, I can make sure that only good circuits are copied. This theory does not work perfectly, but it works rather well, and the number of old errors that are repeated has decreased markedly, recently. Czars have also been appointed in other areas such as: Start-Up Circuits, Zener-Zap Trim Circuits, Proof-reading, and Computer Errors. [2]

After the nominal circuit design is correct, it must be laid out by a mask designer (though sometimes by the Design Engineer himself). Either way, when the layout is complete, it is necessary for another round of detailed criticism, peer review, and close checking. This we denote by the technical phrase, "Beer-check". This differs slightly from the Design Review, as there is opportunity for all the technicians, and mask-designers to review the layout, as well as the engineers. Now, as I mentioned earlier, there may not be as many bad or incorrect ways to lay out a digital circuit as there are bad ways to design it. But for linear circuits, there are a lot of bad things possible to do in a layout, even for a good engineer and a good mask-designer. The kind of flaws that are spotted at almost every beer-check range from the trivial to the disastrous. Fortunately, we often find that the most unlikely persons have the best knack of spotting errors, so we invite the most unlikely people as well as the senior persons. After you try this a few times, it is easy to conclude that the cost of the pizza and potables that are given out as rewards, is 2 or 3 orders of magnitude cheaper than letting the errors go un-detected. An even more important feature of this review, is that every level of worker learns that every person's input is valuable for catching errors, and that nobody is immune to errors. It is especially apparent that when a person has worked on a project for a long time, he becomes incapable of catching his own errors, so that is another reason to bring in outside help, simply to bring new eyes to the task.

Testability

In digital design, when the number of gates rises into the thousands, the importance of testability is well appreciated. So, gates that might fail and go un-detected because of a lack of accessability are given special consideration. Likewise, in linear circuits, the importance of being able to test every important function easily and quickly is recognized - especially in the case of mixed-mode ICs where many analog and digital functions are combined. [3] Also, the techniques for designing testable linear ICs are being expanded. This is true not only for production testing, but also for characterization and for troubleshooting. Every node must be accessible to the troubleshooter with his probes, or else that becomes an invitation for Murphy's Law to strike. [4] Small probe pads are best introduced, to facilitate probing, and vias must bring up nodes from the first layer of metal to the top layer of metal. Fortunately, new instruments are also helpful in aiding the access for probing, but these are tedious and slow.

Computer Simulation

The advantages of the (digital) computer for simulating a circuit's performance are well known and widely trumpeted. The disadvantages are most often swept under a rug or ignored. Many young engineers are slow to catch on that computers lie. Yes, they sometimes lie even when you type in the correct SPICE program. Consequently, wise designers learn to run "sanity checks" on their computer programs, and to blow a whistle if the computer gives a patently absurd answer. Other times, the computer itself blows the whistle and gives up - quits - fails to converge. In every case, the engineer must be prepared to evaluate the results and to decide when to dis-believe the computer.

In a recent case, an operational amplifier was designed with MOSFETs, and all the DC characteristics were simulated and the results were as expected. However, the high-speed response and dynamic stability tests gave no results, due to lack of convergence. Despite pleas, howls, and threats, the Computer Experts were not able to achieve any useful convergence. (Note, this may be partly because the MOSFET models are optimized for accuracy in digital circuits, where the voltages flash from maximum to minimum in less than a nanosecond, but the models often give poor results in linear circuits....) Finally I suggested an analog computer. The designer said that would be difficult as he did not have any suitable kit-parts or test patterns - no suitable samples of MOSFET. I pointed out, that's no problem; replace each MOSFET with a bipolar transistor with a big resistor in its emitter, to simulate the transconductance of the FET. He thought about that, and then objected that the bipolar transistor would have unrealistic frequency response. I replied that he should add in capacitors 1000 times bigger than the actual IC strays. Then the circuit would work at 1/1000 the speed of the real circuit, and the stray capacitances would have negligible effect. He thought about it. He convinced himself. He built the "Analog Computer". It confirmed his theories. He built the IC and it worked just like the Analog Computer. So, when the digital computer in the simulator gives answers that are useless, you are not necessarily stuck.

Recently we had a large non-linear circuit whose simulation gave absurd errors. Fortunately, we were able to snip away 63 of the 65 transistors and the SPICE continued to give absurd errors. By concentrating on the minimal circuit, we were able to convince the Computer Experts that they must

look in their realm for the error, and in only a couple weeks they were able to find 2 or 3 absurd errors. Then we went back and the complete circuit did begin to work correctly.

Breadboards

As an old-time analog-circuit crank, I find that breadboarding is a valuable part of circuit design. For one thing, it does provide an independent check on the results from digital simulation. For another, it can use "kit-part" transistors so if the transistor model is incorrect, the circuit may tell you a different answer than the SPICE, because the model is poorly related to the real world. But, just as I cautioned you that digital computers lie, well, so do Analog ones, and breadboards, also. They lie most easily in high-frequency response, where the stray capacitances of the breadboard cannot relate to the strays in the monolithic circuit. Still, a breadboard is valuable, I find, exactly because it is so touchy, so sensitive, so grouchy. These days, it will normally be foolish to depend only on the breadboard, but in many cases it's not much more foolish than to rely solely on the computer simulation.

Another good use for the breadboard, is to exercise the tester. Several times, I have been able to back up the breadboard to the tester, and confirm that the tester was working OK, even before we had working silicon. So, when you want to run a "sanity check" on the tester, the breadboard can be valuable and reassuring. In some cases, you can see the effect of a sweeping change in the breadboard, quicker than on the computer. So, in many cases, I recommend that you should not abandon the old techniques, if they help you to avoid the occasional hoaxes that the digital computers attempt to foist on you.

More Planning

The way to use PERT charts and Gantt charts to best manage the progress of a project is well documented in the art of program management, and I will not dwell on that. I do want to point out the advantages of using a big Check-list. I have been using Check-lists for a number of years to make sure I didn't forget anything, because on a complex project, there are so many things that could easily be forgotten or neglected.

For example, on a recent project, I started out one evening to write a list of "100 Items to do", and when I was done, I counted 155 items. On another project, I was collaborating with two junior engineers, and each of us wrote down a long list. I had on my list more than a dozen items they had not thought of, and they each had a dozen items I hadn't thought of. So, by combining lists, we came up with a much better list, which will also be useful on other future projects.

Priorities: In these lists, we try to indicate the relative priorities of various aspects of the project. That way, if the noise is the most important feature, and you find yourself compromising the noise to gain other advantages, at least you cannot say that you didn't realize the noise was important. One time I was given a new project, and as I wrote down the list of 100 Items, I tried to decide if it was most important to keep the die size small, or the quiescent power, or the tempco, or, what? (This was because the target specs and guidelines were rather imprecise.) Shortly, I decided that none of these had the highest priority, not even, getting the silicon out in the fastest possible time. Rather, the most important item was, to have a 100% probability of getting the circuit working on the first try. This circuit was to be part of a new library of cells, and the other cells also had to come out on the first try. Some of those cells had 8 components, and others had 18 components, and my circuit had 85 components, so it was not easy to build up the confidence that my circuit would be sure to work on the first try. So, I optimized the circuit in terms of getting each component precisely defined, with no mistakes. And, with a little help from my friends at the Beercheck, we did get it out on the first try, and it worked well.

Now, if it had not worked on the first try, what would be the next most important thing to optimize? In this case, I should optimize the ease of analyzing any errors in the first try, and the ease of fixing them and getting the second silicon to work. So, I also did that. And, even though the silicon did work on the first try, I was able to make a minor metal-mask tweak to make it work even better on the second try.

Why Metal Mask Tweaks? There are a couple advantages in being able to tweak the metal mask. If I have a circuit which works pretty well, I can usually predict how to improve the characteristic nicely by changing a resistor value. So, I could achieve this by changing the base mask. In theory, that is easy to do, but in practice, it's easy to botch the computation, and get resistors that don't do what we want. If, on the other hand, we build in a large group of small resistors in series with the main resistor, and short some of them out with metal links, we can prove that the resistor can be changed to the desired value by opening or shorting links. We can do that on some samples, and then we know that when we change the metal mask, the improvement will be exactly what we want. Further, we can do all sorts of evaluation on these parts, (which would otherwise have to wait until the base mask was changed).

Another reason to make the changes on a metal mask, is that we can more quickly get the changes executed. If we hold some wafers at base (before base mask) then we can slap on a new base mask and get out new wafers in 2 or 3 weeks, but if we hold them

at metal (before metal mask) we can get new wafers in 2 or 3 days.

"Hold At Metal" The other major reason for planning metal mask changes is related to major goof-ups. Many times a new run comes out, and absolutely nothing works. A little study shows an error in the metal mask. We can go back and start new wafers, and put on a new metal mask, at a cost of a few weeks and many thousands of dollars. If instead, we hold 9/10 of that run at metal (before metal mask) and bring out 1/10, we can learn from looking at those few wafers that the metal mask was botched. Then we can expedite the new metal mask, and slap it on the left-over wafers, and cut down on wasted money and time. I have not declared myself the Czar of "Hold at Metal", but I am one of its strongest proponents, and I have convinced everybody I work with that it saves money, as well as time which is more priceless than mere money. Now, it is true that having metal-mask options available on resistors does waste a little die area, but I've seen this cost paid back 10 times over. And the cost advantage of holding wafers at metal could be computed at 100 times or more.

"Bust-proofing"

Ideally, all ICs, all circuits should survive the outputs and the inputs shorted to either supply, and large ElectroStatic Discharges (ESD) to any terminal. But, to be realistic, there are some circuits whose performance would be compromised by any attempt to make the circuit un-bustable. For example, the old UA709 would be destroyed in a few seconds if its output was shorted to either supply, and even faster than that if one input was shorted to a supply. The demise of the UA709 was easy to predict when the LM101 and UA741 came along, as they could survive all these conditions. However, recent requirements that all IC terminals must withstand ESD transients larger than 2000 volts (standard test procedure, applied through 1.5k in series with 100 picofarads) showed that the LM101 had weaknesses when the inputs were pulled toward -80 volts. We were able to add a couple small transistors to clamp the negative-going transients and bring the LM101A's toughness up to an acceptable level, and the user cannot even see the difference. Of course, customers still call us up and ask how many transistors the LM101A has, because they believe MIL-HDBK 217B, which asserts that the circuit will be more reliable if it has fewer transistors. We know better than that. The LM101A has become MORE reliable for the last couple years, not LESS reliable, because it has had 2 transistors added. In general, adding transistors that make the circuit more bust-proof, more forgiving, are good investments in customer satisfaction. In the design of a "high-volume" linear circuit, there are many cases where this is proven wise, to add circuitry in the interest of "foolproofing". But, of course, we don't call it that.

Testing Strategies

Sometimes a test engineer tells me, "Here's a test that almost never has a failure, so I want to delete it." Oh, what is that test? It's a quiescent power drain test at a 6-volt supply. But, that's odd, we know there are a lot of parts that are dead and would have to fail a quiescent current test. It turns out that the 6-volt quiescent current test is run after the 33-volt quiescent current test. Consequently, that test lumps together the dead parts and the ones that break down when the supply is increased from 6 to 33 volts. By swapping the two tests, we can learn quite a bit from these tests. By putting the tests in a foolish sequence, you fail to get your money's worth from these tests.

Recently a product engineer asked me why we were getting such poor yields at final test. I checked into it and discovered that the wafer-sort test had much wider limits than the final test. By being too busy to set the limits in a rational fashion, we were condemned to spend the costs of packaging up bad dice and throwing them away. Normally, it's rational to test and throw away the bad dice before assembling them.

Other times when a test engineer complains that a batch is having a terrible yield on one test. I ask him, "Are they failing any other tests, too?" Often, they reply, " I don't know, after this test, the parts failed and weren't tested any further". I have to explain why they must test these parts in "ALL" Mode, rather than "PASS" mode, so we can learn about the results of the other tests. Usually on a well-designed part, you can learn more from studying bad parts than good ones. I usually spend a minumum amount of time studying tests that give good results, but a lot of time studying parts that fail tests. Sometimes I even have a hundred "inked" dice bonded up. You can learn more from studying the sinners than the saints.

Summary

The world is a complicated (and non-linear) place.

References

[1] R. A. Pease, "How To Get The Right Information From A Datasheet", EE Times, April 29, 1985, pg 31.

[2] R. A. Pease, "Preside Over Power Components With Design Expertise", EDN, Oct. 12, 1989 p. 134

[3] P. Fasang, D. Mullins, and T. Wong, "Design For Testability For Mixed Analog/Digital ASICS" presented at IEEE Custom Integrated Circuits Conference, Rochester NY, pp 16.5.1-16.5.4, May 16-19, 1988

[4] R. A. Pease, "Troubleshooting is More Effective With The Right Philosophy", EDN, Jan 5 1989 pp. 147-156.

THE DESIGN OF BAND-GAP REFERENCE CIRCUITS: TRIALS AND TRIBULATIONS

Robert A. Pease,
Staff Scientist, MS C2500A
National Semiconductor Corp.
2900 Semiconductor Drive
Santa Clara, CA 95051

ABSTRACT

This tutorial will briefly discuss the designs of various band-gap references, with an emphasis on technical problems that caused serious troubles. Practical solutions are shown for problems, and a methodology is shown for solving problems.

TUTORIAL

The band-gap reference has been a popular analog circuit for many years. In 1971, Robert Widlar introduced the LM113, the first band-gap reference. [1] It used conventional junction-isolated bipolar-IC technology to make a stable low-voltage (1.220 V) reference. This type of reference became popular as a stable voltage reference for low-voltage circuits, such as in 5-volt data acquisition systems where zener diodes are not suitable. Band-gaps are also used in digital ICs such as ECL, to provide a local bias that is not adversely affected by ambient noises or transients.

The principle of the band-gap circuit is well known and will be mentioned here in the briefest terms. The circuit relies on two groups of transistors running at different emitter current densities. The rich transistor will typically run at 10 times the density of the lean ones, and a factor of 10 will cause a 60 millivolt delta between the base-emitter voltages of the two groups. This delta voltage is usually amplified by a factor of about 10 and added to a Vbe voltage. The total of these two voltages adds up to 1.25 volts, typically, and that is approximately the band-gap of silicon.

In figure 1, the LM113 schematic diagram shows a basic band-gap circuit. Q1 runs at a relatively high density, about 150 microamperes per square mil. Q2 is operated at a low density, about 10 microamperes per square mil, and so its Vbe is much less, about 70 millivolts. Now, Let's ASSUME that the circuit is at balance and the output is near 1.22 volts. Then the 70 millivolts across R5 is magnified by the ratio of R4 to R5, about 8:1, up to a level of 600 millivolts. This voltage is added to the Vbe of Q4 (about 620 millivolts at room temperature) to make a total of 1.22 volts, as required. Q4 then amplifies the error signal through Q7 and Q9, which provide enough gain to hold the V+ bus at 1.22 volts. The beauty of the band-gap reference is the summation of the Vbe term, which decreases at the rate of about -2 millivolts /°C, and the (delta-Vbe term) which grows at about + 2 millivolts /°C, to achieve an overall Temperature Coefficient (Tempco) that is substantially zero. All band-gaps employ this summation of a growing and a shrinking voltage, to make a stable low-tempco voltage. Further, it has been shown [2] that when a circuit has been trimmed to the correct voltage, the correct tempco will follow, despite process variations in parameters such as Vbe, beta, sheet resistivity, etc. Consequently, band-gap circuits are often trimmed to their ideal voltage so as to provide also a low tempco.

There are many other circuits that have been used for band-gap references, and each one has its own set of advantages and disadvantages. Figure 2 shows a simple brute-force scheme: the stack of D1-D7 run at a rich current, and D8-D13 at a lean current, and the resultant output of 1.2 v has a low tempco. Unfortunately, this circuit is not very suitable for operation on low supply voltages, but its concept is clear!

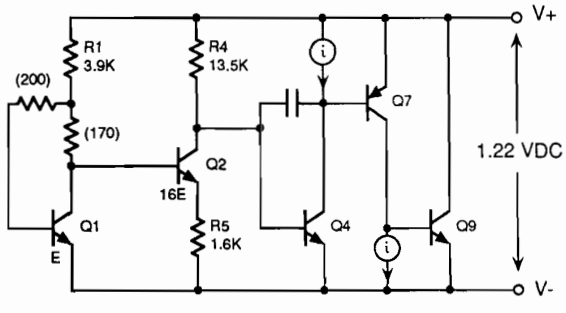

FIGURE 1
SCHEMATIC DIAGRAM, LM113 (SIMPLIFIED)

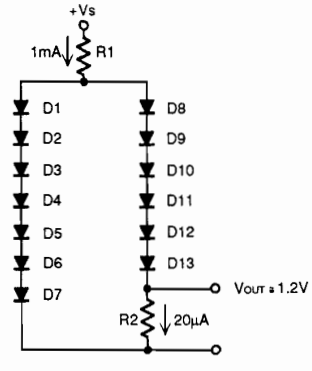

FIGURE 2
SCHEMATIC DIAGRAM of SIMPLE BANDGAP

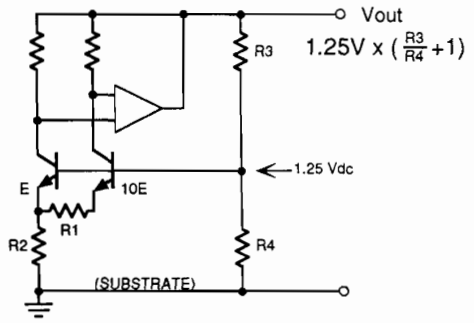

FIGURE 3
SCHEMATIC DIAGRAM, AD580 (SIMPLIFIED)

The Brokaw cell (Figure 3), attributed to Paul Brokaw, [3] is often used for output voltages larger than 1.2 volts, as its Vout can be scaled by the ratio of two resistors, R3 and R4. A similar circuit is used in the LM117, an adjustable medium-power voltage regulator, which can be trimmed by two external resistors to any voltage in the range 1.25 to 57 volts. (Figure 4)

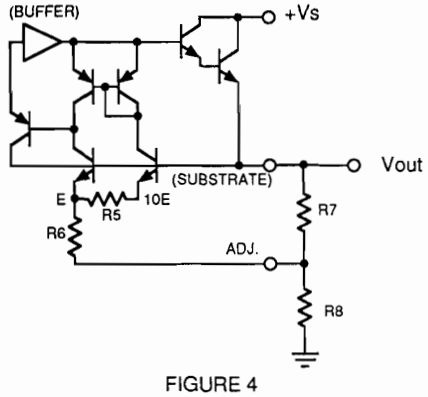

FIGURE 4
SCHEMATIC DIAGRAM, LM117 (SIMPLIFIED)

A different approach is shown in Figure 5, representing the LM185/LM136. The version shown is suitable for operation as a 1.2-volt shunt regulator. The circuit of Figure 6 is for the reference section of the LM10, which provides a reference voltage of 0.200 volts and operates on a supply voltage as low as 1.0 volts.

One of the major drawbacks of all these simple band-gaps is the curvature of their tempco. Ordinarily, all band-gaps have a negative tempco when hot, and a positive one when cold. See figure 7. Around room temperature, tempcos as good as 20 or 30 ppm/°C are typically attainable, but over a wide range, -1% shift is normal at hot and cold temperatures. Several techniques have been devised to neutralize this curvature. Figure 8a shows a little add-on circuit which can do a surprisingly effective improvement at warm temperatures. Sometimes 2 or more transistors are connected with slightly different biasses, for improved curvature correction. A circuit which can improve the tempco curvature at cold temperatures is shown in figure 8b.

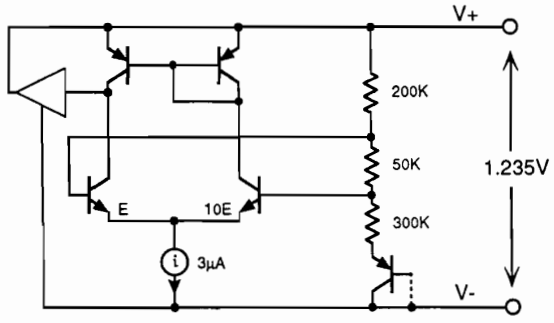

FIGURE 5
SCHEMATIC DIAGRAM, LM185 (SIMPLIFIED)

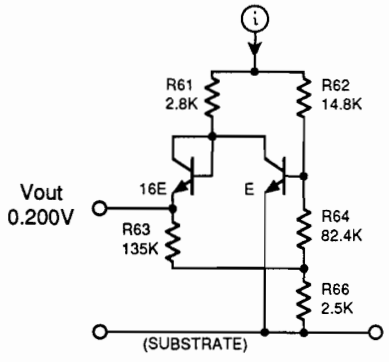

FIGURE 6
SCHEMATIC DIAGRAM, LM10 (SIMPLIFIED)

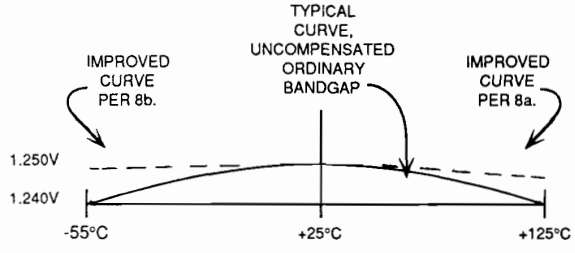

FIGURE 7
V_{REF} VS. TEMPERATURE
BANDGAP REFERENCE

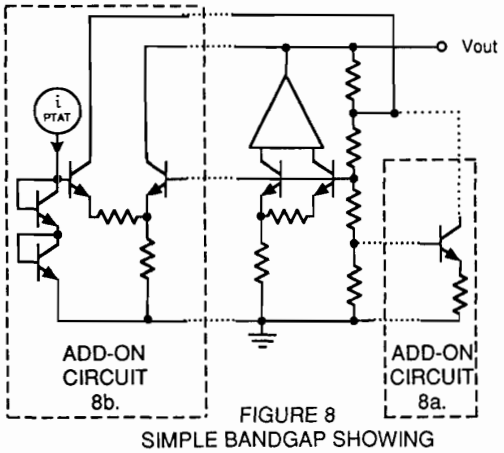

FIGURE 8
SIMPLE BANDGAP SHOWING
CURVATURE COMPENSATION

Other techniques rely on smoothly non-linear techniques to cancel out the parabolic curvature. One example is shown in Figure 9. It is able to make an improvement of about 8:1. Other proprietary techniques and circuits are also used.

Band-gap references are not normally thought of as low-noise devices, because of the high gains needed. Every 1.25 volt of output is made up of 0.6 volts of Vbe, plus a 60 mv signal magnified by a gain of 10. Consequently a 5-volt band-gap includes a 60 millivolt signal amplified by a gain of about 40. If you allow for a transistor running at 1 UA to have a theoretical amount of noise equal to 14 nv per square-root Hz, the output will have at least 800nv per square-root Hz, and for a bandwidth of 10 kHz, about 80 Uv rms, or 500 uv p-p. Obviously, a band-gap reference with 100 ppm of noise will not give good results with a 12-bit DAC. So although some band-gaps may claim advantages of low-power operation of only 11 uA total, for high-performance designs the circuits are more popular when biassed with 30 or 60 uA or more, for each group of emitters.

Many systems-on-a-chip depend on a band-gap to start up and bias all other circuits. Thus, the band-gap must be sure to start promptly under all conditions. It is not trivial to do this, as a band-gap that has not started may not have any bias currents to make sure it does start. Even the leakage of a junction capacitor may be sufficient to prevent the circuit from starting. A good solid sure-start circuit is advisable. If I can think of one, I'll present it. Figure 10, reserved.

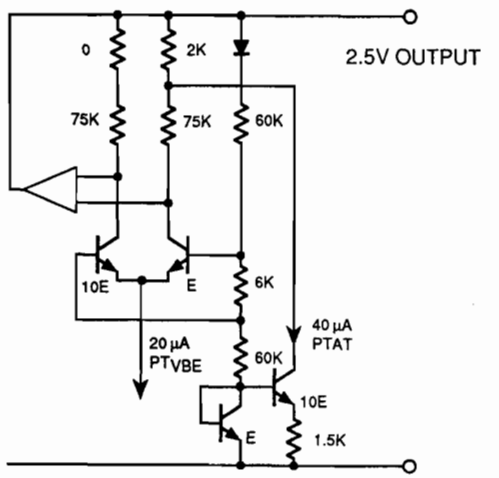

FIGURE 9
SMOOTH CURVATURE-CORRECTION CIRCUIT

There are many other ways to design a band-gap reference that does not work well. There are two basic choices: with a computer, and without a computer.

If you rely on breadboards, you may trim to get a good tempco but the resistors may not be matched in their tempco if you just use RN55D or RN55Cs. If you actually get film resistors with a matched tempco, the tempco of your circuit may still not agree with your final circuit, due to the gross tempco of the actual diffused (or film) resistors in your circuit. Example, diffused resistors have 1600 ppm per ^{o}C; implanted silicon resistors are even steeper. Some SiChrome resistors can have +300 ppm/^{o}C, and they will give a completely different tempco than 50 ppm/^{o}C. Also, diffused resistors sometimes have a different tempco depending on how many volts they are biassed below their tub's voltage, so the tempco of a voltage divider may not be as good as theoretical, unless each resistor sees about the same amount of tub bias; this can be done (in theory, if you have enough space) by giving each resistor its own tub.

FIGURE 10
SURE-START CIRCUIT for BANDGAP REFERENCE

The dynamic stability of a band-gap IC is often inferior to its breadboard due to stray capacitances in the breadboard. Consequently the use of SPICE and similar computer simulation schemes are becoming popular, as the capacitances can be well defined without false influences or strays. However, as with any other computer usage, a "sanity-check" is recommended to insure rational results in a known situation, before the computer can be trusted in unknown cases.

However, there are many ways to get false results in SPICE and other simulation schemes. Most transistor models do not accurately model the shape of the curve of Vbe vs. temperature. It is sometimes possible to tweak the characteristics of the model of a transistor until the tempco of the breadboard or analog model matches that of the computer model. However, after you have done this, it is not safe to assume that reasonable changes in the operating points will cause reasonable changes in the tempco. The actual changes may be different from the computed changes, even for a minor tweak. For example, a minor change of a resistor, which actually gives a change of +10 ppm per ^{o}C may be predicted by the computer to give a change of +5, with no plausible reason for the discrepancy. Often, a SPICE run will simply fail to converge at cold temperatures. If you take the observed operating points from a run at -34^{o}C and insert them as the node-set voltages for a run at -35^{o}C, don't be surprised if you still get no convergence. In a circuit, being close to the right answer helps, but in SPICE, convergence depends on the matrix yielding a valid answer, and that has almost nothing to do with the circuit or reality. For example, on a recent SPICE run, I had a resistor connected only to ground and to a capacitor which also went to ground. If I used an * to comment-out the resistor and capacitor, the circuit refused to converge, but if I put them back into the circuit, they somehow helped the matrix converge. So, in

the future we may be able to insert useless meaningless resistors around the circuit, whose sole function is to help the matrix give good convergence.

Other practical examples will be given of how to make a band-gap work badly, and how to give it a chance of working well.

Advice to the Engineer

While it is not practical or suitable to show examples of How to Design a Good Band-gap Reference, it is possible to show examples of circuits that do not work well. By avoiding the thicket of bad design practices, a good design may be seen to be feasible. A list of examples of bad design will be given here, grouped in categories.

Layout Problems

- Bad cross-coupling. As in a good op-amp, the critical transistors should be laid out with cross-coupling - example, 1/2 of Q1 on one side of Q2 and the other half on the other side of Q2, so the centroid of Q1 and of Q2 will be at the same place. Likewise, critical resistors should be arranged so that 1/2 of R1 is on one side of R2, and the other half on the other side - again, common centroid. This is especially important in a power regulator where large temperature gradients can be expected. Also it is very advantageous to reject gradients in sheet rho, beta, etc. In a typical band-gap, this cross-coupling should be done for the delta-Vbe transistors and also for the PNP transistors that serve as their collector loads, unless you can prove it to be unnecessary.
- Layout of Delta-Vbe circuit vs. Vbe. In some circuits, the transistors that form the delta-Vbe are not the same ones that make the Vbe. These must be laid out to reject gradients from all expected directions.
- Thermal regulation errors. When a band-gap regulator or reference has a step change of dissipation, the thermal gradients across the die may cause significant errors. A good layout with adequate cross-coupling and attention to detail is normally required to keep these errors to acceptable levels. A good power regulator can do 0.005%/W; a good precision reference can do 0.002%/W. Thoughtful layout techniques as mentioned above are especially important for a library cell that will be used in an ASIC, because when it is used, you can never guess where the thermal gradients will come from; you have to lay the circuit out to reject gradients from all directions.
- Rejection of I x R drops in power busses. If some parts of a circuit are connected to a power bus at one point, and other parts at another point, significant errors can be caused when current flows through the bus. In some cases, you can specify that no such current can flow through the bus; in cases where the bus current must be expected to fluctuate, the connections to the power bus must be arranged to reject the I x R drops.
- Thermal stress in corners of the die. When a precision circuit is laid out far off the center-line of the die, near a corner, the thermal stresses can cause errors. For best results, avoid putting precision circuits in corners or far off-axis. This applies to ASIC cells, of course.

Start-up Circuit Problems

- Bad start-up. Normally applies only to series-mode circuits, not shunt-mode.
- Start-up circuit too weak. This often happens when the start-up current (high-value resistor or EPI-FET) puts out too little current. This problem is often exacerbated by high temperatures, leaky diodes, leaky capacitors, or excessive substrate currents if one of the terminals is pulled below the substrate. You can never absolutely be free of this, but you can avoid it if you have a good start-up test. Do not expect the computer to be of much help.
- Start-up circuit too strong. This can happen when the EPI-FET puts out too much current and overwhelms the start-up circuit. It's not easy to model this in SPICE, but you can think about this as a worst-case to be avoided.
- Dynamic start-up with no dc start-up. The circuit can start on the dv/dt of the input, but may not start if dv/dt is small. Use a start-up test.
- Start-up too slow. You may avoid this by good worst-case design and good modelling in SPICE or breadboarding.

Oscillations

- Oscillation due to capacitive load. You avoid this by good circuit design. Sometimes the condition can be helped by pre-load currents, or by a series R-C damper network to ground. Use Pease's Principle to make sure that ringing is not lurking nearby which will turn into oscillation when you turn your back. Check for ringing at all relevant temperatures.
- Oscillations at some temperatures and not others. Check for this by watching the Vout for ringing as the part's temperature is SWEPT from one extreme to the other. Note, monitoring the dc output voltage is not necessarily sufficient to insure freedom from oscillation.
- Oscillations due to improper start-up circuit. Make sure the start-up circuit is well-designed and well-behaved in all worst cases.
- Oscillations because the breadboard had enough strays but the IC does not. This is a matter of good modelling. The breadboard can be expected to lie about this. SPICE may be helpful if applied thoughtfully.
- Obscure oscillations. The computer often lies badly about these. It may refuse to admit that they happen, and refuse to show them happening.

DC Output Voltage Errors (Room temp)

- Excessively broad distribution of Vref. When you expect a tolerance of $\pm$ 3% and the observed distribution is excessive, the problem is usually either badly-matched resistors or transistors. The geometries must be identical as drawn and as masked. If your small resistor is short, your large resistor should be made of a group of short resistors. Sometimes it is advantageous to draw the Rs and Qs as cells, so that the mask-making process acts identically on each cell.
- Parts cannot be trimmed. Make sure that every voltage can be trimmed by at least one combination of trims; avoid any possible "holes" in your trim scheme.

- Interaction of trims. It is a good idea to make sure that in your plan, the size of each trim is (substantially) invariant of whether any of the previous trims have been done.
- Dependence of pre-trim V ref on beta. In general, a higher beta transistor has a lower Vbe. In some designs, a pinch resistor (whose resistance is a linear function of beta) is used to compensate for the shift of Vbe and improve the room-temp accuracy. In other cases, a pinch resistor is used to compensate for tempco, even as it degrades the room-temp accuracy.
- Band-gap "narrowing". With high-speed processes, the band-gap voltage (and the voltage for zero tempco) is decreased vs. ordinary transistors. A good breadboard can help determine the right place to operate. The computer cannot.

Tempco Errors

- Wrong "V-magic". Normally there is one value of Vref, which (if you trim to that voltage) gives the best tempco. But there may be circuit problems, masking problems, or process problems that can cause tempco to vary even when the V ref is well trimmed. The deviations are typically + or - 5 or 10 ppm/^{o}C, but in a poor circuit can be considerably worse.
- Inability to trim to Vmagic. If there are "holes" in your trim scheme and you cannot get the Vref to trim to Vmagic with good precision, the tempco will be degraded by about 3 Uv/^{o}C per millivolt of error.
- Tempco curvature-correction circuit does not work well? MOST tempco curvature correction circuits do not work well the first time. Some are still bad after several tries. Computer models are of little help, or usually of much harm. Breadboards often are no better. Trial-and-error is usually needed.
- Leakages cause high-temp errors. You have to watch out for device leakages, tub leakages to substrate, and capacitor leakages. These are usually hard to model.
- Tempco does not agree with breadboard. If you used discrete low-tempco resistors in your breadboard instead of the monolithic (diffused or implanted) resistors, you can expect bad results.

Computer Modelling Problems

- Bad model of Vbe versus temp. Most transistor models are poor. It is possible to tweak various parameters to get decent match of Vbe vs. temp, but the derivatives may not be realistic.
- Mis-typed data input. Be extremely meticulous about typing errors and copying errors. Strange results may occur otherwise.
- Failure to converge at cold temperatures. Yes, that is a problem....
- Failure to converge despite accurate "Node-set". Sometimes if Nodeset is TOO accurate, convergence is worse than if Nodeset is approximate.
- Failure to converge due to 35 volts across a diode (Vf) which does not conduct any current. Maybe you need a bigger diode....
- False current due to dv/dt across a hidden capacitance. When the collector voltage has a dv/dt, the "collector current" can show the current through the c-b junction, but the current through the c-substrate junction may not be included.
- False dv/dt due to .PLOT command truncating the .TRAN command. This has been observed to happen, when the end of the .PLOT is not at the same time as the end of .TRAN.

Miscellaneous Problems

- Low-voltage-lock-out problems. You must engineer carefully to get good results at low voltages, as the system may require good behaviour at very low supply voltages, and the reference may not want to give it. The breadboard works better than SPICE, here. Check at all temperatures.
- Thermal limit circuit works badly, poor errors, soft knee. Try to use a circuit that has been successful in the past. Hysteresis is often a good feature.
- High noise due to starvation. Be sure to check the breadboard. (Unless you are sure the computer gives good answers.)
- High noise due to high resistance. Check the breadboard carefully, and use realistic transistor samples. High-beta parts will have higher rbb' than low-beta ones.
- Bad matching due to buried layer. If the actual buried layer causes crystal growth to fall in the middle of a critical transistor, Vbe matching may suffer. Note, the crystal growth at the surface is shifted from where the buried layer appears to be.
- Saturation due to insufficient buried layer. Transistors do not run well at warm temperatures when asked to run near saturation, especially at high temp, when the buried layer has been omitted.
- Assembly shift. This is usually caused by stress sensitivity. A good layout can help minimize this.

Czardom

- At National, we appointed a Czar to oversee the design of all band-gap circuits, and to monitor and to log all good and bad results. This has helped cut down on the number of repetitive foolish errors.

Bibliography

1. LM113 Data Sheet, National Semiconductor Linear Data Book, 1972, 1976.

2. R. J. Widlar, "An Exact Expression for the Thermal Variation of the Emitter-Base Voltage of Bipolar Transistors", Proc. IEEE, Jan. 1967.

3. AD580 Data Sheet, Analog Devices, Inc., 1975.

A DUAL HIGH-CURRENT HIGH-VOLTAGE DRIVER

R. Shields and R. A. Pease

National Semiconductor (UK) Ltd., Greenock, Scotland, and
National Semiconductor Corporation, Santa Clara

ABSTRACT

A high-current driver which can control two independent 1.6-ampere, 20-volt loads is described. Inductively induced flyback voltage transients are clamped internally to safe voltages, for inductance values up to 80 millihenries. Short-circuit load conditions are tolerated and controlled.

INTRODUCTION

The design of circuits utilising high-current NPN transistors to switch high current resistive loads to ground is easy. However if the load has an inductive component in addition to an ohmic value, when the load switching transistor is switched OFF, the load voltage may rise to very high voltages causing breakdown and likely damage to the transistor.

In order to prevent this situation from occuring, the load must be voltage clamped to a safe, tolerable value for the output transistor. Frequently, a clamp diode connected between the output and the power supply is a simple solution utilised to protect the output transistor. However for high-performance applications where the stored energy of the load needs to be dissipated quickly, a high voltage clamp is required so that a large V can cause a large di/dt.

A new circuit is described which is able to switch a load current of 1.6 amperes from a 20-volt supply. At turnoff, the stored energy of the inductive load is dissipated internally, clamping the flyback voltage to 72v. During this repetitive turn-OFF transient, the peak dissipation in the clamping circuit can be as high as 116 watts for 2.4 milliseconds typical. This monolithic circuit includes protection features for added robustness and comprehensive system interface capabilities.

DESIGN METHODOLOGY

Figure 1 shows the block diagram of one channel of the device. The control input provides the stimulus to drive the output transistor via the control and linear current amplifiers. The output transistor is surrounded by protection circuitry. In case the load is shorted from output to supply, a current limiter circuit promptly decreases the current flowing in the output transistor to 2.1 amperes typical by way of a current regulation loop. After an internally - set time delay of 200 microseconds, typical, the short-circuit detector sets the short circuit flag indicating to the monitoring media that an overload condition exists. Additionally, the output transistor and associated drive circuitry are switched OFF for the remainder of the control input ON cycle. At the next rising edge of the control input, all internal monitoring circuitry is reset and the output is switched ON.

If the output transistor is switched ON and no significant load current is flowing after 200 microseconds, an open-circuit flag is set indicating to the external monitoring medium that an open-circuit load condition exists.

An overvoltage detector circuit senses when the output voltage exceeds 28 volts nominal. If this situation occurs, the output transistor and associated control circuitry are switched OFF promptly and remain so until the overvoltage condition is removed.

There are some additional situations which arise in the operation of an integrated circuit of this type in real systems. An example is the situation when the load short circuit involves a long length of wire, perhaps in an automotive wireloom. A wire as long as 2-3 meters can represent an inductance of 50 - 150 microhenries in series with less than 0.5 ohms. If such a load is connected to the output of this circuit, the power transistor can turn ON and pull the load to ground momentarily. When the small inductor permits the output current to build up to its maximum value, typically 3 amperes, the output transistor comes out of saturation and the output voltage rises up towards the 20 volt supply. However, when the current limiter circuit causes the output current to decrease from 3.0 to 2.1 amperes, or, if

CH3315-9/93/0000(0076)$1.00 © 1993 IEEE

any other transient causes a momentary decrease in the output current, the L di/dt can cause the output to rise past the overvoltage detector threshold. If this occurs, the output transistor is turned OFF promptly and the output voltage rises rapidly to the clamp level of 72 volts. Momentarily, the dv/dt of the output voltage causes the output transistor to conduct even although its base drive has been actively removed. Within a few microseconds, the output voltage falls back again, through the overvoltage detector threshold towards 20 volts, allowing the output transistor to switch ON again. The repetitive cycling of this effect causes severe overheating of the output transistor and destruction in less than 100 microseconds.

To prevent this mode of failure, improved protection circuits have been designed (U.S. Patent Pending). These circuits sense the normally illogical condition of an over-voltage being detected at the output whilst the control input is high, instructing the output to be ON. (The logical conditions are: The output voltage is LOW when the control input is HIGH. The output is clamped HIGH momentarily when the control input returns to LOW.) If output HIGH and control input HIGH occur simultaneously, a gate detects this condition, and the output transistor is switched off promptly and latched OFF until the next cycle of the control input signal. In this way, the severe overheating condition is considerably reduced and no damage to the output transistor occurs.

As previously discussed, the inductive flyback energy of the load (or long cable) is clamped internally. Due to processing constrictions, the output NPN transistor itself cannot dissipate these large energy surges. Instead, high-power, deep-base vertical PNP structures have been designed in a high-voltage clamp circuit. Figure 2 shows the schematic diagram of such a clamp circuit. The clamp voltage is derived as:

$$VCLAMP = \{Vz1 + i1R5 + i1R4 + Vbe(Q4)\}\frac{R1 + R2 + R3}{R3}$$

Here, i1 is kept relatively constant by the use of Q5 and Q6 which dump excess current to ground, hence:

$$VCLAMP = \{Vz1 + 2VBE(Q5,Q6) + VBE(Q4)\}\frac{R1 + R2 + R3}{R3}$$

A degree of thermal coefficient cancellation is employed to maintain the clamp voltage within 2 - 3 volts over a 150 degree temperature range. The positive tempco of Vz is addressed by the negative tempco of VBE(Q4) and (Q5). The effective tempco of V(R5) is positive which is addressed in part by the negative tempco of VBE(Q4).

The resulting current flowing in Q3 provides drive to the Q1, Q2 power Darlington resulting in clamp voltages of 72 volts at 1.6 amperes for milliseconds or 3 amperes for several microseconds.

Other device features include reverse supply protection and voltage transient protection on the supply rails for both integrated circuit power supply and load voltage supply rails. A thermal limit circuit detects excessive die temperature and can shut down the whole circuit.

CIRCUIT PERFORMANCE

Figure 3 shows the Input/Output (I/O) response of one channel with a load of 18 mH and 15 ohms. The output transistor has a saturation resistance of 400 milliohm typically. The output clamp level at turnoff is 73 volts for approximately 120 microseconds. If this load is shunted by a low value resistance, the short-circuit current threshold is exceeded at around 1.8 amperes typical. Figure 4 shows I/O waveforms where the short circuit threshold current has been exceeded. The output remains ON for the duration of the time delay of 200 microseconds, at which time the output is switched OFF and the short circuit flag is switched active LOW.

Figure 5 shows a partial short circuit load condition of 36 microhenries in series with 4 ohms. Initially, the output transistor is allowed to saturate; however as the current builds in the load, the output transistor starts to come out of saturation. As the current regulation loop pulls the current down to the "short circuit current" value, the output voltage pops upwards in response. In this situation, the output voltage does not exceed the overvoltage detector threshold voltage, hence the output remains ON for the duration of the time delay period.

Figure 6 shows a partial short circuit load condition of 120 microhenries in series with 0.5 ohms. In this situation, the over-voltage detector threshold voltage is exceeded. As the control input signal is also HIGH, the internal protection circuits promptly latch the output transistor OFF

until the next rising edge of the control input. Some typical device characteristics are shown below:

Characteristics:

Supply voltage range - 7 - 20V
Supply current (quiescent) - 4mA typical
Supply current (both channels ON) - - 140mA typical
Clamp energy rating (repetitive) - 138 mJ at 80 mHy

Control inputs:

Input Impedance - 80K ohm typical
Input switch threshold - standard CMOS/TTL

Output current diagnostic thresholds:

Open circuit threshold - 120 mA typical
Short circuit threshold - 1.8A typical
Diagnostic flags sinking current - 5 mA typ.

CONCLUSION

This dual driver circuit with comprehensive protection circuitry is ideal for driving inductive loads such as electromechanical fluid valves where the time lag between electrical control signal and mechanical actuation is important to system accuracy. Its tolerance to short circuits and partial short circuits coupled with internal diagnostics gives it advantages over other solutions currently available.

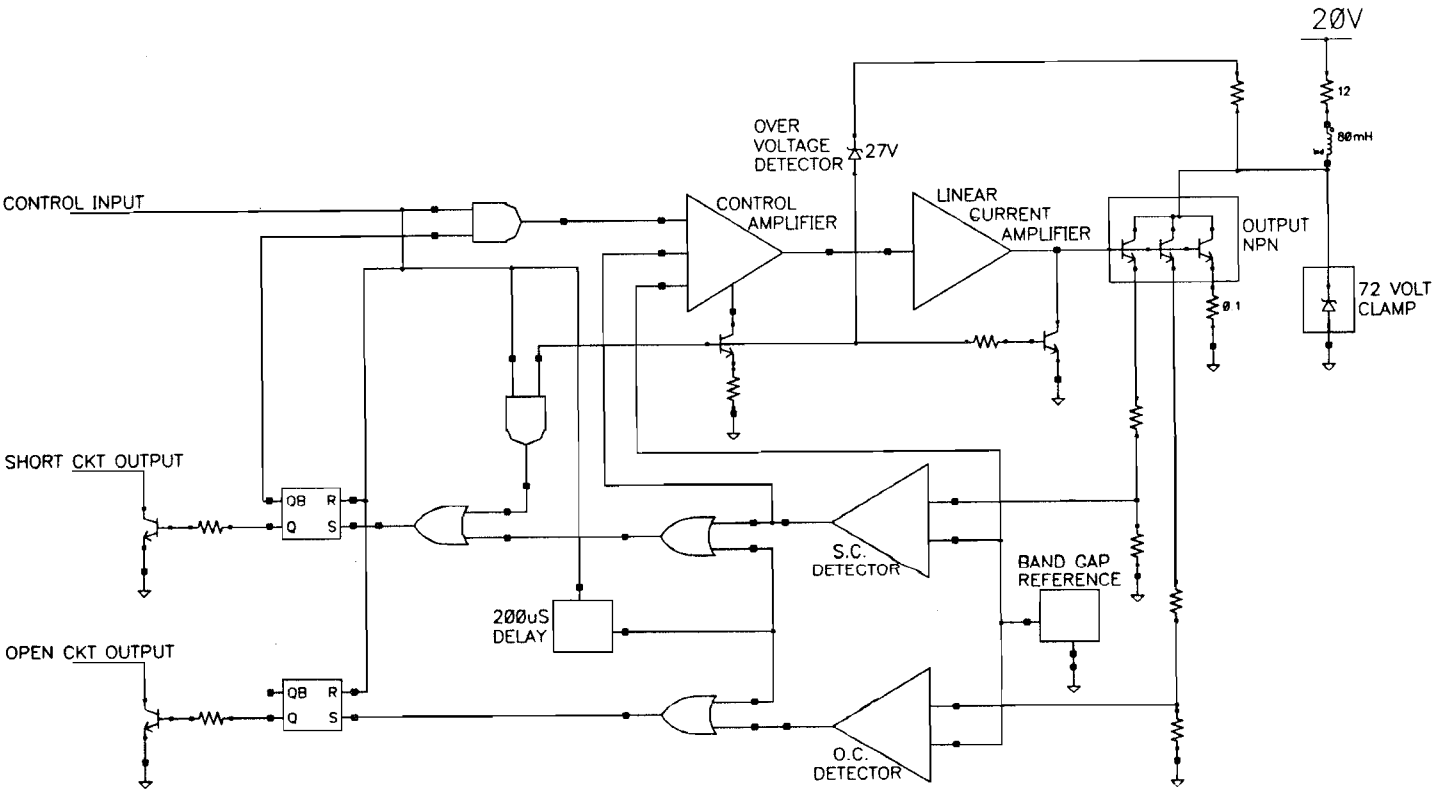

FIGURE 1, HIGH-VOLTAGE HIGH-CURRENT DRIVER

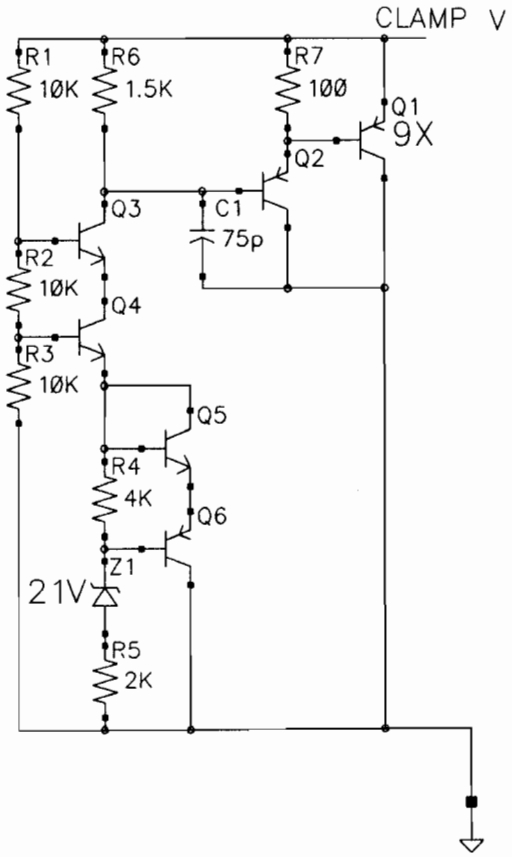

FIGURE 2, 72 VOLT CLAMP CIRCUIT

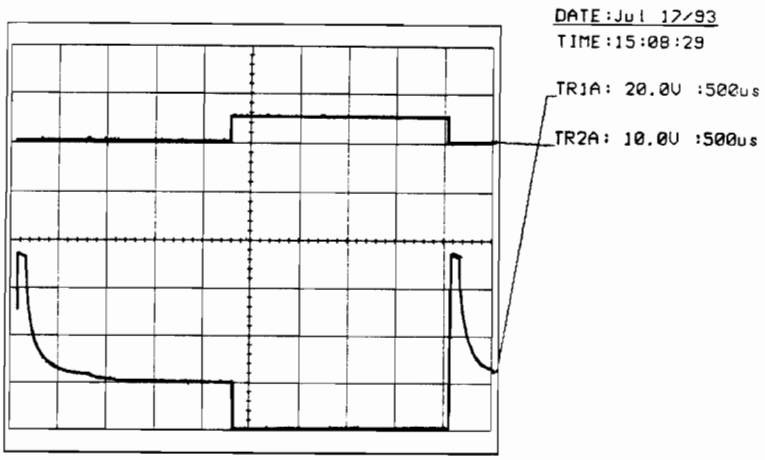

Figure 3. Input/Output Response
TR2A = Vinput, TR1A = Vout

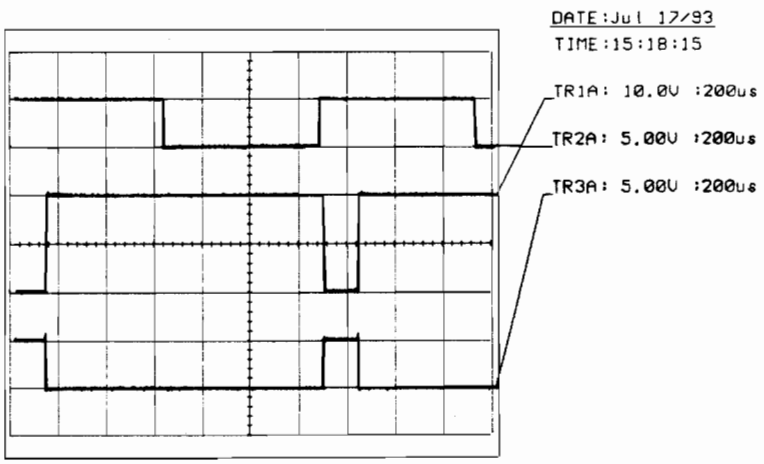

Figure 4. Input/Output Response
TR2A = Vinput, TR1A = Vout,
TR3A = Short Circuit Flag

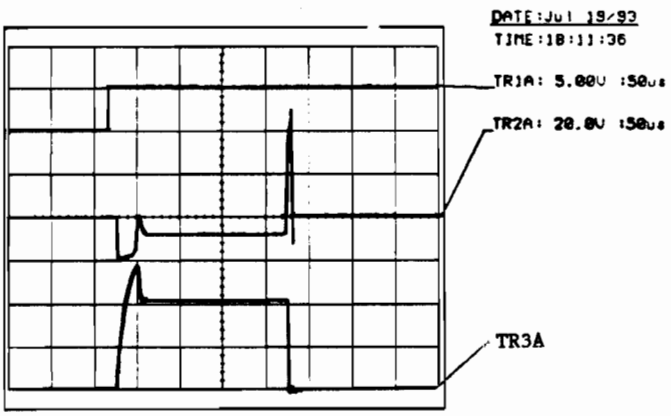

Figure 5. Input/Output Response
TR1A = Vinput, TR2A = Vout,
TR3A = Output Current (1A/div)

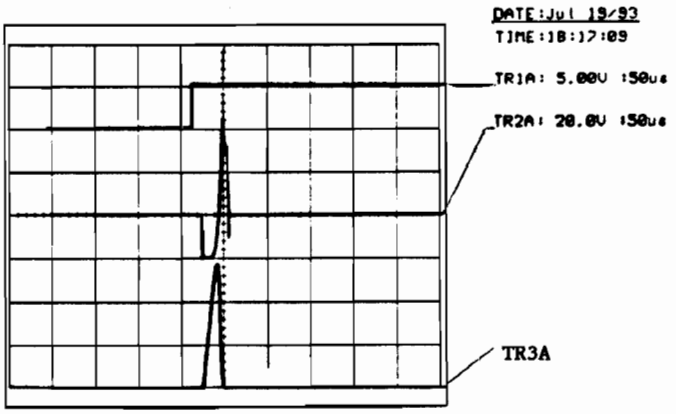

Figure 6. Input/Output Response
TR1A = Vinput, TR2A = V out,
TR3A = Output Current (1A/div)

A 3 V Thermostat Circuit

Sumer Can, Mark S. Richards, Robert A. Pease
National Semiconductor Corporation
2900 Semiconductor Drive, Santa Clara, CA 95052

Abstract

An Integrated Thermostat Circuit is described in this paper. It contains a bandgap reference, a temperature sensor and two hysteresis comparators. Two temperature trip points are generated by dividing down the bandgap voltage reference using three external resistors or an internal resistive network. The temperature sensor circuit generates a voltage proportional to temperature. Two comparators are designed for a 3 degree C internal hysteresis, and their outputs support CMOS/TTL logic levels. The circuit works with a supply voltage of 2.7V to 10V. It shows a trip point accuracy better than 2 degree C at room temperature and 3 degree C over the temperature ranging from -40 degree C to 125 degree C. The temperature sensor output voltage slope is 6.25mV per degree C. The thermostat is fabricated using a standard bipolar process, and packaged in an 8-pin SOT package.

Introduction

Temperature is a critical parameter that needs to be monitored in many electronic systems. Thermistors have long provided a cost-effective solution to electronic temperature sensing needs. A main drawback of thermistors is their non-linearity over temperature requiring some method of linearization. IC thermometers get around this problem by using on-board circuitry to create a linear output that is proportional to temperature in either degrees C or degrees F. The addition of comparators and an adjustable reference allow for an integrated thermostat to be created on one chip. While this straightforward solution has long been available, this paper describes an analog bipolar thermostat that operates off of 2.7V from -40 to +125 degrees C.

Analysis and Design

The sub-circuits of the proposed thermostat IC are analyzed and design equations are given in the following sections.

Start-up Circuit and Biasing *

A simplified schematic of the new start-up and bias circuit is shown in Figure 1. Its unique feature is that it generates supply dependent current through the main bias string consisting of an epi FET J1, Q1, Q2 and R1 until the bandgap reference is turned on. At that point, the startup circuit senses the bandgap and uses its voltage as its reference point for generating the system bias. The bias current flowing through Q4 is given by

$$I_4 = (V_{bg} - V_{be7}) / (R_2 + R_3) \qquad (1)$$

Once the circuit starts, Q2 turns off. The total base current of the current source PNP's flowing through Q7 are compensated by the current sense circuit formed by Q5, Q8, Q9 and Q10. (* *Patent Pending*).

Temperature Sensor and Bandgap Reference

The simplified schematic shown in Figure 2 outlines the low voltage sensor and the bandgap reference. In Figure 2, Q1, Q2 and Q3 form a feedback amplifier. Due to different current densities of the input differential pair Q1 and Q2, a PTAT (Proportional To Absolute Temperature) voltage is created across R1. The matched resistors R2 and R3 are used to multiply the PTAT voltage produced across R1, creating two PTAT voltages; Vp1 across R3, and Vp2 at the emitter of Q3. Voltages Vp1 and Vp2 are expressed in Equations 2 and 3 respectively;

$$V_{p1} = (R_3 / R_1)\, V_t \ln(N) \qquad (2)$$

$$V_{p2} = \left[\frac{R_1 + R_2 + R_3}{R_1} \right] V_t \ln(N) \qquad (3)$$

where N is the emitter ratio of Q1and Q2. Vt is the thermal voltage.

The bandgap voltage is produced by adding the base-emitter voltage of Q4 to Vp1 as given in Equation 4:

$$V_{bg} = V_{be4} + (R_3 / R_1)\, V_t \ln(N) \qquad (4)$$

After producing the initial PTAT voltage, Vp2 is dc level shifted by Q6 and further amplified by R4 and R5. An additional function of Q6 is that it adds approximately 2.2 mV/degree C thermal slope to the Vtmp versus Temperature characteristic. The final expression for Vtmp is given in Equation 5:

$$V_{tmp} = \left(1 + \frac{R_4}{R_5}\right)\left(\left[\frac{R_1 + R_2 + R_3}{R_1} \right] V_t \ln(N) - V_{be6} \right)$$

Eqn. (5)

However a slight nonlinearity is also introduced into the characteristic due to the curvature of the base emitter voltage of Q6 at very high and low values of temperature.
In Figure 2, I3, Q10, Q11, R6 set the collector current for Q4 , and I4, I5, Q8, Q9, R7, R8 set the collector current for Q6.

Hysteresis Comparator

The single-output comparator circuit is shown in Figure 3a. The circuit is formed by joining the outputs of two differential-input folded-cascode comparators in a cross-coupled PNP latch subcircuit. The differential current output of the latch is further converted into a single open collector output through the final stage.

0-7803-3516-3/96/$5.00 ©1996 IEEE

When the voltage Vtmp<Vhi and also Vtmp<Vlo, the transistors Q3, Q6, Q10 and Q7 are off. The transistors Q4, Q11, Q14 are on and Q12 is saturated. Also Q15, Q16, Q17 are off. As Q14 is on and Q17 is off, the comparator output is high. Now consider the case where Vtmp moved high and is between Vlo and Vhi .Since Vtmp > Vlo, Q9 is off, Q10 and Q7 are on. Therefore, the comparator output is still high. It stays high until Vtmp =Vhi. At that point, current through Q5 and Q6 are equal. Then,

$$I_7 = 2\,I_4 \tag{6}$$

which defines the switching point. As soon as Vtmp crosses Vhi , the comparator output switches to low. The output switches back to high when the Vtmp starts falling and it crosses Vlo. The hysteresis is given by:

$$V_{hys} = V_{hi} - V_{lo} \tag{7}$$

The half circuit of the dual output comparator scheme is shown in Figure 3b.The operation of this circuit is similar to the one in Figure 3a. Initially, let us consider the case where Vtmp < Vtset . At this state, Q5, Q3, Q7, Q9 are on and Q6, Q4, Q10, Q11 are off. The comparator output is high. At the switching point:

$$I_4 = 2\,I_3 \tag{8}$$

$$I_{R1} = I_{R2} = I_{R3} = 2I \tag{9}$$

Let $I_3 = I_x$. Then:

$$I_6 = 2I - I_x \tag{10}$$

$$I_5 = I_x \tag{11}$$

Thus,

$$I_4 = 2I - I_x = I_6 \tag{12}$$

From Equations (8) and (12):

$$I_6 = 2\,I_x \tag{13}$$

We can also write:

$$V_{TMP} - V_{Tset} = V_{BE6} - V_{BE5} = V_T \ln(I_6 / I_5) \tag{14}$$

Therefore, from Equations (11), (13) and (14):

$$V_{HYS} = V_{TMP} - V_{Tset} = V_T \ln(2) \tag{15}$$

Set Point Selection Circuit

The Set Point Selection Circuit is a resistive divider network as shown in Figure 4. Two Voltage levels Vhi and Vlo are generated by dividing down the bandgap voltage Vref .In Figure 4, Ra and Rb are almost binary weighted resistors with metal links L1 to L8 . The resistors Rc1 to Rc12 are equal valued resistors with taps Node 1 through Node 13. The resistor Rp is also an almost binary weighted resistor string with metal links L9 to L14 . Node 1 to 13 are the possible connection nodes for nodes i , j, and k. An optimum short or open circuit combination of links L1 to L14 and Node 1 to 13 connection to i , j, and k nodes is determined for a given Vhi and Vlo using a computer program. The voltages Vhi and Vlo are given as:

$$V_{HI} = \left(\frac{[R_B + (13 - k)R + (j - i)R + (R_P / [k - j]R)]}{R_T} \right) V_{REF} \tag{16}$$

$$V_{LO} = \{ [R_A + (13 - k) R] / R_T \} V_{REF} \tag{17}$$

where

$$R_A = R_O + R_1L_1 + R_2L_2 + R_3L_3 + R_4L_4 \tag{18}$$

$$R_B = R_O + R_5L_5 + R_6L_6 + R_7L_7 + R_8L_8 \tag{19}$$

$$R_P = R_9L_9 + \ldots\ldots + R_{14}\,L_{14} \tag{20}$$

$$R_T = R_A + R_B + (j-1)R + (13-k)R + \left(\frac{R_P}{(k - j)R} \right) \tag{21}$$

Experimental Results

The Thermostat circuit described in this article is fabricated using a standard Bipolar process.

The bandgap characteristic of a typical device is shown in Figure 5. The bandgap circuit is trimmed using Zener Zap Trimming techniques. The temperature coefficient is typically 20ppm/degree C.

The Temperature sensor characteristic of a typical device is shown in Fig. 6. It demonstrates a 6.25mV/degree C thermal slope with very little curvature at the high and low end of the operating temperature range.

The comparators show typically 18mV of hysteresis.

Conclusions

A low-voltage thermostat circuit is designed and fabricated using standard bipolar process. The circuit is packaged in an 8-pin SO package. The circuit works from a single 2.7V to 10V power supply. It has a trip point accuracy of 2 degree C at room and 3 degree C over the temperature range of -40 to 125 degree C. The temperature sense output voltage has a 6.25mV/degree C thermal slope. The temperature coefficient of the bandgap is typically 20ppm/degree C.

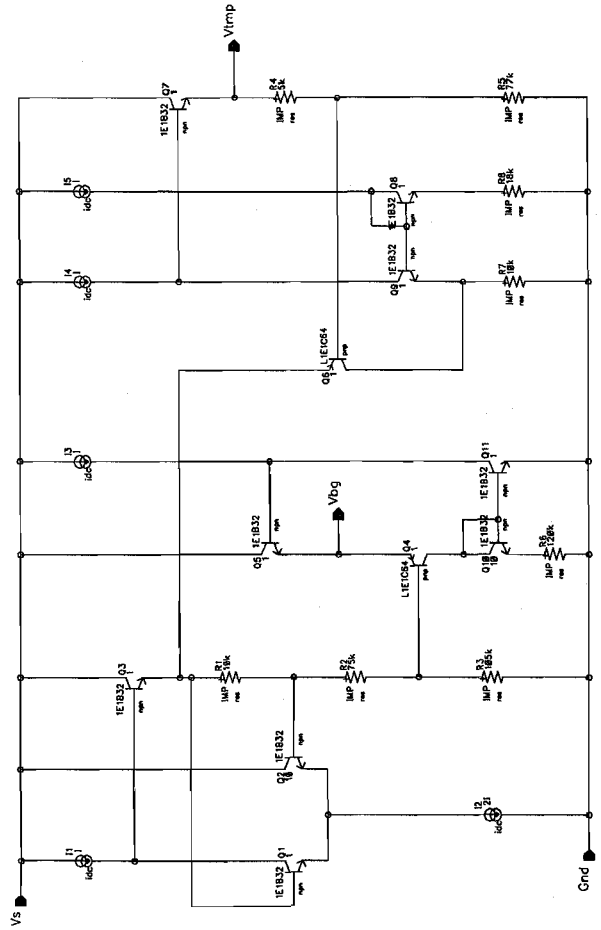

Figure 2

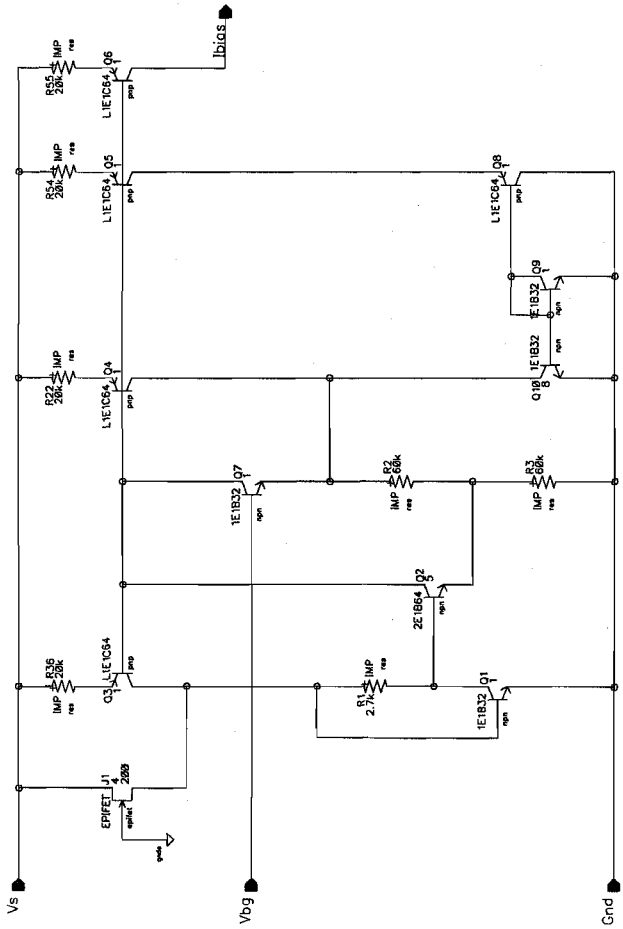

Figure 1

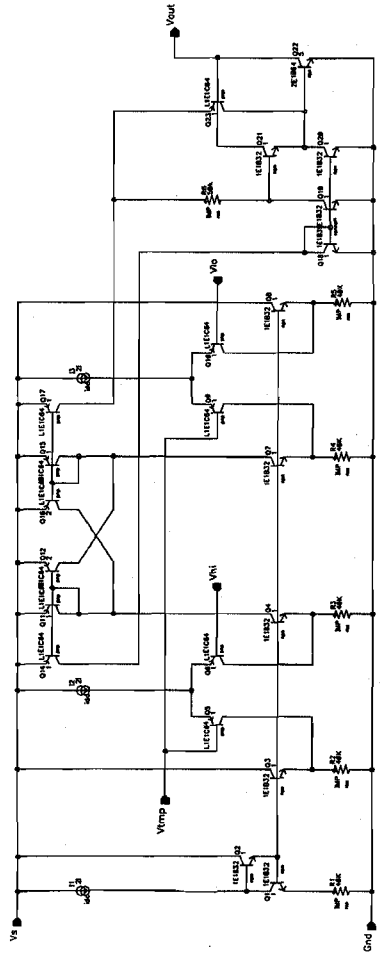

Figure 3b

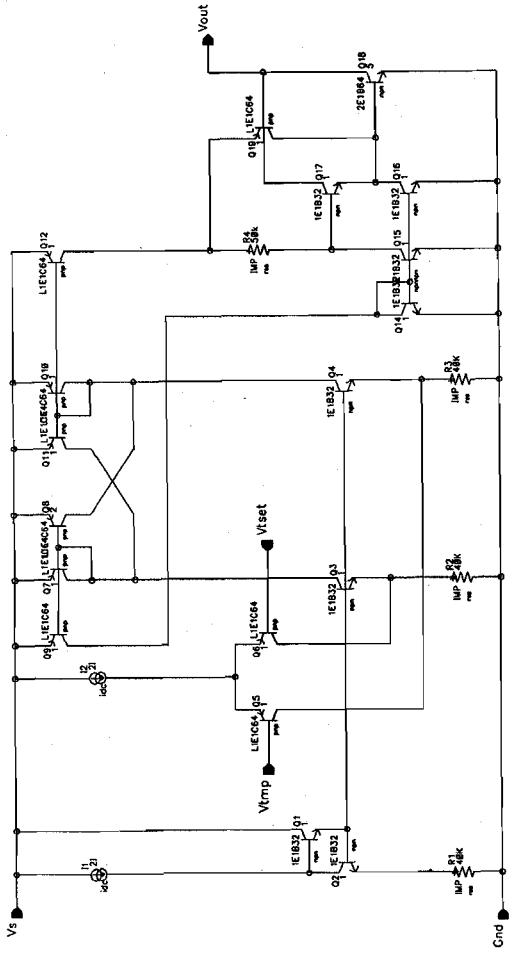

Figure 3a

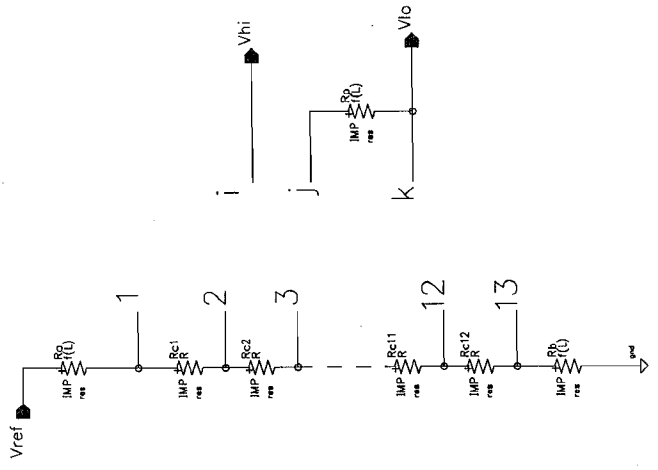

Figure 4

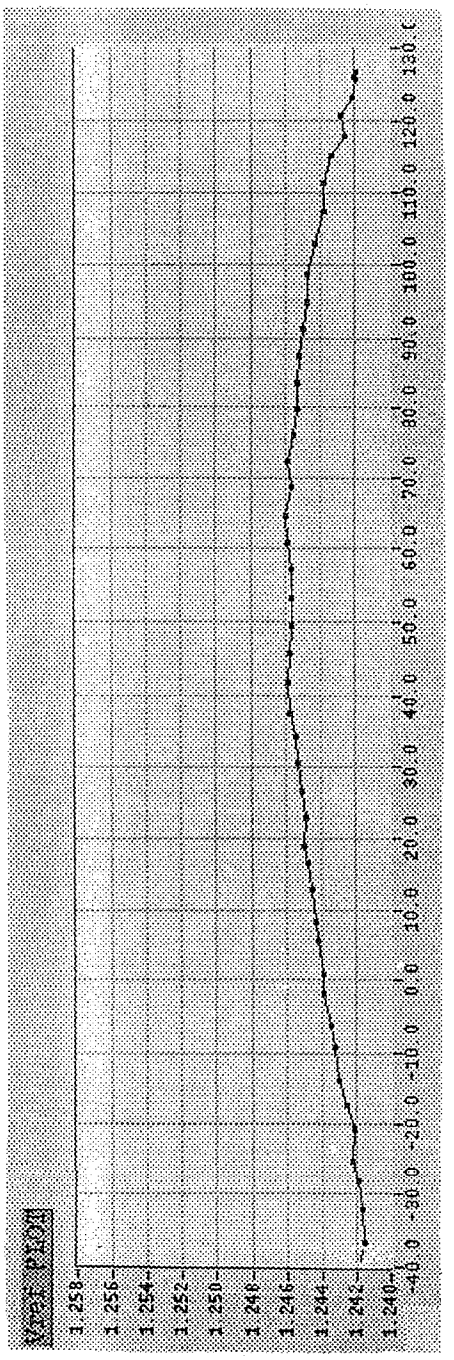

Figure 5

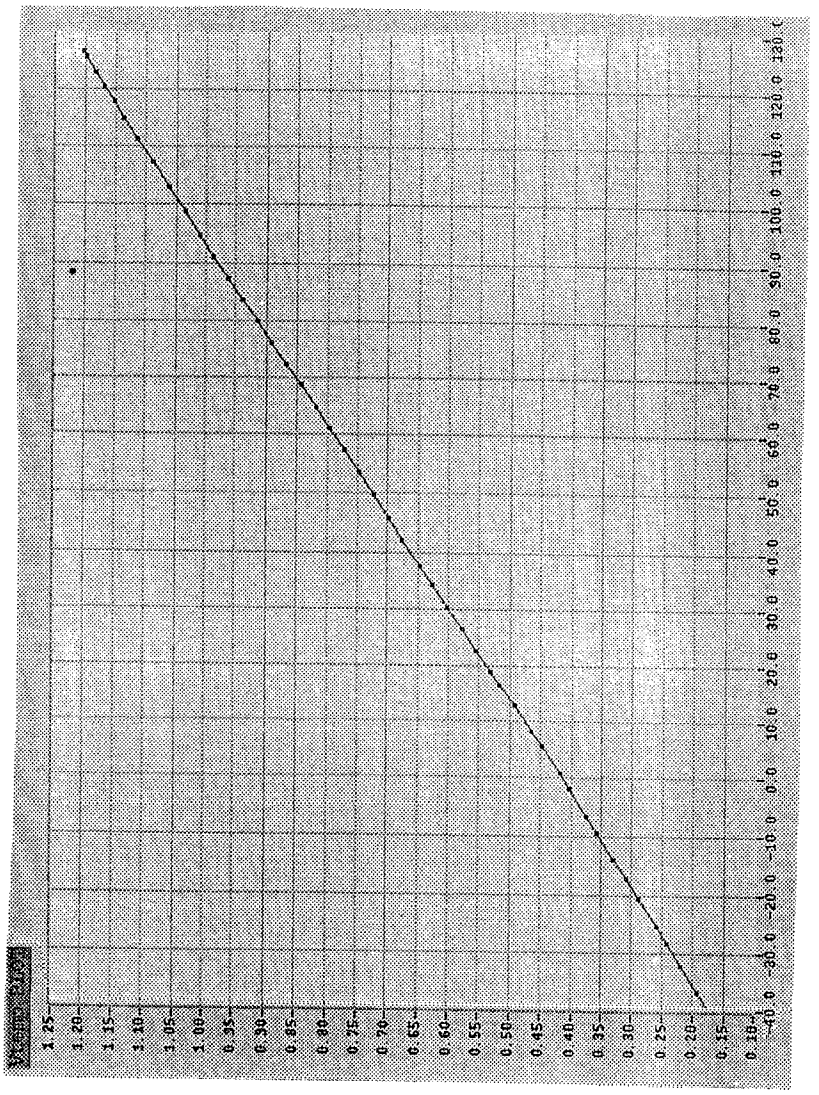

Figure 6

Correspondence

Comments on "Analog Layout Using ALAS!"

Robert A. Pease

The paper mentioned above[1] is quite amusing, yet horrifying. The authors talk of the advantages of laying out matched pairs of transistors, resistors, capacitors, etc., with interdigitated geometries and common centroid. Then they show six examples of "matched devices," and not one has a common centroid.

In Fig. 1, Q1's centroid is about 1/2 inch left of Q2's. Likewise, in Fig. 3, the transistor locations are mismatched by 0.07 inch, and in Fig. 4, by 0.1 inch. In Fig. 5, the capacitors are mismatched by 0.2 inch, and in Fig. 6, the resistors are mismatched by 1/4 inch. In Fig. 7 the 150/4 n-channels have their centroids offset by 0.13 inch and the 60/4 p-channels by 0.4 inch. The computer program "ALAS!" seems to have a remarkable consistency in making layouts where the centroids are never in common.

The authors seem to have the notion that arranging elements as ABAB provides common centroid. It does not. On the contrary, ABBA does provide common centroid. Likewise, AABBAABB as in Fig. 1 does not, but ABBAABBA does, or preferably, ABABBABA. Likewise, in an array of 12 transistors such as Fig. 4,

ABAB		ABBA
BABA	does not provide a common centroid but	BAAB does.
ABAB		ABBA

It is even better to use 16 transistors, to make a pair of transistors each with eight emitters or sources:

ABBA		ABAB
BAAB	or preferably	BABA
BAAB		ABAB
ABBA		BABA

There are many ways of making pairs or sets of transistors with common centroid, and all of them work better than the examples shown by the authors in their paper. The primary advantage of common-centroid layout is to reject linear gradients in threshold, oxide thickness, resistivity, temperature, etc. The layouts provided by this computer are not optimum in any respect. As an example of a fairly good layout which does provide common-centroid advantages, I recommend looking at my old U.S. patent, 3 995 304, which locates groups of 8, 4, 2, 1, and 2 emitters all at common centroids. But I did not use a computer to arrange those.

The authors claim that they obtain superior symmetry and complete parasitic balance. Of course, there are many kinds of symmetry, mirror-image as well as rotational. I prefer the former. If one looks at Fig. 6, the amount of interconnect metal over R1 is not at all the same as the metal over R2. If one uses ABBAABBA instead of ABABABAB, one gets further advantages as there is no need to put interconnect metal over the resistors. I have always found this a superior form of layout. If one lays out npn transistors with ABBA, one can merge the two transistors in the middle into a single tub, thus decreasing the output capacitance of the B devices. This provides worse balance, but may be advantageous in some cases.

In Fig. 7, the p-channel transistors tied to Vbias 1 have a mismatch of their centroids of about 0.4″, and the n-channel transistors tied to Vbias3 are spaced even worse, about 0.6″, although their matching is less critical to the circuit's performance. But when I looked at the n-channels tied to Vbias4, I found that the ALAS! computer had provided a match to within about 0.028″—not perfect, but a lot better than some of the other geometries that are claimed to be so good. It is amusing to contemplate why the computer provides various layouts that vary from bad to fair to good—but never perfect. Why accept layouts anything worse than perfect?

But I really struck pay-dirt when I kept on looking at devices that are "not critical"—maybe they were not inspected very carefully. Look at the "50/4" n-channel device that is connected to Vbias4 and is used to provide a current-source to the input stage. The schematic says it should be 50/4. Actually, it is 25/4, because half of the sources fell off and were never connected.

The authors claimed that the layout of the amplifier in Fig. 7 was completed in less than 20 minutes. I guess they thought they did not have to check their computer's results very carefully and sent in this paper without spotting that the source never got connected. A computer that makes a layout like that in 20 minutes is not fast—it is just half-fast. I prefer to design my circuits by hand, and check them with my eyes and my brain.

The layout of Fig. 7 also suffers from a gain reversal, as the labelling of the output polarity is reversed, compared to the schematic diagram. Likewise in Figs. 1 and 3, "Gate 2" serves "Drain 1," and "Gate 1" serves "Drain 2." This is a sure way to get an unintended gain reversal! The authors' (misplaced) trust and overconfidence in their computers, and their apparent disdain for checking their results, and the resulting layout errors, will get them (and keep them) in deep trouble!

As for this computer program—what can I say but—ALAS!

Manuscript received July 10, 1996

The author is with National Semiconductor Corporation, Santa Clara, CA 95052 USA.

Publisher Item Identifier S 0018-9200(96)07343-X.

[1] J. D. Bruce, H. W. Li, and R. J. Baker, *IEEE J. Solid-State Circuits*, vol. 31, no. 2, pp. 271–274, Feb. 1996.

Authors' Reply

J. D. Bruce, H. W. Li, and R. J. Baker

We would like to thank Bob Pease for his colorful comments on our paper and respond to his concerns. First of all, we would like to clarify the issues regarding common-centroid versus interdigitated layout. We assumed the reader would know that an interdigitated pair of

Manuscript received July 10, 1996.

J. D. Bruce is with Micron Technology, Boise, ID 83706-0006 USA.

H. W. Li is with the Department of Electrical Engineering, University of Idaho, Moscow, ID 83840-1023 USA.

R. J. Baker is with College of Engineering, Engineering Education in Boise, Boise, ID 83712 USA.

Publisher Item Identifier S 0018-9200(96)07342-8.

devices is *not* common centroid. We presented Fig. 1 as the building block for generating common centroid matched pairs. Only the figures in which there are more than one row should be construed as common centroid. Figs. 1 and 6 are examples of interdigitated layouts only. ALAS! produces both interdigitated and common-centroid layout. We failed to make that distinction and apologize for any confusion regarding this issue.

Mr. Pease's analysis of our common centroid examples (Figs. 3, 4, and 5) is correct—insofar as his suggestion that the *reproductions* of our figures are not common centroid. However, what Mr. Pease failed to understand was that through the use of modern computer technology, our figures were screen captured, scaled, cut, pasted, resized and imported (by us and the editors of JSSC) to ensure compatibility with IEEE's journal format. This explains why Mr. Pease's incredibly precise ruler measurements produced such varied results. We would like to ensure the readers of JSSC that the layouts generated by ALAS! are truly common centroid. The reader is also invited to refer to [1], for the definition of common centroid used in the development of ALAS!.

Mr. Pease also cited his patent on a common centroid scheme, patent number 3 995 304 filed January 10, 1972. Upon closer inspection, Mr. Pease's configuration would not have sufficed in our situation because it requires more metal routing as well as different amounts of metal routing for each device. In developing ALAS!, we assigned a high priority to symmetrical routing while minimizing the metal used for each device.

Lastly, Mr. Pease is correct in pointing out the missing source connections and the mislabeled polarities in Fig. 7. However, both mistakes were a result of human error and not the result of the computer program. The original layout generated by ALAS! was verified to be correct; we failed to select the entire layout in preparing the graphic for the article. The mislabeled polarities were a typo. Perhaps, in the future, we will write a program to verify that our computer graphics and layouts are identical....

References

[1] M. Ismail and T. Fiez, *Analog VLSI: Signal and Information Processing*. New York: McGraw-Hill, 1994.

Correction to "Large Bandwidth BiCMOS Operational Amplifiers for SC Video Applications"

Gerhard Nebel, Ulrich Kleine, and Hans-Jörg Pfleiderer

In the above paper[1], Fig. 9(a) was mistakenly reproduced twice instead of Fig. 9(a) and (b). The correct figure appears below.

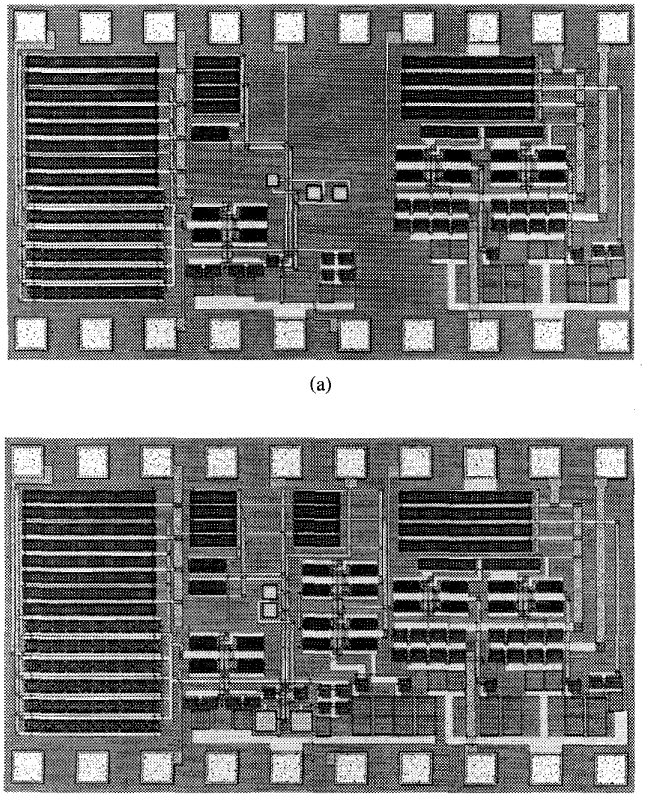

(a)

(b)

Fig. 9. Photomicrographs of (a) the single-ended and (b) the differential amplifier.

Manuscript received July 1, 1996.

G. Nebel is with Siemens AG, Munich, Germany and with the Department of Electronics and Microelectronics, University of Ulm, Ulm, Germany.

U. Kleine is with Siemens AG, Munich, Germany.

H.-J. Pfleiderer is with the Department of Electronics and Microelectronics, University of Ulm, Ulm, Germany.

Publisher Item Identifier S 0018-9200(96)07180-6.

[1] G. Nebel, U. Kleine, and H.-J. Pfleiderer, *IEEE J. Solid-State Circuits*, vol. 31, no. 6, pp. 828–834, June 1996.

29. A Tale of Voltage-to-Frequency Converters

Ancient History

Once upon a time, there weren't any voltage-to-frequency converters (V/F converters) or voltage-controlled oscillators (VCOs). I couldn't tell you exactly when that was, but back in the 1950s and 1960s, very few people ever heard about an oscillator whose frequency could be controlled by a voltage. In those days, when you wanted to change an oscillator's frequency, you changed a pot or a resistor or a capacitor, or maybe an inductor.

I checked up on this, because I spent a couple hours searching in the old M.I.T. *Radiation Lab Series*, published in 1949. There were no oscillators or multivibrators whose frequency could be controlled by a voltage—no VCOs, as far as you could learn by looking at the "Bible" of that time. (See also the last section, A Final Note.) It's true that FM radio transmitters used frequency modulated oscillators back as early as the 1930s, and these were modulated by voltages, but they only covered a relatively narrow frequency; when I refer to a VCO, I am talking about oscillators whose frequency could be controlled over a range of 10:1 or 100:1 or 1000:1 or more. In general, this type of oscillator is expected to have a pulsed or square-wave output, *not* a sine wave.

Less Ancient History

In 1961, when I graduated from M.I.T. and joined up with George A. Philbrick Researches (127–129 Clarendon Street, Boston 16, Massachusetts), I joined a company that made operational amplifiers, analog multipliers, integrators, and all sorts of analog computing modules. And just about everything was made with vacuum tubes. There were applications notes and applications manuals to tell you how to apply operational amplifiers (in those days we would never say "op amp"). And there was the big old Palimpsest, a sort of collection of old stories on things you could do with operational amplifiers and analog computing equipment. But, there were no digital voltmeters, and no voltage-to-frequency converters.

About 1963, I became aware of a high-performance operational amplifier—the 6043 project. This chopper-stabilized amplifier had been designed by Bruce Seddon, one of our senior engineers and one of our vice-presidents. The customer for this amplifier was DYMEC, a subsidiary of Hewlett-Packard, and this amplifier was to be the front end of a voltage-to-frequency converter instrument. The amplifiers were functioning pretty well, and they were meeting just about every specification, but they had a problem with undesired and unpleasant little offset shifts and jumps.

As the project went on, more and more engineers were looking over Bruce's shoulder, trying to help him solve this problem of the jumpy offset. Some people

suspected that the amplifier stages might be rectifying out signals that were caused by—?—perhaps—the transmitter of a local taxicab company? I do not know if it was ever resolved what was the cause of this drift or wobble—it was hard to resolve, just a few dozen microvolts—but, I have the impression that the amplifier project was not a success, and we never went into full production. If there was any circuitry of the actual V/F converter, that was never discussed at Philbrick—that was proprietary within DYMEC, and it was only our job to make an operational amplifier with low offset.

Of course, this amplifier had several vacuum tubes and a chopper. The first amplifier tubes, 6CW4 nuvistors, ran with their 6 V heaters connected in series, so as to run on a 12 V DC bus and avoid the noises inherent in using 60 cps heaters. Then there were two more gain stages, DC-coupled. There was a chopper, Airpax 172 or Bristol or similar, and a chopper amplifier based on a 12AX7. The whole circuit was similar to the Philbrick "USA-3" amplifier; it did, however, run on ± 150 V instead of ± 300 V. Dan McKenna, the senior technician who worked on the project, said he always suspected that it was the blame of the heaters that were connected in series, because the data sheet on the 6CW4s said, "do not connect heaters in series." But I realized later, connecting two heaters in series was surely okay; putting 10 or 20 tubes' heaters in series, as in a TV set, was the procedure that was prohibited. So, even though Dan grouched about it, this series stacking of two heaters was probably quite wise, not risky, as it would force the tubes to run at about the same number of watts in each heater.

Anyhow, that is ancient history, a design with vacuum tubes. Even then, even though we were trying to design part of a V/F Converter, we didn't have one in our lab. We didn't have a digital voltmeter (DVM)—we had the old Fluke 805 differential voltmeters. Now, these meters have many elements of accuracy that good DVMs have these days—good resolution and stability—but they were big and heavy and slow. If you wanted to read 5.032 V, for example, you could do it, but it was a tedious deal. You had to turn the first knob to 5, and then the next one to 0, and the next one to 3, and then turn up the gain and look at the analog meter to see if the residue looked like a "2." That was how you learned the voltage was 5.032 V. If you have ever spent a few hours twisting the knobs of one of those old Fluke meters, you may remember it with nostalgia, but you must admit, it was awfully boring.

When the first DVMs came, from HP and from Non-Linear Systems (NLS), they were slow and (in the case of the NLS) clunky and noisy, and they did not have excellent accuracy or features compared to the old Fluke differential meters. But they sure were faster and easier.

And there was another way to do a DVM—you could buy a voltage-to-frequency converter from DYMEC and feed its output frequency into an "event counter" (an EPUT meter—Events Per Unit Time—from Beckman Berkeley), and its neon-discharge display tubes would glow and tell you what the frequency was, in terms of pulses per unit of time, and that was supposed to be linearly proportional to the input voltage.

This new DYMEC V/F converter had several solid-state circuits and tricky pulse generators. To this day I do not know how those proprietary, secret pulse circuits were supposed to work. It had a pulse generator based on some special pulse transformers and blocking-oscillator circuits. The instrument had a pretty good temperature coefficient (TC), *but*, only because it had a little oven to hold all the transistorized circuits (which *really* had a rotten TC) at a constant temperature of +65 °C. Consequently, when you turned it on, you had to wait at least half an hour before the accuracy would finally settle out to its ultimate value—waiting for the oven to settle its temperature. It could handle a full-scale voltage of + and −1.0 V. It works,

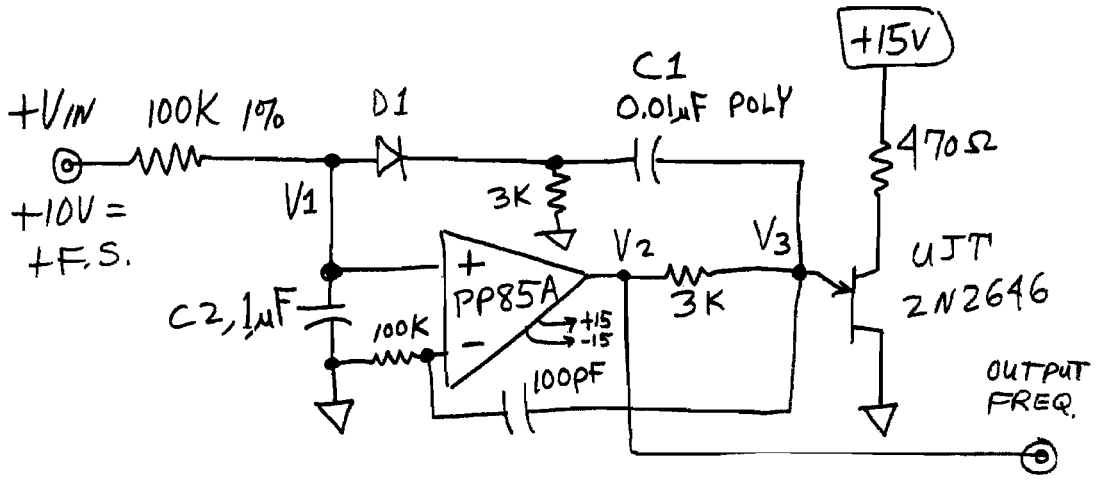

Figure 29-1. The Swedish voltage-to-frequency converter circuit (drawn from memory, approximate schematic).

pretty well. It was claimed, originally, to have better than 0.02% of linearity. I measured it once, and it had some nonlinear errors around 0.024%—not bad, but apparently something had drifted slightly out of spec. It cost $1600, back when $1600 would buy you a Volkswagen. I still have one of the DYMEC V/F Converters, Model DY-2211B, and the book on it.

Now let's move up to about 1967. We engineers at Philbrick were working mostly on solid-state operational amplifiers—amplifiers made of 6 or 8 or 10 discrete transistors. The integrated circuit amplifiers were arriving, but most of them were pretty crude in performance and features.

One day Bill Bernardi, one of the senior applications engineers, told me that a customer in Sweden had made a linear voltage-to-frequency converter using a PP85A (one of our standard operational amplifiers) and a UniJunction Transistor (UJT). And the nonlinearity was, he said, about 0.1%. When I heard this, I got *very* curious, because everybody knows that UJTs are the crudest, dumbest, most imprecise oscillator you can find. Just about every student learned that a UJT looks very cute because it can oscillate with a minimum amount of external parts, *but* it's an awfully junky circuit. You could gold-plate the sow's ear, and it was still a junky circuit. So when I heard that a UJT was involved with a V/F converter of very good linearity, I was impressed, but I was suspicious, and I looked into what they were doing. I didn't know anything about V/F converters, but I was curious. I found that the PP85A was used as a comparator, and the UJT was mostly used to provide some "negative resistance" or positive feedback, to make a pulse whose amplitude or width are not critical and thus did not hurt the accuracy of the V/F converter. Ah, but how is that? How is there a V/F converter that uses one simple comparator and a crude UJT pulser, and no other obvious precision components, and yet provides a 0.1% linear V/F converter?

As near as I can recall and reconstruct, the circuit was basically that in Figure 29-1. The principle of operation is that when the current from V_{in} causes the main integrating capacitor C2 to rise up to 0 V, and the op amp starts to swing its output positive, it triggers the UJT, which then puts out a crude pulse which kicks the minus input of the comparator; and the output *also* kicks a certain amount of charge through a charge-dispensing capacitor, C1, back to the integrating capacitor, to reset it. This amount of charge must be constant and invariant of anything, especially invariant of the repetition rate. If you can get that, you get excellent linearity. Apparently the Swedish engineers had stumbled onto this crude but functional circuit.

Now that I understood the principles, I figured out that there was room for a good

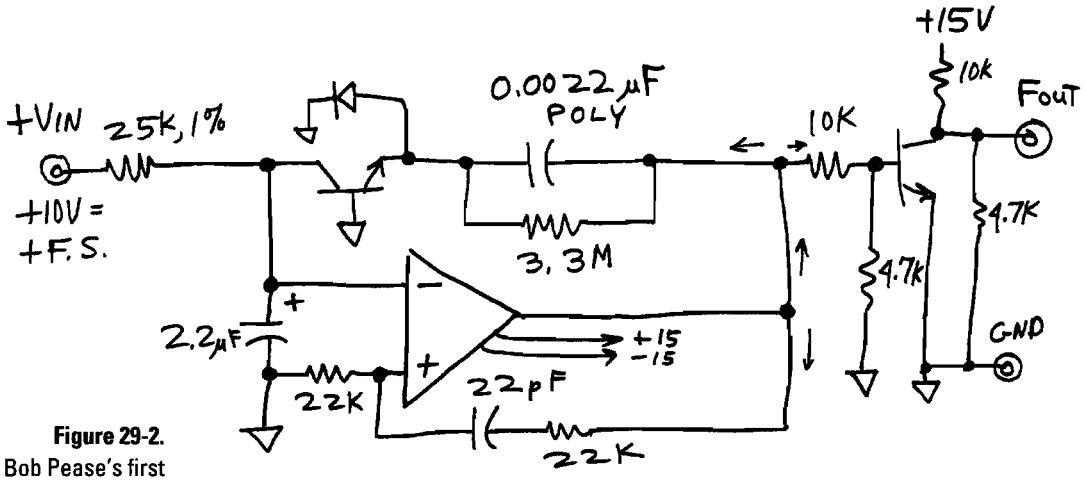

Figure 29-2. Bob Pease's first attempt at designing a voltage-to-frequency converter (A1-Philbrick T82 AH ≅ Amelco 805BE. Full scale is 10 kHz. Good linearity, poor PSRR and TC).

bit of improvement. I started fooling around with some breadboards. I got the full-scale frequency up to 10 kHz (the Swedish circuit worked well up to just 1 kHz, which is not nearly as useful as 10 kHz), and got the nonlinearity down to 0.03%. And I invented a scheme so the operational amplifier's output could be capacitively coupled back to its positive input, causing enough regeneration or positive feedback, that the UJT was no longer needed. I used an Amelco 805BE integrated-circuit operational amplifier as the comparator. Now, the whole thing would fit into a 1.5 in. square package, just 0.5 in. high—a small epoxy-potted module that was rather smaller than the PP85A amplifier and associated parts as shown in Figure 29-2. We built up a prototype and we tested it, and it worked pretty well. We potted it in our usual hard black epoxy and shipped it to a customer in Connecticut—a customer of Larry Plante, who was our Sales Engineer for that region. Also, I sent in a patent application to our patent attorneys. I forget exactly who it was—was it Mr. X in New York, or Mr. Y in Waltham? No matter.

That must have been a busy year, because by the time I got off the other hot projects I was set to work on, for a number of high-priority customers, I realized I had not heard anything from this customer in Connecticut. I got in touch with Larry Plante. All he knew was, the customer didn't like it. Worse yet, a whole year had elapsed since I had sent the part in interstate commerce, and the patent attorney had done nothing, so the patent application was now worthless. I was quite cross, and I read the riot act to these attorneys. Then I set in at the work-bench with a vengeance.

I realized the old circuit had depended on the power supply stability for its gain factor, so it had no power supply rejection, or, to be more nearly correct, a PSRR of about 0 dB. I added a zener in a bridge, to give a PSRR of perhaps 45 dB. (Note, that was the first time I had ever seen that zener bridge (see Figure 29-3)—was I one of the earliest inventors of that circuit? It is a *neat* and useful circuit.) I added improved features around the amplifier, to improve the start-up and the accuracy. I replaced the (sole-sourced) 805BE with the more popular and inexpensive LM301A. Refer to Figure 29-3; a description of how it works is provided nearby. I gave it to my technician, Dick Robie, to oven it and graph out the temperature coefficient (TC) from the temperature data. That night, on the way out the door, I asked Dick what had he seen for the TC. He replied, about zero. I asked, "Whatd'ya mean, zero? You mean, down near 100 parts per million per degrees C?" He replied, "Oh, much better than 100—less than 10 ppm per °C." I was shocked. How could it be that good? The next day, I figured out the fortuitous situation: of course,

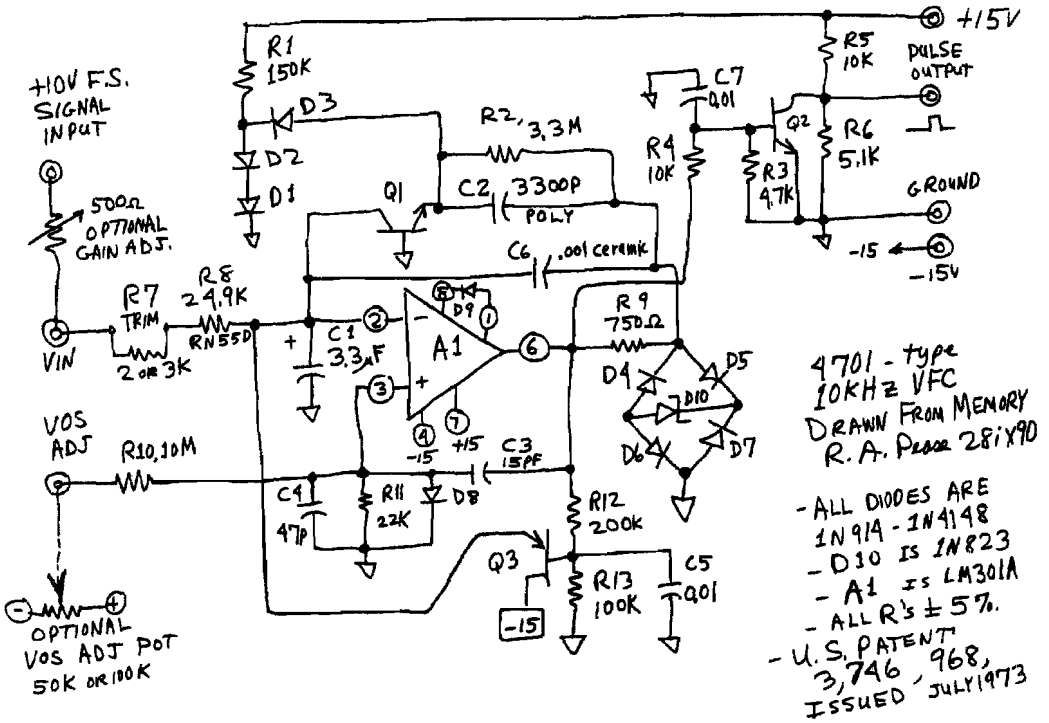

Figure 29-3. The legendary 4701 voltage to frequency converter, designed and drawn by Bob Pease. (All diodes are in 914-1N4148; D10 is IN823; A1 is LM301A; all R's ±5%.)

for the "charge-dispensing capacitor, C2," I had used one of the best capacitors in the house, the most precise and stable ones, which were polystyrenes with a TC of −110 ppm/°C. This TC was just about cancelled out by the TC of the entire group of diodes in the rest of the circuit. Namely, the driven end of the capacitor moves about 12.4 V p-p, plus the V_f of four diodes. These four diodes run rich at about 6 mA, and have a TC of about −2.0 mV/°C. The p-p voltage of these four diodes is approximately cancelled by that of the other four V_{be}'s at the other end of the poly capacitor, but those diodes run at about 0.1 mA, and their TC is about −2.3 mV/ degree. The difference of these is about 4 × 0.3 mV/°C, or 1.2 mV/°C, which is about big enough to cancel out the −110 ppm/°C of the capacitor. Now, there were several things I could have done to fix it if the TCs had not come out even—I could have used 3 diodes, or 5, or 4½ or 3½, but, if 4 was the right answer, I'd go with it.

I got my boss, Dave Ludwig, to approve the layout of a printed-circuit board, and I think Wayne Norwood was the guy who laid it out. We built up a few dozen and evaluated them thoroughly. I wrote up a draft of a data sheet, and I negotiated with Skip Osgood and Bill Bernardi to get it printed properly. I got some test equipment going, and a silkscreen, and we called it the 4701. We were in the voltage-to-frequency converter business.

I don't recall exactly how we got these V/F converters to be so popular. I can't recall how we found so many customers, or how we got out publicity. I asked Frank Goodenough—now the senior editor for analog circuits at *Electronic Design* magazine—I knew he had been involved. He recalled how he had gotten involved: He had looked at some of the characteristics of this 4701, and he suspected that a good V/F converter might be useful at the Foxboro company, the big process-control experts. Indeed, he did find some interest there. They were very interested—but they never bought very many, because Foxboro was very concerned about buying only parts available from multiple sources.

The 4701 became popular, with many customers at many companies. It became profitable. It became good business. If I had all the time and space in the world, I would tell you how the 4701 led to the 4702 (a 10 kHz frequency-to-voltage converter, using the same charge-dispensing principles) and the 4703 (100 kHz full scale V/F converter) and the 4705 and 4707 (1 MHz and 5 MHz V/F converters). Also, the 4709, 4711, 4721, and 4715 V/F converters, and 4704, 4706, 4708, 4710, and 4722 frequency-to-voltage converters. Some of these had a moderate TC of 150 or 44 ppm/°C, but some were well-trimmed and guaranteed to 6 ppm/°C—as good as the best DVMs of the day.

But it all started with that crummy little 4701—and the principle that one cheap operational amplifier driving a little charge-dispenser could make a very linear data converter. This came from an understanding that you could build an oscillator with a UJT and an operational amplifier to help improve the linearity, and *then* throw out the UJT! I was able to do that because I was receptive to the concepts that would make a good V/F Converter, even though I had never seen a V/F Converter! I was able to make accurate measurements, to throw in precision components—zener references, capacitors, resistors—and to get inexpensive IC amplifiers, and to optimize the whole little system. What is the underlying principle?

I like to think of the words of Branch Rickey, manager of the St. Louis Cardinals in the 1920s (and later manager of the Brooklyn Dodgers). One of his favorite sayings was "Good luck is the residue of design." When he referred to his baseball teams, he meant that his players were able to handle any play, field the ball, throw to the right base, etc., no matter how surprising the bounce of the ball might be. If his players had learned the fundamentals, if they were prepared for any situation, they could take advantage of game situations and *win*. Other people might say his players were just lucky, but he knew darned well that he had trained them and drilled them so they would instinctively know to do the right thing.

I, too, was in a situation where I was prepared to take advantage of the opportunity, and I didn't drop the ball when it came my way. (Well, I fumbled it for a while but then I got it right.) We built up that V/F Converter business to about one-tenth of all the business at Teledyne Philbrick, and if you look at the schematic of that little V/F converter, you can tell it was pretty profitable when we sold it for $59. But as I mentioned in another chapter, when a guy had spent his $59, he really got a lot of satisfaction. We had some very loyal customers, because the performance per dollar was right, even if the parts list would have looked pretty skimpy, if the epoxy were not so opaque and obscuring. As with other circuits like the P2, the list of parts looked like not much of a big deal, but the way you put them together was what made the value to the customer. To some extent, this is like the integrated circuit business, where the price of an IC often is related to the value it provides to the customer and is not related at all to the cost of manufacturing.

The V/F Converter business eventually became just popular enough that other analog-circuit manufacturers decided to get into the business. Several competitors such as Dynamics Measurements Corp., Intech, and a couple others started making units that competed. But these companies were all followers—none of them was making or selling V/F converter modules until the 4701 came along and popularized the concept. Finally Raytheon started making an integrated-circuit V/F converter—the RC4151. But it did not use charge-dispensing techniques—it used a timer, similar to the 555. It was inexpensive but did not have any guaranteed TC, and only poor linearity (0.15%).

In 1976 I left Teledyne Philbrick and went to work for National Semiconductor. Bob Dobkin put me to work designing a V/F converter integrated circuit that would have all the advantages of previous designs but avoid the disadvantages. I came up

with a design that was nothing like a 4701—nothing like any of the Philbrick V/F converters. This LM331 would run on any power supply from 4 to 40 V (whereas the Philbrick ones needed ± 15 V, and ± 12 to ± 18 V was about as wide as they would accept). It's been pretty popular, and to this day, 13 years after it was introduced, the LM331 is still getting designed in to new circuits. People tell me they like the way it is versatile and offers good precision.

At Philbrick, after I had refined the circuits of the 4701, I realized that this went far beyond the original patent application I had filed. So, I wrote up a new application to take into account the new schemes, and we filed that one in about 1970, and eventually it issued, in July of 1973. The number of that U.S. Patent is 3,746,968. After 17 years, that patent expired in July 1990, and consequently I don't feel bad at all about talking about the 4701. After all, one of the functions of the patent is to teach the reader how to do something. The patentee holds a monopoly right for 17 years, but 17 years is the limit.

Vignettes—Little Stories...

In those early days of the Philbrick 4701 family, our arch-rival was Analog Devices. We were kind of nervous about how AD would bring out competing modules. Would they steal our circuits? Would they infringe on our patents? Year after year we waited for the shoe to drop. Finally, after just about every other competitor was selling V/F converters and F/V converters, Analog Devices brought out its modules. They did *not* infringe, they did *not* steal our circuits. We also thought they did not have very impressive performance or features, as they had designed their circuits to do things the hard way. But mostly, they were late in the marketplace. Later, we found out why:

At Analog Devices, the engineers always designed what the marketing managers told them to. This was rather different from Philbrick, where the engineers often designed things that marketing people could barely comprehend—but on a good day, these unrequested products made a lot of good business, and a lot of friends, and good profits, too. But the marketing people at AD had looked at the marketplace, and they decided there was "no market for V/F or F/V converters." Year after year, they decreed that, and V/F converters were studied, but no products were introduced. Finally, several of the AD salesmen presented evidence that even though the marketing people could prove there was "no market" for V/F converters, there really were *customers* for V/F converters, and they twisted the marketers' arms until they finally agreed to get into that business.

That reminds me of a funny story. When we had just three V/F converter products at Philbrick, and I was designing several more, one of the marketing managers at Philbrick—I recall that it was Maurice Klapfish, who had recently arrived from Analog Devices—decided to commission a market survey about V/F and F/V converters. The engineers were summoned to hear the report and conclusions of the survey. This fellow had asked many people, "Is there any market for good V/F Converters?" In each case, the answer was, "no, no significant market." We asked this fellow, whom had he asked? He said that he asked all the manufacturers who now made V/F converter instruments. Well, what would you expect *them* to say? Of course they would not want to encourage competitors!

At this point, I stood up and asked if he realized that Philbrick's entire building—the complete facilities and every concrete block in every wall—had been paid for by the profits on a product (the Philbrick P2 amplifier, see Chapter 9) that nobody had ever asked for, and that marketing never said they wanted, and in fact marketing

had never even planned or specified or run a market survey on it—it was just thrust into the marketplace by engineers—did he realize that? The market-survey fellow said, no, he did not believe that. Well, of course I walked out of the meeting.

After a while, our management wisely decided that a few new V/F converter products might just be okay, if the customers continued to show good acceptance of these V/F converter products. I won't say that marketing people never approved of the designs, and features, and specs. Let's just say that I put on my marketing hat, and decided what kind of features a customer might like to buy, in a marketplace where nothing of this sort had ever been seen before.

Frank Goodenough reminded me that at one sales reps' meeting in Chicago, where all the 50-odd reps were gathered to hear the Philbrick marketing experts explain what good new things we had to sell, Frank had stood up to tell the reps about the new V/F converters and F/V converters. For some reason, the Marketing Manager, John Benson, had not gotten his full agreement on how the V/F converter line was to be promoted, so just when Frank (who, if you have ever met him, is a *very* enthusiastic guy) was beginning to say glowing things about the V/F converters, John told him to "shut up and sit down and not talk about those things." Well, if you want to make sure a person pays a lot of attention to what a speaker is saying, you just tell the speaker to "shut up and sit down." Frank did what he was told, but he sure got the reps' interest that day!

In 1988, I was interviewed for a biographical page in *EDN* magazine. When the biographical material was shown to me at the time it was ready for publication, it stated that I had designed a V/F converter using vacuum tubes, and thus I had proved that one could have made good voltage-to-frequency converters back in the 1950s and 1940s. However, even though I had been threatening for 17 years to design and build such a machine, I had never actually done it. I could either tell the editors to drop out that phrase, or I would have to actually design such a machine. I was a little rusty in my art of designing with tubes, but I took an old hi-fi amplifier and rewired a group of tubes—some 6SN7s and 6SL7s and an old OB2 reference tube, and with a minimum of tweaking, I got a good V/F converter running with the same basic design as the 4701. The linearity was down near 0.08% on the first try. So, if had a time machine, I could go back to the year 1940 and be a hero by inventing V/F converters and F/V converters and digital voltmeters using V/F converters. Unfortunately, in the year 1940, I was not even 1 year old, and if anybody was going to invent V/F converters in 1940, it was going to be somebody other than me! Still, in concept, there was nothing to prevent this design from being marvelously useful, back 40 or even 50 or 60 years ago. The V/F converter could have teamed up with the frequency counter or "EPUT Meter," to make digital voltmeters and analog-to-digital converters, 20 or 30 years before the first DVMs were sold commercially...

In 1971 I was invited to do a lecture tour, talking about Philbrick's new products in a 14-day tour of Japan. I talked about many other products as well as the V/F converters, with a slide show and various lectures and data sheets. One of the sequences I made up was to show the linearity of the model 4705. After we trimmed the gain-adjust pot to deliver 1,000,000 Hz at 10.00000 V input, I showed the test results—a 5-digit DVM sat on top of a frequency counter. For example, when you put in 5.00000 V, the frequency might typically be 500,008 Hz—a pretty good match. I showed the sequence we actually followed in our final test sequence—the input was decreased down to 2 V, 1 V, 0.1 V, 0.010 V, 0.001 V, and 0.0001 V. Finally I showed the last slide where the input was 0.00001 V, and the frequency was 1 Hz—showing a dynamic range of 1 million to 1. To my astonishment, the room full of Japanese engineers burst into spontaneous applause! I must say, this happened in Osaka, where the Japanese often show enthusiasm. In Tokyo, I noticed

the engineers in the audience were impressed, but in Tokyo, people are more reserved, and applause did not seem appropriate to them. Still, that applause was the high point of my tour—and perhaps of my entire career!

Notes on "Markets"

Why did the 4701 and its companions become so popular? I think there are several good applications which made for a lot of the popularity, and most of these were not really obvious, so it is not silly to say that marketing people might be fooled.

A major application was analog-to-digital conversion. The 4701 could cover a 10-bit range, from 10 Hz to 10 or 11 kHz, and at the price it was a little better than a comparable 8-bit ADC. Further, you could get the output isolated by 1000 V or more, just by driving the output into an opto-coupler. That made it *very* attractive.

An additional variation on this theme: this ADC could serve as an integrator. Instead of just feeding the output to a 10-bit counter, you could feed it to a 20-bit counter, or to a 16-bit or 24-bit counter with a 10 or 16 or 20-bit prescaler. So, you could integrate signals, just as with a wattmeter, by adding a couple of inexpensive counters, CD4020 or 4040 or similar. I think that made a lot of friends, because in the early 1970s, people were not interested in screwing around with analog computers or analog integrators. They didn't want to reset them, or have anything to do with them. The 4701 let them get the job done and all the user had to add was a couple of DIP ICs in the digital domain. Some of our big orders came because a major Japanese instrument maker was integrating signals related to air pollution, integrating them all day long. . . .

The other major application was isolation of analog signals, over a large AC or DC voltage offset. These days you can run over to Analog Devices or Burr Brown and get some cute isolators that stand off 1000 V or more, *but* those circuits only came along because the 4701 and 4702 pioneered the isolation business. I recall the first order we ever got for 1000 4701 V/F converters *plus* 1000 4702 frequency-to-voltage converters. The customer was in Japan. Many Philbrick people got *really* curious: what was he going to do with them? What industry was he in? Oh, he was in the shipbuilding industry.

But a ship is a solid slab of steel, and why would anybody be interested in isolation in such a case? Finally, some knowledgeable person pointed out, almost everything on a ship might be "grounded," but there are often dozens and hundreds of volts of AC and DC and transients and spikes, between any two "ground" on the ship. Maybe our customers weren't so stupid, after all. Maybe there *was* a "market" for V/F converters! I am still not an expert at "marketing," but I find that when I listen to the customers, as a class, they can figure out a bunch of things I never could imagine. So, even if there is no obvious place for these circuits to go, well, just stand around and see who shows up with money. Shucks—I nearly forgot: the words of Jay Last, one of the founders of Teledyne. He liked to say, "The only valid market survey is a signed purchase order." That's the last word.

How *Does* that Kluge Work?

The best way for me to explain how the 4701-type circuit works is to suggest that you assume that it *really does* work; and after we analyze each section of it, then you will agree that was a reasonable assumption. Figure 29-3 is my schematic of the circuit.

First, let's assume that the negative input of the op amp A1 is a few millivolts more negative than the positive input, and that V_{IN} is some positive voltage. Then rather soon, the voltage at the negative input will cross over and exceed that of the

TELEDYNE PHILBRICK

10kHz High Performance Voltage to Frequency Converters

4701 4713 4725

The 4701, 4713, and 4725 are high performance, low cost, voltage to frequency converters capable of producing a 10Hz to 10kHz output pulse train from a +10mV to +10V input signal. Twenty percent overrange, up to 13 bit resolution and low noise feedthrough are some of the inherent features of these general purpose devices. They are available to three different guraranteed nonlinearity specifications: ±0.1%FS (4713), ±0.05%FS (4705) and ±0.015%FS plus ±0.015% signal (4725). Full scale and offset errors, ±0.75%FS and ±0.03%FS respectively, are the same for the three units. Applications include FM telemetry, magnetic tape recording and digital to frequency conversion.

Applications Information

Precalibrated to meet all published specifications, these devices provide the user with optional trimming for applications requiring greater accuracies (see figure below). Input offset voltage is trimmed by applying a 100mV signal to the input terminal and adjusting R2 for a 100Hz output. Full scale is then trimmed by applying 10V to the input terminals and adjusting R1 for a 10kHz output. Repeat above procedure for precise calibration.

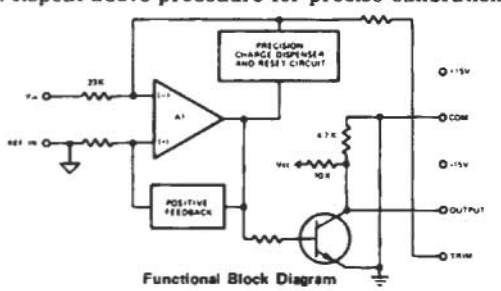

Functional Block Diagram

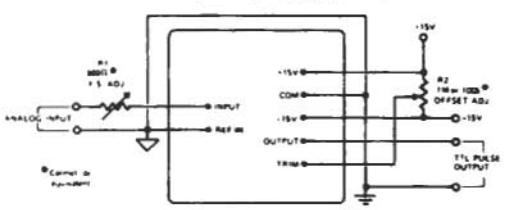

Connections Required for Operation, plus optional Input Offset and Full Scale Adjustments

FEATURES

- ±0.008%FS Nonlinearity
- 20% Overrange
- 13 Bit Resolution
- High Noise Rejection
- Low Cost

APPLICATIONS

- FM Telemetry
- Precision Integrators
- Common Mode Voltage Isolation
- Digital to Frequency Conversion

TELEDYNE PHILBRICK Allied Drive @ Rte. 128, Dedham, Massachusetts 02026
Tel: (617) 329-1600, TWX: (710) 348-6726, Tlx: 92-4439

Figure 29-4.
The data sheet for the 4701 family.

positive input. Now, all this time, A1's output has been at the positive limit, near +13 V; but when the inputs cross, the output will rapidly go to –13 V. What are the interesting things happening then?

1. There is positive feedback through C3 = 15 pF, so the positive input is driven a few volts negative. This ensures that there will be a good wide healthy pulse—at least 17 or 21 μsec.
2. The right-hand end of C2 is driven from about +7.5 V to –7.5 V, and the left-hand end of C2 is discharged through the emitter of Q1. That is a specific amount of charge, $Q = C \times \Delta V$. The ΔV is not just the 15 V p-p at the right-hand end of C2, but rather it is that 15 V p-p *less* the 2.5 V p-p excursion at the left-hand end. When this charge is driven into the emitter of Q1, it comes out the collector (well, at least 99.5% of it does) and pulls the voltage on the 3.3 μF down by about 12 mV. All this time, the voltage at the positive input of A1 is much more negative than that, so this step of –12 mV is just a minor jump. But, that charge is the magic element of precision. The size of the 12-mV jump is not important, but the *charge* is.

Note, in most conventional V/F Converters, the charge is dispensed by a timer circuit such as an LM555, which gates a current ON and OFF, so $Q = I \times T$. However, you need several precision resistors, and even then, the timer is subject to drift and noise, so that is a rather crude, drifty, noisy, unstable kind of charge dispenser. In the 4701, the gain depends almost entirely on just three elements—the zener voltage, the capacitor C2, and the input resistor R8. It's true that the diodes enter in, but since the V_fs of D1, D2, D3, and Q1 cancel out the V_fs of D4, 5, 6, and 7, then there is not a lot of error or drift likely there.

3. Now that the output is staying at a negative 13 V and most of the charge on C2 has flowed through Q1, there are two more details going on:
   a. The voltage at pin 3 of A1 is tailing up gradually to be more positive than at the pin 2. After all, pin 2 was pushed down 12 mV. Soon, after about a total of 20 μsec, V_3 will cross V_2, and the output will bang back up to +13 V.
   b. During that time, the current through R2 pulls at the emitter of Q1 and makes sure that Q1's emitter settles at a stable voltage. It makes sure that Q1's emitter voltage does not tail off to some drifty voltage. Even though R2 looks like it will dump in current that would hurt precision, it actually helps the precision.
4. Okay, now finally V_3 crosses V_2 and the output goes positive. Now we have to wait for the current through R8 to pull V_2 up those 12 mV that it was pushed down. That time will of course depend (inversely) on the size of the signal input; the bigger, the faster. That means the time between pulses could be anything between 70 μsec and 9 or 90 msec. Are we forgetting something? Yes. The p-p voltage at the left end of C2 must be stable and constant and invariant of rep rate. But the diodes there might give a long tail—the voltage might settle quite gradually and give a different p-p value at different speeds. By definition, that would hurt the linearity. What's the fix? The current through R2 is the fix. That current flows through D1, D2, and D3, and forces the left end of C2 to settle to within a millivolt or two in just 50 or 60 μsec. Without R2, the linearity goes to pot. Now, it looks *really stupid* to have a circuit like this where the "precision capacitor" C2 has a resistor across it that obviously makes so much "leakage." But that controlled "leakage" turns out to be *exactly* the reason for the precision and excellent linearity. The Swedish design didn't have this, and while their circuit had good linearity at 1 kHz, it

could not be made to work well at 10 kHz. But this basic charge dispenser, when driven with suitably fast transistors, works well at 100, 1000, and even 5,000 kHz.

What else is in the circuit? D9 is needed between pins 1 and 8 of the LM301A to keep it from wasting excessive current in its negative limit. D8 is a good idea to protect the positive input from overdrive in the positive direction. Q3 functions only when you overdrive the input—let's say—pull V_{IN} up to 50 V, and put in so much current that the V/F converter stops. Then it stops with pin 2 of A1 at +1 V, and pin 6 at −13 V. It would never put out another pulse—it would never restart, even if V_{IN} falls to a legal level such as +1 or +5 V—except that after a lag, C5 gets pulled minus, and Q3 turns on and pulls pin 2 so low that operation does start again. In normal operation, Q3 just sits there and does nothing, biased OFF.

C7 acts as a pulse-stretcher. The pulse width at the output of A1 is about 22 μsec. But we had a companion F/V converter, the 4702, that could only accept 20 μsec (or wider) pulses. If A1's output pulse got any narrower than 20, the 4702 would lose accuracy. We solved the problem by putting in C7 so that when A1 makes a 20 μsec pulse, the base of Q2 would be held off a little longer than that, due to the RC lag—about 15 μsec extra. Then a 4701's pulse was always plenty wide enough to drive a 4702.

The little capacitor C6 was added to make the p-p voltage at V_2 a little bigger, so when some LM301's were a little slow, there was now a bigger signal between V_2 and V_3, and the amplifier would not balk. After all, the LM301 is characterized as an operational amplifier, and if some are a little slower than others when you run them as a comparator, you really can't complain. . .

As you can see, the 4701 circuit did get a couple Band-aids®, but not excessively many, and we never really did get stumped or stuck in production. Our customers loved the linearity, the TC was pretty good, and the frequency output would never quit. They figured they really got their money's worth, and I certainly couldn't disagree with a satisfied customer.

A Final Note

Now, in July 1988 I did read a letter, which a Mr. Sidney Bertram of San Luis Obispo, California had written to the *IEEE Spectrum* (July 1988) about how he had worked on frequency-modulated oscillators in 1941. To quote from the letter: "When I joined the sonar group at the University of California's Division of War Research in 1941, I was told about three frequency-modulated oscillators they had had developed under contract—one by Brush Development, one by Bell Laboratories, one by Hewlett-Packard Co. The Hewlett-Packard oscillator, a positive-bias multivibrator, was the simplest and most satisfactory. It became the heart of the subsequent work, and one of my jobs was to give it a reproducible linear voltage-frequency characteristic."

I wrote off to Mr. Bertram and he was very helpful. He sent me a copy of a paper he had written, published in the *Proceedings of the IRE*, February 1948. This paper showed Mr. Bertram's work (in 1941) to take a basic two-tube multivibrator with six resistors and optimize it to a linearity of ± 200 Hz over a range of 37 kHz to 70 kHz, or about 0.6%. Not bad for a 1941 design!

So, we cannot say there were *no* VCOs before 1950, but they were not common knowledge, as you could not find them unless you looked outside the "Bible"—the 3000 pages of the Rad Lab Series.

9. The Story of the P2—The First Successful Solid-State Operational Amplifier with Picoampere Input Currents

First, let us start with—

A Fable

Once upon a time there were two wizards who decided they wanted to play golf. The first wizard stepped up to the tee, addressed the ball, and drove the ball right down the middle of the fairway; the ball then bounced twice, and rolled, and rolled, and rolled, and rolled right into the cup. The second wizard looked at the first wizard. Then he stepped up to the ball and drove a wicked screaming slice off to the right, which hit a tree, bounced back toward the green, ricocheted off a rock, and plopped into the cup. The first wizard turned to the second wizard and said, "Okay, now let's play golf."

End of Fable

Once upon a time, back in the ancient days of the electronics art, about 1958, there were two wizards, George A. Philbrick and Robert H. Malter, and they enjoyed designing operational amplifiers. In those days, that's what they called them—operational amplifiers, not "op-amps." George had the idea to use some of those new "transistors" to amplify the error signal from a balanced bridge, up to a good level where it could then be demodulated and amplified some more and used to form an operational amplifier. Ah, but what kind of balanced bridge would this be? A ring of conducting diodes? Heavens, no—George proposed a bridge made of 100-pF varactor diodes, so that when the bridge was driven with perhaps 100 mV of RF drive, the diodes would not really conduct very much and would still look like a high impedance—perhaps 10,000 MΩ. Then just a few millivolts of DC signal could imbalance the bridge and permit many microvolts of radio frequency signals to be fed to the AC amplifier. Now, back in 1958, just about the only available transistors at any reasonable price were leaky germaniums, and you certainly could not build a decent operational amplifier out of those. But George got some of the new 2N344 "drift" transistors that still had some decent current gain at 5 Mcps. He ran his oscillator at 5 Mcps, and after running his signal through the whole path of the modulator and then four stages of 2N344 RF amplifier, and a demodulator, he fed it into a DC amplifier stage with push–pull drive to a class AB output. And it was all built as a quasi-cordwood assembly, with seven or eight little PC (Printed Circuit) boards strung between two long PC boards. Since each little PC board had about six wires connecting into the long PC boards, this was a kluge of a *very high order*, and *not* fun to assemble or test, or to evaluate, or to experiment with, or to troubleshoot. George called this amplifier the P7. Please refer to the schematic in Figure 9-1. I know this is the right schematic because I still have a P7. Also, see the photograph of a P7's inner workings in Figure 9-2.

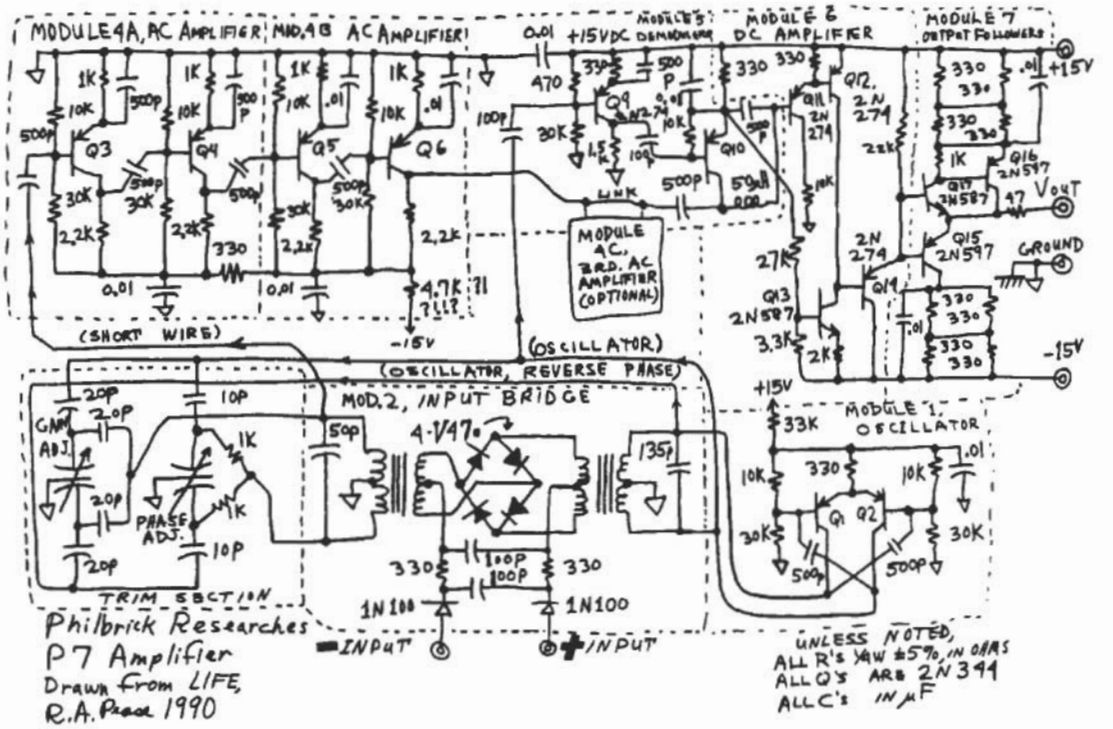

Figure 9-1. Bob Pease's rendition of the schematic for the P7 amplifier.

I never talked very much about the P7 with George, but obviously it had many problems. I don't think it was ever tested successfully as an operational amplifier, not as a working amplifier, nor as a product. Still, it had some promise. After all, if you could get it to work, this little circuit that used about 15 inexpensive germanium transistors could (it was hoped) have a better input current than even the better vacuum-tube amplifiers of that era—less than 10^{-9} A—better than a nA. (Heck, even a 12AX7 had 10 nA of grid current, and that was sort of the standard for ordinary operational amplifiers.) So George Philbrick concentrated on the P7 principle and the P7 design. Some people would use the word *obsessed*. He spent most of his

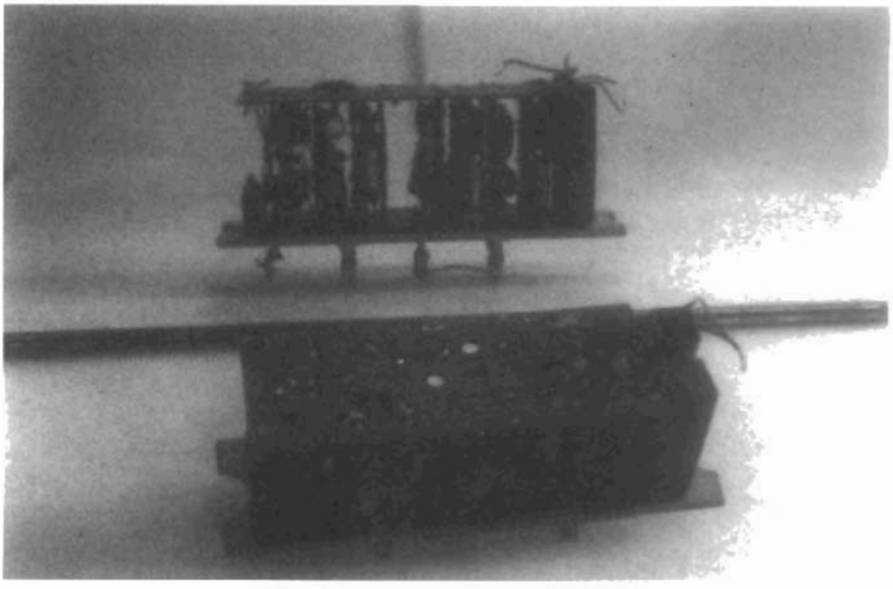

Figure 9-2. The guts of the P7 amplifier. It was photographed in front of a mirror so both sides could be seen.

time for a couple years, and a lot of the company's resources, trying to get the P7 working.

All this experimentation was going on at George A. Philbrick Researches, first at 230 Congress Street and then at 127/129 Clarendon Street, and then at 221 Columbus Avenue and 285 Columbus Avenue, in Boston, Massachusetts, back about 32 years ago. George had started a business to sell analog computers, but even in the 1960 era, the business in operational amplifiers (such as the K2-W) was starting to grow and overshadow the analog computer business. Imagine that—people actually buying op amps so they could build their own instruments!

When Bob Malter arrived at Philbrick in 1957, he was already a smart and accomplished engineer. He was a native of Chicago, and he had served in the army at Dugway Proving Ground. After designing several analog computer modules (which were the flagships of the Philbrick catalog), he became intrigued with the concept of the varactor amplifier, about the time that George was getting frustrated. Now, Bob Malter was a very pragmatic, hard-headed engineer. You would *not* want to bet him that he could not do something, because he would determinedly go out and do it, and prove that he was right—that you were wrong. Bob had his own ideas on how to simplify the P7, down to a level that would be practical. I do not know how many false starts and wild experiments Bob made on what he called the P2, but when I arrived at Philbrick as a green kid engineer in 1960, Bob was just getting the P2 into production.

Instead of George's 10 PC boards, Bob had put his circuits all on 2 PC boards that lay back-to-back. Instead of 14 transistors, he had a basic circuit of 7 transistors—just one more device than the little 6-transistor AM radios of the day. He actually had two little transformers—one to do the coupling from the oscillator down into the bridge and one to couple out of the balanced bridge into the first RF amplifier. A third inductance was connected in the emitter of the output transistor, to help tailor the frequency response. Please refer to the schematic diagram of the P2 in Figure 9-3. I mean, just because *everybody else* used only capacitors to roll off the frequency response of *their* operational amplifiers—well, that did not scare

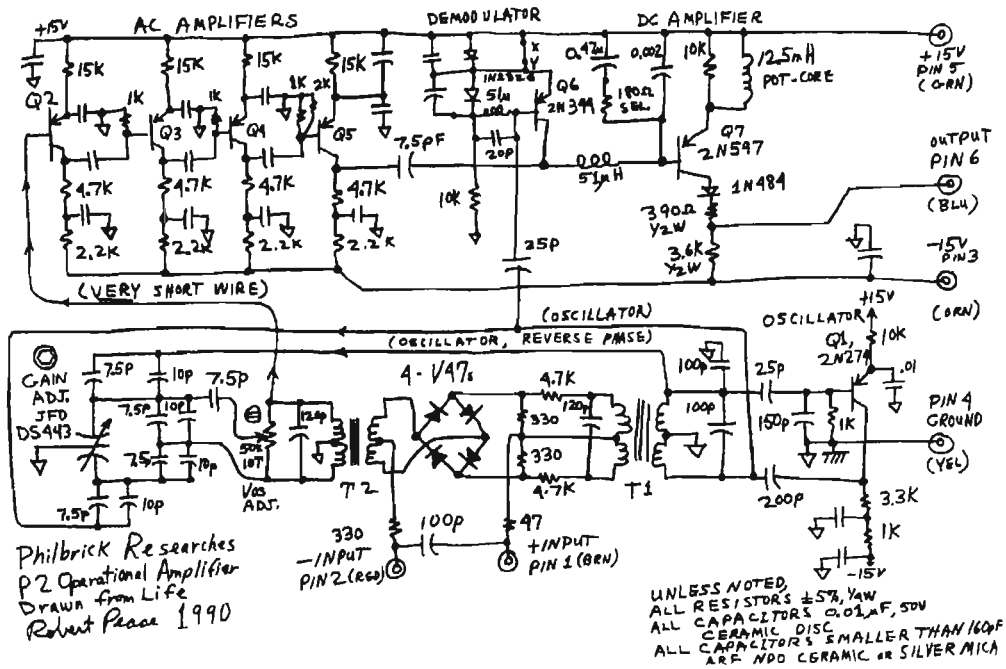

Figure 9-3. The schematic for the P2, as drawn by Bob Pease.

or impress Bob. He had to tailor the response of this operational amplifier with about 75 kcps of unity-gain bandwidth, and he had to roll it off at about 11 dB per octave or else have no output swing past 10 cps. As it was, he got the full 20 V p-p out to 500 cps, and even that was a struggle to accomplish. So Bob used the inductors and anything else he could think of that would help, in addition to various capacitive damping circuits. And he got it all to work. He got it to work quite well.

So, what's the big deal? Here's a pretty crude operational amplifier with a voltage gain of 10,000, and an output of ±1 mA at ±10 V, with a *vicious* slew rate of 0.03 V/μsec. Who would buy an amplifier like that? It turned out that thousands and *thousands* and *thousands* of people bought this amplifier, because the input bias current at either input was just a few *picoamperes*. What the heck is a *picoampere*? Most electrical engineers in 1960 didn't even know what a picofarad was, not to mention a picoampere, but they figured out it was a heck of a small fraction of a microampere—at 10^{-12} A, a picoampere is only 1 millionth of a microampere—and for many high-impedance instrumentation applications, the P2 was by far the only amplifier you could buy that would do the job. And it had this low bias current, only a few picoamperes, because all those germanium transistors were running at 5 Mcps, and their 5 or 10 μA of DC base current had no effect on the precision of the input current. The input current was low, thanks to a well-matched bridge of four V47 varicaps. These were sold by Pacific Semiconductor, Inc. (PSI) for use as varactors in parametric amplifiers, up in the hundreds of "megacycles," in low-noise communications receivers, mixers, and front-end amplifiers—parametric amplifiers. The "V47" designation meant that they had a nominal capacitance of 47 pF at 4 V reverse bias, which is where most RF engineers would bias them. But Bob Malter biased them right around 0 V DC, with a minuscule ±60 mV of AC drive.

At this level of drive, each diode would leak only 20 or 40 pA. But Bob had a gang of technicians working day and night to match up the forward conduction characteristics and the reverse capacitance voltage coefficients, and he was able to make sets of four varactors that would cancel out their offset drift versus temperature, and also their reverse leakage. Of course, there was plenty of experimenting and hacking around, plenty of experiments that didn't work, but eventually a lot of things that worked okay. After all, when you buy 10,000 V47s, *some* of them have to match pretty well.

So, here's a little do-hickey, a little circuit made up of just about as much parts as a cheap $12 transistor radio, but there was quite a lot of demand for this kind of precision. How much demand? Would you believe $227 of demand? Yes! The P2 originally started out selling for $185, but when the supply/demand situation heated up, it was obvious that at $185, the P2 was underpriced, so the price was pushed up to $227 to ensure that the people who got them were people who really *wanted* and *needed* them. So, the people who really wanted a P2 had to pay a price that was more than ⅛ the price of a Volkswagen Beetle—that was back when $227 was a real chunk of money!

Meanwhile, what other kinds of "transistorized" op amps could you buy? Well, by 1963, for $70 to $100, you could buy a 6- or 8-transistor amplifier, with I_{bias} in the ball-park of 60,000 to 150,000 pA, and a common-mode range of 11 V. The P2 had a quiet stable input current guaranteed less than 100 pA (5 or 10 pA, typical), and a common-mode range of ±200 V. (After all, with transformer coupling, the actual DC level at the balanced bridge could be at any DC level, so there was no reason the common mode rejection ratio could not be infinite.)

Wow. A $227 *gouge*. (You couldn't call it a "rip-off," because the phrase hadn't been invented, but perhaps that is the only reason.) Obviously, this must be a very

profitable circuit. Every competitor—and many customers—realized that the P2 must cost a rather small amount to build, even allowing a few hours of work for some special grading and matching and testing. So, people would invest their $227 and buy a P2 and take it home and pull it apart and try to figure out how it worked. The story I heard (it might be partly apocryphal, but most of it probably has a lot of truth) was that Burr Brown hired a bright engineer, handed him a P2 and told him, "Figure out how they do this—figure out how we can do it, too." In a few days the engineer had dismantled the circuit, traced it out, and had drawn up the schematic. Then he analyzed it and began experiments to be able to meet or exceed the P2's performance. But he couldn't get it to work well. He tried every approach, but he never could. After a full year, Burr Brown gave up and put the engineer to work on some other project. Burr Brown never did get into the varactor-input amplifier business, and I believe there is truth behind that story. Let me tell you why.

The P2 had an offset-adjust trim that was a little 20-turn trim-pot—that's not a surprise. But it also had a "gain adjust." This was not any ordinary gain adjust. This was a 20-turn variable trim capacitor—a differential piston capacitor—which the user could trim, per the instruction sheet, with a tiny little Allen wrench or hexagonal key. But it did not just have a linear control over the gain. If you trimmed the pot over to one end, the gain might be at 300 or 500, and then as you trimmed it closer to the center, the gain might rise to 700 or 900—and then, suddenly, the gain would pop up to 7,000 and then to 10,000 before the nonlinearity made the gain fall off again, when you turned the adjustment too far. The test techs called this, "going into mode." I used to wonder what they meant by that.

Several years later, George Philbrick brought me in to help him on an up-dated, up-graded version of the P2—the "P2A". We had to redesign the P2 because Philco was stopping the manufacture of those old 2N344s, and we couldn't buy any more, so we had to redesign it to use the more modern (cheaper) Silicon Mesa transistors, such as 2N706 or 2N760 or whatever. When we bought them from Texas Instruments, they were labelled "SM0387." I had a new circuit working pretty well, with the help and advice of George Philbrick, because Bob Malter had passed away, sadly, after a long bout with multiple sclerosis, about 1966. Anyhow, I was getting some results with the silicon-transistor version, but the improvements weren't coming along as well or as fast as I expected, so I went back to fool around with several real P2s, and to study them.

The original P2 had an apparent imbalance at the output of its demodulator. Well, that looked kind of dumb, that the first DC transistor would be turned off unless there was a pretty big signal coming out of the demodulator. To turn on the DC transistor, you had to have a considerable imbalance of the RF. So I took one unit and modified it to bias the demodulator about 1 V_{be} down from the positive supply, so it would not have to handle a great amount of signal just to drive the DC transistor Q7. Refer to the schematic of the P2 (Figure 9-3). Normally, Q6's emitter was connected directly to the +15 V bus. I disconnected it by removing link X–Y and connected it to a bias diode. Yes, the RF amplifier ran with less RF signal at balance—but the gain refused to "come into mode." So that "improvement" scheme was unusable. Now, what was *that* trying to tell me?

After some more study, I planned a few more experiments, and then I tried pulling apart the two PC boards so I could access some of the signals down in between the boards. As I eased the two boards apart (with power ON), the gain "jumped out of mode." I gradually realized the P2 amplifier was running, all these years, as a reflex amplifier. The "gain adjust" consisted of changing the phase between the oscillator and the bridge, so when the amplified signal came down to the end of the RF ampli-

fier (four stages, remember) and was patched back to the other PC board, it would be able to regeneratively amplify even more than the honest gain of the RF amplifier. *That* was why the demodulator wouldn't work right unless a certain constant minimum amount of 5 Mcps signal was always flowing through the amplifier. *That* was why the gain would "pop into mode" (and when it wouldn't "pop into mode," that explained why not). *That* was why the engineer down at Burr Brown couldn't figure out how to get it working right—the gain depended on the two PC boards being spaced just the right distance apart! *That* was the trick that Bob Malter had accidentally built into the P2, and that he had figured out how to take advantage of. To this day, I am not sure if Bob Malter knew exactly what a tiger he had by the tail. But I would never dare to underestimate Malter's tough and pragmatic brilliance, so I guess he probably did know and understand it. (I never did have the brass to ask him exactly how he thought it worked. I bet if I had had the brass to ask him, he would have told me.) I must say, if any engineer was bright enough to grasp and take advantage of a strange interaction like this, well, Bob Malter was that sharp guy.

Now, since my P2A was designed on a single board, with the demodulator far away from the inputs and oscillator, we wouldn't have any "mode" to help us. But that was okay—now that I understood the "mode" business, I could engineer the rest of it okay without any "mode," and I did. But that explained why none of our competitors ever second-sourced the P2. And the P2A and SP2A remained profitable and popular even when the new FET-input amplifiers came along at much lower prices. It was years before these costly and complex parametric amplifiers were truly and finally made obsolete by the inexpensive monolithic BIFET™ (a trademark of National Semiconductor Corporation) amplifiers from National Semiconductor and other IC makers. Even then, the FET amplifiers could not compete when your instrument called for an op amp with a common-mode range of 50 or 200 V.

A friend pointed out that in 1966, Analog Devices came out with a "Model 301," which had a varactor input stage. It did work over a wider temperature range, but it did not use the same package or the same pin-out as the P2.

Still, it is an amazing piece of history, that the old P2 amplifier did so many things right—it manufactured its gain out of thin air, when just throwing more transistors at it would probably have done more harm than good. And it had low noise and extremely good input current errors—traits that made it a lot of friends. The profits from that P2 were big enough to buy us a whole new building down in Dedham, Massachusetts, where Teledyne Philbrick is located to this day. The popularity of the P2 made a lot of friends, who (after they had paid the steep price) were amazed and delighted with the performance of the P2. And the men of Philbrick continued to sell those high-priced operational amplifiers and popularize the whole concept of the op amp as a versatile building block. Then, when good low-cost amplifiers like the μA741 and LM301A came along, they were accepted by most engineers. *Their* popularity swept right along the path that had been paved by those expensive amplifiers from Philbrick. If George Philbrick and Bob Malter and Dan Sheingold and Henry Paynter and Bruce Seddon hadn't written all those applications notes and all those books and stories, heck, Bob Widlar might not have been able to *give* away his μA709s and LM301s! And the P2—the little junk-box made up virtually of parts left over from making cheap transistor radios—that was the profit-engine that enabled and drove and powered the whole operational-amplifier industry.

Since George Philbrick passed away about 1974, and Bob Malter had died earlier, I figured I had an obligation to tell this story as there was nobody else left to tell it. Even though I was not in on the design of the P7 or the P2, I understood their designs better than just about anybody else. So, I just have to express my appreciation to Jim Williams for leading and editing this book. I know he will want to read about

the P2, because I know he has one in his lab. (Meanwhile, I agreed to write a chapter for Jim's book, and to support and encourage the book, because I want to read all the *other* stories that will be in here.)

Vignettes: Additional Little Stories about the P2

One time Bob Malter came back from the big WESCON show in Los Angeles. He said, "I made a good bargain for a new spec on the piston capacitors. I got the price down from $2.15 to $1.65. That savings will pay for my trip and then some." It sure did.

One time, there were some P2s that had a lousy tempco. Most of the units had a drift much better than 6 mV from 20 to 45 °C. But this time a couple batches had a lousy yield for drift. So Bob figured out where to install some little thermistors—across one of the legs of the 50 k pot—and his wizardly technicians delved away like mad, and trimmed and tweaked and tested, and sure enough they got the drift to improve enough to meet specs. I said "delved," because they had to dig through the room-temperature-vulcanizing (RTV) potting material to access the places they needed.

One time, just a couple days after Bob went on vacation, the frequency response began to go to pot, and none of the usual tricks would fix it. So the senior technician, Tom Milligan, got on the phone to Bob (who had given him his vacation phone number), and Bob figured out a tweak, and by the time Bob got back from his vacation, the problem was completely cured.

One time, I was standing around in front of the Philbrick booth at the big IEEE show in New York City. A couple engineers were hiking past the booth, and one said to the other, nodding his head toward the booth, ". . . and there's the company that makes a *big* f------ profit." Well, at that time George A. Philbrick Researches was indeed making big profits from the P2. Can't deny it.

On various occasions, customers would ask about how to get the best long-term stability of the offset voltage. It turns out that most parts, if held at a constant temperature, could hold an offset voltage better than 100 μV/hour, and some were as good as 20 μV/hour. We had our little Rustrak meter to prove it. Heavens, we used

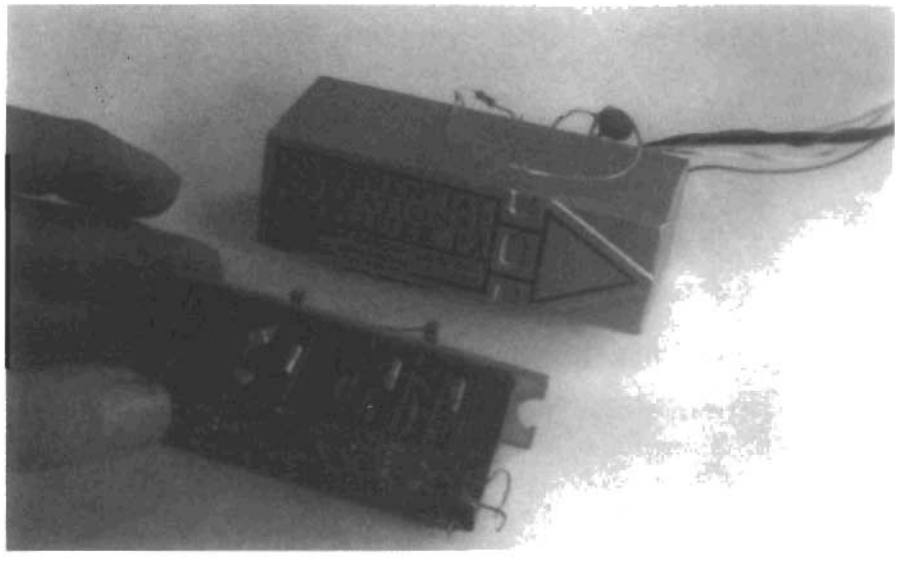

Figure 9-4. It was what was inside that was important!

miles of that Rustrak paper. When the P2A came along, it was able to do as good as 1 or 2 μV p-p for an hour. But Bruce Seddon, one of the senior engineers, was always skeptical about the possibility of a P2 having V_{OS} stability that good. He computed that a single microvolt was worth about 600 electrons on each varactor. Since a varactor diode had really a rather shallow slope, you could compute that a 1 mV DC input would cause a 0.03 pF imbalance in a 200 pF bridge. And a 1 μV DC imbalance would cause a 0.00003 pF imbalance. Needless to say, that was a preposterous situation. You could compute that even a couple of *atoms* of shift on the components nearest to the varactors would cause worse imbalance than that. But we measured a lot of P2s, and a lot of P2As, and some of them would hold better than 1 μV p-p for an hour or two. Bruce always was incredulous about that.

Now, if you trimmed the offset voltage to zero, the input current was pretty small, about 5 or 10 pA typical, and 100 pA guaranteed max. Some people would pay a surcharge for selected units with extra-low input current. But many people would just crank the input offset pot over to the side—perhaps a few millivolts—and get the input current down to less than 1 pA. It wasn't perfectly stable if somebody suddenly turned on the airconditioner, but under constant ambient conditions, it was better than all but the best electrometer-tube amplifiers.

In addition to having low DC errors, the P2 had fairly decent low noise. The P2A was guaranteed to be better than 1 μV rms in the bandwidth 1–100 Hz, and many P2s were almost that good. Now, how can an operational amplifier have noise as good as that, right where most solid-state amplifiers have many microvolts of noise p-p? The fact is that the varactors transform, or down-convert, the noise of the first RF transistor at (5 Mcps to 5.0001 Mcps), down to input noise at the inputs of the P2, in the frequency range (0–100 cps). Those varactors really did provide the advantages of a parametric amplifier. And those germaniums weren't bad at 5 Mcps, so the P2 did a respectably good job for low noise. It took many years before its performance was matched by FET amplifiers.

The P2 was assembled with its two little PC boards rivetted securely together and then installed in a cast aluminum case. Then the whole cavity was filled with room-temperature-vulcanizing (RTV) silastic material. It did seem to keep things at a constant temperature, and if there was much moisture, the RTV did seem to help keep it off the boards. Still, on some moist days, they could not get the P2s to pass a 100 pA final test, so they would just set them aside and wait for drier weather. When we built the P2A, we did not use RTV, because at +85 °C, the RTV would expand and *pop* the P2 right out of its case. We just used several heavy coats of Humiseal, and that gave very good results. I don't think moisture gave us much problem on the P2A.

According to some of Bob's friends, Bob said that he could tell when the women assembling the P2s were menstruating. He thought it was the amount of sweat that would cause corrosion or leaky printed-circuit boards. He could check out failure rates versus serial numbers versus the initials of the assembler, and the yield would go up and down every 28 days. I know I was impressed that there were always two inspectors, inspecting the PC boards after they were hand soldered. They could spot badly soldered joints and cold-soldered joints, and mark them with a red pen, to go back and get resoldered and touched up, because a P2 would sometimes run really badly, noisy and flaky, if there were cold-soldered joints on the board.

To this day, I still have the dismantled carcasses of a few P2s. That is because Bob Malter decreed that if you could not get a P2 to meet spec, after you had tried everything, the technicians would pull off the valuable components—the trim-pot and the piston capacitor—for re-use. Then the transistors and transformers would be removed, so that even a competitor who wanted to raid our trash cans would not

learn anything much. And in retrospect, well, Bob had a lot of good hunches, and he probably had a good hunch in this respect also.

Of course, if you wait long enough, any good thing can become obsolete. As of 1989, you could buy low-leakage amplifiers such as the NSC LMC660, with input currents normally less than 0.004 pA, for about 50 cents per amplifier ($2.00 for a quad). But what do you expect after a 30-year wait?

Notes on George Philbrick's P7 Circuit

1. The AC amplifiers are all supplied through a single 4.7 kΩ (?!) resistor! George wanted to run all 4 AC amplifier stages on barely 1 mA total! In the P2, Bob Malter was willing to spend 4 mA. I could never understand why George was so unwilling to spend just 15 mW for the entire four-stage AC amplifier when he spent 30 mW to bias up the output stage (27 k on Module 6). Maybe if he fed any more current through those AC amplifiers, they would break into song and oscillate hopelessly?! Because in the layout, the output of the fourth stage is right next to the input of the first AC stage?! Those of you who are not chess players might like to know that in the notation of chess, "?!" signifies a *blunder*.
2. Likewise, the P7 oscillator was intended to run on less than 0.3 mA (?!), whereas the P2 spent 1.5 mA. I checked the actual P7 circuit to see if these values represented a typo error—but they didn't.
3. The P7 I have uses two AC amplifier modules—four stages of transistor—but the arrangement of the upper and lower (mother and father?) boards left room for three amplifier modules—six stages of AC gain. You can see the gap in the middle of the assembled unit, where another two stages could have fitted in. But if you had six stages, then the possibility of oscillation would become hopelessly bad. No wonder George backed up to go with just four stages.
4. The germanium 1N100 diodes in series with the inputs are intended to act as low-impedances near null but to act as current-limiters (just a few microamperes) when overdriven. That is what George intended—a neat concept. But in actuality, I bet they made bad errors when they rectified any ambient noise, and I bet they had *awful* leakages when operated in ambient light. Furthermore, if you ever got around to running this amplifier with feedback, you would find they add a lot of phase shift. At room temperature, they would cause a lag of perhaps 25 kΩ and 600 pF, or a 10 kHz roll-off. If you get it cold, the break at +5°C would be at 2.5 kHz. This confirms my suspicions that George never really got the DC operation working okay, so he never even *began* to think seriously about AC response. The circuit shows no evidence of a big filter around the output stage to give 6 dB-per-octave rolloff.

Comments on Bob Malter's P2 Circuit

1. Obviously, 30 years ago this was a high-level industrial secret. But as I mentioned, even if I gave you the schematic, that would not help you make a P2 that works. Since the P2 has been out of production for more than 20 years, this is more like industrial archaeology than espionage.
2. The doubled capacitors (several places where you see [7.5 pF in parallel with 10 pF]) were arranged so that the test technicians could do some coarse trims by snipping out one or another of the caps. Much judgment and experience were needed. There were many other places where discretionary trim resistors and

capacitors could be added. To improve the temperature coefficent of V_{OS}, for example, you could install a thermistor from the wiper of the 50 k pot over to one side *or* the other.

3. Unlike the P7, the P2 had a lot of AC roll-off, provided by the 15 millihenry inductor and the 0.47 μF capacitor. It rolled the DC gain off at a steep 10 dB-per-octave rate down to about 15 kcps and then there was a lead (selected resistor in series with the 0.47 μF) so it could cross over at a unity gain frequency of about 75 kcps at about 6 or 7 dB per octave. The frequency response was trimmed and fitted on each individual unit.

   However, it is fair to note that the roll-off did not use any Miller integrator around the output transistor. Consequently, the high-frequency open-loop output impedance of the P2 was not a whole lot lower than 3 kΩ. If you combine that statement with the fact that the P2's input capacitance is just about 600 pF, you can see that the output impedance, just trying to drive the input capacitance, gives you an extra phase shift of about 40 degrees. No wonder each unit was hand-fitted for response!
4. The demodulator (Q6) would put out a voltage right near the +15-V bus if you did not feed in any amplitude from the AC amplifiers, and then the DC transistor (Q7) would not turn on. To get the output transistor on, you had to have a minimum amount (perhaps 400 mV p-p) of 5 Mcps signal coming through. And it was the interaction of that signal that talked back from one board to the other and let the gain come "into mode." Look at the coupling capacitor from the fourth AC amplifier into the demodulator! The P7 had a reasonable value—500 pF. But Bob Malter found something magic about the 7.5 pF, probably because it was the right way to get the amplifier into "mode." Surely, Bob Malter was the embodiment of "The Lightning Empiricist."

Comments on Rustrak data

I set up a P2—Jim Williams loaned me his old P2—at a gain of 20. I followed this with a gain of 200 (or 100 or 400) to get the offset voltage's drift up to a decent level, and fed it through 10 kΩ into an old 1–mA Rustrak meter—the kind that goes

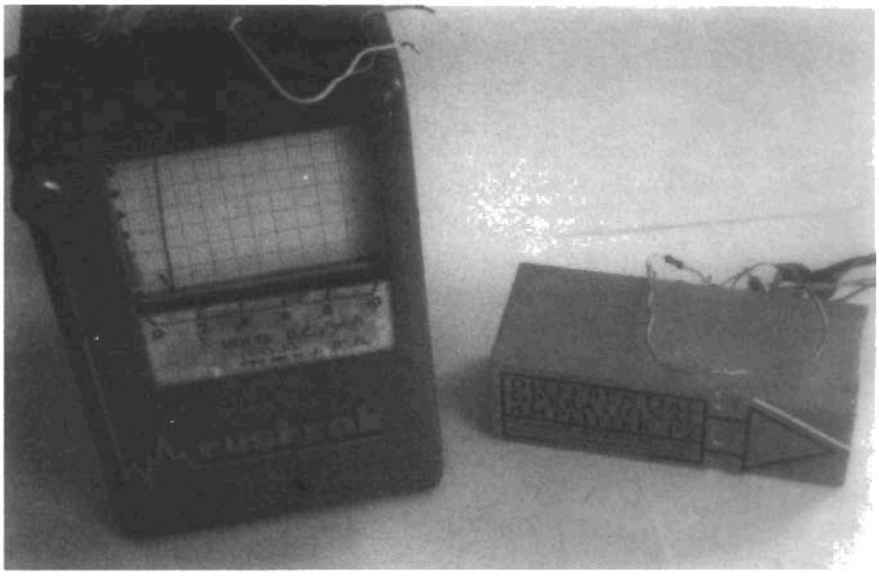

Figure 9-5. A Rustrak strip recorder tracking the offset drift of a P2.

B. Pease

GEORGE A.
PHILBRICK
RESEARCHES, INC.

DIFFERENTIAL
OPERATIONAL AMPLIFIER
SOLID STATE

GAP/R
TECHNICAL
BULLETIN
P2

(Preliminary and Tentative)

All Solid State!

NO CHOPPERS
NO TUBES
NO COMMON
MODE ERROR

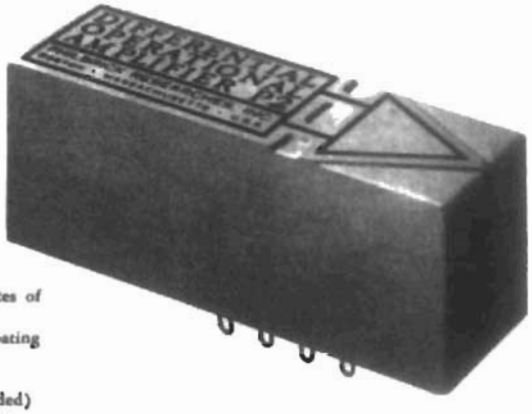

If the rigors of your applications (or the tastes of your customers) demand:

- FULL DIFFERENTIAL INPUT (Truly floating in respect to ground or output)
- COOLNESS (less than 330 milliwatt dissipation, fully loaded)
- COMPACTNESS (4" L x 1¼" W x 1 11/16" H)
- LOW INPUT CURRENT (Typically less than one-tenth milli-microampere (10^{-10} amp))
- NO HEATER SUPPLY OR TUBE REPLACEMENT
- LOW COST **$185**

you owe it to yourself (and your Company) to consider the virtues of the Philbrick P2.

The Philbrick philosophy of more value per dollar that has led to the fame of the K2-W, USA-3, and USA-4 amplifiers and the R-100B Power Supply, has scored another triumph with Model P2, which, at its price of $185 (1-9 units), has more performance per dollar than any solid state operational amplifier now on the market.

GENERAL DESCRIPTION

Model P2 is a dc amplifier with differential inputs designed for applications far beyond just analog computing. They include high reliability process control, electrometer applications, and instrumentation in general. Besides the usual operations of addition, integration, scaling, and inversion, its differential inputs make possible high impedance voltage following and amplification, subtraction, and precise current driving to grounded loads. It even allows itself such dogmatic luxuries as multivibration, flipping, flopping, clipping, tripping, and the like, but with precision. Philbrick's famous 24-page Applications Manual applies (but more simply) to our solid state amplifiers as well as to the K2's. Moreover, you can ignore the section on biasing techniques.

Unlike most solid state operational amplifiers, P2's high input impedance and low input leakage current allow it to be used with the same resistors and capacitors normally used in vacuum tube operational circuitry (but at lower voltage ratings and dissipations, hence with greater compactness and longer life).

Input current is so low, that its typical drift when connected as a 1 second integrator using a *1.0 microfarad* capacitor, is less than 500 microvolts per second. 24-hour amplifier offset drifts in a standard computor installation will be typically below 500 microvolts.

The P2 requires about 11 ms at ±15 volts, and it will operate continuously for 90 hours in the field when energized by two pairs of mercury batteries (TR-136).

Write to Philbrick for information regarding Solid-State Power supplies for use with P2.

SPECIFICATIONS AND CHARACTERISTICS

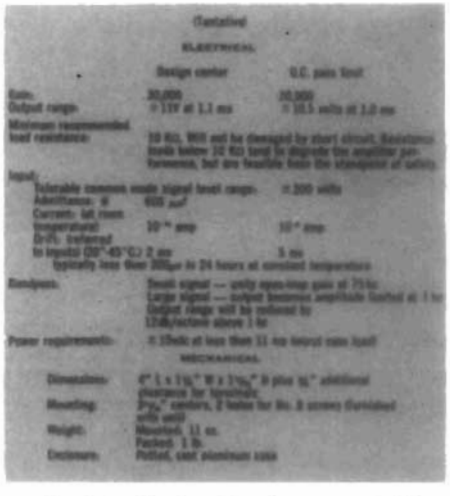

(Tentative)

ELECTRICAL

	Design center	0°C–60°C limit
Gain	30,000	20,000
Output range	±11V at 1.1 ma	±10.5 volts at 1.0 ma
Minimum recommended load resistance	10 KΩ. Will not be damaged by short circuit. Resistive loads below 10 KΩ tend to degrade the amplifier performance, but are feasible from the standpoint of safety.	
Input		
Tolerable common mode signal level range		±200 volts
Admittance, @	600 μμf*	
Current, (at room temperature)	10^{-10} amp	10^{-9} amp
Drift, referred to input (20°-40°C) 2 hr		5 ma
	typically less than 200μv in 24 hours at constant temperature	
Bandpass	Small signal — unity open-loop gain at 75 kc Large signal — output becomes amplitude limited at 1 kc Output range will be reduced by 12db/octave above 1 kc	
Power requirements	±15vdc at less than 11 ma output rate load	

MECHANICAL

Dimensions	4" L x 1¼" W x 1 11/16" H plus ¼" additional clearance for terminals
Mounting	3½" centers, 2 holes for No. 8 screws furnished with unit
Weight	Mounted, 11 oz. Packed, 1 lb.
Enclosure	Potted, cast aluminum case

* 600 μμf, differential, NOT 2-M

Figure 9-6.
The original data sheet from Philbrick describing the P2.

"—tick—tick—tick—." After some warming up and some clearing away the cobwebs, this P2 began giving pretty good stability. In some hours when the temperature wasn't changing much, it would hold 20 or 60 μV p-p—not bad for a unit with perhaps 200 μV/°C. Also not too bad, considering that the trim pot had an end-to-end range of 100 mV, so that asking it to hold 100 μV—the equivalent of 0.1% of span—was about as optimistic a task as anybody ever demanded of a carbon pot. But sometimes it did a lot better than that.

However, the offset kept drifting to the left. Could it be some kind of chemical interaction, where the RTV is changing slightly after all these years of inactivity? After all, the P2 really is *not* a well-balanced circuit. Maybe the drift rate will slow down if I do some warm-temperature burn-in?!

You never can tell. . . .

Preface

Comments on "World-Class" Analog Design

Achieving excellence in analog circuit design has always been challenging. These days it is still not always easy, so we want to help with some general advice. All the authors of these chapters have presented their best ideas as the kinds of things a good analog circuit designer must know to consistently accomplish very good circuits.

These days so much of analog circuit design can be done using operational amplifiers (op-amps) with a small number of discrete resistors and capacitors. It is often very easy to slap in resistors and the circuit works well. However, this is still not trivial. You might have to pick sets of matched resistors or add a trimpot. Even these days some young engineers have to ask, "So, should I make a 1-ohm/1-ohm unity-gain inverter?" Some kids really don't know how to pick appropriate resistor values; they have never done any practical work or lab work. So we have to teach them about practical circuits. We have to teach them about error budgets. Sometimes 1% resistors are quite appropriate; other times 5%, 10%, 0.1%, or 0.01% might be right. Richard Burwen has good comments on resistors. More on error budgets later.

Recently a guy showed me his design with eight precision op-amps and sixteen precision resistors. After I did some whittling out, we got it down to two precision resistors and one precision op-amp and a greatly improved error budget. More on error budgets later.

Once upon a time, in the 1950s, there were no operational amplifiers that you could buy. The engineers at Philbrick Researches wrote a twenty-eight-page *Applications Manual for Octal Plug-In Computing Amplifiers* (such as the K2-W, see **Figure P-1**). With a little advice from this pamphlet, you could design analog computing circuits and some simple instrumentation, too. I came to work at Philbrick about that time (1960). I studied operational amplifiers based on vacuum tubes and then high-performance solid-state amplifiers.

Applications Notes

Then about 1965, the new arts and applications using transistorized op-amps showed the need for a comprehensive *Applications Manual for Computing Amplifiers for Modeling,*

Figure P-1: Philbrick K2-W, 1952 to 1992.

Measuring, Manipulating, and Much Else. Dan Sheingold, George Philbrick, Bruce Seddon, and several others wrote a lot. I contributed a small bit. This book was *very* useful. My theory is that when Bob Widlar brought out the μA709, he couldn't have *given* it away, but Philbrick had sold and given away many thousands of these books, which made it reasonably easy to apply those IC op-amps. This book was sold for several years for $3. Recently, a good copy sold on e-bay for $300+. It's darn near worth it. Can you get the basic info off the Web? I'll have to look it up on Analog Devices' Website.

Other companies such as Burr Brown, Analog Devices, and TI wrote lots of App Notes and books on op-amp applications. I was never very impressed with them; they were not good explainers. NSC published lots of App Notes. Not all were well documented, but they were pretty good circuits.

Which Op-Amp?

Even for experienced engineers, this can be a bewildering question. There are many low-voltage and high-voltage op-amps; low-voltage noise and high-voltage noise; low-power and high-speed amplifiers; and cheap and expensive ones. Let's see what insights Bonnie Baker can offer.

Precision Capacitors?

How many kid engineers know the price of 1% capacitors? Precision capacitors are rarely justifiable. Yet not all 1% capacitors are really high priced. Sometimes a dime will get you that; other times, it could take a dollar or two. And sometimes a circuit really does

need 1% capacitors. I just got a *thick* Digikey Catalog the other day, and it has 2% and 1% tolerance polypropylene capacitors at surprisingly reasonable prices, even in small quantities!

Inductors?

Inductors are specialized animals that may be required for filters and for switch-mode regulators. Usually the designer of the switcher provides detailed advice on what to buy. If not, then designing with inductors, or redesigning to adjust the inductor type or values, is a special advanced area of expertise. Most schools don't teach much of this. The design of switchers can be either a high-tech specialty or a monkey-see, monkey-do exercise. The latter might not be as cheap, but it usually does work well.

Diodes

Diodes can be a truly bewildering field. Some can carry small milliamperes; some can leak less than a picoampere; some rectifiers can carry amperes without overheating. But the big ones (such as 1N4005) often cannot be used at high frequencies. The 2N5819 Schottky rectifier can carry a couple amperes, but it is somewhat leaky. Still, it can rectify up to 1 MHz without misbehaving. Who's going to teach everybody about diodes?

Especially tricky is the fact that some good, fast small-signal diodes (1N4148/1N914) do turn on and off quickly—faster than 1 nanosecond sometimes—but at low rep rates, some of them sort of forget how to turn on and have a bad overshoot. That's annoying.

Transistors and Designing With Them

Now, when you get to transistors, this becomes complicated. Designing with transistors is a whole 'nother game. Even experienced analog designers try to minimize that when they can. But sometimes you have to use transistors. Sometimes the transistor's inherent log characteristics are very important. Can you buy a logger? Yes, several companies make and sell loggers. But loggers can be designed for special cases, which a store-bought logger cannot handle, such as low voltage. I've done a couple of these in the last year. I still design low-noise and high-speed amplifiers occasionally using selected transistors, such as 2N3904 and LM394. I often use the curves from "What's All This V_{BE} Stuff, Anyhow?" Or you might merely need to use a transistor as a switch—a crude one or a precision switch.

Filters

When you need a filter, it might not be hard to figure out what is needed; other times more research is needed. Can you avoid inductors? Can you avoid expensive op-amps?

Can you avoid high impedances *or* large capacitors? As with all of analog design, this covers a *huge* dynamic range, and there is usually nothing simple about it. Yet it gets done.

SPICE

I usually try to avoid using SPICE. I use pen and paper; I call it "back-of-envelope SPICE." I do mostly hand computations, and good approximations, using my slide rule or by doing the math in my head. You might say I am in agreement with Dick Burwen's chapter, "How to Design Analog Circuits Without a Computer." Other people think that SPICE is acceptable over a wide range of applications. That makes me nervous. I find that you can use SPICE to save an hour of computation every day for a month and then discover that SPICE has made a costly mistake that wastes all the time you thought you saved. Some people agree with me on that.

Also, when people use monolithic transistor models (such as the ones in the monolithic array, LM3046), that is different from designing with discrete transistors. I mean, who will give you a free model of a 2N3904 that is worth what you paid for it? And in what regimes do you trust it? I would trust it for *only* the crudest noncritical applications.

Some people say they like to trust SPICE. If they get good models and they know what they are doing, good luck to them.

I will mention a few particular places where SPICE models do not usually work well:

- At low values of *Vce* (or *Vds*), where the transistors are starting to saturate.
- At high frequencies at low values of *Vce* (or *Vds*), where the frequency response of the transistor does not ring true.
- Monolithic transistors are *often* badly modeled where they saturate (or start to saturate) since the substrate currents get large.
- Sometimes when an op-amp's inputs get reversed, it will still appear to work like an op-amp without saturating. *Some* kinds of SPICE do work right in this situation, but not all.
- If somebody gives you a bad model, you might have problems. Even when you make your own model, it could have problems.
- Sometimes SPICE fails to converge and wastes a lot of your time.
- Sometimes SPICE gives an absurd answer, such as saying that a 10 exp-25 ampere current step has a real risetime. How can a "current" that consists of 1 electron per day show a "risetime"?

- Usually in a band-gap reference, the fine details of a temperature characteristic do not go in the right direction. SPICE cannot lead you to a better answer. My old LM131 from 1977 had (and still has) a good tempco because it was based on good breadboards. When I tried to run it in SPICE many years later, SPICE said it did not work and could not be made to work. It's a good thing I didn't try it in SPICE in 1977. SPICE was wrong.
- In any circuit where transistors are heated or self-heated, the temp rise of the transistors is *very* hard to model, especially in a distributed layout.
- And sometimes SPICE just *lies*. Sometimes it just gives incorrect answers.

I've had debates with many "SPICE experts" and they try to tell me I am wrong. But I have seen too many cases where I was right and SPICE was wrong. I say this because people bring me their problems when their circuit does not work. I can see through the errors of SPICE; I use special test techniques (mostly in the time domain or in thought experiments) to show why a circuit is misbehaving. SPICE is not only *no help*, it leads to "computer-hindered design."

How Many?

How many are you going to build? If you are going to build large numbers or small, it makes a difference how you engineer it, for minimum overall cost and maximum output.

Low Noise?

Many general-purpose op-amps are pretty quiet, but some that are quiet at low impedance are noisy at high impedances. Others that are quiet at high impedance are noisy at low impedance. Let's see what comments Bonnie Baker has on this topic.

Troubleshooting?

Once you get your circuit built, you apply power and then it does (or does not) work correctly. How do you do the troubleshooting? Better yet, how do you plan in advance a way that you can easily do the needed troubleshooting?

Check out the Bob Pease book *Troubleshooting Analog Circuits*. With 39,000 copies in print in six languages, it has *legs*—and that's because analog circuit troubles do not go away by wishing and sometimes not even by engineering. Sometimes they are solved only by real troubleshooting. But planning ahead can help. See www.national.com/rap/Book/0,1565,0,00.html.

The Future?

People often ask, "Would you encourage your son or daughter to go into engineering?" I reply, "Yes, if it is analog circuit design." They say, "Explain!"

I respond, "My friends and I know many analog design techniques, tricks, and secrets. They cannot be learned from SPICE. Every year there are 200,000 Chinese engineering graduates, and they don't know what we know. We can solve problems they cannot."

I rest my case. /rap

—Robert A. Pease
Staff Scientist, NSC
Santa Clara, CA
August 2007
rap@galaxy.nsc.com

P.S. One of the authors of a chapter in this book said that he took a "well-designed" system and put a good model of it into SPICE. When he ran it, he was surprised to find a sneaky sampling error. So we should not say that SPICE cannot be helpful. We just have to be cautious about trusting SPICE—in any positive or negative way.

CHAPTER 16

What's All This Error Budget Stuff, Anyhow?

Robert A. Pease

Well, I stated at the start of the story the reason it's important to do an error budget on even a simple circuit—and then I showed the size of trouble you can get into if you don't. I rest my case. /rap

I was just on the phone explaining to a young engineer how to do an error budget analysis on some fairly simple circuits. Later, I mentioned this while I was visiting my friend Martin, and he said he had been quite surprised when he found that many engineers in Europe were quite unfamiliar with the concept of an error budget. How can you design a good circuit without being aware of which components will hurt your accuracy?

When I was a kid engineer back in 1962, my boss George Philbrick gave me a book on differential amplifiers by Dr. R. David Middlebrook, and he asked me to do a book review. I studied the book, and it was full of hundreds of partial differential equations. If you wanted the output of a circuit with fourteen components, you could see a complete analysis of how each component would affect the output offset and gain. Each equation filled up a whole page. It did this several times.

Yet the book didn't offer any insights into what's important. I mean, is $\beta \times d(R1)$ more important than $R1 \times d(\beta)$? In retrospect, I'm glad I didn't submit any critique of that book. I would've done more harm than good. Such a mess! Even now, it would be hard to write a critique on a book that was so true but so unhelpful.

Things are much simpler now that people are mostly (but not entirely) designing with op-amps. The best thing is that the output offset and DC gain and AC gain errors are largely orthogonal. An "operational" amplifier does perform, largely, an "operation" based on what task you ask it to perform when you "program it" with *R*s and *C*s. If the offset

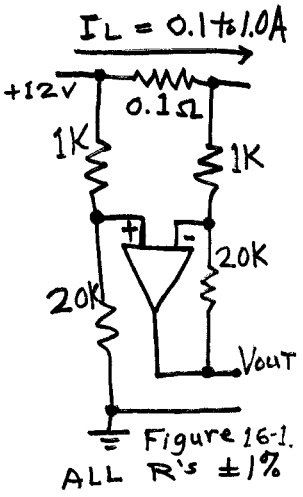

Figure 16-1: Conventional differential amplifier.

varies, the gain does not, and vice versa. We all agree that it's very helpful that you can compute what the performance will be with almost no interaction. No partial derivatives.

Now let's take a look at a couple of applications—real circuits—and their tolerances within an error budget. Here is an amplifier to magnify the $I \times R$ drop of current through a 0.1-Ω resistor and bring it back down to ground. **Figure 16-1** shows a conventional differential amplifier, with the common mode up at +12 V. The gain of −20 will bring the 1.0 A × 0.1 Ω signal down to a ground level. If the current is 0.1 A, the output will be 0.2 V, small scale. A full-scale current of 1 A will bring the output up to 2.0 V, which is suitable to send to a detector or analog-to-digital converter.

Let's select an op-amp like the LMC6482A, with low offset voltage less than 1.0 mV. (There are other versions of this amplifier with less than 0.35 mV, but let's select an intermediate model.) This 1 mV does cause 21 mV of output error. This op-amp has less than 20 pA of I_B at all temperatures, so at least that's negligible. (Bipolar op-amps might have small I_B errors, but you'd have to check it.)

Now let's see what the resistors add. Assuming that all *R*s have a 1% tolerance, the gain of (2.0 V per A) has a tolerance of ±3%. This would cause ±60 mV at full scale, but only ±6 mV at small scale (0.1 A). This might be acceptable.

Then let's consider the common-mode errors. If *R*4 has a 1% tolerance and it has 11.4 V across it, the 1% tolerance could cause a 114-mV error. By symmetry, a 1% error of each of *R*1, *R*2, *R*3 can cause another 114 mV! Added together, the common mode could cause an output error of 456 mV! That's about ±1/4 of full scale—even for small signals. That doesn't look so good to me!

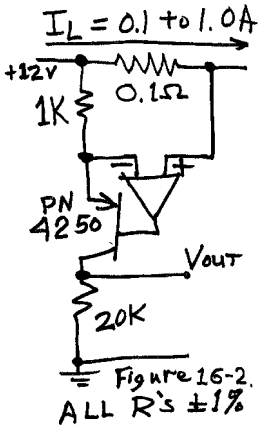

Figure 16-2: Alternative circuit.

It's true that if adjacent 1-kΩ resistors are inserted, they're likely to match within ±1/2%, so the probable error between the pair might cause ±60 mV and the ±1/2% matching between the 20 kΩ would cause another 60 mV. That added to the 21 mV from the V_{OS} would add to 141 mV.

Some textbooks teach you that you should add these errors arithmetically to 141 mV. Others point out that they could be added in an RMS way, so that 60 + 60 + 21 mV = 87 mV. Typically, this might be true. But the worst case of 141 or 456 mV might be more realistic. I mean, if you're going to build 1000 circuits and most of them are better than 141 mV, what are you going to do with the 400 circuits that are worse than 141 mV? And that's still 7% of full scale.…

You could go shopping for 0.1% resistors, but they aren't cheap. You could put in a trimpot to trim the error (to no offset error) for small signals. But as you might have noticed, a trimpot has to be properly trimmed. And if that pot is accessible, it could someday be mistrimmed and it would have to be corrected in some awkward calibration cycle. Most people want to avoid that trimpot. Before we decide that this 141 mV is unacceptable, let's look at another circuit.

Figure 16-2 shows an alternative circuit with the same gain, 2.0 V per A, using a PN4250 or 2N4250, a high-beta pnp transistor. What does the error budget look like? The same op-amp causes just 20 mV of output error. The 1% resistor tolerances cause the same gain error, 60 mV at full scale, or 6 mV at small scale. The newly added transistor adds (?1/3%) max from its alpha, or less than 7 mV, at full scale.

What is the offset error due to common-mode rejection ratio (CMRR) or due to resistor mismatch? Nothing. Zero. The transistor doesn't care about the voltage across it. There are no resistors with 12 V across them.

So the offset error is ±20 mV, due primarily to the amplifier's VOS (which could be reduced), not ±400 mV. This little circuit has greatly reduced errors compared to Figure 16-1, even if Figure 16-1 had a couple bucks of 0.1% resistors. This might be acceptable. Even the offset errors could be reduced to 7 mV by selecting the LMV841 or LMC7701.

So we have seen that circuits with similar functions can have completely different error budgets. I love to recommend amplifiers with high CMRR. But depending on cheap 1% resistors can hurt your error budget a lot more than you'd suspect.

This article was first printed in *Electronic Design Magazine*, June 8, 2006, and is reprinted with the kind permission of Penton Publishing.

CHAPTER 16

What's All This Error Budget Stuff, Anyhow?

Robert A. Pease

Well, I stated at the start of the story the reason it's important to do an error budget on even a simple circuit—and then I showed the size of trouble you can get into if you don't. I rest my case. /rap

I was just on the phone explaining to a young engineer how to do an error budget analysis on some fairly simple circuits. Later, I mentioned this while I was visiting my friend Martin, and he said he had been quite surprised when he found that many engineers in Europe were quite unfamiliar with the concept of an error budget. How can you design a good circuit without being aware of which components will hurt your accuracy?

When I was a kid engineer back in 1962, my boss George Philbrick gave me a book on differential amplifiers by Dr. R. David Middlebrook, and he asked me to do a book review. I studied the book, and it was full of hundreds of partial differential equations. If you wanted the output of a circuit with fourteen components, you could see a complete analysis of how each component would affect the output offset and gain. Each equation filled up a whole page. It did this several times.

Yet the book didn't offer any insights into what's important. I mean, is $\beta \times d(R1)$ more important than $R1 \times d(\beta)$? In retrospect, I'm glad I didn't submit any critique of that book. I would've done more harm than good. Such a mess! Even now, it would be hard to write a critique on a book that was so true but so unhelpful.

Things are much simpler now that people are mostly (but not entirely) designing with op-amps. The best thing is that the output offset and DC gain and AC gain errors are largely orthogonal. An "operational" amplifier does perform, largely, an "operation" based on what task you ask it to perform when you "program it" with *R*s and *C*s. If the offset

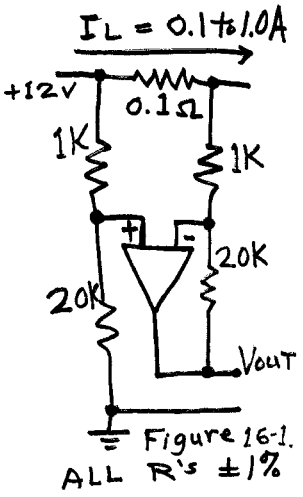

Figure 16-1: Conventional differential amplifier.

varies, the gain does not, and vice versa. We all agree that it's very helpful that you can compute what the performance will be with almost no interaction. No partial derivatives.

Now let's take a look at a couple of applications—real circuits—and their tolerances within an error budget. Here is an amplifier to magnify the $I \times R$ drop of current through a 0.1-Ω resistor and bring it back down to ground. **Figure 16-1** shows a conventional differential amplifier, with the common mode up at +12 V. The gain of −20 will bring the 1.0 A × 0.1 Ω signal down to a ground level. If the current is 0.1 A, the output will be 0.2 V, small scale. A full-scale current of 1 A will bring the output up to 2.0 V, which is suitable to send to a detector or analog-to-digital converter.

Let's select an op-amp like the LMC6482A, with low offset voltage less than 1.0 mV. (There are other versions of this amplifier with less than 0.35 mV, but let's select an intermediate model.) This 1 mV does cause 21 mV of output error. This op-amp has less than 20 pA of I_B at all temperatures, so at least that's negligible. (Bipolar op-amps might have small I_B errors, but you'd have to check it.)

Now let's see what the resistors add. Assuming that all *R*s have a 1% tolerance, the gain of (2.0 V per A) has a tolerance of ±3%. This would cause ±60 mV at full scale, but only ±6 mV at small scale (0.1 A). This might be acceptable.

Then let's consider the common-mode errors. If *R*4 has a 1% tolerance and it has 11.4 V across it, the 1% tolerance could cause a 114-mV error. By symmetry, a 1% error of each of *R*1, *R*2, *R*3 can cause another 114 mV! Added together, the common mode could cause an output error of 456 mV! That's about ±1/4 of full scale—even for small signals. That doesn't look so good to me!

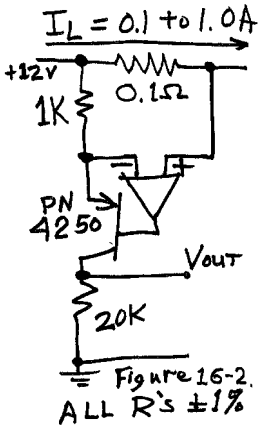

Figure 16-2: Alternative circuit.

It's true that if adjacent 1-kΩ resistors are inserted, they're likely to match within ±1/2%, so the probable error between the pair might cause ±60 mV and the ±1/2% matching between the 20 kΩ would cause another 60 mV. That added to the 21 mV from the V_{OS} would add to 141 mV.

Some textbooks teach you that you should add these errors arithmetically to 141 mV. Others point out that they could be added in an RMS way, so that 60 + 60 + 21 mV = 87 mV. Typically, this might be true. But the worst case of 141 or 456 mV might be more realistic. I mean, if you're going to build 1000 circuits and most of them are better than 141 mV, what are you going to do with the 400 circuits that are worse than 141 mV? And that's still 7% of full scale.…

You could go shopping for 0.1% resistors, but they aren't cheap. You could put in a trimpot to trim the error (to no offset error) for small signals. But as you might have noticed, a trimpot has to be properly trimmed. And if that pot is accessible, it could someday be mistrimmed and it would have to be corrected in some awkward calibration cycle. Most people want to avoid that trimpot. Before we decide that this 141 mV is unacceptable, let's look at another circuit.

Figure 16-2 shows an alternative circuit with the same gain, 2.0 V per A, using a PN4250 or 2N4250, a high-beta pnp transistor. What does the error budget look like? The same op-amp causes just 20 mV of output error. The 1% resistor tolerances cause the same gain error, 60 mV at full scale, or 6 mV at small scale. The newly added transistor adds (?1/3%) max from its alpha, or less than 7 mV, at full scale.

What is the offset error due to common-mode rejection ratio (CMRR) or due to resistor mismatch? Nothing. Zero. The transistor doesn't care about the voltage across it. There are no resistors with 12 V across them.

So the offset error is ±20 mV, due primarily to the amplifier's VOS (which could be reduced), not ±400 mV. This little circuit has greatly reduced errors compared to Figure 16-1, even if Figure 16-1 had a couple bucks of 0.1% resistors. This might be acceptable. Even the offset errors could be reduced to 7 mV by selecting the LMV841 or LMC7701.

So we have seen that circuits with similar functions can have completely different error budgets. I love to recommend amplifiers with high CMRR. But depending on cheap 1% resistors can hurt your error budget a lot more than you'd suspect.

This article was first printed in *Electronic Design Magazine*, June 8, 2006, and is reprinted with the kind permission of Penton Publishing.

CHAPTER 17

What's All This V_{BE} Stuff, Anyhow?

Robert A. Pease

Part 1

This V_{BE} topic has come up many, many times since junction transistors were introduced in the 1950s. Usually scholars like to use lots of exponential equations, and they seem to pretend that I_S is constant. Then they show the old, trite curve where Ic bends like a hockey stick at $V_{BE} = 0.6$ volts on a linear-linear scale. They pretend the transistor has no collector current below $V_{BE} = 0.4$ volts. This does not help a user or engineer understand how things change versus Ic or temperature. They ignore the way that the transistor's current shrinks exponentially, all the way down to just a few millivolts of V_{BE}, and does not magically stop below a certain "threshold." This analysis helps me a lot; how about you? Bob Widlar used graphical techniques to design transistor circuits that ran some transistors on a small number of nanoamperes or of millivolts. /rap

The other day, I was walking past the applications engineering area when I heard a grouchy debate between a couple of guys over in the corner. As they saw me walk by, they called out, "Bob, come on over here, and maybe you can solve this problem for us." I looked at their problem.

"Bob, we were trying to use the standard diode equation to compute the tempco of a transistor's V_{BE}, and it doesn't seem to make any sense." I looked at their standard equation:

$$I_C = i_S \times e^{(qV_{BE}/kt)} \qquad [17\text{-}1]$$

Yes, there was a term for temperature, t, in there, but it wasn't a very prominent term. Obviously they had tried to see how this equation responded to temperature. They were puzzled because it does *not* respond properly to temperature. It doesn't give anything *like* −2 mV/°C. I began assisting them by explaining, "When they give you this equation in school, they neglect to tell you that the i_S isn't a constant but rather a very wild function

of temperature. This function is so wild that they won't tell it to you, because it's not very useful. You can't successfully differentiate it versus temperature. So you're better off *not* having such an unusable equation."

They responded, "Okay, what are we supposed to use?" I replied, "Ah, let's do a graphical approach. Let me make up a couple of sketches." First I scribbled out **Figure 17-1**, showing the log of collector current versus V_{BE}.

I went on to explain, "That schoolbook linear plot of V_{BE} versus I_c isn't very useful, because it just shows a severe knee. I never use that one. Look at the middle line of this plot. It shows that at room temperature, the slope of the log of I_c versus V_{BE} is quite linear over seven, eight, or nine decades of current. Only at high currents does the curve bend, due to emitter resistance. And only at very small currents do you get errors due to leakages. So, in the whole midrange, you get a wide range of conformity to the slope of 60 mV/decade." The two guys agreed with what I had said.

After this, I pointed to the upper line. "At a hot temperature such as +127°C, the curve is very similar. But at a shallower slope, the millivolts per decade is *worse*, very close to 80 mV/decade. Indeed, this number of millivolts per decade is predicted by the diode equation." They further agreed that my explanations seemed correct. Plus I showed the guys that the lower sloping line is sort of like the curve for −73°C, but it's at a slope of 40 mV/decade—a rather higher gain, with a higher gm. Fine.

Also, it's possible to see that all the curves tend to converge or extrapolate to a single high point at a *very* high base-emitter voltage, perhaps +1.24 V, at a *very* high current, maybe 10,000 Amperes. Based on this outrageously high theoretical current at an absurd voltage, one could (theoretically) compute what the V_{BE} is really doing—not very accurately, or usefully.

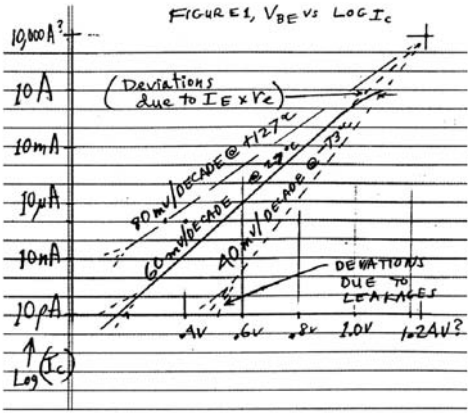

Figure 17-1

But I pointed out that this curve is just good for giving a ballpark overview of what goes on. Yes, in concept this could be used for computing the actual V_{BE} of the transistor, at various currents and temperatures. But it's too crude and too broad to be useful. What we want to use is closer to **Figure 17-2.**

I sketched away madly to get this figure, showing the plot of V_{BE} versus temperature. This illustrates the bias of transistors at various constant currents versus temperature. "*This*," I said, "is *useful*—and let me show you where and why." I stated that it was based on the real data for a real standard transistor, and it's what I use to compute biases for real precision linear circuits, such as band-gap references or temperature sensors. This and a slide rule (or a little handheld scientific calculator) lets me compute the operating points I need.

I pointed out the middle, solid, sloping line. "This line is based on some measured data. This transistor, when used in a band-gap reference, has a Magic Voltage of about 1.240 V. That's where the band-gap runs flattest. So this line is drawn in order to go through 1.240 V DC at absolute zero temperature. That's where the V_{BE} extrapolates to—if the transistor were cooled off—and that is not real data.

"The other point of calibration is where it goes through 0.640 V of V_{BE} at 10 µA at room temperature, about +27°C. That's a simple, factual, measured data point." Then one guy asked, "But why +27°C? Why not +25?" I replied that +27°C is, with an accuracy better than 0.2°C, exactly 300°Kelvin. Therefore, it makes the math much easier to work with, at +200, +300, and +400°Kelvin. They agreed.

Furthermore, I pointed out that the voltage represented by this line is just the nominal V_{BE} of the transistor versus temperature at a constant emitter current. This has a nominal

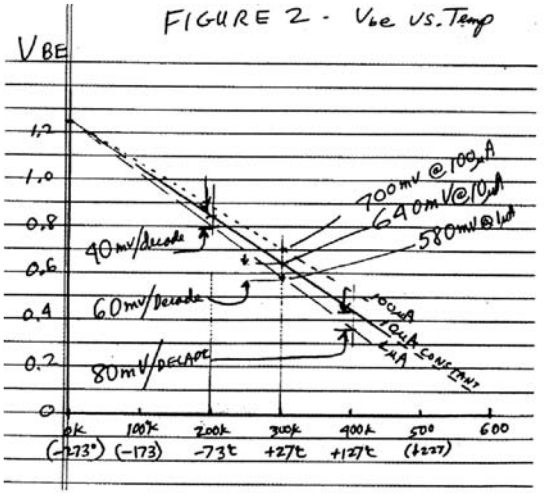

Figure 17-2

slope which is quite close to −2.00 mV/°C. *This* is a very useful thing to know—the bias at which the transistor runs at −2.00 mV/°C—because we will soon see that at many other operating currents, the tempco is *not* −2.00.

Next I stated to them that *if* the voltage between the solid, slanting line and the horizontal line at 1.240 V is studied, you can see that it's a voltage proportional to absolute temperature (VPTAT). Therefore, when we want to build a band-gap reference that's 1.240 V, all we have to do is *add* to the V_{BE} a voltage that's VPTAT. Then we can make a band-gap reference.

This is all you have to do: If you have a VPTAT that's 60 mV at room temperature, and you can amplify this with a gain of 10, you can add that onto a V_{BE} to make a band-gap reference—as Mr. Widlar proved about 30 years ago. They agreed that made sense as well.

I had to admit that the solid, sloping line appears to be nominally linear, and I drew it as more or less linear—but it's *not* truly linear. The V_{BE} curve actually is bowed downward at both hot and cold temperatures, perhaps as much as 2 to 4 mV. But for many uses, that's a negligible error, which is easy to make corrections for later.

I explained further: "Let's take a look at the upper, dotted line of Figure 17-2 (and **Figure 17-3**). This is for the transistor running at 100 µA. It, too, extrapolates back toward that point at absolute zero. This line does *not* have a slope of −2.000 mV/°C, but instead −1.800. This line isn't parallel to the other line. It's set above it by 60 mV/decade at room temp, by 80 mV at +127°C, and by 40 mV at −73°C. This difference is very accurately a VPTAT."

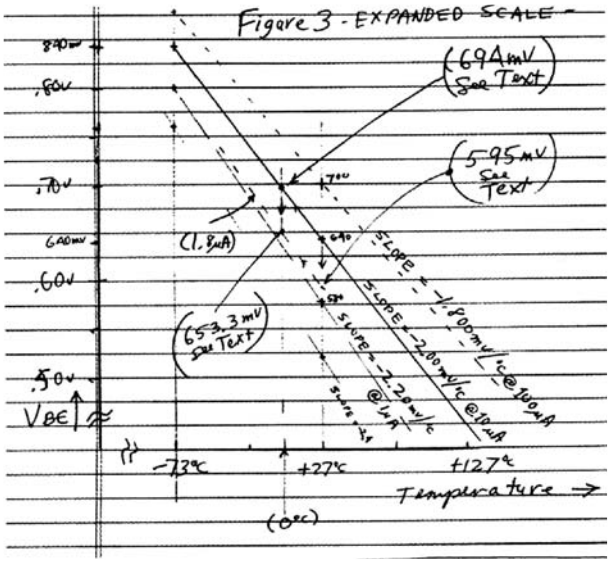

Figure 17-3

The *lower* dashed line is the line for a bias of 1 μA. It has a slope of −2.200 mV/°C. It's offset by 60 mV/decade at room temperature, more when hot and less when cold. The tiny segment of the line is at 0.1 μA and has a slope of −2.40 mV/°C (Figure 17-3).

So, all lines for V_{BE} at a constant current are all fanned out, radiating from that point at absolute zero. THIS is the curve from which it's easy to compute temperature coefficients and operating points. Now, an expanded plot, Figure 17-3, depicts only the central portion of Figure 17-2.

Let's say we want to estimate a V_{BE} at some other biases. I'll take you through some examples. The main point is, though, that you can fairly easily compute the bias for any normal situation.

Okay, we agree that we know the V_{BE} at those specified conditions: 640 mV at 10 μA and +27°C. Let's say I want to compute the V_{BE} of the transistor at the same current but at a different temperature, such as 0°C. In the example shown, the temperature coefficient of V_{BE} is −2.000 mV/°C. A shift of −27°C will cause the V_{BE} to increase by (−27) × (−2.0) = +54 mV, up to 694.0 mV. That's not very hard. For any change of temperature, at a constant bias current, simply multiply the change in temperature by the tempco of V_{BE}. But the tempco of −2.0 mV/°C only applies at 10 μA in this example. At any other current, the tempco will be different. (More on this later.)

What if we want to start from our initial conditions and move to a different current, such as 1.8 μA at +27°C? For this case, where things are at a constant temperature, you can use the diode equation:

$$IC1 = I_S \times e^{qV_{BE}/kt} 2 \qquad [17\text{-}2]$$

or its inverse:

$$V_{BE}1 - V_{BE}2 = kt/q \ \ln(IC1)/(IC2) \qquad [17\text{-}3]$$

The ratio of currents is 0.18, and the natural log of 0.18 is −1.7148. At +27°C, the factor kt/q = 26.06 mV per factor of *e*, which is the same factor as 60.0 mV/decade.

Therefore, the delta V_{BE} will be −1. 7148 × 26.06 mV, or −44.7 mV. The V_{BE} will decrease from 640 mV to (640 − 44.7) = 595.3 mV. This isn't a surprise. Any time the collector current of a transistor changes at a constant temperature, the V_{BE} changes in a nice logarithmic way. But that 26.06 mV is only at that value at +27°C. At all other temperatures, it's different, as a linear function of absolute temperature.

Another useful way to look at it is that any time you change the current by a factor of 10 at room temp (about +27°C), the V_{BE} will shift by 60 mV, up or down, as appropriate.

For many cases where decades of current are the important factor, the multiples of 60 mV make calculations simple. No computers or calculators are required.

Now let's consider the case where you want to compute the V_{BE} when both the current and the temperature are changed. There are two ways to compute this. And both of these computations had better give the same answer.

Let's say we want to compute the V_{BE} at 1.8 μA at 0°C. You could first change the temperature of the 10-μA transistor to 0° at constant current and then change the current at a constant temperature.

Let's do that: We just agreed that the V_{BE} would be 694 mV at 10 μA at 0°C. How much will V_{BE} change if we then go to 1.8 μA? At 0°C, kt/q isn't 26.06 mV, but 273/300 × 26.06 mV, or 23.712 mV, as the temperature has decreased by that factor. Therefore, as we decrease the current by a factor of 0. 18, the VBE changes by −1.7148 × 23.712 mV, or −40.7 mV, so the V_{BE} decreases to 653.3 mV.

What if we arrive at this point by the other route of first decreasing the current, *then* decreasing the temperature? We just computed that the V_{BE} at +27°C and at 1.8 μA was 595.3 mV. What is the tempco of V_{BE} at *this* current? It isn't −2.000 mV/°C, as it is at 10 μA. And, and it isn't −2.200 mV/°C, as it is at 1 μA. It's at an intermediate value. These slopes are all proportional to absolute temperature, as they intercept absolute zero at 1240 mV. So the slope of (1240 mV − 595.3 mV)/300°C is 644.7 mV/300°C, or −2.149 mV/°C. If you multiply this tempco by a −27°C change, the shift will be 58.02 mV. When you add this to 595.3 mV, the answer is 653.32 mV. So, fortunately, we get the correct answer when we compute it either way.

If you need to know the tempco of V_{BE}, it normally changes −200 μV/°C every time the current is reduced by a factor of 10. Thus while the transistor of this example had −2.000 mV/°C at 10 μA, it has −2.200 mV/°C at 1 μA, −2.400 mV/°C at 0.1 μA, −2.6 mV/°C at 10 nA, and −3.0 mV/°C at 100 PA. Although most people don't bias transistors down there, that does *not* mean that the tempco isn't surprisingly well defined down there, and it's a *lot* bigger than just −2.0 mV/°C!

What other factors should we take into account when we want to compute V_{BE}? With monolithic npn transistors, it's fairly safe to assume that the transistors' V_{BE}s are fairly well matched and predictable. We need to only take into account a difference of about 5 or 10 mV if the transistors are designed with similar geometries. That's even if no special care is taken to match them perfectly. With discrete transistors from the same batch, the matching might be similar, or it might be *poor* if the transistors came from different batches. There could be a lot of deviations, but you can't count on that.

As mentioned earlier, the curvature of V_{BE} versus temperature will cause the V_{BE} to be 1 or 2 mV smaller, at 0°C and also at +70°C, compared to the linear predictions. It could easily be 3 or 4 mV lower at −55°C or +150°C—it really is quite close to a parabolic error.

Additionally, Early Effect will normally cause a low-beta transistor (beta = 50 or 100) to run 1 or 2 mV lower in V_{BE} if the VCE is as high as 20 V rather than 0.6 V. On high-beta transistors (beta = 200 or 400), the decrease in V_{BE} may easily run 3 or 4 mV. (At another time, we can discuss the complete ramifications of this Early Effect. Suffice it to say here, transistors with high beta might have smaller *current* errors, but they tend to have correspondingly poorer *voltage* errors.)

Of course, if you run the transistor at high currents where $V = IE \times RE$ is significant, that effect can be additive (approximately) and is usually fairly linear and predictable (not to mention self-heating). If the IC or IE are small, the leakages could cause significant deviations. Also, if IE becomes quite small, some transistors could have a rapid fall-off of beta, so you cannot be sure the base current is negligible any more! And if you ever let the transistor saturate, the V_{BE} can rise or fall considerably, depending on how the transistor was made. Still, these graphical techniques can do a pretty good job of helping you to estimate the V_{BE} of a bipolar npn transistor—and of a discrete pnp, too.

Now you could write a fancy equation to compute all this, but I prefer a graphical approach. That way, I get good insights into what's going on, and I don't get fooled by computational mistakes.

Part 2

For a given V_{BE}, the collector current of a transistor is well behaved for many values of V_{CE}, extending nicely plus and minus a couple hundred millivolts above and below V_{BE} = 0.7 volts. But "The Technical Books" all talk about the "forward active region" that only has V_{CE} greater than V_{BE}. Not so! This essay just points out the obvious fact that this region extends quite a good way down from $V_{CE} = V_{BE}$, down well below 0.7 volts, down toward V_{CE} = 0.25 volts. It also corrects a couple other myths. /rap

I've been debating with a guy who argues that a transistor won't work as a transistor unless its V_{CE} is bigger than its V_{BE} (see **Figure 17-4**). He keeps reading this in books. Also, he points out that if the base and collector are nominally tied together to make a diode, you might think that it's okay. But actually, he says, the $I \times R$ in the collector path makes the V_{CE} lower than the base voltage, so it won't work. Well, I've been looking in some of those books, and they sometimes do say that. But when they do, they're wrong.

When a transistor's V_{CE} is slightly less than its V_{BE}, it keeps right on working like a transistor. Can I prove this? Sure. Look in the NSC linear Databook at circuits such as

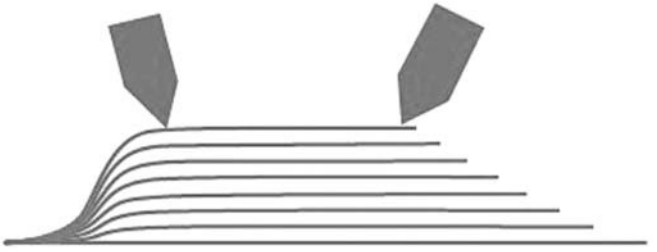

The upper curve represents a 2N3904 biased up with I_B = 7 μA, V_{BE} = 700 mV, and I_C = 1.05 mA. As V_{CE} decreases from 0.7 V (at the right-hand arrow) to V_{CE} = 0.25 V (at the left-hand arrow), this transistor obviously is acting like a transistor in its active region, even though V_{CE} is well below V_{BE}. Of course!

Figure 17-4

the LM10. The LM10 wouldn't work on a 1.1-V power supply if the transistors aren't working well with V_{CE} as low as 350 or 250 or even 150 mV, which is far below V_{BE}. Of course, you have to be a good engineer to make these circuits work well.

NSC guys (like Bob Widlar) have been doing this for 40 years. Look at the V_{CE} curves of any transistor. When V_{CE} falls below V_{BE}, it's not a disaster. Put a transistor on a curve tracer. Apply a bias like 1 μA per step to the base. When you change the V_{CE} from +1.0 V to 0.6 V to 0.5 or 0.4 V, I_C doesn't change much, does it?

Okay, maybe when you get V_{CE} down to 0.35 V the gain starts to degrade some. But above that, at room temperature, it's not a big deal. There is no demarcation between $V_{CE} > V_{BE}$ and $V_{CE} < V_{BE}$. No inflection. The beta doesn't even change more than perhaps 2% per volt, and it does so smoothly.

Now run the temperature up to 125°C. Can you design a circuit that works up there? It's not easy. But if you don't need a lot of swing, some specialized circuits work just fine. Look at the LM4041-1.2 or the LM185. Many of their V_{CE}s are about 0.3 V, yet they work hot and cold.

How about 160°C? How about 260°C? I can't, but Widlar could, and did, in the LM12. After all, in the old days, a pentode could run with a very low V_{PLATE}—much lower than V_{SCREEN}. It's hard to comprehend this, but after a while you get to understand and believe it. It's an analogous situation that the output voltage is so low, you can't believe it will work. But it's true. It does work.

Next Topic

When V_{BE} = approximately zero, changes in V_{BE} certainly have no effect on IC, right? Wrong. In Widlar's LM12, some of the transistors are so biased, when $V_{BE} = 0$, the

V_{BE} can still influence the collector current (and vice versa) whether the V_{BE} is a few millivolts positive or negative.

Admittedly, you can't see this easily in a silicon transistor at room temperature. But you can see this in a silicon transistor at 220°C or in a germanium transistor at room temperature, which is about the same idea. Go ahead and measure it. When I did, I was impressed by Bob Widlar's brilliance.

Also, the beta of a transistor can still be important, even when V_{BE} is about zero. That's because as V_{BE} moves up and down a few millivolts compared to zero, the base current needed might be small but finite—not negligible. The base current and its changes are necessary. And if you start at $I_B = 0$ and pull the base negative, the collector current can decrease.

I must remind you that high-beta transistors (300 and up) still have disadvantages in terms of voltage gain or mu. When the beta gets too high, and because mu is inversely proportional to beta, the voltage gain is hurt. I remember a test that asked how much voltage gain a particular amplifier design has. The answer was supposed to be 20,000. But the gain was really 9000, as the betas were too high and the Early Effect was too strong. I passed the test after I explained my solution.

As a rule of thumb, I use mu $\times$ beta $=$ 2 million. On some devices, that product is only as good as 1 million, or even 4 million on LM194. If the beta gets better, the mu $= 1/h_{RB}$ gets worse and the voltage gain suffers. Be careful not to allow in transistors with too high beta in circuits where poor mu could cause poor performance. Beta is often important. Too little of it can do harm; so can TOO MUCH.

Part 1: This article was printed in *Electronic Design's Analog Applications Supplement*, June 26, 2000, and is reprinted with the kind permission of Penton Publishing.

Part 2: This article was printed in *Electronic Design Magazine*, June 21, 2007, and is reprinted with the kind permission of Penton Publishing.

Technology Edge

Power Supply Design and the *Thinking* Engineer

Wanda Garrett, Technical Communications Manager
Robert A. Pease, Staff Scientist

Over the last 15 years, the design of switching power supplies has entered the mainstream. Now, no longer solely the responsibility of an expert, almost anyone involved in system design may end up putting together the switching power supply. Fortunately, people of all levels of expertise can use National's WEBENCH® online tools to design switching regulators. Because experts usually want to tailor their designs for specific components or targeted performance, these tools allow a great deal of customization of the designs. This means it is possible to take a design that works perfectly well and "optimize" it into oscillation. The online tools can be used to discover design problems and correct them - as long as thinking is applied as well.

Taking the First Step

The first step in online power supply design is defining the power supply requirements, including voltage range, output voltage(s) and load current(s). Possible solutions are automatically evaluated and one or two recommendations are presented to the user. This is the first place a designer can get in trouble: if the requirements are not correctly stated (for example, if the input voltage range is really higher or lower than the value entered), an unsuitable solution will be shown. It is easy to try several sets of requirements, but the user must have a clear understanding of the system needs. After a regulator-based solution is selected, components are identified for the circuit. Wherever possible, real components with part numbers are shown. The user has the option of changing the components to predefined alternates or to enter a custom component. There are guidelines given for the component values and any critical parasitic values. If a custom component is used that is quite different from the recommendations, be prepared for poorer performance.

Evaluating Performance

Once circuit components are selected, it is time to evaluate the performance. In general, look for frequency response values (crossover frequency and phase margin), peak currents and voltages, and thermal values (efficiency, junction temperature, and component temperatures). Even though these calculations are based on models, the results match well with the bench data.

Electrical and Thermal Simulation

Some solutions are supported by electrical simulation. These simulators will display the schematic and allow the user to further change components and run tests on the regulator circuit. Available tests include Bode plot, steady-state, line transient, load transient, and startup. (Note that Bode plots can be run only for those circuits using fixed-frequency regulator ICs.) For the online tests to be useful, it is critical that the user check all of the conditions of the test. Input voltage and load current may be varied for each test, and the defaults may not match the situation in the user's system. The user must think about what results should be expected and if the simulated results are different, to figure out why. Thermal simulation can be run for many solutions. The online tools evaluate the regulator circuit as implemented on a PCB using a reference design layout. Results showing component and board temperature are presented in a full color image as well as in a table. Because the thermal simulation is designed to run fairly quickly (giving results within a few minutes), the accuracy is not as high as it would be if a detailed CFD (computational fluid dynamics) multi-hour simulation were run. However, the temperature estimates are typically within 20°C of actual. This is close enough to identify if there are any hot spots on the board or components in danger of exceeding temperature ratings.

Testing a Protoype

The final step in the switching regulator design process before preparing for production is to build a prototype to test on the bench. Some solutions are supported with customizable evaluation boards, while other solutions have reference design boards available. With the power of the online tools, it may be tempting to skip this step - don't! Most designs will work fine, but some need careful layout for best performance. The real components may not exactly match the simulations, especially considering their parasitics, so the actual performance (including the effects of board layout) will be slightly different from the simulated results. Again, the key is to determine what is expected and to analyze any differences. Using lab tools including a good, high-speed oscilloscope, a DVM, a current probe, and thinking will give the designer the best chance of success.

National Semiconductor

THE BEST OF BOB PEASE

The Design of Band-Gap Reference Circuits: Trials and Tribulations

ABSTRACT

This tutorial will briefly discuss the designs of various band-gap references, with an emphasis on technical problems that caused serious troubles. Practical solutions are shown for problems, and a methodology is shown for solving problems.

TUTORIAL

The band-gap reference has been a popular analog circuit for many years. In 1971, Robert Widlar introduced the LM113, the first band-gap reference.[1] It used conventional junction-isolated bipolar-IC technology to make a stable low-voltage (1.220 V) reference. This type of reference became popular as a stable voltage reference for low-voltage circuits, such as in 5-volt data acquisition systems where zener diodes are not suitable. Band-gaps are also used in digital ICs such as ECL, to provide a local bias that is not adversely affected by ambient noises or transients.

The principle of the band-gap circuit is well known and will be mentioned here in the briefest terms. The circuit relies on two groups of transistors running at different emitter current densities. The rich transistor will typically run at 10 times the density of the lean ones, and a factor of 10 will cause a 60 millivolt delta between the base-emitter voltages of the two groups. This delta voltage is usually amplified by a factor of about 10 and added to a Vbe voltage. The total of these two voltages adds up to 1.25 volts, typically, and that is approximately the band-gap of silicon.

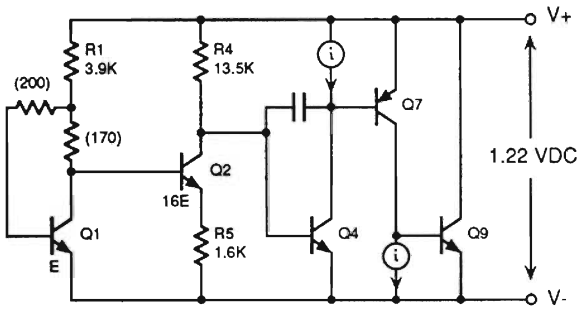

FIGURE 1
SCHEMATIC DIAGRAM, LM113 (SIMPLIFIED)

In figure 1, the LM113 schematic diagram shows a basic band-gap circuit. Q1 runs at a relatively high density, about 150 microamperes per square mil. Q2 is operated at a low density, about 10 microamperes per square mil, and so its Vbe is much less, about 70 millivolts. Now, Let's ASSUME that the circuit is at balance and the output is near 1.22 volts. Then the 70 millivolts across R5 is magnified by the ratio of R4 to R5, about 8:1, up to a level of 600 millivolts. This voltage is added to the Vbe of Q4 (about 620 millivolts at room temperature) to make a total of 1.22 volts, as required. Q4 then amplifies the error signal through Q7 and Q9, which provide enough gain to hold the V+ bus at 1.22 volts. The beauty of the band-gap reference is the summation of the Vbe term, which decreases at the rate of about -2 millivolts /°C, and the (delta-Vbe term) which grows at about + 2 millivolts /°C, to achieve an overall Temperature Coefficient (Tempco) that is substantially zero. All band-gaps employ this summation of a growing and a shrinking voltage, to make a stable low-tempco voltage. Further, it has been shown[2] that when a circuit has been trimmed to the correct voltage, the correct tempco will follow, despite process variations in parameters such as Vbe, beta, sheet resistivity, etc. Consequently, band-gap circuits are often trimmed to their ideal voltage so as to provide also a low tempco.

There are many other circuits that have been used for band-gap references, and each one has its own set of advantages and disadvantages. Figure 2 shows a simple brute-force scheme: the stack of DI-D7 run at a rich current, and D8-DI3 at a lean current, and the resultant output of 1.2 v has a low tempco. Unfortunately, this circuit is not very suitable for operation on low supply voltages, but its concept is clear!

FIGURE 2
SCHEMATIC DIAGRAM
of SIMPLE BANDGAP

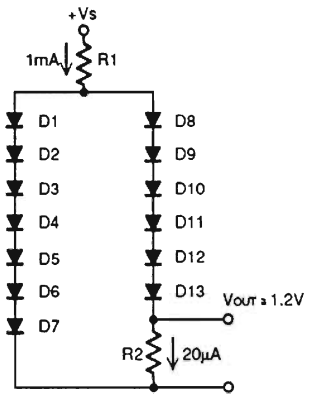

The Brokaw cell (Figure 3), attributed to Paul Brokaw,[3] is often used for output voltages larger than 1.2 volts, as its Vout can be scaled by the ratio of two resistors, R3 and R4. A similar circuit is used in the LM117, an adjustable medium-power voltage regulator, which can be trimmed by two external resistors to any voltage in the range 1.25 to 57 volts. (Figure 4)

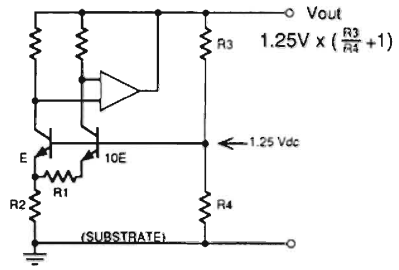

FIGURE 3
SCHEMATIC DIAGRAM, AD580 (SIMPLIFIED)

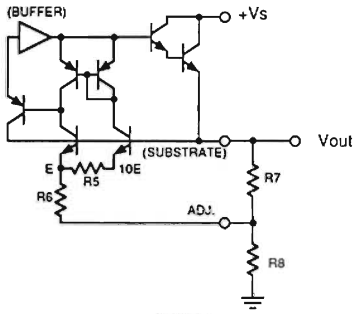

FIGURE 4
SCHEMATIC DIAGRAM, LM117 (SIMPLIFIED)

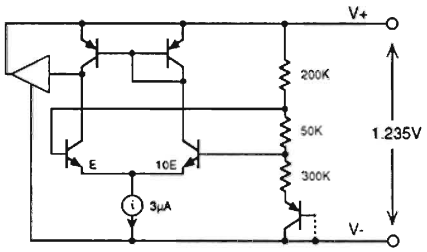

FIGURE 5
SCHEMATIC DIAGRAM, LM185 (SIMPLIFIED)

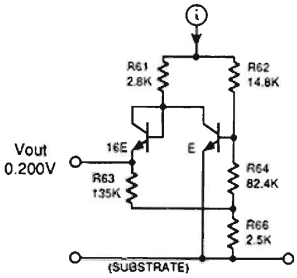

FIGURE 6
SCHEMATIC DIAGRAM, LM10 (SIMPLIFIED)

A different approach is shown in Figure 5, representing the LM136. The version shown is suitable for operation as a 1.2-volt shunt regulator. The circuit of Figure 6 is for the reference section of the LM10, which provides a reference voltage of 0.200 volts and operates on a supply voltage as low as 1.0 volts.

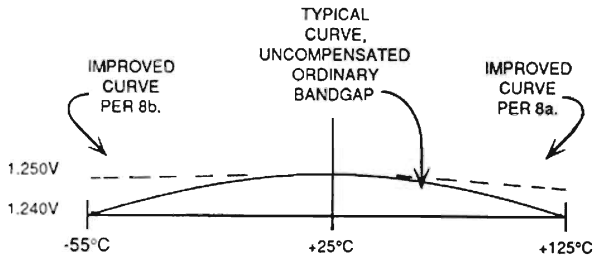

FIGURE 7
V_{REF} VS. TEMPERATURE
BANDGAP REFERENCE

One of the major drawbacks of all these simple band-gaps is the curvature of their tempco. Ordinarily, all band-gaps have a negative tempco when hot, and a positive one when cold. See figure 7. Around room temperature, tempcos as good as 20 or 30 ppm/°C are typically attainable, but over a wide range, -1% shift is normal at hot and cold temperatures. Several techniques have been devised to neutralize this curvature. Figure 8a shows a little add-on circuit which can do a surprisingly effective improvement at warm temperatures. Sometimes 2 or more transistors are connected with slightly different biasses, for improved curvature correction. A circuit which can improve the tempco curvature at cold temperatures is shown in figure 8b.

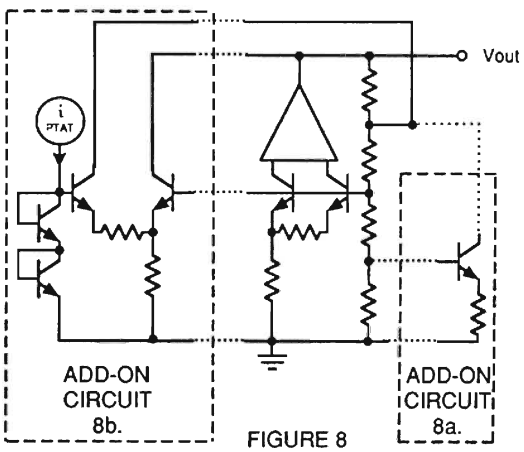

FIGURE 8
SIMPLE BANDGAP SHOWING
CURVATURE COMPENSATION

Other techniques rely on smoothly nonlinear techniques to cancel out the parabolic curvature. One example is shown in Figure 9. It is able to make an improvement of about 8:1. Other proprietary techniques and circuits are also used.

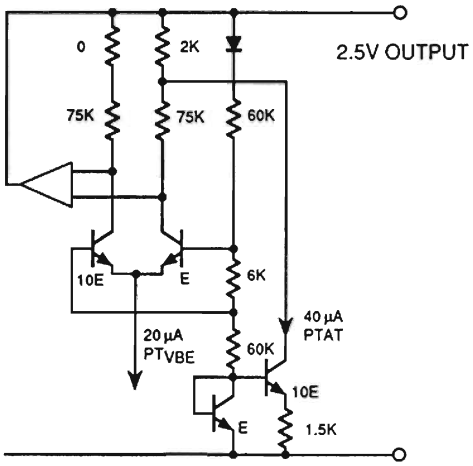

FIGURE 9
SMOOTH CURVATURE-CORRECTION CIRCUIT

Band-gap references are not normally thought of as low-noise devices, because of the high gains needed. Every 1.25 volt of output is made up of 0.6 volts of Vbe, plus a 60 mv signal magnified by a gain of 10. Consequently a 5-volt band-gap includes a 60 millivolt signal amplified by a gain of about 40. If you allow for a transistor running at 1 µA to have a theoretical amount of noise equal to 14 nv per square-root Hz, the output will have at least 800 nv per square-root Hz, and for a bandwidth of 10 kHz, about 80 µv rms, or 500 µv p-p. Obviously, a band-gap reference with 100 ppm of noise will not give good results with a 12-bit DAC. So although some band-gaps may claim advantages of low-power operation of only 11 µA total, for high-performance designs the circuits are more popular when biassed with 30 or 60 µA or more, for each group of emitters.

Many systems-on-a-chip depend on a bandgap to start up and bias all other circuits. Thus, the band-gap must be sure to start promptly under all conditions. It is not trivial to do this, as a band-gap that has not started may not have any bias currents to make sure it does start. Even the leakage of a junction capacitor may be sufficient to prevent the circuit from starting. A good solid sure-start circuit is advisable. If I can think of one, I'll present it. Figure 10, reserved.

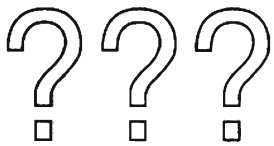

FIGURE 10
SURE-START CIRCUIT for BANDGAP REFERENCE

There are many other ways to design a band-gap reference that does not work well. There are two basic choices: with a computer, and without a computer.

If you rely on breadboards, you may trim to get a good tempco but the resistors may not be matched in their tempco if you just use RN55D or RN55Cs. If you actually get film resistors with a matched tempco, the tempco of your circuit may still not agree with your final circuit, due to the gross tempco of the actual diffused (or film) resistors in your circuit. Example, diffused resistors have 1600 ppm per °C; implanted silicon resistors are even steeper. Some SiChrome resistors can have +300 ppm/°C, and they will give a completely different tempco than 50 ppm/°C. Also, diffused resistors sometimes have a different tempco depending on how many volts they are biassed below their tub's voltage, so the tempco of a voltage divider may not be as good as theoretical, unless each resistor sees about the same amount of tub bias; this can be done (in theory, if you have enough space) by giving each resistor its own tub.

The dynamic stability of a band-gap IC is often inferior to its breadboard due to stray capacitances in the breadboard. Consequently the use of SPICE and similar computer simulation schemes are becoming popular, as the capacitances can be well defined without false influences or strays. However, as with any other computer usage, a "sanity-check" is recommended to insure rational results in a known situation, before the computer can be trusted in unknown cases.

However, there are many ways to get false results in SPICE and other simulation schemes. Most transistor models do not accurately model the shape of the curve of Vbe vs. temperature. It is sometimes possible to tweak the characteristics of the model of a transistor until the tempco of the breadboard or analog model matches that of the computer model. However, after you have done this, it is not safe to assume that reasonable changes in the operating points will cause reasonable changes in the tempco. The actual changes may be different from the computed changes, even for a minor tweak. For example, a minor change of a resistor, which actually gives a change of +10 ppm per °C may be predicted by the computer to give a change of +5, with no plausible reason for the discrepancy. Often, a SPICE run will simply fail to converge at cold temperatures. If you take the observed operating points from a run at -34°C and insert them as the node-set voltages for a run at -35°C, don't be surprised if you still get no convergence. In a circuit, being close to the right answer helps, but in SPICE, convergence depends on the matrix yielding a valid answer, and that has almost nothing to do with the circuit or reality. For example, on a recent SPICE run, I had a resistor connected only to ground and to a capacitor which also went to ground. If I used an * to comment-out the resistor and capacitor, the circuit refused to converge, but if I put them back into the circuit, they somehow helped the matrix converge. So, in the future we may be able to insert useless meaningless resistors around the circuit, whose sole function is to help the matrix give good convergence.

Other practical examples will be given of how to make a band-gap work badly, and how to give it a chance of working well.

Advice to the Engineer

While it is not practical or suitable to show examples of How to Design a Good Band-gap Reference, it is possible to show examples of circuits that do not work well. By avoiding the thicket of bad design practices, a good design may be seen to be feasible. A list of examples of bad design will be given here, grouped in categories.

Layout Problems

- Bad cross-coupling. As in a good op-amp, the critical transistors should be laid out with cross-coupling -- example, 1/2 of Q1 on one side of Q2 and the other half on the other side of Q2, so the centroid of Q1 and of Q2 will be at the same place. Likewise, critical resistors should be arranged so that 1/2 of R1 is on one side of R2, and the other half on the other side -- again, common centroid. This is especially important in a power regulator where large temperature gradients can be expected. Also it is very advantageous to reject gradients in sheet rho, beta, etc. In a typical band-gap, this cross-coupling should be done for the delta-Vbe transistors and also for the PNP transistors that serve as their collector loads, unless you can prove it to be unnecessary.
- Layout of Delta-Vbe circuit vs. Vbe. In some circuits, the transistors that form the delta-Vbe are not the same ones that make the Vbe. These must be laid out to reject gradients from all expected directions.
- Thermal regulation errors. When a band-gap regulator or reference has a step change of dissipation, the thermal gradients across the die may cause significant errors. A good layout with adequate cross-coupling and attention to detail is normally required to keep these errors to acceptable levels. A good power regulator can do 0.005%/W; a good precision reference can do 0.002%/W. Thoughtful layout techniques as mentioned above are especially important for a library cell that will be used in an ASIC, because when it is used, you can never guess where the thermal gradients will come from; you have to lay the circuit out to reject gradients from all directions.
- Rejection of I x R drops in power busses. If some parts of a circuit are connected to a power bus at one point, and other parts at another point, significant errors can be caused when current flows through the bus. In some cases, you can specify that no such current can flow through the bus; in cases where the bus current must be expected to fluctuate, the connections to the power bus must be arranged to reject the I x R drops.
- Thermal stress in corners of the die. When a precision circuit is laid out far off the center-line of the die, near a corner, the thermal stresses can cause errors. For best results, avoid putting precision circuits in corners or far off-axis. This applies to ASIC cells, of course.

Start-up Circuit Problems

- Bad start-up. Normally applies only to series-mode circuits, not shunt-mode.
- Start-up circuit too weak. This often happens when the start-up current (high-value resistor or EPI-FET) puts out too little current. This problem is often exacerbated by

high temperatures, leaky diodes, leaky capacitors, or excessive substrate currents if one of the terminals is pulled below the substrate. You can never absolutely be free of this, but you can avoid it if you have a good start-up test. Do not expect the computer to be of much help.
- Start-up circuit too strong. This can happen when the EPI-FET puts out too much current and overwhelms the start-up circuit. It's not easy to model this in SPICE, but you can think about this as a worst-case to be avoided.
- Dynamic start-up with no dc start-up. The circuit can start on the dv/dt of the input, but may not start if dv/dt is small. Use a start-up test.
- Start-up too slow. You may avoid this by good worst-case design and good modelling in SPICE or breadboarding.

Oscillations

- Oscillation due to capacitive load. You avoid this by good circuit design. Sometimes the condition can be helped by pre-load currents, or by a series R-C damper network to ground. Use Pease's Principle to make sure that ringing is not lurking nearby which will turn into oscillation when you turn your back. Check for ringing at all relevant temperatures.
- Oscillations at some temperatures and not others. Check for this by watching the Vout for ringing as the part's temperature is SWEPT from one extreme to the other. Note, monitoring the dc output voltage is not necessarily sufficient to insure freedom from oscillation.
- Oscillations due to improper start-up circuit. Make sure the start-up circuit is well-designed and well-behaved in all worst cases.
- Oscillations because the breadboard had enough strays but the IC does not. This is a matter of good modelling. The breadboard can be expected to lie about this. SPICE may be helpful if applied thoughtfully.
- Obscure oscillations. The computer often lies badly about these. It may refuse to admit that they happen, and refuse to show them happening.

DC Output Voltage Errors (Room temp)

- Excessively broad distribution of Vref. When you expect a tolerance of ±3% and the observed distribution is excessive, the problem is usually either badly-matched resistors or transistors. The geometries must be identical as drawn and as masked. If your small resistor is short, your large resistor should be made of a group of short resistors. Sometimes it is advantageous to draw the Rs and Qs as cells, so that the mask-making process acts identically on each cell.
- Parts cannot be trimmed. Make sure that every voltage can be trimmed by at least one combination of trims; avoid any possible "holes" in your trim scheme.
- Interaction of trims. It is a good idea to make sure that in your plan, the size of each trim is (substantially) invariant of whether any of the previous trims have been done.
- Dependence of pre-trim Vref on beta. In general, a higher beta transistor has a lower Vbe. In some designs, a pinch resistor (whose resistance is a linear function of beta) is used to compensate for the shift of Vbe and improve the room-temp accuracy. In other cases, a pinch resistor is used to compensate for tempco, even as it degrades the room-temp accuracy.
- Band-gap "narrowing". With high-speed processes, the band-gap voltage (and the voltage for zero tempco) is decreased vs. ordinary transistors. A good breadboard can help determine the right place to operate. The computer cannot.

Tempco Errors

- Wrong "V-magic". Normally there is one value of Vref, which (if you trim to that voltage) gives the best tempco. But there may be circuit problems, masking problems, or process problems that can cause tempco to vary even when the Vref is well trimmed. The deviations are typically + or - 5 or 10 ppm/°C, but in a poor circuit can be considerably worse.
- Inability to trim to Vmagic. If there are "holes" in your trim scheme and you cannot get the Vref to trim to Vmagic with good precision, the tempco will be degraded by about 3 µv/°C per millivolt of error.
- Tempco curvature-correction circuit does not work well? MOST tempco curvature correction circuits do not work well the first time. Some are still bad after several tries. Computer models are of little help, or usually of much harm. Breadboards often are no better. Trial-and-error is usually needed.
- Leakages cause high-temp errors. You have to watch out for device leakages, tub leakages to substrate, and capacitor leakages. These are usually hard to model.
- Tempco does not agree with breadboard. If you used discrete low-tempco resistors in your breadboard instead of the monolithic (diffused or implanted) resistors, you can expect bad results.

Computer Modelling Problems

- Bad model of the Vbe versus temp. Most transistor models are poor. It is possible to tweak various parameters to get decent match of Vbe vs. temp, but the derivatives may not be realistic.
- Mis-typed data input. Be extremely meticulous about typing errors and copying errors. Strange results may occur otherwise.
- Failure to converge at cold temperatures. Yes, that is a problem...
- Failure to converge despite accurate "Node-set". Sometimes if Nodeset is TOO accurate, convergence is worse than if Nodeset is approximate.
- Failure to converge due to 35 volts across a diode (Vf) which does not conduct any current. Maybe you need a bigger diode...
- False current due to dv/dt across a hidden capacitance. When the collector voltage has a dv/dt, the "collector current" can show the current through the c-b junction, but the current through the c-substrate junction may not be included.
- False dv/dt due to .PLOT command truncating the .TRAN command. This has been observed to happen, when the end of the .PLOT is not at the same time as the end of .TRAN.

Miscellaneous Problems

- Low-voltage-lock-out problems. YOU must engineer carefully to get good results at low voltages, as the system may require good behaviour at very low supply voltages, and the reference may not want to give it. The breadboard works better than SPICE, here. Check at all temperatures.
- Thermal limit circuit works badly, poor errors, soft knee. Try to use a circuit that has been successful in the past. Hysteresis is often a good feature.
- High noise due to starvation. Be sure to check the breadboard. (Unless you are sure the computer gives good answers.)
- High noise due to high resistance. Check the breadboard carefully, and use realistic transistor samples. High-beta parts will have higher rbb' than low-beta ones.
- Bad matching due to buried layer. If the actual buried layer causes crystal growth to fall in the middle of a critical transistor, Vbe matching may suffer. Note, the crystal growth at the surface is shifted from where the buried layer appears to be.
- Saturation due to insufficient buried layer. Transistors do not run well at warm temperatures when asked to run near saturation, especially at high temp, when the buried layer has been omitted.
- Assembly shift. This is usually caused by stress sensitivity. A good layout can help minimize this.

Czardom

At National, we appointed a Czar to oversee the design of all band-gap circuits, and to monitor and to log all good and bad results. This has helped cut down on the number of repetitive foolish errors.

Bibliography

1. LM113 Data Sheet, National Semiconductor Linear Data Book, 1972, 1976.
2. R. J. Widlar, "An Exact Expression for the Thermal Variation of the Emitter-Base Voltage of Bipolar Transistors", Proc. IEEE, Jan. 1967.
3. AD580 Data Sheet, Analog Devices, Inc., 1975.

Originally published in the IEEE Proceedings of the 1990 Bipolar Circuits and Technology Meeting, September 17-18 1990 in Minneapolis, Minnesota.

Robert A. Pease
Mail Stop D2597A
National Semiconductor
P.O. Box 58090
Santa Clara, CA 95052-8090

National Semiconductor

THE BEST OF BOB PEASE

What's All This Neatness Stuff, Anyhow?

Once upon a time, there was a rapidly converging conflict: My boss thought my office was getting messier and messier, and he wanted me to make it neater. Now, this was just a year or so after my desk had won a $500 prize for being the "Ugliest Desk in Northern California." So I guess he thought he was justified in pressuring me to clean up my act. He solved that problem by making it one of my goals to get my office to an acceptable (whatever that meant) level of neatness. Well, we never found out what that meant. Every time he would ask me how I was coming on the neatness campaign, I would tell him all of the other things I was doing to help our customers.

What if I came in on a Saturday with good intentions of neatening up some of my office, and the phone rang. Should I tell the customer, "No, I won't help you, I have more important things to do"?

So every year he would mark me down points in my review for not fulfilling my goals. He finally got so discouraged that he left the company. The poor guy. He just wasn't devious enough! He could have waited until the next earthquake and told one of the guys to knock over a couple of my piles of papers. Then he could then explain that I had to get it at least to a reasonable level of *safety*. But he never figured that out, and I didn't tell him until after he left.

Some people keep their desk neat because that's what feels good to them. I find that neatness is not a priority compared to a number of other things, such as answering the phone when a guy needs help, or volunteering advice when a customer has a problem. Some people find it easy to keep a neat desk because they throw out things that make it look messy. I just don't operate that way.

One time I was working on a Saturday after being at National just a few months. My desk was already stack up pretty high. Another guy was at his desk, which had just a few dozen things on it. He was picking them up, one by one, studying them, and then throwing most of them in the wastebasket. I commented, "You sure do keep your desk neat." He said, "Yeah, if I find something I don't need, I throw it out." I said, "Doesn't your wife ever get nervous?" He replied, "It's my third wife ..." No, I don't operate that way.

One day, an engineer stepped gingerly into the entry way of my office and asked, "Bob, do you have a Siliconix catalog?" I replied, "Sean ... you're standing on it." He looked down and, indeed, he was. He was impressed. But I knew right where it was, because I had recently tossed it over by the doorway so I could then put it in the bookcase by the door. Sean just happened to walk in before I put it in the bookcase.

More recently, I inherited a couple of filing cabinets and a huge 7 ft. x 3 ft. x 7 ft. cabinet from a Fairchild laboratory. Our secretary explained that I would have to junk it unless I could find a use for it. I said, "Well, I could always put it in my office." She looked at this huge ark and said, aghast, "No, you couldn't do *that*."

I thought about it. I got a yardstick, and I figured out that, with an inch to spare, I *could* do that. My technician and I spent nearly all morning reassembling that cabinet and easing it into the corner of my office. I put about 1/3 of a million cubic inches of my paperwork into that, and into the other file cabinets, and improved the appearance of my office so much that our senior secretary admitted that I qualified for an "Enviros Award." In the past, the various departments would vie to achieve cleaner clean rooms and higher-yield fab lines by having better cleanliness. A whole department of 20 or 30 people would work real hard to cut down the number of *particles* in their area and win an Enviros Award. But I got my Enviros Award single-handedly. I hate to guess how many particles I straightened up.

Right now, my office seems to be in the getting-messier-again phase. When I have to review a mask set, with precision down to the last tenth of a micron, I get my head in the right mood to do that. And when I'm done, in sheer rebellion I guess, I abandon the neatness for a while. I save what seems to me to be of value. Often that includes documents and papers and notes that other people would think aren't very valuable - until they come to see me years later, hoping I might have the information they need. Often I do. Go ahead, call me *retentive*. See if I care.

Now that the NBS has changed its name to the "NIST" or "National Institute of Standards and Technology," I have figured out the next way to enhance the neatness of my office. I'm going to buy a big dresser with 6 big drawers and a mirror and everything. I'm going to put it right at the entrance of my office, and put our ultra-precision resistors and capacitors in those dresser drawers. And I'm going to call it "The National Bureau of Standards."

Comments invited! / RAP
Robert A. Pease / Engineer

Originally published in Electronic Design,
October 25, 1990.

RAP Update: *This was the fourth column I wrote, the fourth to be published. And suddenly I began to get some Fan Mail. A lot of people said they had desks that were pretty close behind mine in sheer messiness. My desk still is a Federal Disaster Area. I tried to put a recent picture of my desk in this web page, but it isn't even recognizable as a desk. Is it?*

National Semiconductor

THE BEST OF BOB PEASE

What's All This Splicing Stuff, Anyhow?

Several months ago, a reader wrote in to one of the local newspapers: "If I want to move my speakers a few feet further from my amplifiers, can I splice in a few more feet of speaker cable, or should I buy all new cable? My brother-in-law claims that splicing would hamper the sound." The resident expert at the paper stated that the brother-in-law was wise, as the spliced wire would give inferior audio results.

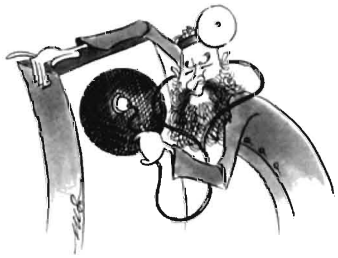

I promptly wrote in to the resident expert, asking him on what basis he could say this. Was he claiming that *he* could *hear* the difference? I demanded that he show us readers how the spliced wire could possibly make any difference. I challenged him to listen to any music, under any audio conditions, and I would swap in various pieces of speaker wire (enclosed in boxes, on a double-blind basis) that had 0 or 1 or 2 or 6 or 12 splices. How, short of clairvoyance, could he tell which wire had the splices, using ordinary audio-frequency signals? Of course, if you used an impedance analyzer with a bandwidth of several gigahertz, you could "see" some of the splices. But, for good high-fidelity audio, there's no way you could discern this, especially as a splice may make the wire's impedance lower or higher or unchanged.

The expert, with his "golden ears" and all, never wrote back. So, I sent my criticism to one of the local skeptics' groups called "BASIS," the Bay Area Skeptics Information Sheet.

They edited it lightly, and in their newsletter, they printed my complaint, which amounted to this: If a person claims to talk to the dead, or summon spirits, or show extrasensory perception, then we must apply some skepticism so as not to encourage gullible persons to invest their money in these hoaxes.

But if a person who is endorsed as the "high-fidelity expert" says that you can hear the difference between spliced and unspliced wires, then we, as technical people, have an obligation to express our doubts and our skepticism. Why should a hi-fi salesman be able to sell a bright-eyed yuppie a $50 hank of speaker wire, (or $100 or $200 or $400 or more, which is where the really high-end speaker wire is priced these days - believe it or not) just because an "expert" says it's *better* to buy new wires rather than splice on a few extra feet? Obviously, ethics in technical electronics and science is involved here.

Many hi-fi experts, with their "golden ears," claim that they can hear differences in sophisticated speakers, expensive amplifiers, or just fancy wires, that I can't possibly discern or detect. It might take many thousands of dollars to just buy the equipment and duplicate the experiment. And, their ears might be correct - much more discerning than mine, more than I could imagine. But, when the "expert" talks about wire and splices, then I find myself compelled to comment and raise doubts. There are some experiments that even I can propose and that I could conduct, that would be decisive, if the "expert" did not duck the challenge.

Now, there are many persons who have *golden* ears and will claim that they can easily distinguish between *good, better,* and *best-quality* speaker cables. However, when these persons are invited to a double-blind test, they usually have a strong tendency to demur. Some people like to call this *the shyness factor.* Other people liken this to the tendency of cockroaches to scuttle into a dark corner when the lights are turned on.

I was only slightly concerned about how to conduct the test, because to do a fair test, you might have to change back and forth from, say, speaker wire #1 to speaker wire #2 or #6. If you do that with screwdrivers and pliers, it might take a long time to make the changes; a critical listener's judgment might be affected by long delays, and it would be unfair to ask for good judgment under those conditions. But if I proposed to use a number of selector switches, the man with the "golden ears" might argue that the switch's impedance would be worse than the splices, so a switch would be suspect! No, you can't use switches when you want to do an A-B comparison!

But in the last few weeks, the hi-fi review column of this "expert" was discussing how he compares different speakers: He said to change from one set of speakers to another, he uses switches! I just hope the switches don't cloud his judgment, as if they were (God forbid) *splices.*

Comments invited! / RAP
Robert A. Pease / Engineer

Originally published in Electronic Design,
December 27, 1990.

RAP Update: *This was my ninth column, and I got a LOT of fan mail on this, dozens and dozens of letters. People sent me all sorts of advertising pages for outrageous claims on speaker cables. Some of this material went into my 1994 column on Hoaxes, and some of the ideas were brought into the recent one on Speaker Cables. Now that Tom Nousaine has established that nobody can tell any difference between speaker cables, the conceit of this fellow, Harry Somerfield, that he could hear the difference between spliced and un-spliced cables, is seen to be just hilarious - as it was at the time. I sent Harry Somerfield a challenge, to tell how many splices were in the wires in each of 10 sealed boxes. Are you surprised that he didn't ever reply?*

National Semiconductor

THE BEST OF BOB PEASE

Third Thoughts on Fuzzy Logic

An ardent skeptic of fuzzy logic, Robert A. Pease presents a viewpoint opposing that of the theme articles in this issue. Pease is a staff scientist at National Semiconductor and writes the column "Pease Porridge." He has a BSEE from MIT. -- Ed.

I became an instant "expert" on fuzzy logic (FL) in May of 1993, after questioning[1] the advantages claimed for it -- too many were preposterous and made no sense. Now after writing five more columns on FL[2-6], lecturing, and receiving many thoughtful letters, I am not much more of an expert on doing FL than in 1993.

But, no problem: most people publishing papers touting FL as having great advantages over conventional systems know less about good conventional systems than I do about FL. And a lot of my bold statements about FL been confirmed. For example, when an FL system claims to show great advantages over a conventional system, the advantages arise because the conventional system was badly done, or because it was heavily nonlinear.

So-called experts make silly claims for FL: "An elevator not based on FL runs at constant speed until it reaches the destination, then comes to a rough stop."[7] Absurd! Then an FL "expert" on the radio[8] claimed that the 1993 Saturn automobile's automatic transmission was redesigned with FL. Daniel McNeill, author of a collection of platitudes and puffery about FL,[9] stated that on steep upgrades, the transmission used FL to cut out "hunting" and excessive shifting up and down. He said that ordinary automatic transmissions "always shift at the same speed," but by using FL, they can be made to shift at different speeds -- poppycock!! Well, the Saturn's transmission was improved, but it did not use FL to improve its behavior on upgrades.[10] Designers used FL only to provide downshifts on downgrades. Every example of FL included bad thinking and untrue statements -- which made me very suspicious!

Several people recommended the Sendai train as an excellent example of FL, providing smoother acceleration, faster speed, and better energy economy in a practical system. The Sendai train actually provides worse energy economy;[4] it only appeared to give better economy due to a flawed 1985 computer analysis. In addition, the only published research[11] about the train appeared before it began operation. So all FL enthusiasts must stop claiming that the Sendai train uses less energy. It is a very good, smooth train, and very well engineered, but it cannot travel faster between stations and still consume less energy. This is not because FL is bad, but just because the train provides smoother acceleration.

Several other people have pointed out that, even today, most technical papers on FL show trivial or "toy" examples, not real-world examples. Such papers still set up "straw men." Also, they are written in esoteric symbols and couched in obscure, scholarly phraseology -- and thus are incomprehensible to most serious students or practitioners of FL.

Most FL controllers respond only to proportional and derivative terms[5] and do not have integral response. So a true PID (proportional-integral-derivative) controller, whether analog, digital, or FL, can achieve better accuracy and better dynamic response. These days, few engineers know what PID control[6] is or why it is useful.

Some authorities (notably Bart Kosko) argue that when we go from two parameters to three, and use seven rules in each dimension, the number of fuzzy rules increases grossly, from 49 (7^2) to 343 (7^3), for example. In most cases, that are more or less linear, however, we can use three or five rules, rather than seven. if we add an integral term to the main (proportional) parameter, we may add a small number of new rules, while easily accommodating a third or fourth parameter. So, conversion of an FL to a full-PID controller may require only a few more rules, leading to greatly improved performance.

Dave Brubaker, a serious practitioner of FL, showed[12] the advantages of using PID control in FL. And Adaptive Logic (San Jose, California,) has new software that facilitates the design of PID FL controllers. This can be powerful: a full-PID controller that adds FL's ability to handle nonlinearities.

So what else are we learning about FL? For one, there is still a lot of hype. Ads insist that we'll prefer a $249.95 FL electric razor or a $99.99 FL electric toothbrush. Maybe; I'll wait to see what Consumer Reports says. Meanwhile, there are still many impractical promoters of FL -- balanced by a number of excellent, practical ones.

Such a case comes from Constantin von Altrock of Inform[13] (Oak Park, Illinois). A team from Inform built an autonomous racing car with a 1-hp electric motor for its 10 lbs. of weight, with a complete computer on board, plus an FL pattern recognition system. They put tape down on a parking lot to indicate a race track, so the car could sense its location on the track and when it was coming up to curves.

After a couple of hours writing software, the team got the car to move around the track -- but it performed poorly, skidding badly and driving stupidly. Yet, two weeks later, the software improved, and the car could really drift the corners and race competently. Impressive! Next, the team threw boxes (with tape on the edges, for visibility) onto the track, and the car dodged the boxes and kept racing! Bravo! This may not yet be useful, but it is a virtuoso performance, nontrivial, and not feasible with conventional computers. It can't even be done with op amps! Surely such capabilities will soon prove useful.

Enthusiasts make positive claims for process controllers using FL added to or instead of conventional PID controllers. However, people write saying that, despite these positive claims, FL controllers did not work as well as a good PID controller, in their particular application. Still, a marketeer at Omron Electronics, (Schaumburg, Illinois) pointed out that a good controller does cost extra with added FL; and these augmented controllers are selling pretty well. Why would people pay more? Because, in some applications, the ones with FL really do work better. This fact is compelling. As Jay Last (a Teledyne founder) says, "The only valid market survey is a signed purchase order."

At the 1994 Fuzzy Logic conference, the most impressive demonstration had a ping-pong ball floating atop a column of moving air in a vertical plastic cylinder. An FL controller moved the ball between three levels. This problem is fairly nonlinear and not easily controlled, even with a good detector to indicate the ball's location. (Most FL examples never mention the sensors used, or the sample rate, which can be very important.) Two of the movements were smooth, but the third overshot badly and oscillated. After a few hours, the third transition was alright, but the second one was poor. They never got all three transitions working correctly. So, while it is possible to optimize an FL system, doing so is difficult, despite some FL enthusiasts' claims that optimizing an FL system is easy. Claims that PID is difficult to use are similarly untrue.[14]

An article about balancing a ball on a tilting beam[15] claimed that an FL controller, using only a 486-based computer, was faster and settled better (compared to an untrained human operator). The published plot of the ball's motion was unbelievable because it showed that the ball sometimes did not accelerate after the beam tilted and sometimes accelerated when the beam did not tilt. In addition, the ball had a 100-percent overshoot five times before settling. And despite these poor results, the authors insisted that they got better performance using FL and did not have to provide a mathematical model for the system. When confronted,[16] the authors insisted their design was good,[17] but did not explain the erratic data.

I built my own ball-on-a-beam balancer and used a mathematical model. The ball's velocity is the integral of the tilt, and when a ball rolls, its position is the integral of its velocity. Thus, the ball's position is a double integration on the beam's tilt. I designed the controller in a couple of hours using scrap paper and pen -- not a computer. The design uses four op amps (LMC660) as integrators and differentiators.

This design ran on the first try. While its motion is not yet optimal, the ball still settles promptly without even 5 percent of overshoot, in about 6 seconds for a 20-in. motion. Thus, four op amps (costing less than $1) can, in a thoughtful application, outperform a $1,500 personal computer. This works because a good model exists for the rolling ball (rather than pretending to be better off without any model).

A while back, Lotfi Zadeh and I discussed ball-on-beam balancing. He thought he could make a ball roll on a lossy tilting surface (such as carpet) with humps to hinder the ball's free roll, and could even make it arrive on target at a particular time. I said that was possible with op amps, too, but he could surely do it better, because FL is very good for nonlinear cases.

Then Lotfi spoke of a colleague who still dislikes FL. This person said, "Lotfi, I hope I live long enough to see you invited to the White House, where the President will present you with a medal 'For fooling the Japanese into thinking that fuzzy logic is a good idea'." It was really charming that Lotfi had such a fine sense of humor, and could tell a joke on himself, on such a serious topic!

I invite your comments.

Robert A. Pease
Mail Stop D2597A
National Semiconductor
P.O. Box 58090
Santa Clara, CA 95052-8090
E-mail: czar44@me.com

References

References 1-6 and 15-17 available from the author -- Ed.

2. R.A. Pease, "What's All This Fuzzy Logic Stuff, Anyhow?," Electronic Design, May 13, 1993, pp. 77-79.
3. R.A. Pease, "What's All This Fuzzy Logic Stuff, Anyhow (Part II)?," Electronic Design, Nov. 1, 1993, pp. 95-98.
4. R.A. Pease, "What's All This Fuzzy Logic Stuff, Anyhow (Part III)?," Electronic Design Nov. 11, 1993, pp. 105-108.
5. R.A. Pease, "What's All This Acceleration Stuff, Anyhow?," Electronic Design, Nov. 7, 1994, pp. 63-64.
6. R.A. Pease, "What's All This Refrigerator Stuff, Anyhow?," Electronic Design, Dec. 16, 1994, pp.111- 113.
7. R.A. Pease, "What's All This PID Stuff, Anyhow?," Electronic Design, June 26, 1995 (special issue on Analog Applications).
8. Electronic Design, June 25,1992, p. 37.
9. Interview with Daniel McNeill, "Marketplace," on NPR's KQED-FM, 6:49 p.m., Mar. 23, 1993, Marketplace Radio, Los Angeles, Calif.
10. D. McNeill and P. Freiberger, Fuzzy Logic, Simon & Schuster, Old Tappan, N.J., 1993.
11. Automotive Industries, Feb. 1993, pp. 149-150.
12. S. Yasunobu and S. Miyamoto, "Automatic Train Operation System by Predictive Fuzzy Control," Industrial Applications of Fuzzy Control, Elsevier Science Publishers, New York, 1985, pp. 1-18.
13. D. Brubaker, "Design a Fuzzy Set-Point Controller," EDN, Jan. 5, pp. 167-170; Feb. 16, pp. 163-166; and Mar. 16, 1995, pp. 133-136.
14. INFORM GmbH, "Advanced Fuzzy Logic Control of a Model Car," J. Fuzzy Sets and Systems, Elsevier Science Publishers, New York, 1994.
15. QuickMotion, Pro-Log, Monterey, Calif.
16. H. Li and Y. Ji, "Fuzzy-Logic Tools on Tap for IC Wafers," IEEE Circuits and Devices, Mar. 1994, pp. 30-35,
17. R. Pease, "Letters to the Editor," IEEE Circuits and Devices, Sept. 1994, pp. 4-6.
18. H. Li, "Letters to the Editor," IEEE Circuits and Devices, Sept. 1994, pp. 6, 19.

Originally published in the August 1995 IEEE *Micro Magazine*.

National Semiconductor

THE BEST OF BOB PEASE

Understand Capacitor Soakage to Optimize Analog Systems

Dielectric absorption can cause subtle errors in analog applications such as those employing S/H circuits, integrating ADCs and active filters. But knowing how to measure this soakage and compensate for it helps you minimize its effects.

Veteran circuit designers often got a shocking introduction to dielectric absorption when supposedly discharged high-voltage oil-filled paper capacitors reached out and bit them. Indeed, the old oil-filled paper capacitors were notorious for what was once called soakage -- a capacitor's propensity to regain some charge after removal of a momentary short. Today, you won't find very many of these capacitors in use, but you will still encounter soakage. Do you know how to deal with it?

Nowadays, you're more likely to notice the effects of dielectric absorption in some more subtle way, perhaps in the performance of an integrator that can't be reset to zero or a sample/hold that refuses to work correctly. But whether you literally feel its effects or merely observe them in a circuit's behavior, dielectric absorption is an undesirable characteristic that every capacitor possesses to some degree. This characteristic is inherent in the dielectric material itself, although a poor manufacturing procedure or inferior foil electrodes can contribute to the problem.

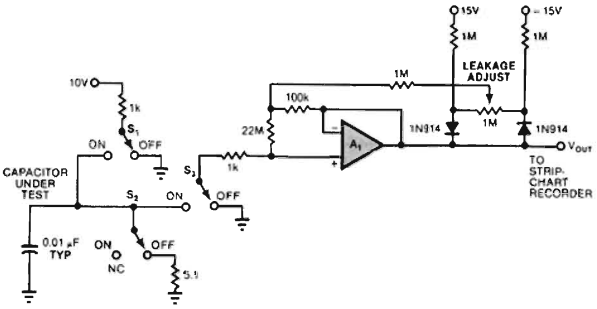

Fig 1 - A simple test fixture lets you evaluate dielectric absorption at low speeds. To use the one shown here, start with all switches off and throw S_1 and S_2 on for 1 min; throw S_1 and S_2 off and wait 6 sec, throwing S_3 on during the wait period. Next, turn S_2 on and watch V_{OUT} for 1 min. To compensate for leakage, leave all switches off for 1 min and then throw S_2 and S_3 on. Monitor V_{OUT} for 1 min and subtract this value from the V_{OUT} value obtained earlier. (View a larger version of the image.)

Indeed, soakage seems an apt term for dielectric absorption when you note what the capacitor seems to be doing. Consider a typical example: A capacitor charges to 10V for a long time T and then discharges through a small-value resistor for a short time t. If you remove the short circuit and monitor the capacitor terminals with a high-impedance voltmeter, you see the capacitor charge back to 0.1%, 1% or as much as 10% of the original voltage. For example, a 1-µF Mylar capacitor charged to 10V for 60 sec (T_{CHARGE}) and discharged for 6 sec ($T_{DISCHARGE}$) charges to 20 or 30 mV after 1 min (T_{HOLD}). Fig 1 shows a simple evaluation circuit for measuring this characteristic.

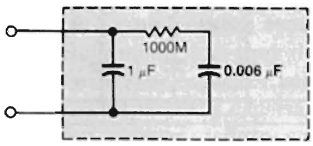

Fig 2 - To model the soakage characteristic of a 1-µF Mylar capacitor, consider a circuit that incorporates a 0.006-µF capacitor to represent the dielectric's charge-storage characteristics.

A capacitor exhibiting dielectric absorption acts as if during its long precharge time the dielectric material has soaked up some charge that remains in the dielectric during the brief discharge period. This charge then bleeds back out of the dielectric during the relaxation period and causes a voltage to appear at the capacitor terminals. Fig 2 depicts a simple model of this capacitor: When 10V is applied for 1 min, the 0.006-µF capacitor gets almost completely charged, but during a 6-sec discharge period it only partially discharges. Then, over the next minute, the charge flows back out of the 0.006-µF and charges the 1-µF capacitor to a couple of dozen millivolts. This example indicates that a longer discharging time reduces soakage error but that discharging for only a small fraction of that time results in a larger error. Illustrating this point, Fig 3 shows the results of conducting Fig 1's basic test sequence for 1-, 6- and 12-sec discharge times. Note that the capacitor tries to remember its old voltage, but the longer you hold it at its new voltage, the better it forgets -- in the Fig 3 case, soakage errors equal 31 mV at $t_{DISCHARGE}$=I sec, 20 mV at $t_{DISCHARGE}$=6 see and 14 mV at $t_{DISCHARGE}$=12 sec.

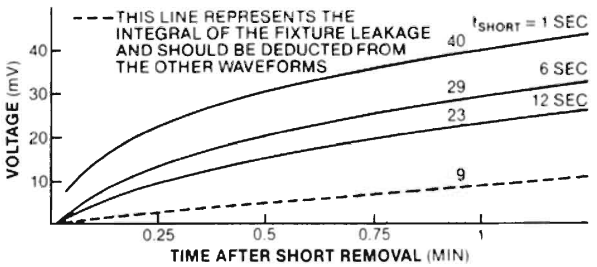

Fig 3 - Obtained using Fig 1's test circuit, these dielectric-absorption-measurement results for a 1-µF capacitor shown that longer $t_{DISCHARGE}$ times reduce soakage-caused errors.

High-speed tests predict S/H performance

You might now ask whether these low-speed tests have any bearing on a capacitor's suitability in fast millisecond or microsecond sample/hold applications. If you repeat the Fig 1 experiment for T_{CHARGE} = T_{HOLD} = 1000 µsec and $t_{DISCHARGE}$ = 100 µsec, you see very similar capacitor-voltage waveforms but with about 10-times-smaller amplitudes. In fact, for a constant T:t ratio, the resulting soakage error decreases only slightly in tests ranging in length from minutes to microseconds.

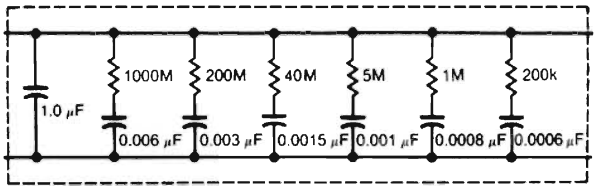

Fig 4 - More precise than Fig 2's equivalent circuit, a capacitor model employing several time constants proves valid for a wide range of charge and discharge times. This model approximates a Mylar capacitor.

Fig 4's circuit approximates this capacitor characteristic, which you can observe on actual capacitors by using Fig 5's test setup. Here, a sample/hold IC exercises the capacitor under test at various speeds and duty cycles, and a limiter amplifier facilitates close study of the small residual waveforms, without over-driving the oscilloscope when the capacitor is charged to full voltage.

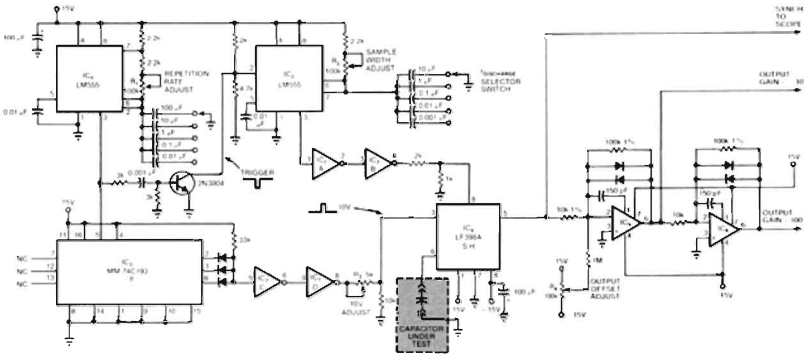

Fig 5 - Capable of automatically sequencing the dielectric-absorption tests, a circuit employing timers, a sample/hold and limiting stages allows you to make measurements for a wide range of T_{CHARGE}, T_{HOLD}, and $t_{DISCHARGE}$ values. Fig 7 shows the results obtained using the circuit shown here. (View a larger version of the image.)

Notes:

1. ALL DIODES = 1N914
2. IC5, IC6 = LM301A
3. IC7 = MM74C04
4. USE R4 OR -10 GAIN TO KEEP SCOPE WAVEFORM BELOW 200mV SO AS TO AVOID DISTORTION OR FALSE ATTENUATIONS

Such experiments illustrate that if you put a certain amount of charge into a less-than-ideal capacitor, you will get out a different amount of charge, depending on how long you wait. Thus, using low-soakage capacitors proves important in applications such as those involving high-resolution dual-slope integrating ADCs. And sure enough, many top-of-the-line digital voltmeters do use polypropylene (a low-soakage dielectric) devices for their main integrating capacitors.

But dielectric-absorption characteristics are most obviously detrimental in applications involving sample/holds. Manufacturers guarantee how fast these devices can charge a capacitor in their Sample mode and how much their circuits' leakage causes capacitor-voltage droop during the Hold mode, but they don't give any warning about how much the capacitor voltage changes because of soakage. This factor is especially important in a data-acquisition system, where some channels might handle small voltages while others operate near full scale. Even with a good dielectric, a sample/hold can hurt your accuracy, especially if the sample time is a small fraction of T_{HOLD}. For example, although a good polypropylene device can have only 1-mV hysteresis per 10V step if T/t=100 msec/10 msec, this figure increases to 6 mV if the T/t ratio equals 100 msec/0.5 msec. Because most sample/hold data sheets don't warn you of such factors, you should evaluate capacitors in a circuit such as Fig 5's, using time scaling suited to your application.

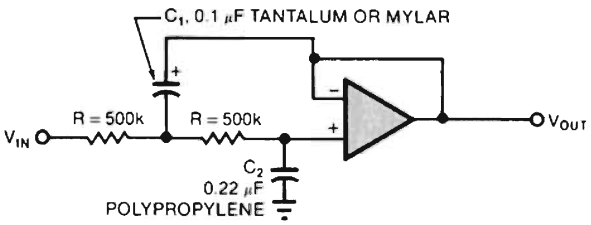

Fig 6 - Soakage can present problems when you're designing a fast-settling amplifier or filter. In the circuit shown here, for example, C_1 can be a Mylar or tantalum unit, but making C_2 a polypropylene device improves performance.

Other applications in which soakage can degrade performance are those involving fast-settling ac active filters or ac-coupled amplifiers. In Fig 6's circuit, C_1 can be a Mylar or tantalum unit because it always has 0V dc on it, but making C_2 polypropylene instead of Mylar noticeably improves settling. For example, settling to within -0.2 mV for a 10V step improves from 10 to 1.6 sec with the elimination of Mylar's dielectric absorption. Similarly, voltage-to-frequency converters benefit from low-soakage timing capacitors, which improve V/F linearity.

Some dielectrics are excellent at all speeds

Fortunately, good capacitors such as those employing polystyrene, polypropylene, NPO ceramic and Teflon dielectrics perform well at all speeds. Fig 7 shows the characteristics of capacitors using these dielectrics and others such as silver mica and Mylar. In general, polystyrene, polypropylene or NPO-ceramic capacitors furnish good performance, although polystyrene can't be used at temperatures greater than 80°C. And although NPO ceramic capacitors are expensive and hard to find in values much larger than 0.01 μF, they do achieve a low temperature coefficient (a spec not usually significant for a S/H but one that might prove advantageous for precision integrators or voltage-to-frequency converters). Teflon is rather expensive but definitely the best material to use when high performance is important. Furthermore, only Teflon and NPO ceramic capacitors suit use at 125°C.

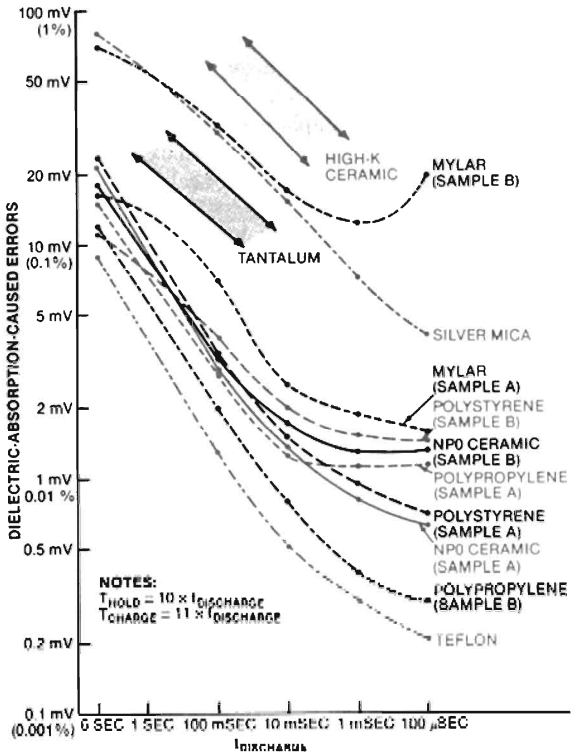

Fig 7 - Soakage-measurement results for a variety of capacitors illustrate the effects of $t_{DISCHARGE}$ values on dielectric-absorption-caused errors. Note the curves for two different samples of NP0 ceramic capacitors intersect.

If you look at Fig 7's dielectric-absorption values, you can see wide differences in performance for a given dielectric material. For example, polypropylene sample A is about as good as B at t=6 sec, but B is four times better at high speeds. Similarly, NP0-ceramic sample A is slightly worse than NP0-ceramic sample B at low speeds, but A is definitely better at high speeds. And some Mylar capacitors (sample A) get better as speed increases from 1000 to 100 µsec, but others (sample B) get worse. So if you want consistently good performance from your capacitors, evaluate and specify them for the speed at which they'll be used in your application. Keep in mind that because most sample/holds are used at much faster speeds than those corresponding to the 1- or 5-min ratings usually given in data sheets, a published specification for dielectric absorption has limited value.

In addition, other dielectrics furnish various levels of performance:

- Because any long word that starts with poly seems to have good dielectric properties, how about polycarbonate or polysulfone? No -- they are about as bad as Mylar.
- Does an air or vacuum capacitor have low soakage? Well, it might, but many standard capacitors of this type are old designs with ceramic spacers, and they might give poor results because of the ceramic's hysteresis.
- If a ceramic capacitor is not an NP0 device, is it any good? Most of the conventional high-K ceramics are just terrible -- 20 to 1000 times worse than NP0 and even worse than tantalum.
- Is silicon dioxide suitable for small capacitances? Although Fig 5's test setup, used in preparing Fig 7's chart, only measures moderate capacitances (500 to 200,000 pF), silicon dioxide appears suitable for the small capacitors needed for fast S/Hs or deglitchers.

Cancellation circuit improves accuracy

A practical method of getting good performance with less-than-perfect capacitors is to use a soakage-cancellation circuit such as one of the form shown in Fig 8, in which a capacitor of the type modeled in Fig 4 serves as an integrator. (Only the first two soakage elements are shown.) The integrator's output is inverted with a scale factor of -0.1, and this voltage is then fed through one or more experimentally chosen RC networks to cancel the equivalent network inherent in the capacitor's dielectric material.

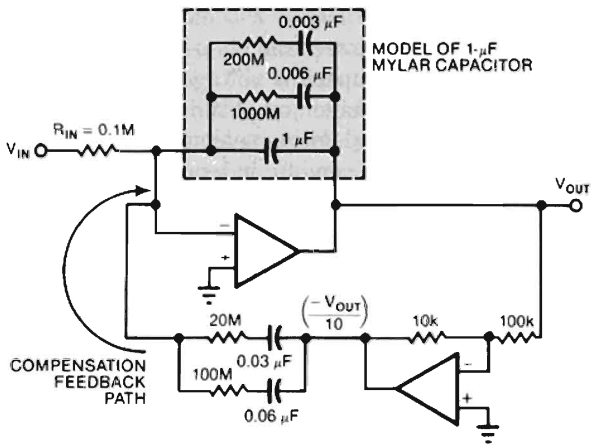

Fig 8 - You can compensate an integrator for dielectric absorption by feeding its inverted output back to the input through one or more experimentally chosen RC networks, which cancel the equivalent network inherent in the capacitor's dielectric material.

Fig 9 shows a practical sample/hold circuit with an easily trimmed compensator. This network provides about a 10-fold improvement for sample times in the 50- to 2000-µsec range (Fig 10). Although this compensation is subject to limitations at very fast or slow speeds, the number of RC sections and trimming pots employed can be extended.

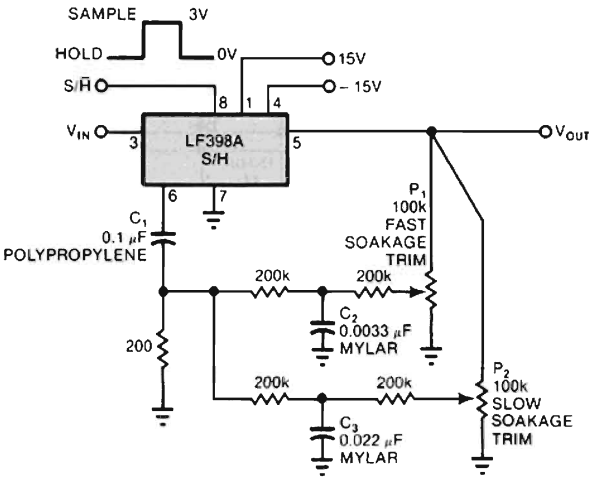

Fig 9 - Adding compensation circuitry to a sample/hold yields better-than-Teflon performance with a polypropylene capacitor. Using Teflon capacitors in such circuits can yield a 15- to 17-bit dynamic range.

Simple circuits similar to Fig 9's or Fig 8's have been used in production to let inexpensive polypropylene capacitors provide better-than-Teflon performance. In turn, using these compensator circuits with a good Teflon capacitor furnishes a dynamic range of 15 to 17 bits.

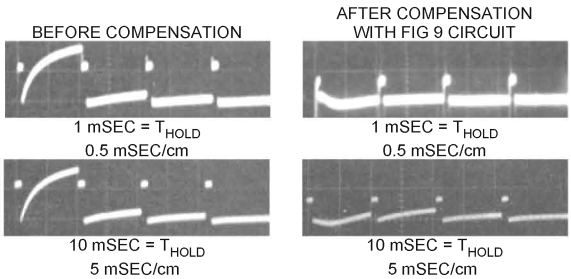

Fig 10 - Adding Fig 9's compensation network to a sample/hold circuit yields a 10-fold performance improvement for sample times of 50 to 2000 µsec; additional RC networks and trimming pots can extend the time range. The short pulses represent normal S/H jumps and occur during the sample time. The exponentially rising waveform during the hold time results from soakage. Note that soakage effects are still visible during the second hold period.

Notes:

1. Dielectric absorption errors with good polypropylene capacitor
2. All waveforms at 1mV/cm; T_{SAMPLE}/T_{HOLD} = 1/10

Originally published in EDN,
October 13, 1982.

RAP Update: *FIFTEEN years ago, back in 1982, when I wrote this, I had never seen any study of capacitors and their "soakage" -- nor of the kind of circuits you could use to shrug off the effects of soakage. Nobody ever talked about this, at high speed.*

To this day, I have not seen any other articles that covered either subject. So this is still about the prime source of info on how to evaluate capacitors for soakage, AND how to build good Sample-and-Hold circuits, so as to NOT get hurt by that soakage.

National Semiconductor

Remembering Bob Pease

Legendary Analog Expert

National Semiconductor honors the accomplishments of the company's legendary design engineer, Bob Pease, who died on June 18. During his 33-year career at National, Bob received 21 patents and designed more than 20 integrated circuits. Bob Pease was loved by the analog community and we celebrate Bob's passion for analog.

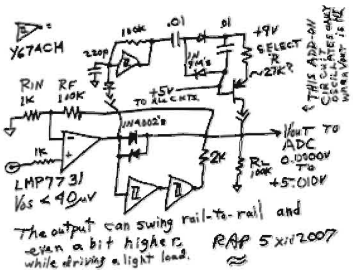

East Coast Beginnings

After attending Northfield Mount Herman High School in Massachusetts, Bob earned a Bachelor of Science degree in electrical engineering from MIT in 1961 and started working at George A. Philbrick Researches. The Boston-based company launched the commercial use of the operational amplifier in 1952. There Bob worked on affordable mass-produced op amps using discrete solid-state components. He passed up an opportunity to work at nearby Analog Devices and instead moved to the San Francisco Bay Area to work at National in 1976, which had rapidly become one of the top three US semiconductor companies on the strength of its analog technology.

Among the semiconductor products Bob designed are temperature-voltage frequency converters used in groundbreaking medical research expeditions to Mt. Everest in the 1980s. He also designed a seismic pre-amplifier chip used to measure lunar ground tremors in the US Apollo moon landing missions. Among his more memorable designs are the LM331 voltage to frequency converter and the LM337 adjustable voltage regulator.

"Bob Pease was one of those analog engineers who spanned the semiconductor industry's early history," said EDN columnist Paul Rako, a former colleague at National. "He started working on vacuum tubes and discrete components, then monolithic analog circuits with the planar process. Later in his career he put all of this accumulated knowledge to use as an applications engineer. That's what gave him such breadth."

Analog Seminars & Pease Porridge

While at National, Bob's reputation grew as he shared the secrets of analog design with engineers around the world via National's Analog Seminars. Bob's passion for sharing information knew no bounds. He worked long hours from home as well as at National answering phone calls and emails from anyone with questions about analog design: customers, students, veteran engineers—it didn't matter.

"Discussing the solutions made one think differently and look at alternative possibilities," said strategic technologist Don Archer. "It was a fun time working with Bob and he loved working through difficult problems to find elegant answers."

During his tenure, he began writing a continuing popular monthly column in Electronic Design magazine entitled "Pease Porridge" about his experiences in the world of electronic design and application. Stories often started with his trademark expression "What's All This [topic] Stuff, Anyhow?" He also wrote for EDN magazine for a time. He authored eight books, his most popular being "Troubleshooting Analog Circuits." He also had his own website at National.com with the subtitle, "What's All This Homepage Stuff, Anyhow?" where you can find "several kinds of useful information" as well as Bob's passion for Nepal and its people.

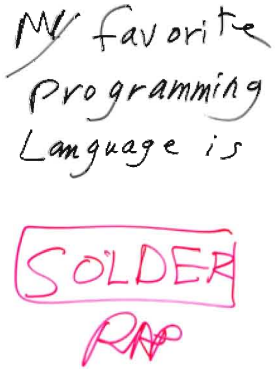

Analog by Design Show

Starting in 2003, Bob hosted the semiconductor industry's first online webcast tailored specifically for analog design engineers. The "Analog By Design Show" was an engineering talk show by engineers for engineers. Early segments of the half-hour talk show included co-hosts Paul Rako and Paul Grohe, plus a guest expert. As Bob's reputation grew and he became a face of National, his persona grew as well. His image was featured in many promotions, often with a humorous slant. Behind the humor, there was always a serious note around engineering excellence. As part of a campaign to help engineers avoid repeating old mistakes in their new band gap reference circuits, Bob donned an outlandish Czar of Band-Gaps uniform.

"Bob Pease goes back to the wild days of analog design," Rako said. "This is when a core group of passionate engineers and scientist would work hard, play hard, and do as they pleased."

The Right Stuff

Recently Bob was working on material for National.com called "How to Choose an Op Amp" featuring sage advice and his inimitable style. He was also updating the popular application note, AN-31, a collection of op amp circuits dating back 1969. He was in the process of writing about the history and evolution of these circuits and the best ways to implement these circuits using today's op amps.

Bob received many awards for his work, including the Lifetime Achievement Award from the Embedded Systems Conference (2010) and Electronic Design magazine's Electronic Engineering Hall of Fame (2002). He was listed as one of the top 10 analog engineers of all time in a 2009 EE Times story.

National Fellow Dennis Monticelli remembers Bob as a helpful colleague and friend. "We go way back to my days as a green engineer when his gregarious personality and sheer knowledge drew me in. Bob was always generous with his time and never forgot what interested you

whether work-related or not. He could multi-task like no other, yet also dive deep and narrow into esoteric areas. While famous for his analog expertise and passion for seat-of-the-pants engineering, his interests were actually quite broad and he would gladly engage on a wide variety of subjects. I will miss him yet take solace in the fact that his legacy lives on in the hearts and minds of true analog engineers everywhere."

"The industry has lost an analog giant," said National CEO Don Macleod. "Bob Pease was an extraordinarily talented engineer who cared deeply that others gained the knowledge they needed to advance their own work. He was a spokesperson for National for many years, with a worldwide following."

Bob Pease is a hero of the analog world. We will miss him greatly.

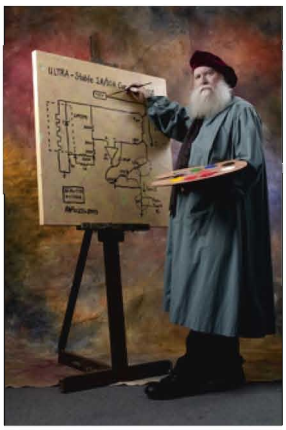

Related Links

Bob's Stuff	Friends' Stuff	Charity
• The Best of Bob Pease (Bob's site on National.com) • Analog by Design Show Videos • Pease Porridge Column at Electronic Design • Troubleshooting Analog Circuits Book (Amazon) 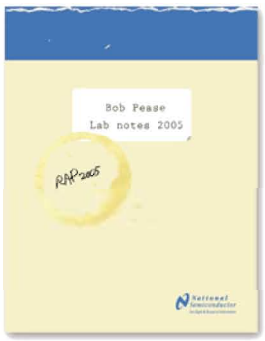 Bob Pease Lab Notes 2005 (What's All This Common Mode Rejection Stuff, Anyhow...)	• EE Times (Bill Schweber) • EDN (Paul Rako) • Analog Engineering Legend Bob Pease Remembered by Friends (Paul Rako, EDN) • Honoring Bob (Paul Rako, EDN) • Bob Pease Remembered for Pease Porridge and a Whole Lot More (Joe Desposito, Electronic Design) • Remembering Bob Pease The Writer (Don Tuite, Electronic Design) • An English Major Remembers an Analog Giant (Richard Gawel, Electronic Design) • EEWeb Interview • Bob Pease Didn't Hate SPICE Models (Paul Rako, EDN) • PCB Design 007 Panel Interview • Bob Pease's Office (YouTube video from Paul Rako) • Walt Jung Remebers Bob Pease (EDN) • A Bob Pease Photo Gallery (Paul Rako, EDN) • Sheingold, Titus, Thompson, Hanrahan, Franco, et al on Bob Pease (Paul Rako, EDN)	NEPAL YOUTH FOUNDATION Bob supported the Nepal Youth Foundation, helping Nepalese children live better lives 

New Phase-Locked-Loops Have Advantages as Frequency to Voltage Converters (and more)

National Semiconductor
Application Note 210
Robert Pease
April 1979

A phase-locked-loop (PLL) is a servo system, or, in other words, a feedback loop that operates with frequencies and phases. PLL's are well known to be quite useful (powerful, in fact) in communications systems, where they can pluck tiny signals out of large noises. Here, however, we will discuss a new kind of PLL which cannot work with low-level signals immersed in noise, but has a new set of advantages, instead. It does require a clean noise-free input frequency such as a square wave or pulse train.

This PLL can operate over a wide frequency range, not just 1 or 2 octaves but over 1 or 2 or 3 decades. It naturally provides a voltage output which responds quickly to frequency changes, yet does not have any inherent ripple. Thus, it can be used as a frequency-to-voltage (F-to-V) converter which does not have any of the classical limitations or compromises of (large ripple) vs (slow response), which most F-to-V converters have (Note 1). The linearity of this F-to-V converter will be as good as the linearity of the V-to-F converter used, and this linearity can easily be better than 0.01%. Other advantages will be apparent as we study the circuit further.

The basic circuit shown in *Figure 1* has all the functional blocks of a standard PLL. The frequency and phase detection do not consist of a quadrature detector, but of a standard dual-D flip-flop. When the frequency input is larger than F_2, Q1 will be forced high a majority of the time, and provide a positive error signal (via CR3, 4, 5, and 6) to the integrator.

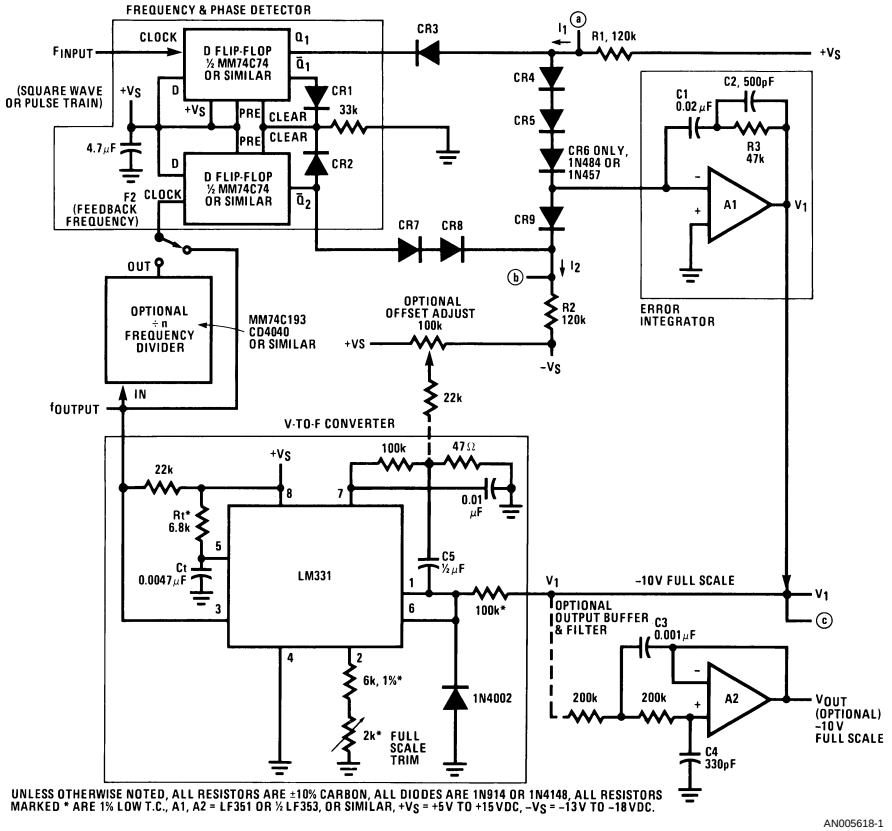

FIGURE 1. Basic Wide-Range Phase-Locked Loop

Note 1: Appendix C, "V/F Converter ICs Handle Frequency-to-Voltage Needs," National Semiconductor Linear Applications book.

New Phase-Locked-Loops Have Advantages as Frequency to Voltage Converters (and more)

AN-210

If F input and F_2 are the same, but the rising edges of F input lead the rising edges of F_2, the duty cycle of Q1=HI will be proportional to the phase error. Thus, the error signal fed to the integrator will decrease to nearly zero, when the loop has achieved phase-lock, and the phase error between F_{IN} and F_2 is zero. Actually, in this condition, Q1 will put out 30 nanosecond positive pulses, at the same time that Q2 puts out 30 nanosecond negative pulses, and the net effect as seen by the integrator is zero net charge. The 30 nanosecond pulses at Q1 and Q2 enable both flip-flops to be CLEARED, and prepared for the next cycle. This phase-detector action is substantially the same as that of an MC4044 Phase-Detector, but the MM74C74 is cheaper and uses less power. It is fast enough for frequencies below 1 MHz. (At higher frequencies, a DM74S74 can be used similarly, with very low delays.)

The error integrator takes in the current from R1 or R2, as gated by the Q1 and $\overline{Q2}$ outputs of the flip-flop. For example, when F_{IN} is higher, and Q1 is HIGH, I_1 will flow through CR4, 5, and 6 and cause the integrator's output to go more negative. This is the direction to make the V-to-F converter run faster, and bring F_2 up to F input. Note that A1 does not merely integrate this current in C1 (a mistake which many amateur PLL designers make). The resistor R3 in series with C1 makes a phase *lead* in the loop response, which is essential to loop stability. The small capacitor C2 across R3 is not essential, but has been observed to offer improved settling at the voltage output.

The output of the integrator, V1, is fed to a voltage-to-frequency (V-to-F) converter. The example shown here utilizes a LM331. This converter runs on a single supply, and responds quickly with nonlinearity better than 0.05% (even though an op-amp is not used nor needed). The output of the VFC is fed back to F_2, as a feedback frequency, either directly or through an (optional) frequency divider. Any number of standard frequency dividers such as MM74C193, CD4029, or CD4018, can be used, subject to reasonable limits. A divider of 2, 3, 10, or 16 is often used. The output voltage of the integrator will be proportional to the F input, as linearly as the V-to-F can make it. Thus, the integrator's output voltage V1 can be used as the output of an ultralinear F-to-V converter. However during the brief pulses when the flip-flop is CLEARing itself, there will be small glitches found on the output of A1. The RMS value of this noise may be very small, typically 0.5 to 5mV, but the peak amplitude, sometimes 10 to 100mV, can be annoying in some systems. And, no additional filtering can be added in the main loop's path, for any further delay in the route to the VFC would cause loop instability. Instead, the output may be obtained from a separate filter and buffer which operates on a branch path. A2 provides a simple 2-pole active filter (as discussed in Reference 1) which cuts the steady-state ripple and noise down below 1mV peak-to-peak, an excellent level for such a quick F-to-V (as we shall see).

What is not obvious about A2 is that its output can settle (within a specified error-band such as ±10 millivolts from the final DC value) earlier and more quickly than A1's output. The waveforms in *Figures 2, 3* show F_{IN} stepping up instantly from 5 kHz to 10 kHz; it also shows F_2 stepping up very quickly. The error signal at Q1 is also shown. The critical waveforms are shown in *Figures 4, 5*, the outputs of A1 and A2. While A1 puts out large spikes (caused by I1 flowing through R3), these large spikes cause the V-to-F converter to jump from 5 kHz to 10 kHz without any delay. There is, as shown in *Figures 2, 3*, a significant phase error between F_{IN} and F_2, but an inspection of these frequencies shows that frequency lock has been substantially instantaneous. Not one cycle has been lost. The phase lock and settling takes longer to achieve. Still, we know that if the frequency out of the VFC is 10 kHz, its input voltage must be −10 VDC. If there is noise on it, all we have to do is filter it in A2. *Figures 4, 5* shows that A2 settles very quickly — actually, in 2.0 milliseconds, which is just 20 cycles of the new frequency. A2's output has settled (i.e., the frequency has settled). While A1's output error (which is indicative of phase error being servo'ed out) continues to settle out for another 12 ms. Thus, this filter permits its output voltage to settle faster than its input, and it is responsible for the remarkable quickness of this circuit as an F-to-V converter. The waveforms of *Figures 4, 5* can be compared to the response (shown in *Figures 6, 7, 8, 9, 10*) of a conventional F-to-V converter. The upper trace is the output of a conventional FVC after a 4-pole filter (Note 2), and the lower trace is the output of the circuit of *Figure 1.* The phase-locked-loop F-to-V converter is quicker yet quieter.

Note 2: AN-207, V-to-F and F-to-V Converter Applications.

Vertical sensitivity=10 V/DIV (CMOS logic levels)
Horizontal sensitivity=0.5 ms/DIV

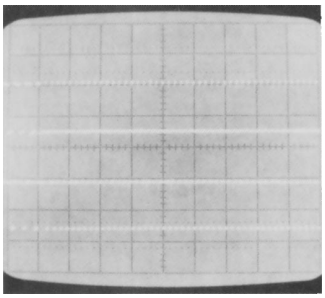

AN005618-2

FIGURE 2. F output steps up from 5 kHz to 10 kHz as quickly as the input, never missing a beat.
Top Trace = input "F_{IN}" to PLL.
Bottom Trace = output "F_{OUT}" from PLL.

Vert=10 V/DIV, Horiz=2 ms/DIV

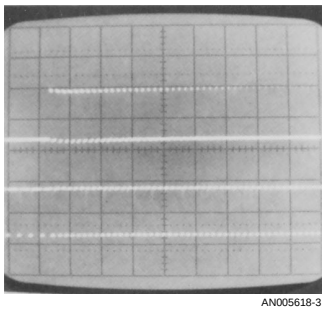

AN005618-3

FIGURE 3. Error Signal. Top Trace = error signal at Q1.
Bottom Trace = output "F_{OUT}" from PLL.

Vert = 2 V/DIV, Horiz = 2 ms/DIV

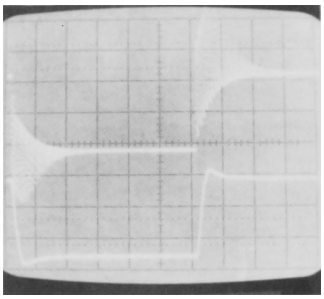

AN005618-4

FIGURE 4. Settling waveforms, as F_{IN} goes from 5 kHz to 10 kHz and back again, using circuit of *Figure 1*. Top Trace = output of integrator (V1). Bottom Trace = output of filter (V_{OUT}).

Vert = 2 V/DIV, Horiz = 0.5 ms/DIV

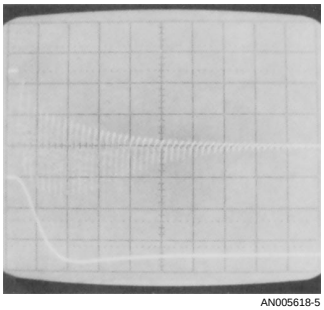

AN005618-5

FIGURE 5. PLL Settling Waveforms. The same waveform as in *Figure 4*, but time base is expanded to 0.5 ms/DIV to show fine detail of settling.

Vert = 2 V/DIV, Horiz = 20 ms/DIV

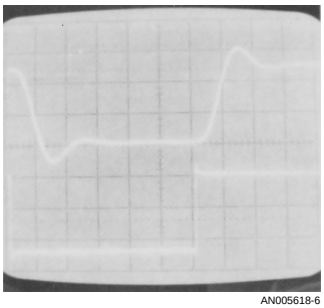

AN005618-6

FIGURE 6. FVC Response vs PLL Response. The PLL can settle rather more quickly than a conventional F-to-V converter. Top Trace = conventional F-to-V converter with 4-pole active filter, responding to a 5 kHz to 10 kHz step. Bottom Trace = PLL FVC, with the same input, circuit of *Figure 1*.

Vert = 2 V/DIV, Horiz = 20 ms/DIV

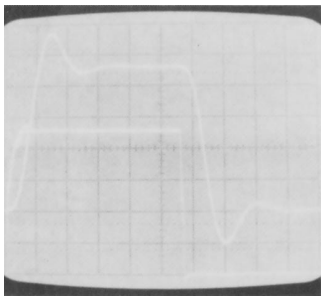

AN005618-7

FIGURE 7. FVC Step Response. This waveform is similar to that in *Figure 6* but the frequency change covers a 10:1 ratio, from 10 kHz to 1 kHz and back to 10 kHz. For this waveform, the adaptive current sources of *Figure 11* connect to *Figure 1* (whereas for *Figure 6* R1 = R2 = 120k).

Vert = 2 V/DIV, Horiz = 5 ms/DIV

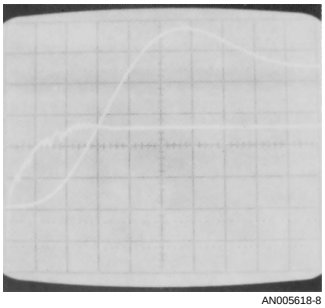

AN005618-8

FIGURE 8. FVC Response. The same as *Figure 7*, but time base expanded to 5 ms/DIV, to show detail of rise time. Top Trace = conventional FVC. Bottom Trace = PLL FVC.

Vert = 2 V/DIV, Horiz = 5 ms/DIV

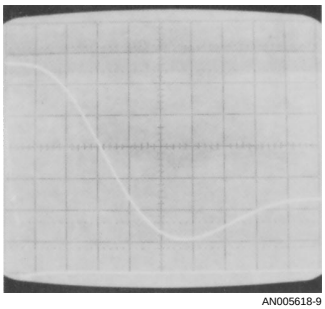

AN005618-9

FIGURE 9. FVC Response The same as *Figure 7*, but expanded to 5 ms/DIV to show details of fall time. Top Trace = conventional FVC. Bottom Trace = PLL FVC.

Vert = 0.2 V/DIV, Horiz = 50 ms/DIV

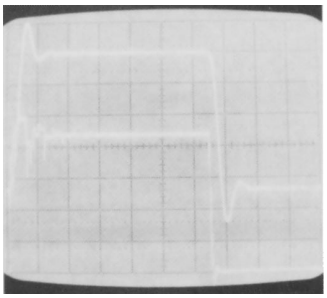

AN005618-10

FIGURE 10. PLL Settling Waveforms at Low Frequencies. The same idea as in *Figure 7*, but 10 x slower, from 1.0 kHz to 100 Hz (and back). The settling to 1 kHz is still distinctly faster for the PLL, but at 100 Hz, it is a bit slower. Still, the PLL is faster than the FVC at all speeds from 200 Hz to 10 kHz.

So far we have shown a PLL which operates nicely over a frequency range of about 3:1. If the frequency is decreased below 3 kHz, the loop gain becomes excessive, and the currents I1 and I2 are large enough to cause loop instability. The loop gain increases at lower frequencies, because a given initial phase error will cause the fixed current from R1 or R2 to be integrated for a *longer* time, causing a *larger* output change at the integrator's output, and a *larger* change of frequency. When the frequency is thus corrected, and the period of one cycle is changed, at a low frequency it may be *over-corrected*, and the phase error on the next cycle may be as large as (or larger than) the initial phase error, but with reversed sign (Note 3). To avoid this and to maintain loop stability at lower frequencies, e.g. 0.5 to 1 kHz, R1 and R2 can be simply raised to 1.5 MΩ. However, response to a step will be proportionally slower. To achieve a wide frequency range (20:1), and optimum quickness at all frequencies, it is necessary to servo I1 and I2 to be *proportional to the frequency*. Fortunately, as V1 is normally proportional to F, it is easy to generate current sources I1' and I2' which are proportional to F. The circuit of *Figure 11* can be connected to the basic PLL, instead of R1 and R2, and provides good, quick loop stability over a 30:1 frequency range, from 330 Hz to 10 kHz. For best results over a 30:1 frequency range, change R3, the damping resistor in *Figure 1*, from 47k to 100k. However, if the frequency range is smaller (such as 2:1 or 3:1), constant resistors for R1 and R2 or very simple current sources may give adequate response in many systems. (To cover wider frequency ranges than 30.1 with optimum response, the circuits in the precision phase-locked-loop, below, are much more suitable.)

Often a frequency multiplier is needed, to provide an output frequency 2 or 3 or 10 or n times higher than the input. By inserting a ÷n frequency divider in the feedback loop, this is easily accomplished. [Of course, a ÷m frequency divider can be inserted ahead of the frequency input, to provide correct scaling, and the output frequency then will be F_{IN}(n/m).]

To obtain good loop stability in a frequency multiplier with n = 2, remember that a 20 kHz V-to-F converter followed by a ÷2 circuit has exactly the same loop response and stability needs as a 10 kHz V-to-F converter, because it *is* a 10 kHz V-to-F converter, even though it provides a useful 20 kHz output. Thus, the frequency of the F_2 (minimum and maximum) will determine what loop gains and loop damping components are needed.

To accommodate a 1 kHz V-to-F loop, simply make C1 and C2 10 times bigger than the values of *Figure 1*; treating C3, C4, C5 and Ct similarly is used. To accommodate a 100 Hz V-to-F, increase them by another factor of 10.

If the PLL is to be used primarily as a frequency multiplier, it may be necessary to use stable, low-temperature-coefficient components, because the accuracy of V_{OUT} will not be important. The parts cost can be cut considerably. (Make sure that the VFC does not run out of range to handle all frequencies of interest.) On the other hand, the damping components will be chosen quite a bit differently if slow, stable jitter-free response is needed or if quick response is required. The circuits shown are just a starting place, to start optimizing your own circuit.

Note 3: Optimize phase-lock loops to meet your needs or determine why you can't. Andrzej B. Przedpelski, Electronic Design, September 13, 1978.

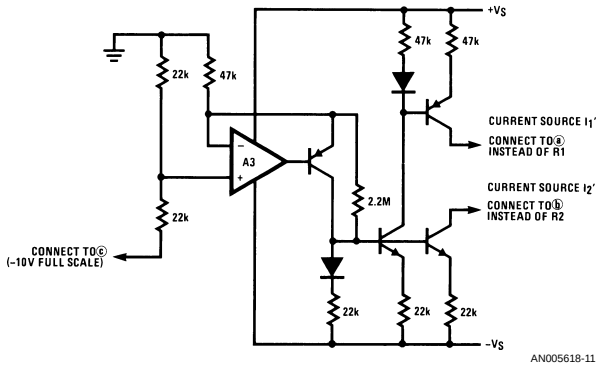

A3 — LF351, LM741 OR ANY NPN TRANSISTOR — 2N3904, 2N2222 OR ANY SILICON NPN ; PNP TRANSISTOR — 2N3906, 2N2907 OR ANY SILICON PNP ; ALL RESISTORS ±10%
ALL DIODES 1N914 OR 1N4148 OR SIMILAR

FIGURE 11. Proportional Current Source for Basic PLL

A Single-Supply PLL

The single-supply PLL is shown in *Figure 12* as an example of a simple circuit which is effective when battery operation or single-supply operation is necessary. This circuit will function accurately over a 10:1 frequency range from 1 kHz to 10 kHz, but will not respond as quickly as the basic PLL of *Figure 1*. The reason is the use of the CD4046 frequency detector. When an F_{IN} edge occurs ahead of a F feedback pulse, pin 13 of the CD4046 pulls up on C1 via R1 = 1 kΩ. This current cannot be controlled or manipulated over as wide a range as "I1" in *Figure 1*. As a consequence, the response of this PLL is never as smooth nor fast-settling as the basic PLL, but it is still better behaved than most F-to-V converters. As with the basic PLL, the detector feeds a current to be integrated in C1 (and R2 provides the necessary "lead"). A1 acts simply as a buffer for the R1, C1 integrator. A3, optional, can provide a nicely filtered output. And A2 servos Q1, drawing a current out of C6 which is proportional to V2. Here the LM331 acts as a current-to-frequency converter, and F output is precisely proportional to the collector current of Q1. As with the basic circuit, this PLL can be used as a quick and/or quiet F-to-V converter, or as a frequency multiplier. One of the most important uses of an F-to-V is to demodulate the frequency of a V-to-F converter, which may be situated at a high common-mode voltage, isolated by photoisolators, or to recover a telemetered signal. An F-to-V converter of this sort can provide good bandwidth for demodulating such a signal.

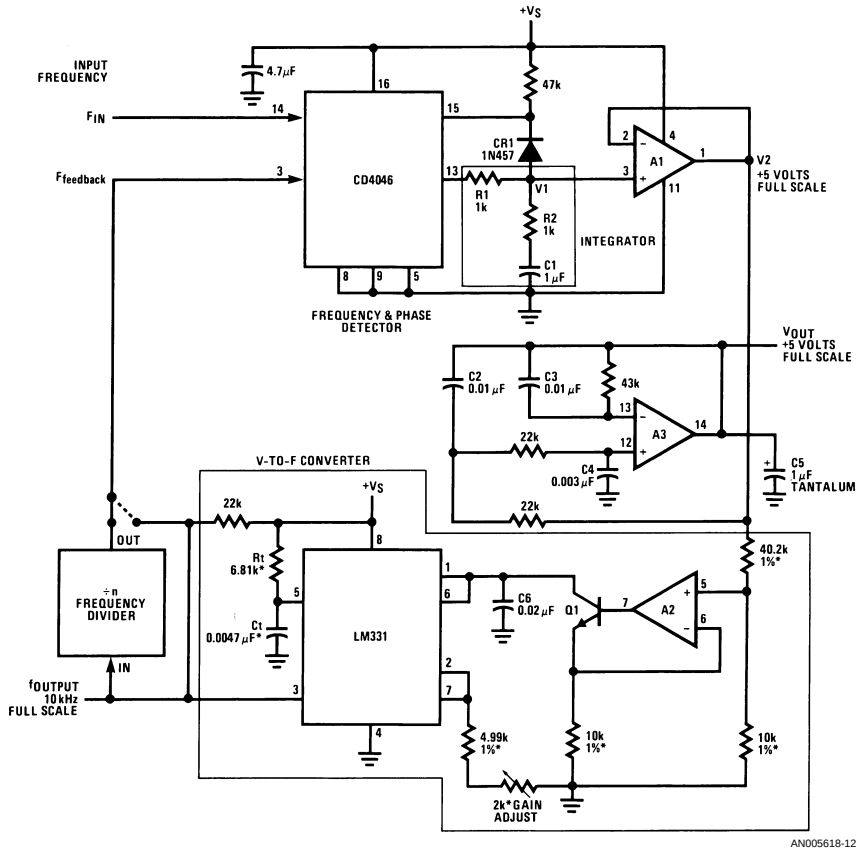

— Q1 = 2N3565 OR 2N3904 HIGH BETA NPN
— A1, A2, A3 = ¼ LM324
— ON CD4046, PINS 1, 2, 4, 6, 7, 10, 11, 12 ARE NO CONNECTION
— USE STABLE, LOW-T.C. PARTS FOR COMPONENTS MARKED*
— $+V_S$ = +7 TO +15 VDC

FIGURE 12. Single Supply Phase Locked Loop

The precision PLL in *Figure 13* acts very much the same as the basic PLL, with refinements in various places.

- The flip-flops in the detector have a gate G1 to CLEAR them, for quicker response.
- The currents which A1 integrates are steered through Q1, Q2 and Q3, Q4 because transistors are quicker then diodes, yet have much lower leakage.
- The V-to-F converter uses A2 as an op-amp integrator, to get better than 0.01% nonlinearity (max).
- G2 is recommended as an inverter, to invert the signal on the LM331's pin 3, avoid a delay, and improve loop stability. (However, we never found any *real* improvement in loop stability, despite theories that insist it must be there. Comments are invited.)

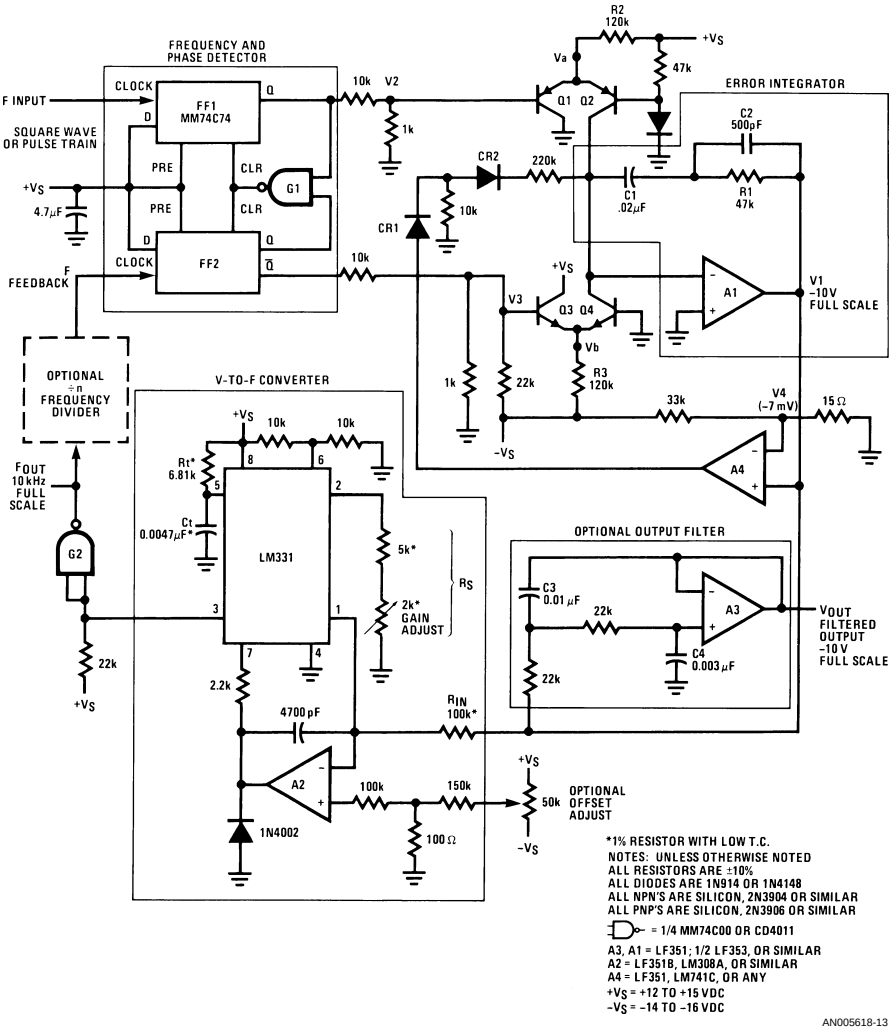

FIGURE 13. Precision PLL

- A4 is included as an (optional) limiter, to prevent V1 from ever going positive. This will facilitate quick startup and recovery from overdrive conditions.

Also, in *Figure 14*, the wide-range current pump for the precision PLL is a "semiprecision" circuit, and provides an output current proportional to −V1, give or take 10 or 15%, over a 3-decade range. The 22 MΩ resistors prevent the current from shutting off in case −V becomes positive (probably unnecessary if A4 is used). For best results over a full 3-decade range (11 kHz to 9 Hz), do use A4, delete the four 22 MΩ resistors, and insert the (diode parallel to the 470 kΩ) in series with the R_G as shown. This will give good stability at all frequencies (although stability cannot be extended below 1/1500 of full scale without extra efforts).

This PLL has been widely used in testing of VFCs, as it can force the LM331 to run at a crystal-controlled frequency (established as the F input), and the output voltage at V_{OUT} is promptly measured by a 6-digit (1 ppm nonlinearity, max) digital voltmeter, with much greater speed and precision than can be obtained by forcing a voltage and trying to read a frequency. While at 10 kHz, the advantages are clearcut; at 50 Hz it is even more obvious. Measuring a 50 Hz signal with ±0.01 Hz resolution cannot be done (even with the most powerful computing counter-timer) as accurately, quickly, and conveniently as the PLL's voltage output settles.

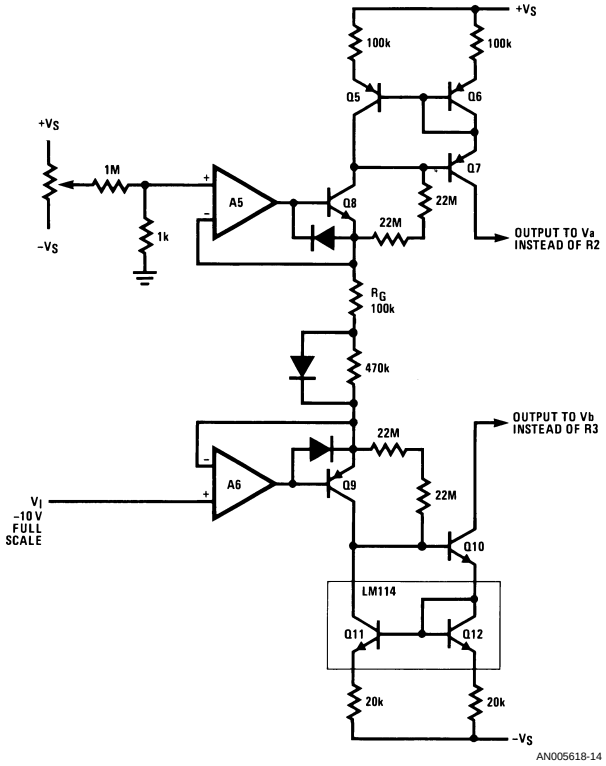

A5, A6, ARE LF351 OR ½LF353 OR SIMILAR
Q5, Q6, Q7, Q9 ARE 2N2907 OR 2N3906 OR SIMILAR
Q8, Q10 ARE Q8 IS 2N3565 OR 2N3904 OR SIMILAR

FIGURE 14. Wide Range Current Pumps for Precision PLL of *Figure 13*

One final application of this PLL is as a wide-range sine generator. The VFC in *Figure 15* puts out an adequate sine-shaped output, but does not have good V-to-F linearity, and its frequency stability is not much better than 0.2%. An LM331 makes an excellent linear stable V-to-F converter, with a pulse output; but it can not make sines. But it can command, via a PLL, to force the sine VFC to run at the correct frequency. Simply connect the sine VFC of *Figure 15* into one of the PLLs, instead of the LM331 VFC circuit. Then use a precise linear low-drift VFC based on the LM331 to establish the F_{IN} to the PLL. If the voltage needed by the sine VFC to put out a given frequency drifts a little, that is okay, as the integrator will servo and make up the error. The use of a controlled sine-wave generator in a test system was the first of many applications for a wide-range phase-locked-loop.

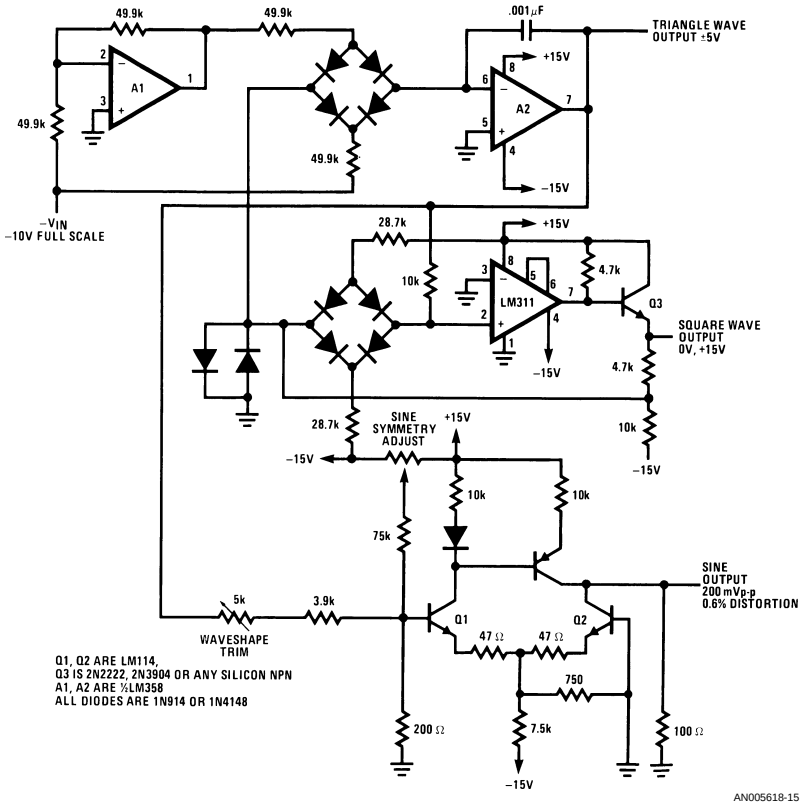

FIGURE 15. Sine-Wave VFC to Use with PLL

Book Extract End

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component in any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
Americas
Tel: 1-800-272-9959
Fax: 1-800-737-7018
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 1 80-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 1 80-530 85 85
English Tel: +49 (0) 1 80-532 78 32
Français Tel: +49 (0) 1 80-532 93 58
Italiano Tel: +49 (0) 1 80-534 16 80

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: sea.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5620-6175
Fax: 81-3-5620-6179

Working with High Impedance Op Amps

National Semiconductor
Application Note 241
Robert J. Widlar
February 1980

Abstract. *New developments have dramatically reduced the error currents of IC op amps, especially at high temperatures. The basic techniques used to obtain this peformance are briefly described. Some of the problems associated with working at the high impedance levels that take advantage of these low error currents are discussed along with their solutions. The areas involved are printed-circuit board leakage, cable leakage and noise generation, semiconductor-switch leakages, large-value resistors and capacitor limitations.*

INTRODUCTION

A new, low cost op amp reduces dc error terms to where the amplifier may no longer be the limiting factor in many practical circuits. FET bias currents are equalled at room temperature; but unlike FETs, the bias current is relatively stable even over a −55˚C to 125˚C temperature range. Offset voltage and drift are low because bipolar inputs and on-wafer trimming are used. The 100 µV offset voltage and 25 pA bias current are expected to advance the state of the art for high impedance sensors and signal conditioners.

BIAS CURRENTS

There has been a continual effort to reduce the bias current of IC op amps ever since the µA709 was introduced in 1965. The LM101A, announced in 1968, dropped this current by an order of magnitude through improved processing that gave better transistor current gain at low operating currents. In 1969, super-gain transistors (see appendix) were applied in the LM108 to beat FET performance when temperatures above 85˚C were involved.

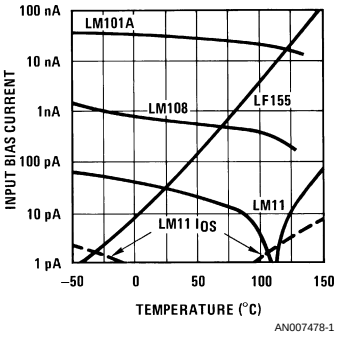

FIGURE 1. Comparison of typical bias currents for various types of IC op amps. New bipolar device not only has lower bias current over practical temperature ranges but also lower drift. Offset current is unusually low with the new design.

In 1974 FETs were integrated with bipolar devices to give the first FET op amp produced in volume, the LF155. These devices were faster than general purpose bipolar op amps and had lower bias current below 70˚C. But FETs exhibit higher offset voltage and drift than bipolars. Long-term stability is also about an order of magnitude worse. Typically, this drift is 100 µV/year, but a small percentage could be as bad as 1 mV. Laser trimming and other process improvements have lowered initial offset but have not eliminated the drift problem.

The new IC is an extension of super-gain bipolar techniques. As can be seen from *Figure 1*, it provides low bias currents over a −55˚C to 125˚C temperature range. The offset current is so low as to be lost in the noise. This level of performance has previously been unavailable for either low-cost industrial designs or high reliability military/space applications.

This low bias current has not been obtained at the expense of offset voltage or drift. Typical offset voltage is under a millivolt and provision is made for on-water trimming to get it below 100 µV. The low drift exhibited in *Figure 2* indicates that the circuit is inherently balanced for exceptionally low drift, typically 1 µV/˚C below 100˚C.

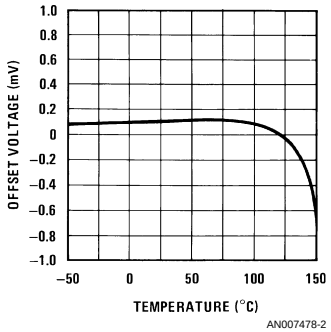

FIGURE 2. Bipolar transistors have inherently low offset voltage and drift. The low drift of the LM11 over a wide temperature range shows that there are no design problems degrading performance.

THE NEW OP AMP

The LM11 is, in essence, a refinement of the LM108. A modified Darlington input stage has been added to reduce bias currents. With a standard Darlington, one transistor is biased with the base current of the other. This degrades dc amplifier performance because base current is noisy, subject to wide variation and generally unpredictable.

Supplying a bleed current greater than the base current, as shown in *Figure 3*, removes this objection. The 60 nA provided is considerably in excess of the 1 nA base current. The bleed current is made to vary as absolute temperature to maintain constant impedance at the emitters of Q1 and Q2. This stabilizes frequency response and also reduces the

AN-241

thermal variation of bias current. Parasitic capacitances of the current generator have been bootstrapped so that the 0.3 V/µs slew rate of the basic amplifier is unaffected.

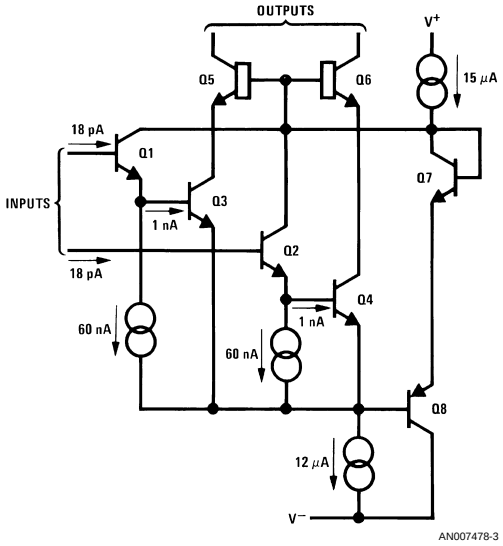

FIGURE 3. Modifying Darlington with bleed current reduces offset voltage, drift and noise. Unique circuitry provides well-controlled current with minimal stray capacitance so that speed of the basic amplifier is unaffected.

Results to date suggest that the base currents of this modified Darlington input are better matched than the simple differential amplifier. In fact, offset current is so low as to be unmeasurable on production test systems. Therefore, guaranteed limits are determined by the test equipment rather than the IC.

NOISE

Operating transistors at very low currents does increase noise. Thus, the LM11 is about a factor of four noisier than the LM108. But the low frequency noise, plotted in *Figure 4*, is still slightly less than that of FET amplifiers. Long-term measurements indicate that the offset voltage shift is under 10 µV.

In contrast to the noise voltage, low frequency noise current is subject to greater unit-to-unit variation. Generally, it is below 1 pA, peak-to-peak, about the same magnitude as the offset current.

With the LM11, both voltage and current related dc errors have been reduced to the point where overall circuit performance could well be noise limited, particularly in limited temperature range applications.

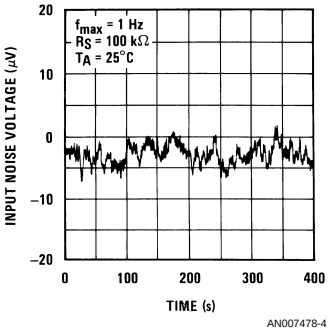

FIGURE 4. Lower operating currents increase noise, but low frequency noise is still slightly lower than IC FET amplifiers. Long-term stability is much improved.

RELIABILITY

The reliability of the LM11 is not expected to be substantially different than the LM108, which has been used extensively in military and space applications. The only significant difference is the input stage. The low current nodes introduced here might possibly be a problem were they not bootstrapped, biased and guarded to be virtually unaffected by both bulk and surface leakages. This opinion is substantiated by preliminary life-test data.

This IC could, in fact, be expected to improve reliability when used to replace discrete or hybrid amplifiers that use selected components and have been trimmed and tweaked to give the required performance.

From an equipment standpoint, reliability analysis of insulating materials, surface contamination, cleaning procedures, surface coating and potting are at least as important as the IC and other components. These factors become more important as impedance levels are raised. But this should not discourage designers. If poor insulation and contamination cause a problem when impedance levels are raised by an order of magnitude, it is best found out and fixed.

Even so, it may not be advisable to take advantage of the full potential of the LM11 in all cases, especially when hostile environments are involved. For example, there should be no great difficulty in finding an LM11 with offset current less than 5 pA over a −55˚C to 125˚C temperature range. But anyone

designing high-reliability equipment that is going to be in trouble if combined leakages are greater than 10 pA at 125°C had best know what he is about.

ELECTRICAL GUARDING

The effects of board leakage can be minimized using an old trick known as guarding. Here the input circuitry is surrounded by a conductive trace that is connected to a low impedance point at the same potential as the inputs. The electrical connection of the guard for the basic op amp configurations is shown in *Figure 5*. The guard absorbs the leakage from other points on the board, drastically reducing that reaching the input circuitry.

To be completely effective, there should be a guard ring on both sides of the printed-circuit board. It is still recommended for single-sided boards, but what happens on the unguarded side is difficult to analyze unless Teflon inserts are used on the input leads. Further, although surface leakage can be virtually eliminated, the reduction in bulk leakage is much less. The reduction in bulk leakage for double-sided guarding is about an order of magnitude, but this depends on board thickness and the width of the guard ring. If there are bulk leakage problems, Teflon inserts on the through holes and Teflon or kel-F standoffs for terminations can be used. These two materials have excellent surface properties without surface treatment even in high-humidity environments.

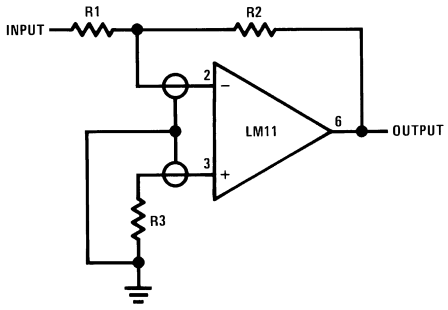

a. inverting amplifier

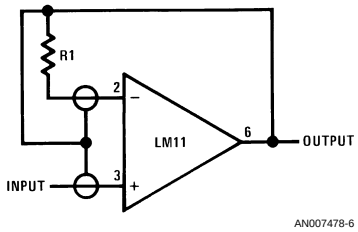

b. follower

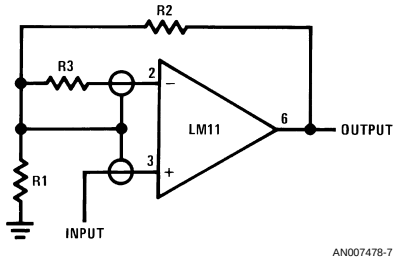

c. non-inverting amplifier

FIGURE 5. Input guarding for various op amp connections. The guard should be connected to a point at the same potential as the inputs with a low enough impedance to absorb board leakage without introducing excessive offset.

An example of a guarded layout for the metal-can package is shown in *Figure 6*. Ceramic and plastic dual-in-line packages are available for critical applications with guard pins adjacent to the inputs both to facilitate board layout and to reduce package leakage. These guard pins are not internally connected.

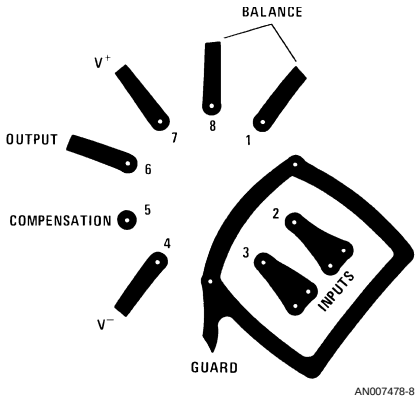

Bottom View

FIGURE 6. Input guarding can drastically reduce surface leakage. Layout for metal can is shown here. Guarding both sides of board is required. Bulk leakage reduction is less and depends on guard ring width.

SIGNAL CABLES

It is advisable to locate high impedance amplifiers as close as possible to the signal source. But sometimes connecting lines cannot be avoided. Coaxially shielded cables with good insulation are recommended. Polyethelene or virgin (not reconstituted) Teflon is best for critical applications.

In addition to potential insulation problems, even short cable runs can reduce bandwidth unacceptably with high source resistances. These problems can be largely avoided by bootstrapping the cable shield. This is shown for the follower connection in *Figure 7*. In a way, bootstrapping is positive feedback; but instability can be avoided with a small capacitor on the input.

CABLE BOOTSTRAPPING

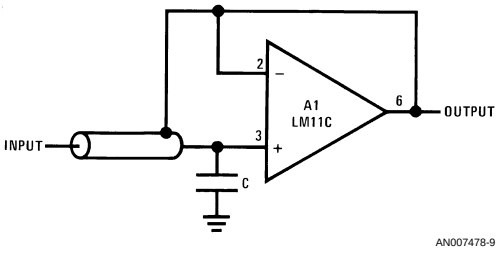

FIGURE 7. Bootstrapping input shield for a follower reduces cable capacitance, leakage and spurious voltages from cable flexing. Instability can be avoided with small capacitor on input.

With the summing amplifier, the cable shield is simply grounded, with the summing node at virtual ground. A small feedback capacitor may be required to insure stability with the added cable capacitance. This is shown in *Figure 8*.

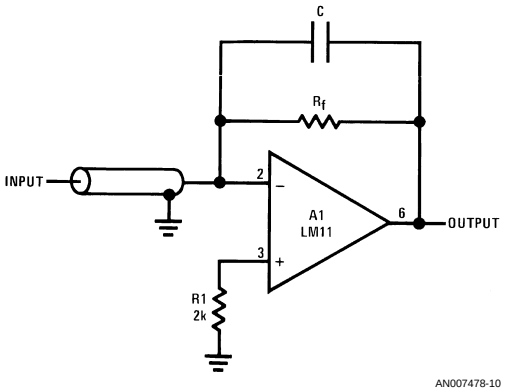

FIGURE 8. With summing amplifier, summing node is at virtual ground so input shield is best grounded. Small feedback capacitor insures stability.

An inverting amplifier with gain may require a separate follower to drive the cable shield if the influence of the capacitance, between shield and ground, on the feedback network cannot be accounted for.

High impedance circuits are also prone to mechanical noise (microphonics) generated by variable stray capacitances. A capacitance variation will generate a noise voltage given by

$$e_n = \frac{\Delta C}{C} V,$$

where V is the dc bias on the capacitor. Therefore, the wiring and components connected to sensitive nodes should be mechanically rigid.

This is also a problem with flexible cables, in that bending the cable can cause a capacitance change. Bootstrapping the shield nearly eliminates dc bias on the cable, minimizing the voltage generated. Another problem is electrostatic charge created by friction. Graphite lubricated Teflon cable will reduce this.

SWITCH LEAKAGE

Semiconductor switches with leakage currents as low as the bias current of the LM11 are not generally available when operation much above 50˚C is involved. The sample-and-hold circuit in *Figure 10* shows a way around this problem. It is arranged so that switch leakage does not reach the storage capacitor.

Isolating leakage current requires that two switches be connected in series. The leakage of the first, Q1, is absorbed by R1 so that the second, Q2, only has the offset voltage of the op amp across its junctions. This can be expected to reduce leakage by at least two orders of magnitude. Adjusting the op

amp offset to zero at the maximum operating temperature will give the ultimate leakage reduction, but this is not usually required with the LM11.

MOS switches with gate-protection diodes are preferred in production situations as they are less sensitive to damage from static charges in handling. If used, D1 and R2 should be included to remove bias from the protection diode during hold. This may not be required in all cases but is advised since leakage from the protection diode depends on the internal geometry of the switch, something the designer does not normally control.

A junction FET could be used for Q1 but not Q2 because there is no equivalent to the enhancement mode MOSFET. The gate of a JFET must be reverse biased to turn it off, and leakage on its output cannot be avoided.

HIGH-VALUE RESISTORS

Using op amps at very high impedance levels can require unusually large resistor values. Standard precision resistors are available up to 10 MΩ. Resistors up to 1 GΩ can be obtained at a significant cost premium. Larger values are quite expensive, physically large and require careful handling to avoid contamination. Accuracy is also a problem. There are techniques for raising effective resistor values in op amp circuits. In theory, performance is degraded; in practice, this may not be the case.

With a buffer amplifier, it is sometimes desirable to put a resistor to ground on the input to keep the output under control when the signal source is disconnected. Otherwise it will saturate. Since this resistor should not load the source, very large values can be required in high-impedance circuits.

Figure 9 shows a voltage follower with a 1 GΩ input resistance built using standard resistor values. With the input disconnected, the input offset voltage is multiplied by the same factor as R2; but the added error is small because the offset voltage of the LM11 is so low. When the input is connected to a source less than 1 GΩ, this error is reduced. For an ac-coupled input, a second 10 MΩ resistor could be connected in series with the inverting input to virtually eliminate bias current error; bypassing it would give minimal noise.

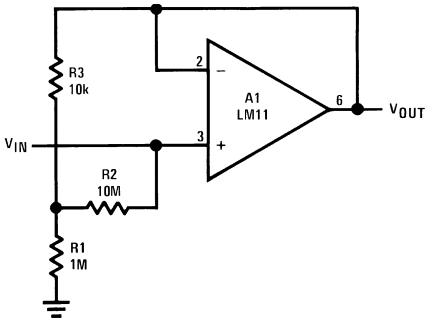

FIGURE 9. Follower input resistance is 1 GΩ. With the input open, offset voltage is multiplied by 100, but the added error is not great because the op amp offset is low.

The voltage-to-current converter in *Figure 11* uses a similar method to obtain the equivalent of a 10 GΩ feedback resistor. Output offset is reduced because the error can be made dependent on offset current rather than bias current. This would not be practical with large value resistors because of cost, particularly for matched resistors, and because the summing node would be offset several hundred millivolts from ground. In *Figure 11*, this offset is limited to several millivolts. In addition, the output can be nulled with the usual balance potentiometer. Further, gain trimming is easily done.

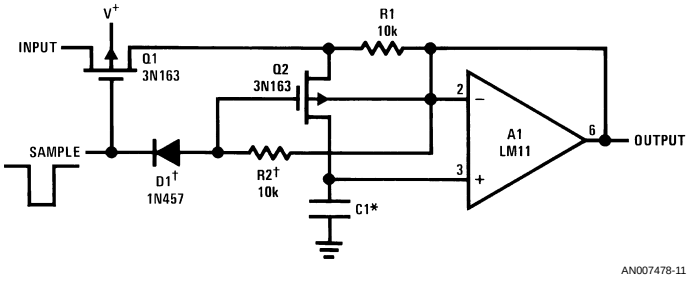

* polystyrene or Teflon

† required if protected-gate switch is used

FIGURE 10. Switch leakage in this sample and hold does not reach storage capacitor. If Q2 has an internal gate-protection diode, D1 and R2 must be included to remove bias from its junction during hold.

Resistance Multiplication

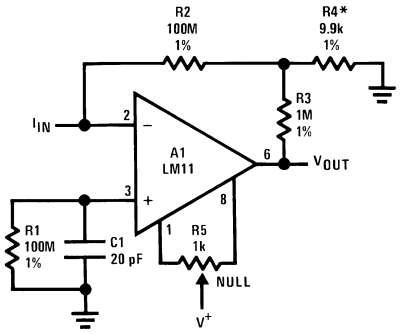

AN007478-13

FIGURE 11. Equivalent feedback resistance is 10 GΩ, but only standard resistors are used. Even though the offset voltage is multiplied by 100, output offset is actually reduced because error is dependent on offset current rather than bias current. Voltage on summing junction is less than 5 mV.

This circuit would benefit from lower offset current than can be tested and guaranteed with automatic test equipment. But there should be no problem in selecting a device for critical applications.

CAPACITORS

Op amp circuits impose added requirements on capacitors, and this is compounded with high-impedance circuitry. Frequency shaping and charge measuring circuits require control of the capacitor tolerance, temperature drift and stability with temperature cycling. For smaller values, NPO ceramic is best while a polystyrene-polycarbonate combination gives good results for larger values over a −10˚C to 85˚C range.

Dielectric absorption can also be a problem. It causes a capacitor that has been quick-charged to drift back toward its previous state over many milliseconds. The effect is most noticeable in sample-and-hold circuits. Polystyrene, Teflon and NPO ceramic capacitors are most satisfactory in this regard. Choice depends mainly on capacitance and temperature range.

Insulation resistance can clearly become a problem with high-impedance circuitry. Best performer is Teflon, with polystyrene being a good substitute below 85˚C. Mylar capacitors should be avoided, especially where higher temperatures are involved.

Temperature changes can also alter the terminal voltage of a capacitor. Because thermal time constants are long, this is only a problem when holding intervals are several minutes or so. The effect is reported to be as high as 10 mV/˚C, but Teflon capacitors that hold it to 0.5 mV/˚C are available*.

An op amp with lower bias current can ease capacitor problems, primarily by reducing size. This is obvious with a sample-and-hold because the capacitor value is determined by the hold interval and the amplifier bias current. The circuit in *Figure 12* is another example. An RC time constant of more than a quarter hour is obtained with standard component values. Even when such long time constants are not required, reducing capacitor size to where NPO ceramics can be used is a great aid in precision work.

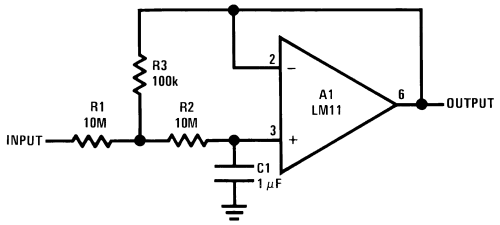

AN007478-14

$$\tau = \frac{R1\,C}{R3}(R2 + R3)$$

$$\Delta V_{OUT} = \frac{R1 + R3}{R3}(I_B\,R2 + V_{OS})$$

FIGURE 12. This circuit multiplies RC time constant to 1000 seconds and provides low output impedance. Cost is lowered because of reduced resistor and capacitor values.

Note: *Component Research Co., Inc., Santa Monica, California.

CONCLUSIONS

A low cost IC op amp has been described that not only has low offset voltage but also advances the state of the art in reducing input current error, particularly at elevated temperatures. Designers of industrial as well as military/space equipment can now work more freely at high impedance levels.

Although high-impedance circuitry is more sensitive to board leakages, wiring capacitances, stray pick-up and leakage in other components, it has been shown how input guarding, bootstrapping, shielding and leakage isolation can largely eliminate these problems.

ACKNOWLEDGMENT

The author would like to acknowledge the assistance of the staff at National Semiconductor in implementing this design and sorting out the application problems. Discussions with Bob Dobkin, Bob Pease, Carl Nelson and Mineo Yamatake have been most helpful.

APPENDIX

SUPER-GAIN TECHNIQUES

Super-gain transistors are not new, having been developed for the LM102/LM110 voltage followers in 1967 and later used on the LM108 general-purpose op amp. They are similar to regular transistors, except that they are diffused for high current gains (2,000–10,000) at the expense of breakdown voltage. A curve-tracer display of a typical device is shown in *Figure 13*. In an IC, super-gain transistors can be

made simultaneously with standard transistors by including a second, light base predeposition that is diffused less deeply.

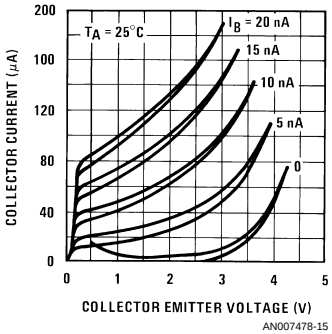

FIGURE 13. Curve tracer display of a super-gain transistor

Super-gain transistors can be connected in cascode with regular transistors to form a composite device with both high gain and high breakdown. The simplified schematic of the LM108 input stage in *Figure 14* shows how it is done. A common base pair, Q3 and Q4, is bootstrapped to the input transistors, Q1 and Q2, so that the latter are operated at nearly zero collector-base voltage, no matter what the input common-mode. The regular NPN transistors are distinguished by drawing them with wider base regions.

Operating the input transistors at very low collector-base voltage has the added advantage of drastically reducing collector-base leakage. In this configuration bipolar transistors are affected little by the leakage currents that limit performance of FET amplifiers.

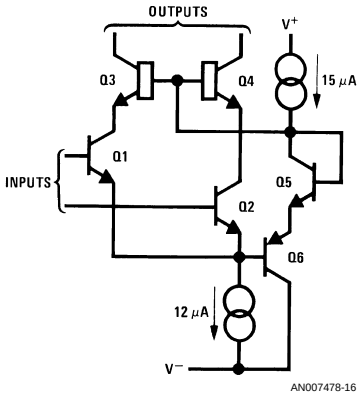

FIGURE 14. A bootstrapped input stage

***See Addendum at the End of Application Note 242.**

Book Extract End

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component in any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
Americas
Tel: 1-800-272-9959
Fax: 1-800-737-7018
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 1 80-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 1 80-530 85 85
English Tel: +49 (0) 1 80-532 78 32
Français Tel: +49 (0) 1 80-532 93 58
Italiano Tel: +49 (0) 1 80-534 16 80

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: sea.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5620-6175
Fax: 81-3-5620-6179

PrintDate=1998/04/24 PrintTime=15:10:20 39739 an007478 Rev. No. 3 cmserv **Proof**

Applying a New Precision Op Amp

National Semiconductor
Application Note 242
Robert J. Widlar
Bob Pease
Mineo Yamatake
April 1980

***Abstract**: A new bipolar op amp design has advanced the state of the art by reducing offset voltage and bias current errors. Its characteristics are described here, indicating an ultimate input resolution of 10 µV and 1 pA under laboratory conditions. Practical circuits for making voltmeters, ammeters, differential instrumentation amplifiers and a variety of other designs that can benefit from the improved performance are covered in detail. Methods of coupling the new device to existing fast amplifiers to take advantage of the best characteristics of both, even in follower applications, are explored.*

INTRODUCTION

A low cost, mass-produced op amp with electrometer-type input currents combined with low offset voltage and drift is now available. Designated the LM11, this IC can minimize production problems by providing accuracy without adjustments, even in high-impedance circuitry. On the other hand, if pushed to its full potential, what has been impossible in the past becomes entirely practical.

Significantly, the LM11 is not restricted to commercial and industrial use. Devices can be completely specified over a −55˚C to +125˚C range. Preliminary data indicates that reliability is the same as standard ICs qualified for military and space applications.

The essential details of the design along with an introduction to the peculiarities of high-impedance circuits have been presented elsewhere.* This will be expanded here. Practical circuitry that reduces effective bias current for those applications where performance cannot be made dependent on offset current are described. In addition, circuits combining the DC characteristics of the new part with the AC performance of existing fast amplifiers will be shown. This will be capped with a number of practical designs to provide some perspective into what might be done.

DC ERRORS

Barring the use of chopper or reset stabilization, the best offset voltage, drift and long-term stability are obtained using bipolar transistors for the op amp input stage. This has been done with the LM11. On-wafer trimming further improves performance. Typically, a 100 µV offset with 1 µV/˚C drift results.

Transistors with typical current gains of 5000 have been used in the manufacture of the LM11. The input stage employs a Darlington connection that has been modified so that offset voltage and drift are not degraded. The typical input currents, plotted in *Figure 1*, demonstrate the value of the approach.

Note: *R. J. Widlar, "Working with High Impedance Op Amps", National Semiconductor AN-241.

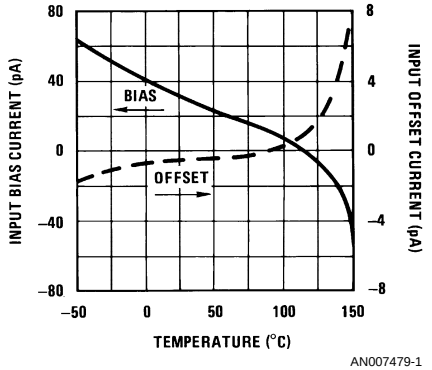

FIGURE 1. Below 100˚C, bias current varies almost linearly with temperature. This means that simple circuitry can be used for compensation. Offset current is unusually low.

The offset current of this op amp is so low that it cannot be measured on existing production test equipment. Therefore, it probably cannot be specified tighter than 10 pA. For critical applications, the user should have little difficulty in selecting to a tighter limit.

The bias current of the LM11 equals that of monolithic FET amplifier at 25˚C. Unlike FETs, it does not double every 10˚C. In fact, the drift over a −55˚C to +125˚C temperature range is about the same as that of a FET op amp during normal warm up.

Other characteristics are summarized in *Table 1*. It can be seen that the common-mode rejection, supply-voltage rejection and voltage gain are high enough to take full advantage of the low offset voltage. The unspectacular 0.3V/µs slew rate is balanced by the 300 µA current drain.

TABLE 1. Typical characteristics of the LM11 for T_j = 25˚C and V_S = ±15V. Operation is specified down to V_S = ±2.5V.

Parameter	Conditions	Value
Input Offset Voltage		100 µV
Input Offset Current		500 fA
Input Bias Current		25 pA
Input Noise Voltage	0.01 Hz ≤ f ≤ 10 Hz	8 µVpp
Input Noise Current	0.01 Hz ≤ f ≤ 10 Hz	1 pApp
Long Term Stability	T_j = 25˚C	10 µV
Offset Voltage Drift	−55˚C ≤ T_j ≤ 125˚C	1 µV/˚C
Offset Current Drift	−55˚C ≤ T_j ≤ 125˚C	20 fA/˚C
Bias Current Drift	−55˚C ≤ T_j ≤ 125˚C	500 fA/˚C

TABLE 1. Typical characteristics of the LM11 for T_j = 25˚C and V_S = ±15V. Operation is specified down to V_S = ±2.5V. (Continued)

Parameter	Conditions	Value
Voltage Gain	V_{OUT} = ±12V, I_{OUT} = ±0.5 mA	1,200V/mV
	V_{OUT} = ±12V, I_{OUT} = ±2 mA	300V/mV
Common-Mode Rejection	−12.5V ≤ V_{CM} ≤ 14V	130 dB
Supply-Voltage Rejection	±2.5V ≤ V_S ≤ ±20V	118 dB
Slew Rate		0.3V/µs
Supply Current		300 µA

As might be expected, the low bias currents were obtained with some sacrifice in noise. But the low frequency noise voltage is still a bit less than a FET amplifier and probably more predictable. The latter is important because this noise cannot be tested in production. Long term measurements have not indicated any drift in excess of the noise. This is not the case for FETs.

It is worthwhile noting that the drift of offset voltage and current is low enough that DC accuracy is noise limited in room-temperature applications.

BIAS CURRENT COMPENSATION

The LM11 can operate from MΩ source resistances with little increase in the equivalent offset voltage, as can be seen in *Figure 2*. This is impressive considering the low initial offset voltage. The situation is much improved if the design can be configured so that the op amp sees equal resistance on the two inputs. However, this cannot be done with all circuits. Examples are integrators, sample and holds, logarithmic converters and signal-conditioning amplifiers. And even though the LM11 bias current is low, there will be those applications where it needs to be lower.

Referring back to *Figure 1*, it can be seen that the bias current drift is essentially linear over a −50˚C to +100˚C range. This is a deliberate consequence of the input stage design. Because of it, relatively simple circuitry can be used to develop a compensating current.

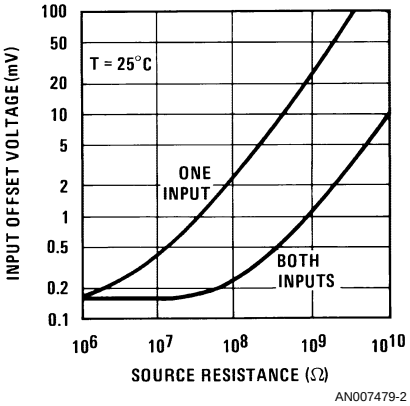

FIGURE 2. The LM11 operates from MΩ source resistances with little DC error. With equal source resistances, accuracy is essentially limited by low frequency current noise.

Bias current compensation is not new, but making it effective with even limited temperature excursions has been a problem. An early circuit suggested for bipolar ICs is shown in *Figure 3(a)*. The compensating current is determined by the diode voltage. This does not vary as rapidly with temperature as bias current nor does it match the usual non-linearities.

With the improved circuit in *Figure 3(b)*, the temperature coefficient can be increased by using a transistor and including R2. The drop across R2 is nearly constant with temperature. The voltage delivered to the potentiometer has a 2.2 mV/˚C drift while its magnitude is determined by R2. Thus, as long as the bias current varies linearly with temperature, a value for R2 can be found to effect compensation.

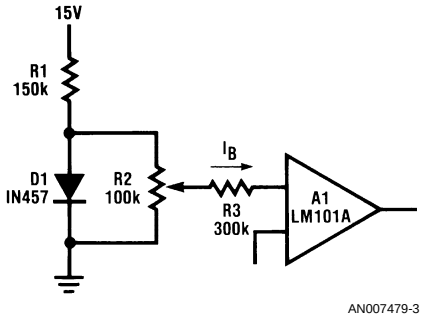

a. original circuit

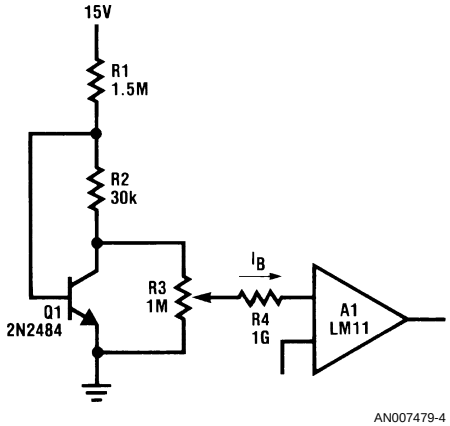

b. improved version

FIGURE 3. Bias-current compensation. With the improved version, the temperature coefficient of the compensating current can be varied with R2. It is effective only if bias current has linear, negative temperature coefficient.

In production, altering resistors based on temperature testing is to be avoided if at all possible. Therefore, the results that can be obtained with simple nulling at room temperature and a fixed value for R2 are of interest. *Figure 4* gives this data for a range of parts with different initial bias currents. This was obtained from pre-production and initial-production

runs. The bias current variations were the result of both h_{FE} variations and changes in internal operating currents and represent the worst as well as best obtained. They are therefore considered a realistic estimate of what would be encountered among various production lots.

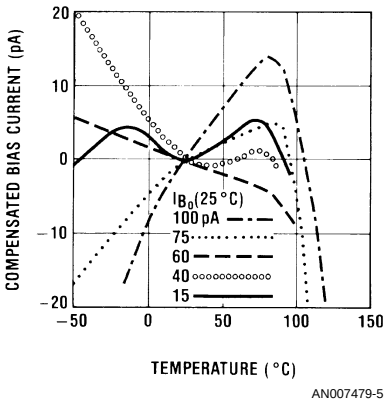

FIGURE 4. Compensated bias current for five representative units with a range of initial bias currents. The circuit in *Figure 3*(b) was used with balancing at 25°C. High drift devices could be improved further by altering R2.

Little comment need be made on these results, except that the method is sufficiently predictable that another factor of five reduction in worst case bias current could be made by altering R2 based on the results of a single temperature run.

One disadvantage of the new circuit is that it is more sensitive to supply variations than the old. This is no problem if the supplies are regulated to 1%. But with worst regulation it suffers because, with R2, the transistor no longer functions as a regulator and because much tighter compensation is obtained.

The circuit in *Figure 5* uses pre-regulation to solve this problem. The added reference diode has a low breakdown so that the minimum operating voltage of the op amp is unrestricted. Because of the low breakdown, the drop across R3 can no longer be considered constant. But it will vary linearly with temperature, so this is of no consequence. The fact that this reference can be used for other functions should not be overlooked because a regulated voltage is frequently required in designs using op amps.

In *Figure 5* a divider is used so that the resistor feeding the compensating current to the op amp can be reduced. There will be an error current developed for any offset voltage change across R6. This should not be a problem with the LM11 because of its low offset voltage. But for tight compensation, mismatch in the temperature characteristics of R4 and R5 must be considered.

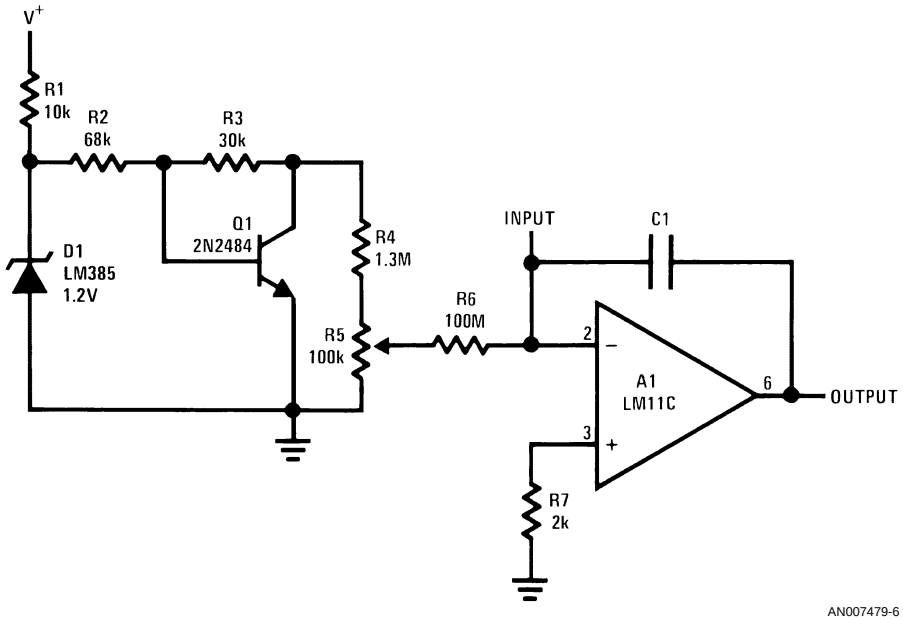

FIGURE 5. Bias current compensation for use with unregulated supplies. Reference voltage is available for other circuitry.

Bias current compensation is more difficult for non-inverting amplifiers because the common-mode voltage varies. With a voltage follower, everything can be bootstrapped to the output and powered by a regulated current source, as shown in *Figure 6*. The LM334 is a temperature sensor. It regulates against voltage changes and its output varies linearly with temperature, so it fits the bill.

Although the LM334 can accommodate voltage changes fast enough to work with the LM11, it is not fast enough for the high-speed circuits to be described. But compensation can still be obtained by using the zener diode pre-regulator bootstrapped to the output and powered by either a resistor or FET current source. The LM385 fits well here because both the breakdown voltage and minimum operating current are low.

With ordinary op amps, the collector base voltage of the input transistors varies with the common-mode voltage. A 50% change in bias current over the common-mode range is not unusual, so compensating the bias current of a follower has limited value. However, the bootstrapped input stage of the LM11 reduces this to about 2 pA for a ±20V common-mode swing, giving a 2 x $10^{13}\Omega$ common-mode input resistance.

FAST AMPLIFIERS

A precision DC amplifier, although slow, can be used to stabilize the offset voltage of a less precise fast amplifier. As shown in *Figure 7*, the slow amplifier senses the voltage across the input terminals and supplies a correction signal to the balance terminals of the fast amplifier. The LM11 is particularly interesting in this respect as it does not degrade the input bias current of the composite even when the fast amplifier has a FET input.

Surprisingly, with the LM11, this will work for both inverting and non-inverting connections because its common-mode slew recovery is a lot faster than that of the main loop. This was accomplished, even with circuitry running under 100 nA, by proper clamping and by bootstrapping of internal stray capacitances.

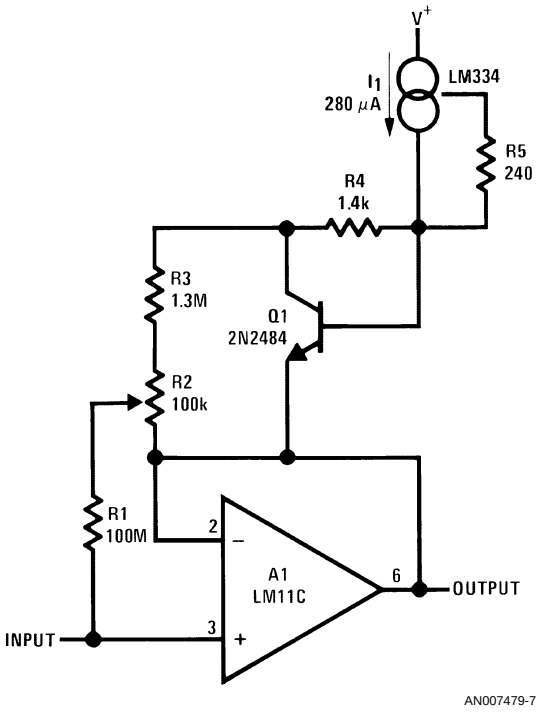

FIGURE 6. This circuit shows how bias current compensation can be used on a voltage follower.

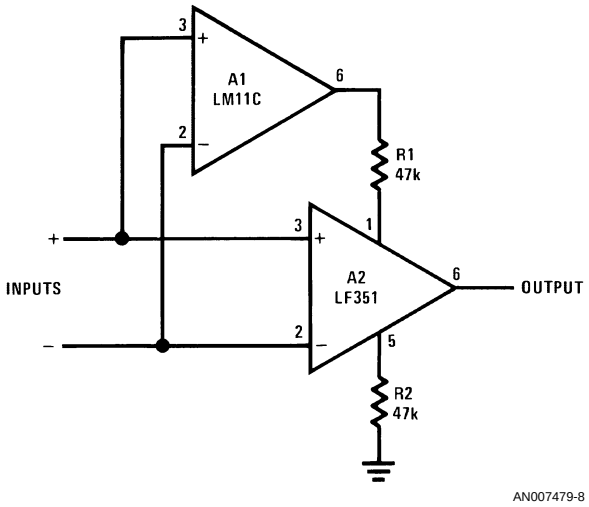

FIGURE 7. A slow amplifier can be used to null the offset of a fast amplifier.

An optimized circuit for the inverting amplifier connection is shown in *Figure 8*. The LM11 is DC coupled to the input and drives the balance terminals of the fast amplifier. The fast amplifier is AC coupled to the input and drives the output. This isolates FET leakage from the input circuitry.

As can be seen, the method of coupling into the balance terminals will vary depending on the internal configuration of the fast amplifier. If the quiescent voltage on the balance terminals is beyond the output swing of the LM11, a differential coupling must be used, as in *Figure 8(a)*. A lead capacitor, C2, reduces the AC swing required at the LM11 output. The clamp diode, D1, insures that the LM11 does not overdrive the fast amplifier in slew.

If the quiescent voltage on the balance terminals is such that the LM11 can drive directly, the circuit in *Figure 8(b)* can be used. A clamp diode from the other balance terminal to internal circuitry of the LM11 keeps the output from swinging too far from the null value, and a resistor may be required in series with its output to insure stability.

Measurements indicate that the slew rate of the fast amplifier is unimpaired, as is the settling time to 1 mV for a 20V output excursion. If the composite amplifier is overdriven so that the output saturates, there will be an added recovery delay because the coupling capacitor to the fast amplifier takes on a charge with the summing node off ground. Therefore, C1 should be made as small as possible. But going below the values given may introduce gain error.

If the bias current of the fast amplifier meets circuit requirements, it can be direct-coupled to the input. In this case, offset voltage is improved, not bias current. But overload recovery can be reduced. The AC coupling to the fast-amplifier input might best be eliminated for limited-temperature-range operation.

This connection also increases the open-loop gain beyond that of the LM11, particularly since two-pole compensation can be effected to reduce AC gain error at moderate frequencies. The DC gains measured showed something in excess of 140 dB.

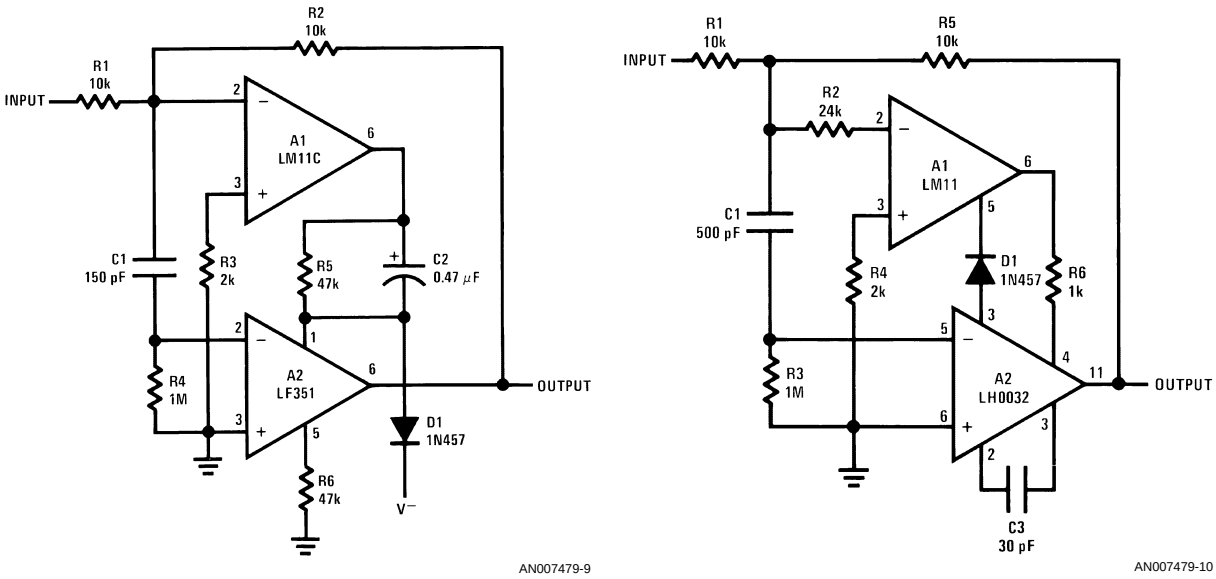

a. with standard BI-FET **b. with fast hybrid**

FIGURE 8. These inverters have bias current and offset voltage of LM11 along with speed of the FET op amps. Open loop gain is about 140 dB and settling time to 1 mV about 8 µs. Excess overload-recovery delay can be eliminated by directly coupling the FET amplifier to summing node.

A voltage-follower connection is given in *Figure 9*. The coupling circuitry is similar, except that R5 was added to eliminate glitches in slew. Overload involves driving the fast amplifier outside its common-mode range and should be avoided by limiting the input. Thus, AC coupling the fast amplifier is less a problem. But the repetition frequency of the input signal must also be limited to 10 kHz for ±10V swing. Higher frequencies produce a DC error, believed to result from rectification of the input signal by the voltage sensitive input capacitance of the FET amplifier used. A fast bipolar amplifier like the LM118 should work out better in this respect. To avoid confusion, it should be emphasized that this problem is related to repetition frequency rather than rise time.

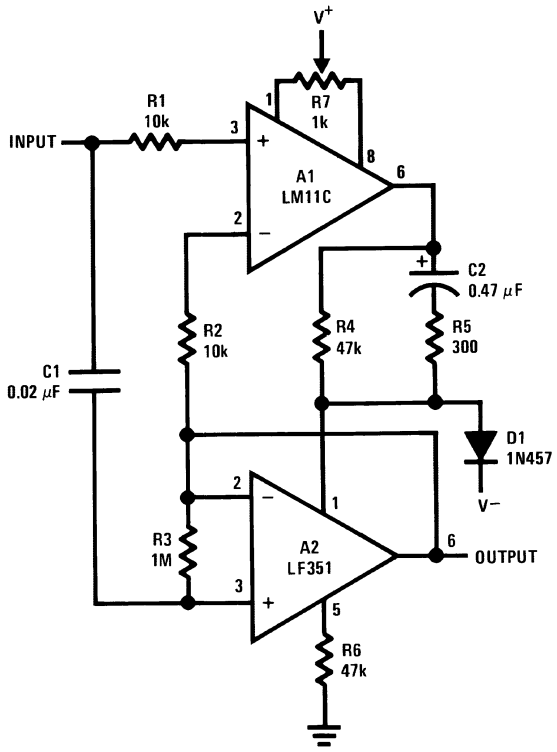

FIGURE 9. Follower has 10 µs settling to 1 mV, but signal repetition frequency should not exceed 10 kHz if the FET amplifier is AC coupled to input. The circuit does not behave well if common-mode range is exceeded.

A precision DC amplifier with a 100 MHz gain-bandwidth product is shown in *Figure 10*. It has reasonable recovery (~7 µs) from a 100% overload; but beyond that, AC coupling to the fast amplifier causes problems. Alone, the gain error and thermal feedback of the LH0032 are about 20 mV, input referred, for ±10V output swing. Adding the LM11 reduces this to microvolts.

PICOAMMETER

Ideally, an ammeter should read zero with no input current and have no voltage drop across its inputs even with full-scale deflection. Neither should spurious indications nor inaccuracy result from connecting it to a low impedance. Meeting all these requirements calls for a DC amplifier, and one in which both bias current and offset voltage are controlled.

The summing amplifier connection is best for measuring current, because it minimizes the voltage drop across the input terminals. However, when the inputs are shorted, the output state is indeterminate because of offset voltage. Adding degeneration as shown in *Figure 11* takes care of this problem. Here, R2 is the feedback resistor for the most sensitive range, while R1 is chosen to get the meter deflection out of the noise with a shorted input. Adding the range resistor, as shown, does not affect the degeneration, so that there is minimal drop across the input for full-scale on all ranges.

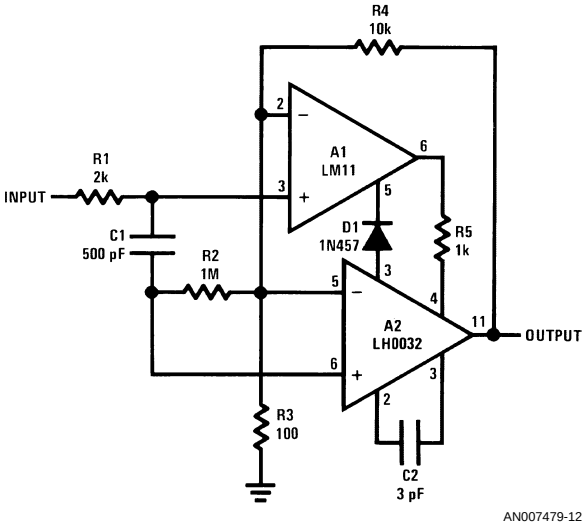

FIGURE 10. This 100X amplifier has small and large signal bandwidth of 1 MHz. The LM11 greatly reduces offset voltage, bias current and gain error. Eliminating long recovery delay for greater than 100% overload requires direct coupling of A2 to input.

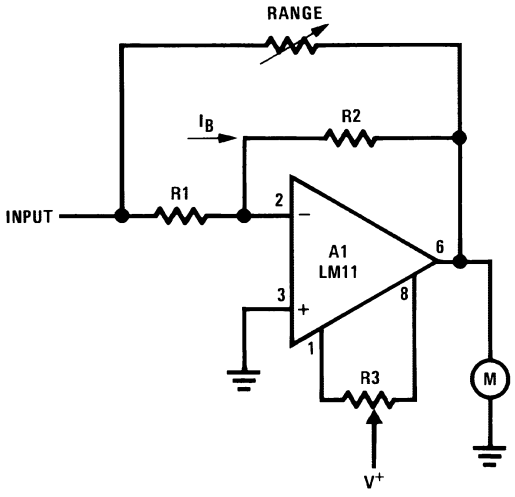

FIGURE 11. An ammeter that has constant voltage drop across its input at full-scale, no matter what the range. It can have a reasonably-behaved output even with shorted inputs, yet a maximum drop of ten times the op amp noise voltage.

The complete meter circuit in *Figure 12* uses a different scheme. A floating supply is available so that the power ground and the signal ground can be separated with R12. At full-scale, the meter current plus the measured current flow

through this resistor, establishing the degeneration. This method has the advantage of allowing even-value range resistors on the lower ranges but increases degeneration as the measured current approaches the meter current.

Bias-current compensation is used to increase the meter sensitivity so there are two zeroing adjustments; current balancing, that is best done on the most sensitive range where it is needed, and voltage balancing that should be done with the inputs shorted on a range below 100 µA, where the degeneration is minimal.

With separate grounds, error could be made dependent on offset current. This would eliminate bias-current compensation at the expense of more complicated range switching.

The op amp input has internal, back-to-back diodes across it, so R6 is added to limit current with overloads. This type of protection does not affect operation and is recommended whenever more than 10 mA is available to the inputs. The output buffers are added so that input overloads cannot drag down the op amp output on the least-sensitive range, giving a false meter indication. These would not be required if the maximum input current did not approach the output current limit of the op amp.

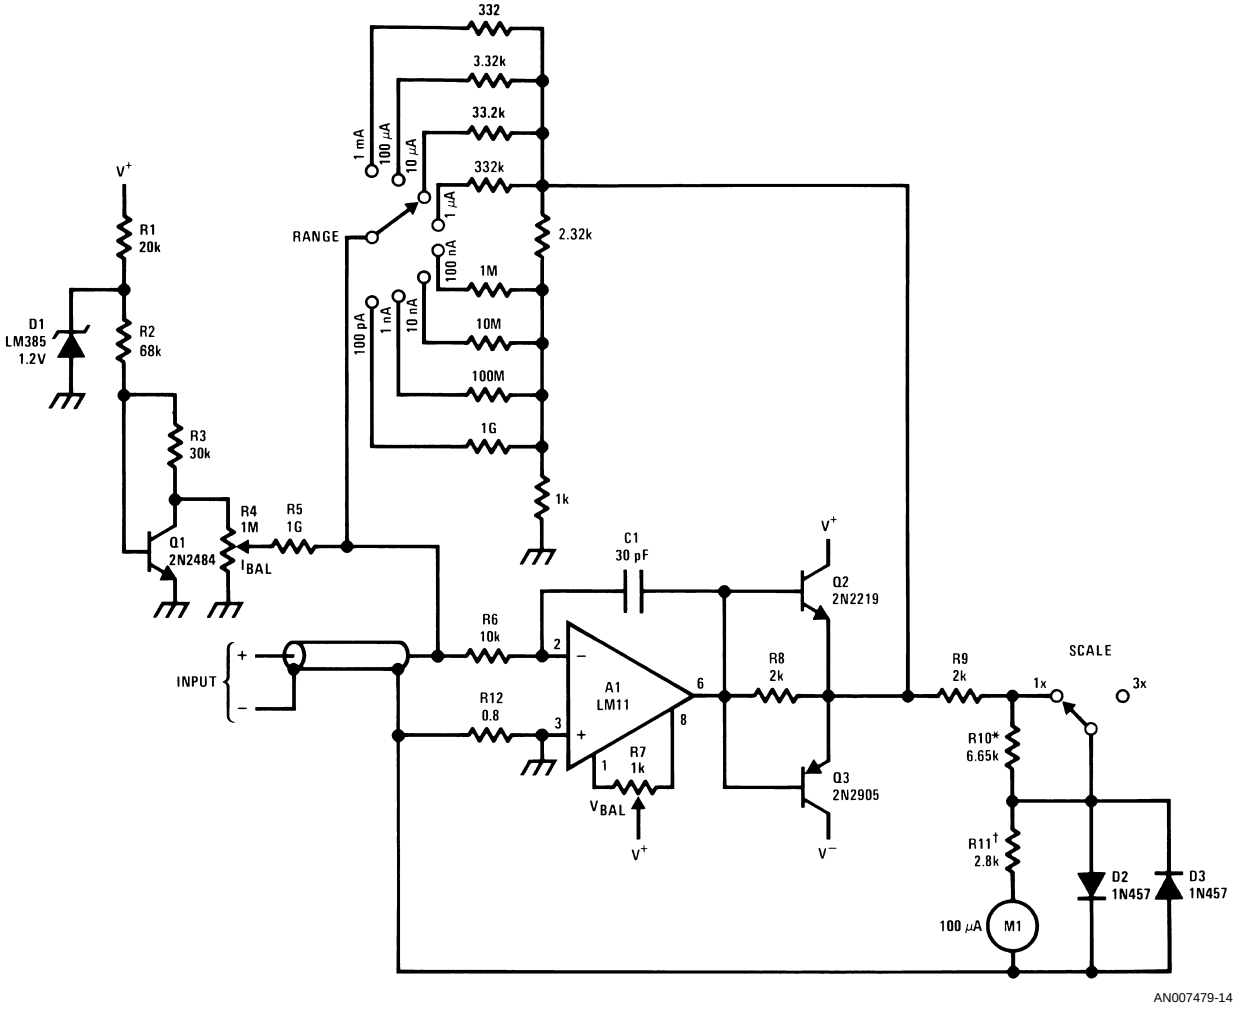

FIGURE 12. Current meter ranges from 100 pA to 3 mA, full-scale. Voltage across input is 100 µV at lower ranges rising to 3 mV at 3 mA. Buffers on op amp are to remove ambiguity with high-current overload. Output can also drive DVM or DPM.

MILLIVOLTMETER

An ideal voltmeter has requirements analogous to those discussed for the ammeter, and *Figure 13* shows a circuit that will satisfy them. In the most-sensitive position, the range resistor is zero and the input resistance equals R1. As voltage measurement is desensitized by increasing the range resistor, the input resistance is also increased, giving the maximum input resistance consistent with zero stability with the input open. Thus, at full-scale, the source will be loaded by whatever multiple of the noise current is required to give the desired open-input zero stability.

This technique is incorporated into the voltmeter circuit in *Figure 14* to give a 100 MΩ input resistance on the 1 mV scale rising to 300 GΩ on the 3V scale. The separation of power and signal grounds has been used here to simplify

bias-current compensation. Otherwise, a separate op amp would be required to bootstrap the compensation to the input.

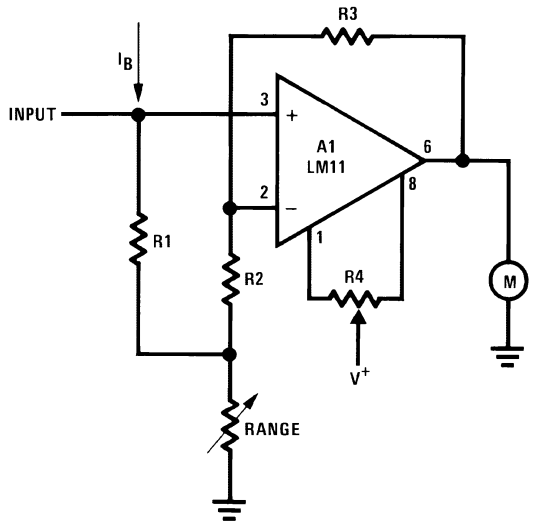

FIGURE 13. This voltmeter has constant full-scale loading independent of range. This can be only ten times the noise current, yet the output will be reasonably behaved for open input.

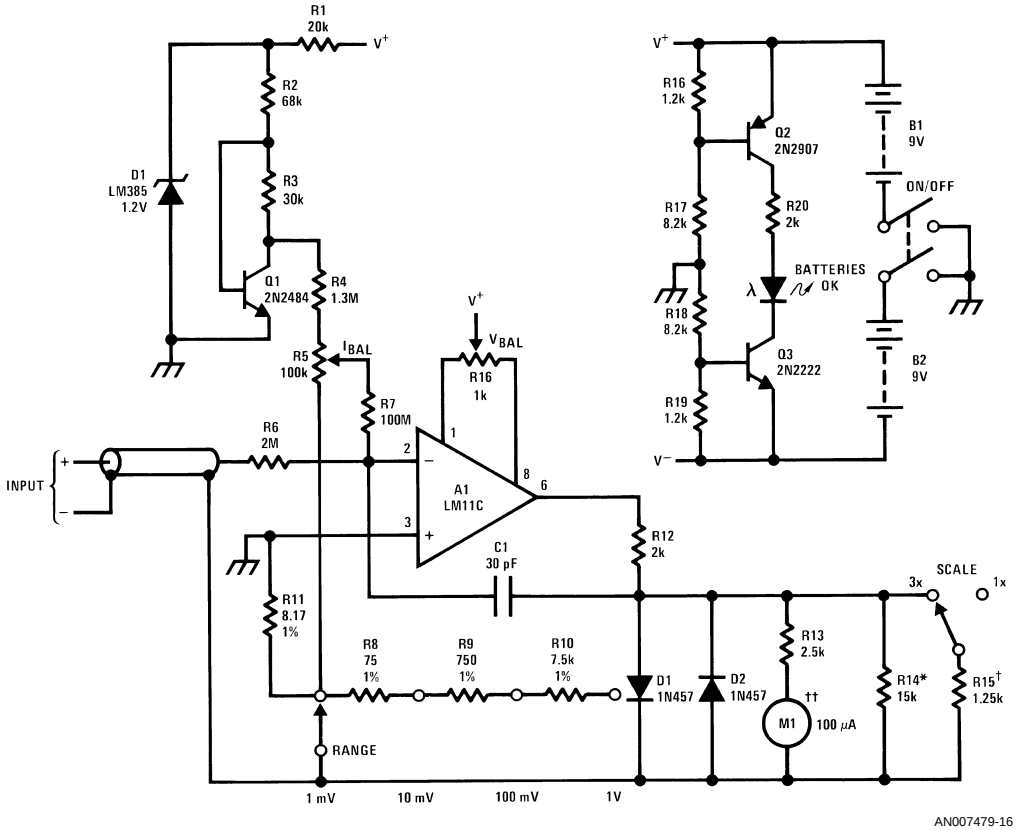

*1x scale calibrate
†3x scale calibrate
††includes reversing switch

FIGURE 14. High input impedance millivoltmeter. Input current is proportional to input voltage, about 10 pA at full-scale. Reference could be used to make direct reading linear ohmmeter.

The input resistor, R6, serves two functions. First, it protects the op amp input in the event of overload. Second, it insures that an overload will not give a false meter indication until it exceeds a couple hundred volts.

Since the reference is bootstrapped to the input, this circuit is easily converted into a linear, direct-reading ohmmeter. A resistor from the top of D1 to the input establishes the measurement current so that the voltage drop is proportional to the resistance connected across the input.

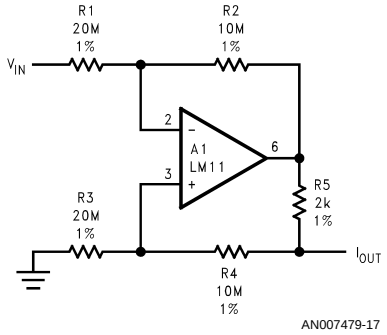

$R1 = R3;\ R2 = R4$

$$I_{OUT} = \frac{R2\ V_{IN}}{R1\ R5}$$

FIGURE 15. Output resistance of this voltage/current converter depends both on high value feedback resistors and their matching.

CURRENT SOURCES

The classical op amp circuit for voltage-to-current conversion is shown in *Figure 15*. It is presented here because the output resistance is determined by both the matching and the value of the feedback resistors. With the LM11, these resistors can be raised while preserving DC stability.

While the circuit in *Figure 15* can provide bipolar operation, better performance can be obtained with fewer problems if a unipolar output is acceptable. A complete, battery-powered current source suitable for laboratory use is given in *Figure 16* to illustrate this approach. The op amp regulates the voltage across the range resistors at a level determined by the voltage on the arm of the calibrated potentiometer, R3. The voltage on the range resistors is established by the current through Q2 and Q3, which is delivered to the output.

The reference diode, D1, determines basic accuracy. Q1 is included to insure that the LM11 inputs are kept within the common-mode range with diminishing battery voltage. A light-emitting diode, D2, is used to indicate output saturation. However, this indication cannot be relied upon for output-current settings below about 20 nA unless the value of R6 is increased. The reason is that very low currents can be supplied to the range resistors through R6 without developing enough voltage drop to turn on the diode.

If the LED illuminates with the output open, there is sufficient battery voltage to operate the circuit. But a battery-test switch is also provided. It is connected to the base of the op amp output stage and forces the output toward V^+.

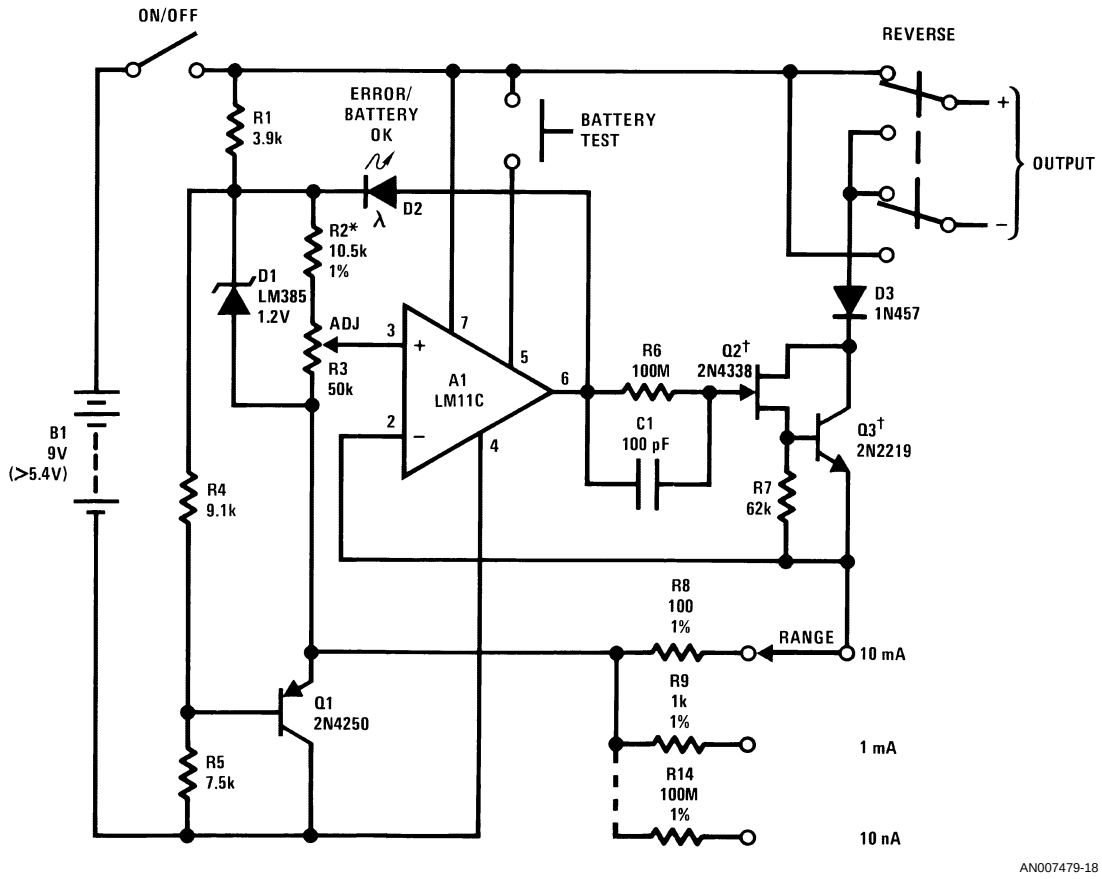

*calibrate range

†select for $I_{CBO} \leq 100$ pA

FIGURE 16. Precision current source has 10 nA to 10 mA ranges with output compliance of 30V to −5V. Output current is fully adjustable on each range with a calibrated, ten-turn potentiometer. Error light indicates output saturation.

Bias current compensation is not used because low-range accuracy is limited by the leakage currents of Q2 and Q3. As it is, these parts must be selected for low leakage. This should not be difficult because the leakage specified is determined by test equipment rather than device characteristics. It should be noted in making substitutions that Q2 was selected for low pinch-off voltage and that Q3 may have to dissipate 300 mW on the high-current range. Heating Q3 on the high range could increase leakage to where the circuit will not function for a while when switched to the low range.

LOGARITHMIC CONVERTER

A logarithmic amplifier that can operate over an eight-decade range is shown in *Figure 17*. Naturally, bias current compensation must be used to pick up the low end of this range. Leakage of the logging transistors is not a problem as long as Q1A is operated at zero collector-base voltage. In the worst case, this may require balancing the offset voltage of A1. Non-standard frequency compensation is used on A1 to obtain fairly uniform response time, at least at the high end of the range. The low end might be improved by optimizing C1. Otherwise, the circuit is standard.

LIGHT METER

This logging circuit is adapted to a battery-powered light meter in *Figure 18*. An LM10, combined op amp and reference, is used for the second amplifier and to provide the regulated voltage for offsetting the logging circuit and powering the bias current compensation. Since a meter is the output indicator, there is no need to optimize frequency compensation. Low-cost single transistors are used for logging since the temperature range is limited. The meter is protected from overloads by clamp diodes D2 and D3.

Silicon photodiodes are more sensitive to infrared than visible light, so an appropriate filter must be used for photography. Alternately, gallium-arsenide-phosphide diodes with suppressed IR response are becoming available.

DIFFERENTIAL AMPLIFIERS

Many instrumentation applications require the measurements of low-level signals in the presence of considerable ground noise. This can be accomplished with a differential amplifier because it responds to the voltage between the inputs and rejects signals between the inputs and ground.

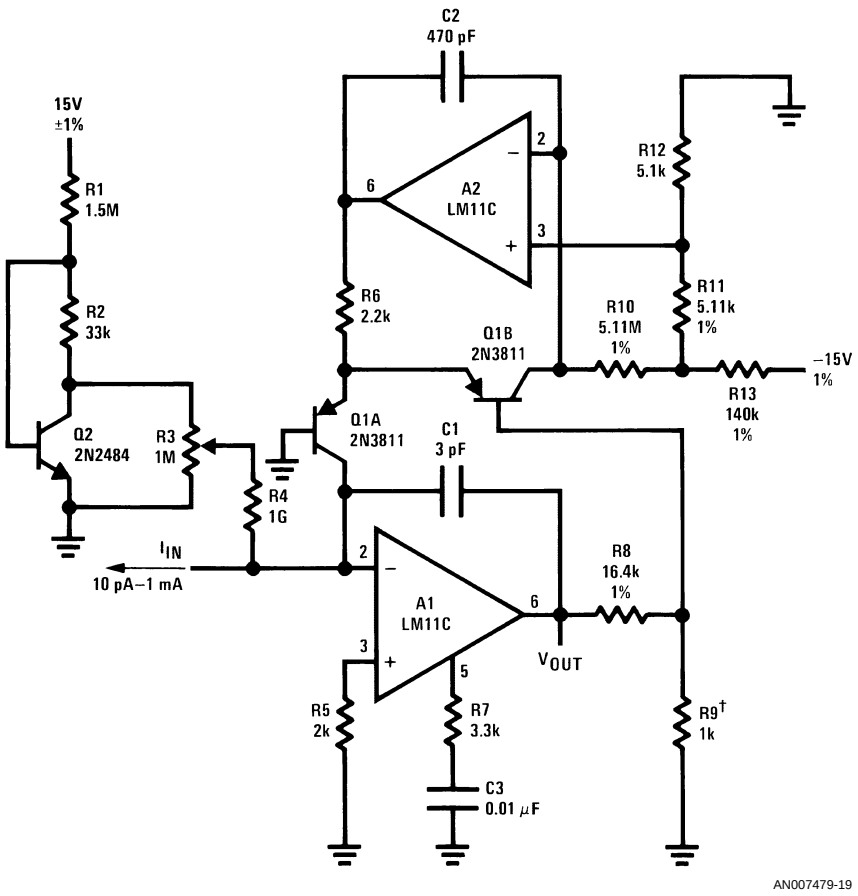

330 ppm/°C. Type Q209 available from Tel Labs, Manchester, N.H.Inc.,
a. set R11 for V_{OUT} = 0 at I_{IN} = 100 nA
b. set R8 for V_{OUT} = 3V at I_{IN} = 100 µA
c. set R3 for V_{OUT} = 4V at I_{IN} = 10 pA

FIGURE 17. Unusual frequency compensation gives this logarithmic converter a 100 µs time constant from 1 mA down to 100 nA, increasing from 200 µs to 200 ms from 10 nA to 10 pA. Optional bias current compensation can give 10 pA resolution from −55°C to +100°C. Scale factor is 1V/decade and temperature compensated.

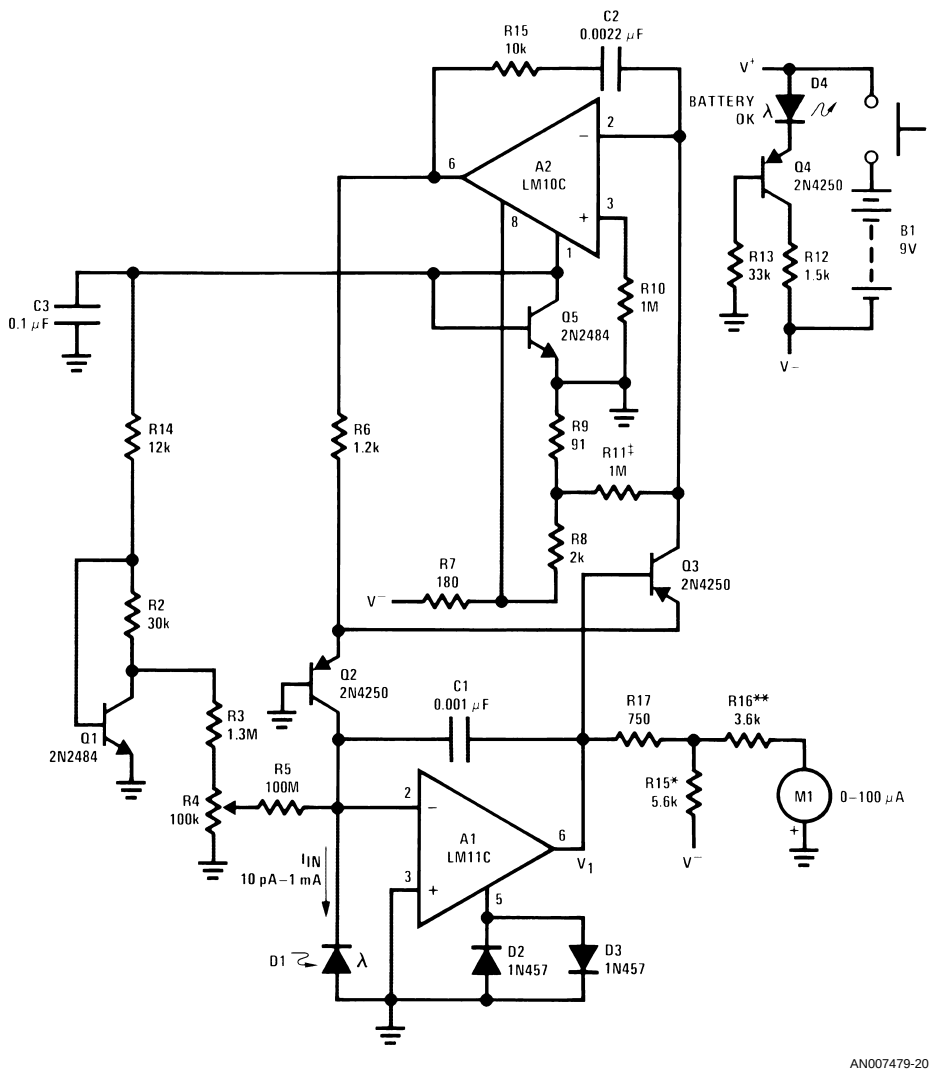

V1 = 0@I_{IN} = 100 nA
†V1 = −0.24V@I_{IN} = 10 pA
*M1 = 0@I_{IN} = 10 pA
**M1 = f.s.@I_{IN} = 1 mA

FIGURE 18. Light meter has eight-decade range. Bias current compensation can give input current resolution of better than ±2 pA over 15˚C to 55˚C.

Figure 19 shows the classic op amp differential amplifier connection. It is not widely used because the input resistance is much lower than alternate methods. But when the input common-mode voltage exceeds the supply voltage for the op amp, this cannot be avoided. At least with the LM11, large feedback resistors can be used to reduce loading without affecting DC accuracy. The impedances looking into the two inputs are not always the same. The values given equalize them for common-mode signals because they are usually larger. With single-ended inputs, the input resistance on the inverting input is R1, while that on the non-inverting input is the sum of R2, R4 and R5.

Provision is made to trim the circuit for maximum DC and AC common-mode rejection. This is advised because well matched high-value resistors are hard to come by and because unbalanced stray capacitances can cause severe deterioration of AC rejection with such large values. Particu-

lar attention should be paid to resistor tracking over temperature as this is more of a problem with high-value resistors. If higher gain or gain trim is required, R6 and R7 can be added.

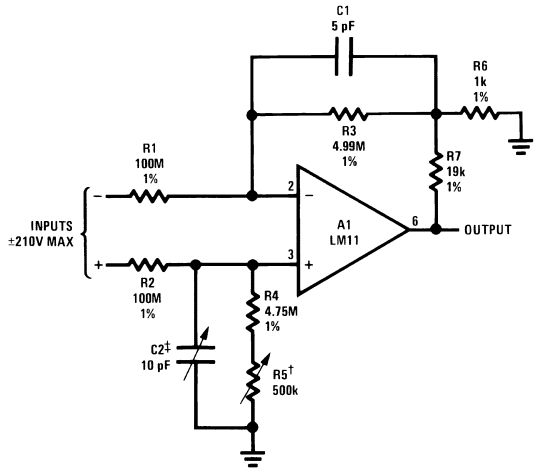

V_S = ±15V

$$V_{CM(MAX)} = \frac{R1}{R3} V_{OUT(MAX)}$$

$$A_V = \frac{R3}{R1}\left(\frac{R6 + R7}{R6}\right)$$

†trim for DC CMRR
trim for AC CMRR

FIGURE 19. This differential amplifier handles high input voltages. Resistor mismatches and stray capacitances should be balanced out for best common-mode rejection.

The simplest connection for making a high-input-impedance differential amplifier using op amps is shown in *Figure 20*. Its main disadvantage is that the common-mode signal on the inverting input is delayed by the response of A1 before being delivered to A2 for cancellation. A selected capacitor across R1 will compensate for this, but AC common-mode rejection will deteriorate as the characteristics of A1 vary with temperature.

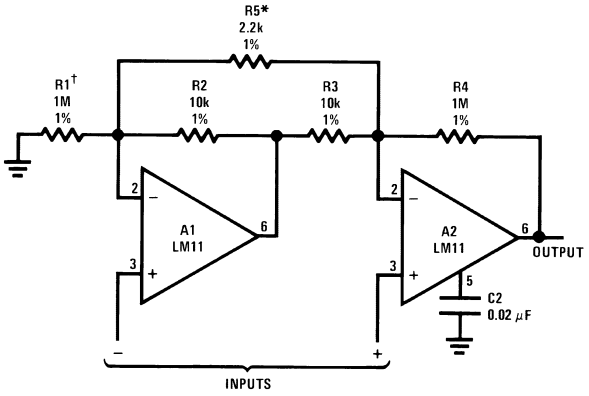

$f_2 \simeq 10$ Hz
*gain set
†trim for DC CMRR

FIGURE 20. Two-op-amp instrumentation amplifier has poor AC common-mode rejection. This can be improved at the expense of differential bandwidth with C2.

When slowly varying differential signals are of interest, the response of A2 can be rolled off with C2 to reduce the sensitivity of the circuit to high frequency common-mode signals. If single-resistor gain setting is desired, R5 can be added. Otherwise, it is unnecessary.

A full-blown differential amplifier with extremely high input impedance is shown in *Figure 21*. Gain is fixed at 1000, but it can be varied with R10. Differential offset balancing is provided on both input amplifiers by R18.

The AC common-mode rejection is dependent on how well the frequency characteristics of A2 and A3 match. This is a far better situation than encountered with the previous circuit. When AC rejection must be optimized, amplifier differences as well as the effects of unbalanced stray capacitances can be compensated for with a capacitor across R13 or R14, depending on which side is slower. Alternately, C1 can be added to control the differential bandwidth and make AC common-mode rejection less dependent on amplifier matching. The value shown gives approximately 100 Hz differential bandwidth, although it will vary with gain setting.

A separate amplifier is used to drive the shields of the input cables. This reduces cable leakage currents and spurious signals generated from cable flexing. It may also be required to neutralize cable capacitance. Even short cables can attenuate low-frequency signals with high enough source resistance. Another balance potentiometer, R8, is included so that resistor mismatches in the drive to the bootstrapping amplifier can be neutralized. Adding the bootstrapping amplifier also provides a connection point, as shown, for bias-current compensation if the ultimate in performance is required.

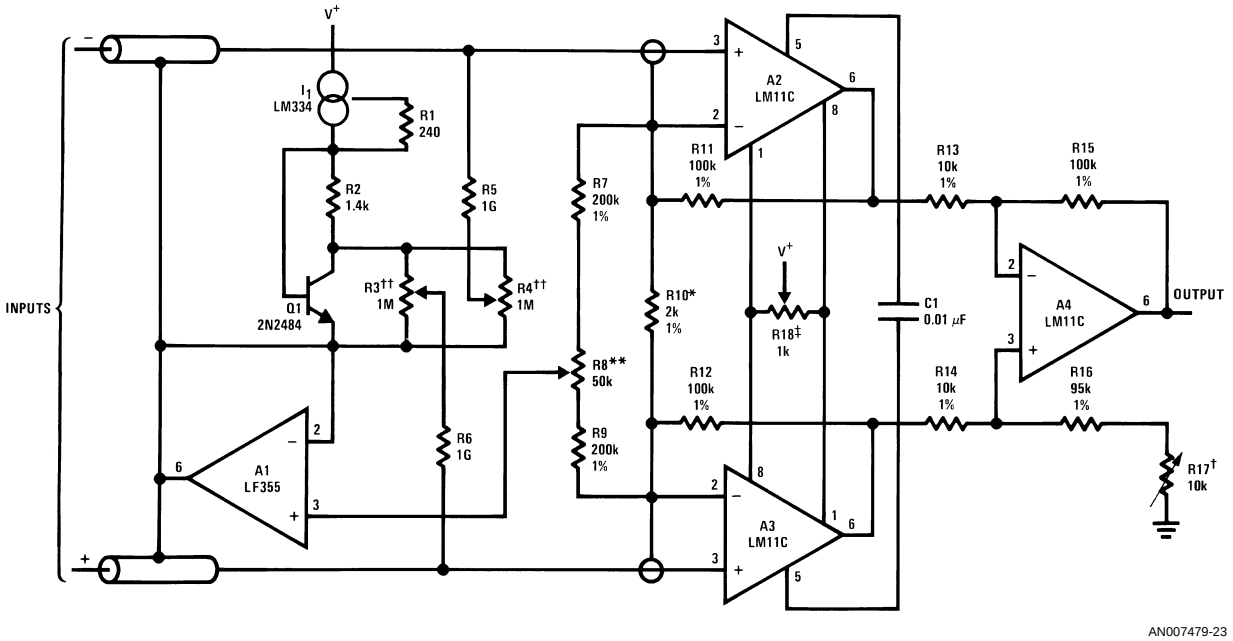

††current zero
voltage balance
*gain
†DC CMRR
**AC CMRR

FIGURE 21. High gain differential instrumentation amplifier includes input guarding, cable bootstrapping and bias current compensation. Differential bandwidth is reduced by C1 which also makes common-mode rejection less dependent on matching of input amplifiers.

As can be seen in *Figure 23*, connecting the input amplifiers as followers simplifies the circuit considerably. But single resistor gain control is no longer available and maximum bandwidth is less with all the gain developed by A3. Resistor matching is more critical for a given common-mode rejection, but AC matching of the input amplifier is less a problem. Another method of trimming AC common-mode rejection is shown here.

INTEGRATOR RESET

When pursuing the ultimate in performance with the LM11, it becomes evident that components other than the op amp can limit performance. This can be the case when semiconductor switches are used. Their leakage easily exceeds the bias current when elevated temperatures are involved.

The integrator with electrical reset in *Figure 22* gives a solution to this problem. Two switches in series are used to shunt the integrating capacitor. In the off state, one switch, Q2, disconnects the output while the other, Q1, isolates the leakage of the first. This leakage is absorbed by R3. Only the op amp offset appears across the junctions of Q1, so its leakage is reduced by two orders of magnitude.

A junction FET could be used for Q1 but not for Q2 because there is no equivalent to the enhancement mode MOSFET. The gate of a JFET must be reverse biased to turn it off and leakage on its output cannot be avoided.

MOS switches with gate-protection diodes are preferred in production situations as they are less sensitive to damage from static charges in handling. If used, D1 and R2 should be inlcuded to remove bias from the internal protection diode when the switch is off.

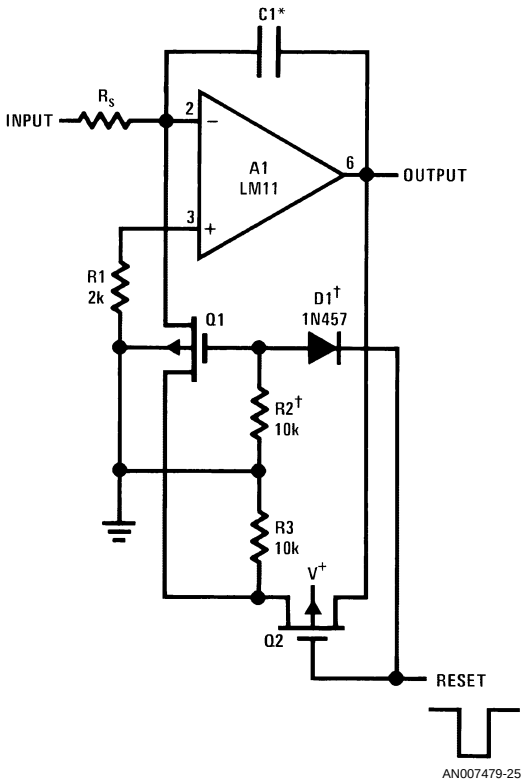

*polystyrene recommended
†required if protected gate switch is used

FIGURE 22. Reset is provided for this integrator and switch leakage is isolated from the summing junction. Greater precision can be provided if bias-current compensation is included.

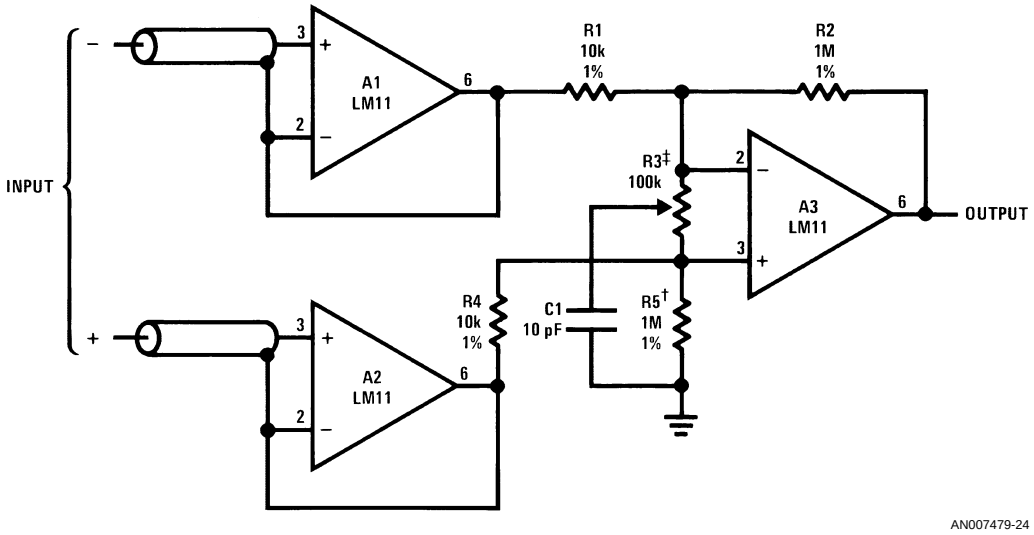

R1 = R3; R2 = R4

$$A_V = \frac{R2}{R1}$$

†trim for DC CMRR
set for AC CMRR

FIGURE 23. For moderate-gain instrumentation amplifiers, input amplifiers can be connected as followers. This simplifies circuitry, but A3 must also have low drift.

PEAK DETECTOR

The peak detector in *Figure 24* expands upon this idea. Isolation is used on both the peak-detecting diode and the reset switch. This particular circuit is designed for a long hold interval so acquisition is not quick. As might be expected from an examination of the figure, frequency compensation of an op amp peak detector is not exactly straightforward.

OVEN CONTROLLER

The LM11 is quite useful with slow servo systems because impedance levels can be raised to where reasonable capacitor values can be used to effect loop stabilization without affecting accuracy. An example of this is shown in *Figure 25*. This is a true proportional controller for a crystal oven.

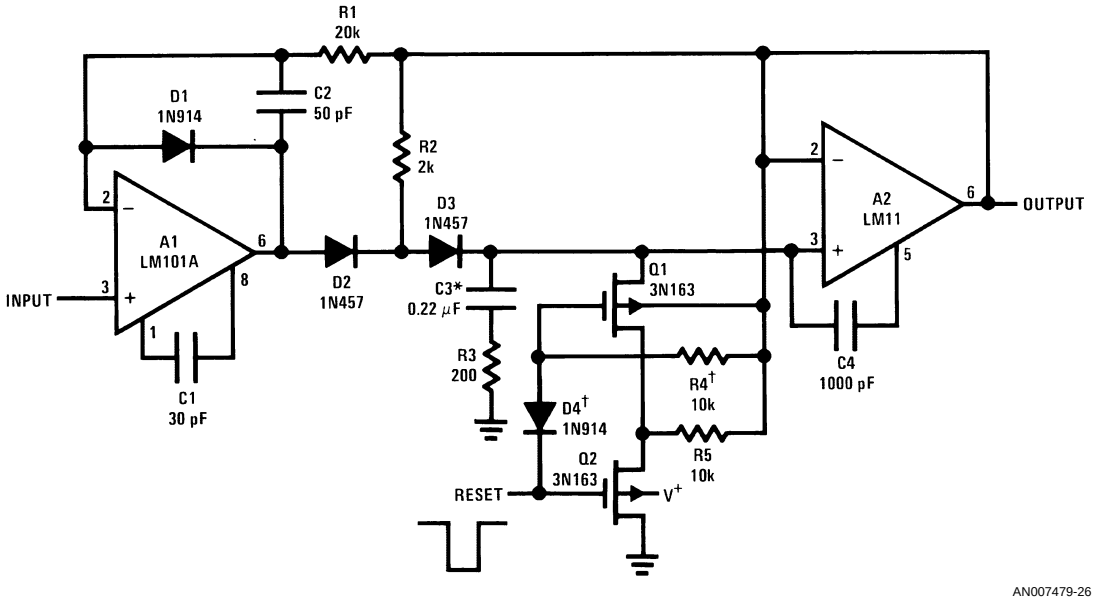

300 μs min single pulse
200 μs min repetitive pulse
300 Hz max sine wave error < 5 mV
†required if Q1 has gate protection diode
*polystyrene or Teflon

FIGURE 24. A peak detector designed for extended hold. Leakage currents of peak-detecting diodes and reset switch are absorbed before reaching storage capacitor.

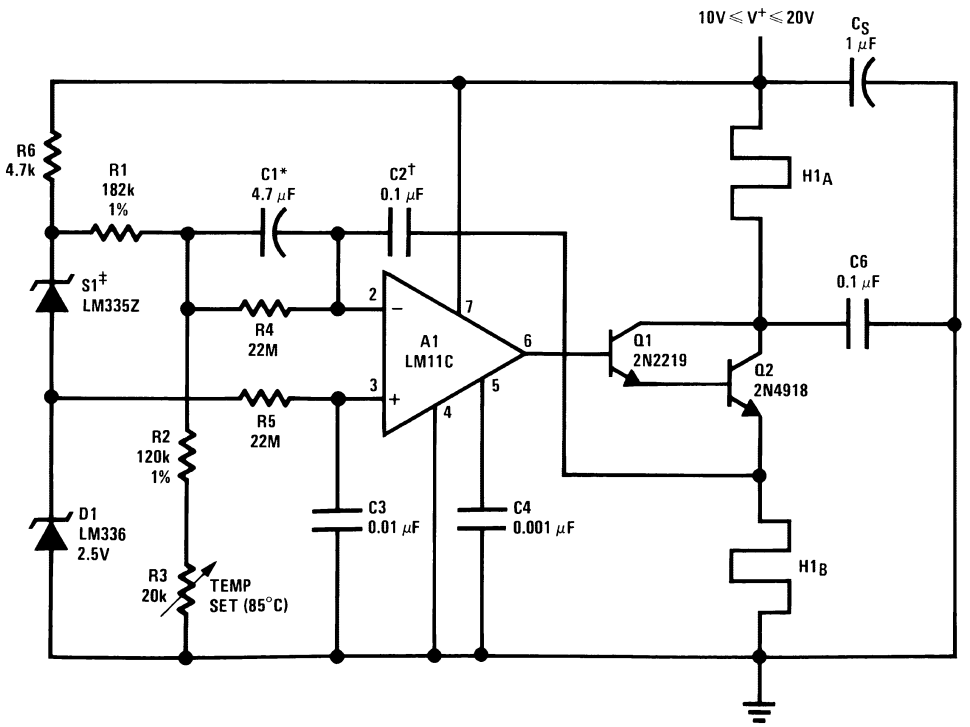

*solid tantalum
†mylar
close thermal coupling between sensor and oven shell is recommended

FIGURE 25. Proportional control crystal oven heater uses lead/lag compensation for fast settling. Time constant is changed with R4 and compensating resistor R5. If Q2 is inside oven, a regulated supply is recommended for 0.1˚C control.

Temperature sensing is done with a bridge, one leg of which is formed by an IC temperature sensor, S1, and a reference diode, D1. Frequency stabilization is done with C2 providing a lag that is finally broken out by C1. If the control transistor, Q2, is put inside the oven for maximum heating efficiency, some level of regulation is suggested for the heater supply when precise control is required. With Q2 in the oven, abrupt supply changes will alter heating, which must be compensated for by the loop. This takes time, causing a small temperature transient.

Because the input bias current of the LM11 does not increase with temperature, it can be installed inside the oven for best performance. In fact, when an oven is available in a piece of equipment, it would be a good idea to put all critical LM11s inside the oven if the temperature is less than 100˚C.

AC AMPLIFIER

Figure 26 shows an op amp used as an AC amplifier. It is unusual in that DC bootstrapping is used to obtain high input resistance without requiring high-value resistors. In theory, this increases the ouput offset because the op amp offset voltage is multiplied by the resistance boost.

But when conventional resistor values are used, it is practical to include R5 to eliminate bias-current error. This gives less output offset than if a single, large resistor were used. C1 is included to reduce noise.

STANDARD CELL BUFFER

The accuracy and lifetime of a standard cell deteriorate with loading. Further, with even a moderate load transient, recovery is measured in minutes, hours or even days. The circuit in *Figure 27* not only buffers the standard cell but also disconnects it in the event of malfunction.

The fault threshold is determined by the gate turn-on voltage of Q1. As the voltage on the gate approaches the threshold either because of low battery voltage or excessive output loading, the MOS switch will begin to turn off. At the turn off threshold, the output voltage can rise because of amplifier bias current flowing through the increasing switch resistance. Therefore, a LED indicator is included that extinguishes as the fault condition is approached. The MOS threshold should be higher than the buffer output so disconnect and error indicator operates before the output saturates.

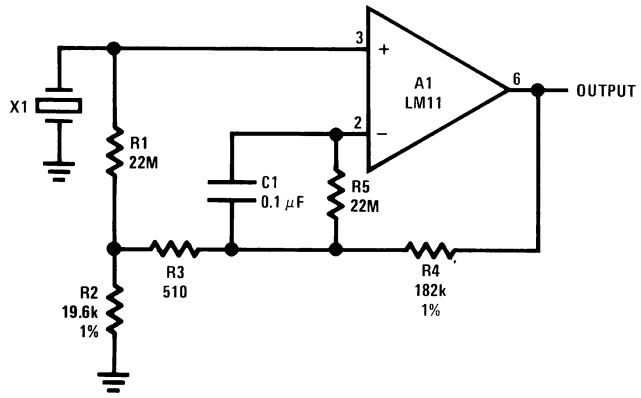

$$R_{IN} = R1\left(1 + \frac{R2}{R3}\right); \quad A_V = \frac{R2 + R3 + R4}{R2 + R3}$$

FIGURE 26. A high input impedance AC amplifier for a piezoelectric transducer. Input resistance of 880 MΩ and gain of 10 is obtained.

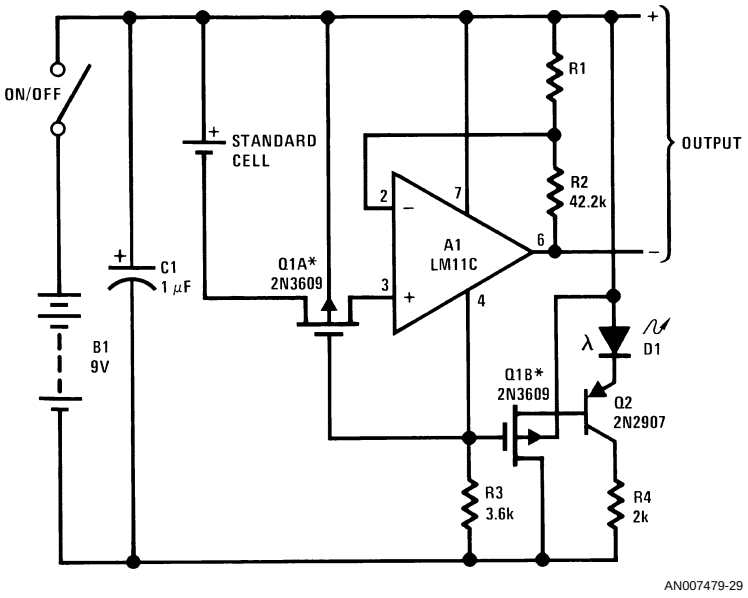

*cannot have gate protection diode; $V_{TH} > V_{OUT}$

FIGURE 27. Battery powered buffer amplifier for standard cell has negligible loading and disconnects cell for low supply voltage or overload on output. Indicator diode extinguishes as disconnect circuitry is activated.

CONCLUSIONS

Although the LM11 does not provide the ultimate in performance in either offset voltage or bias current for nominal room temperature applications, the combination offered is truly noteworthy. With significant temperature excursions, the results presented here are much more impressive. With full-temperature-range operation, this device does represent the state of the art when high-impedance circuitry is involved.

Combining this new amplifier with fast op amps to obtain the best features of both is also interesting, particularly since the composite works well in both the inverting and non-inverting modes. However, making high-impedance circuits fast is no simple task. If higher temperatures are not involved, using the LM11 to reduce the offset voltage of a FET op amp without significantly increasing bias current may be all that is required.

An assortment of measurement and computational circuits making use of the unique capabilities of this IC were presented. These circuits have been checked out and the results should be of some value to those working with high impedances. These applications are by no means all-inclusive, but they do show that an amplifier with low input current can be used in a wide variety of circuits.

Although emphasis was on high-performance circuits requiring adjustments, the LM11 will see widest usage in less demanding applications where its low initial offset voltage and bias current can eliminate adjustments.

ACKNOWLEDGEMENT

The authors would like to thank Dick Wong for his assistance in building and checking out the applications described here.

***See Addendum that follows this Application Note.**

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT AND GENERAL COUNSEL OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
Americas
Tel: 1-800-272-9959
Fax: 1-800-737-7018
Email: support@nsc.com
www.national.com

National Semiconductor Europe
Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5639-7560
Fax: 81-3-5639-7507

Applications of the LM392 Comparator Op Amp IC

National Semiconductor
Application Note 286
September 1981

The LM339 quad comparator and the LM324 op amp are among the most widely used linear ICs today. The combination of low cost, single or dual supply operation and ease of use has contributed to the wide range of applications for these devices.

The LM392, a dual which contains a 324-type op amp and a 339-type comparator, is also available. This device shares all the operating features and economy of 339 and 324 types with the flexibility of both device types in a single 8-pin mini-DIP. This allows applications that are not readily implemented with other devices but retain simplicity and low cost. *Figure 2* provides an example.

Sample-Hold Circuit

The circuit of *Figure 2* is an unusual implementation of the sample-hold function. Although its input-to-output relationship is similar to standard configurations, its operating principle is different. Key advantages include simplicity, no hold step, essentially zero gain error and operation from a single 5V supply. In this circuit the sample-hold command pulse (Trace A, *Figure 1*) is applied to Q3, which turns on, causing current source transistor Q4's collector (Trace B, *Figure 1*) to go to ground potential. Amplifier A1 follows Q4's collector voltage and provides the circuit's output (Trace C, *Figure 1*). When the sample-hold command pulse falls, Q4's collector drives a constant current into the 0.01 µF capacitor. When the capacitor ramp voltage equals the circuit's input voltage, comparator C1 switches, causing Q2 to turn off the current source. At this point the collector voltage of Q4 sits at the circuit's input voltage. Q1 insures that the comparator will not self trigger if the input voltage increases during a "hold" interval. When a DC biased sine wave is applied to the circuit (Trace D, *Figure 1*) the sampled output (Trace E, *Figure 1*) is available at the circuit's output. The ramping action of the Q4 current source during the "sample" states is just visible in the output.

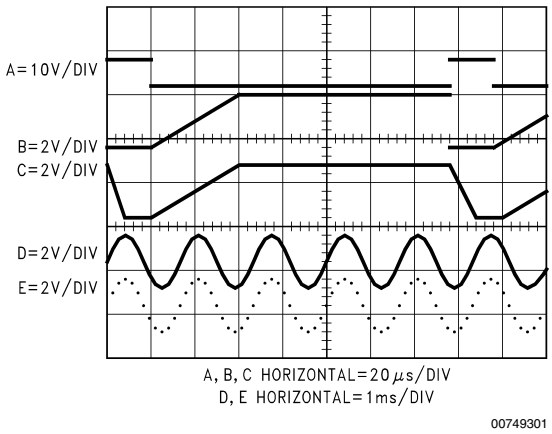

FIGURE 1.

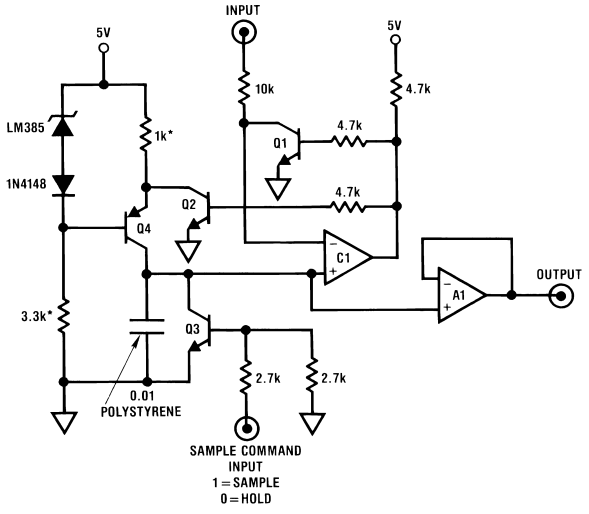

Q1, Q2, Q3 = 2N2369

Q4 = 2N2907

C1, A1 = LM392 amplifier-comparator dual

*1% metal film resistor

FIGURE 2.

"Fed-Forward" Low-Pass Filter

In *Figure 3* the LM392 implements a useful solution to a common filtering problem. This single supply circuit allows a signal to be rapidly acquired to final value but provides a long filtering constant. This characteristic is useful in multiplexed data acquisition systems and has been employed in electronic infant scales where fast, stable readings of infant weight are desired despite motion on the scale platform. When an input step (Trace A, *Figure 4*) is applied, C1's negative input will immediately rise to a voltage determined by the 1k pot-10 kΩ divider. C1's "+" input is biased through the 100 kΩ-0.01 µF time constant and phase lags the input. Under these conditions C1's output will go low, turning on Q1. This causes the capacitor (Waveform B, *Figure 1*) to charge rapidly towards the input value. When the voltage across the capacitor equals the voltage at C1's positive input, C1's output will go high, turning off Q1. Now, the capacitor can only charge through the 100k value and the time constant will be long. Waveform B clearly shows this. The point at which the filter switches from short to long time constant is adjustable with the 1 kΩ potentiometer. Normally, this is adjusted so that switching occurs at 90%–98% of final value, but the photo was taken at a 70% trip point so circuit operation is easily discernible. A1 provides a buffered output. When the input returns to zero the 1N933 diode, a low forward drop type, provides rapid discharge for the capacitor.

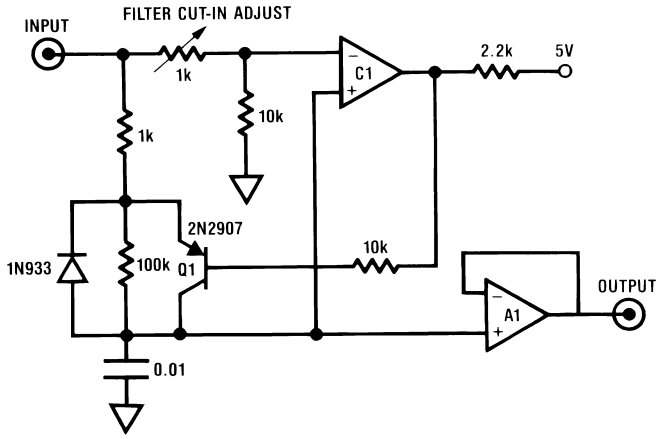

A1, C1 = LM392 amplifier-comparator dual

FIGURE 3.

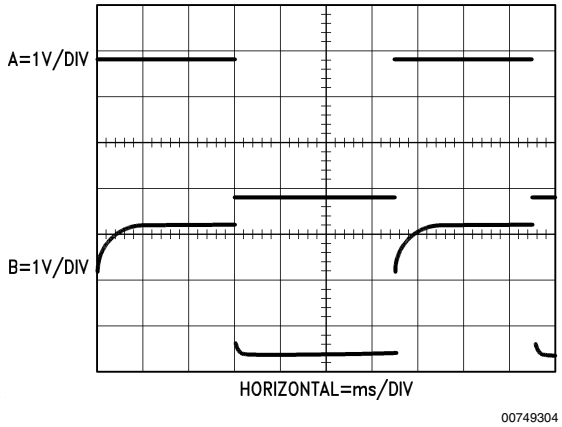

FIGURE 4.

Variable Ratio Digital Divider

In *Figure 5* the circuit allows a digital pulse input to be divided by any number from 1 to 100 with control provided by a single knob. This function is ideal for bench type work where the rapid set-up and flexibility of the division ratio is highly desirable. When the circuit input is low, Q1 and Q3 are off and Q2 is on. This causes the 100 pF capacitor to accumulate a quantity of charge (Q) equal to

$Q = CV$

where C = 100 pF

and V = the LM385 potential (1.2V) minus the $V_{CE(SAT)}$ of Q2.

When the input goes high (Trace A, *Figure 6*) Q2 goes off and Q1 turns on Q3. This causes Q3 to displace the 100 pF capacitor's charge into A1's summing junction. A1's output responds (Waveform B, *Figure 6*) by jumping to the required

Variable Ratio Digital Divider

(Continued)

value to maintain the summing junction at 0V. This sequence is repeated for every input pulse. During this time A1's output will form the staircase shape shown in Trace B as the 0.02 µFfeedback capacitor is pumped up by the charge dispensing action into A1's summing junction. When A1's output is great enough to just bias C1's "+" input below ground, C1's output (Trace C, *Figure 6*) goes low and resets A1 to 0V. Positive feedback to C1's "+" input (Trace D, *Figure 6*) comes through the 300 pF unit, insuring adequate reset time for A1. The 1 MΩ potentiometer, by setting the number of steps in the ramp required to trip C1, controls the circuit input-output division ratio. Traces E–G expand the scale to show circuit detail. When the input (Trace E) goes high, charge is deposited into A1's summing junction (Trace F) and the resultant staircase waveform (Trace G) takes a step.

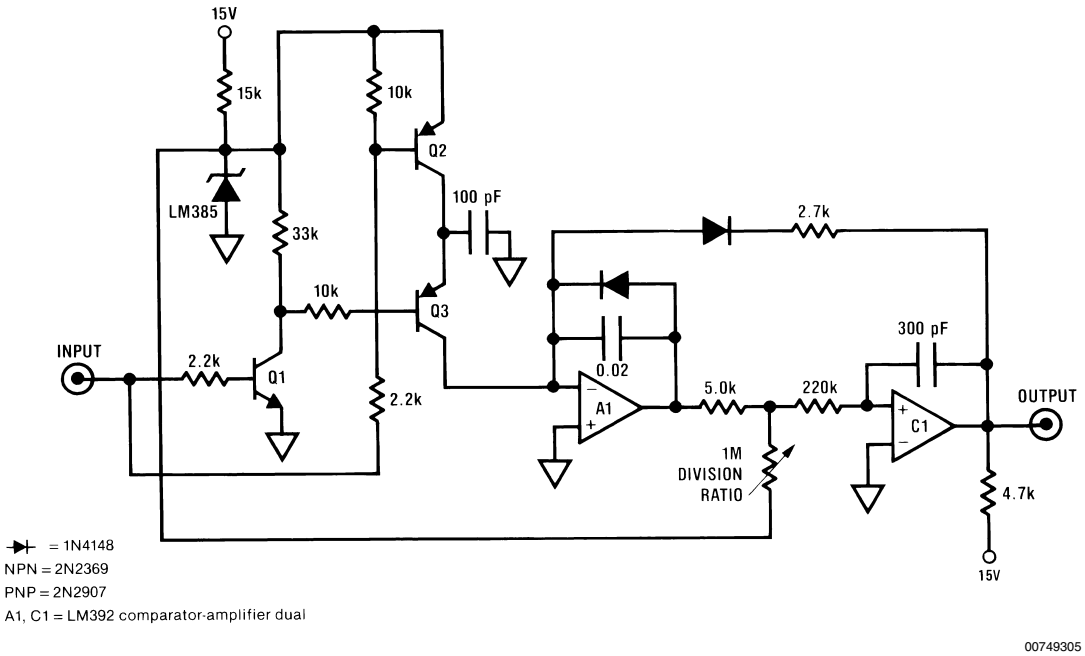

FIGURE 5.

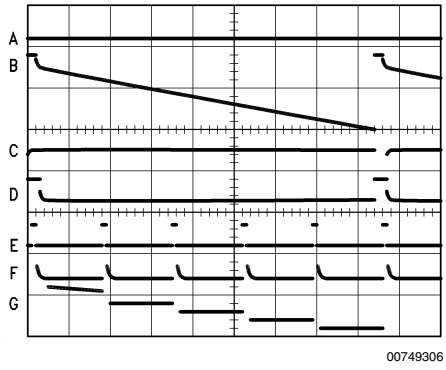

FIGURE 6.

Trace	Vertical	Horizontal
A	10V	500 µs
B	1V	500 µs
C	50V	500 µs
D	50V	500 µs
E	10V	50 µs

Trace	Vertical	Horizontal
F	10 mA	50 µs
G	0.1V	50 µs

Exponential V/F Converter for Electronic Music

Professional grade electronic music synthesizers require voltage controlled frequency generators whose output frequencies are exponentially related to the input voltages. *Figure 7* diagrams a circuit which performs this function with 0.25% exponential conformity over a range from 20 Hz to 15 kHz using a single LM392 and an LM3045 transistor array. The exponential function is generated by Q1, whose collector current will vary exponentially with its base-emitter voltage in accordance with the well known relationship between BE voltage and collector current in bipolar transistors. Normally, this transistor's operating point will vary wildly with temperature and elaborate and expensive compensation is required. Here, Q1 is part of an LM3045 transistor array. Q2 and Q3, located in the array, serve as a heater-sensor pair for A1, which servo controls the temperature of Q2. This causes the entire LM3045 array to be at constant temperature, eliminating thermal drift problems in Q1's operation. Q4 acts as a clamp, preventing servo lock-up during circuit start-up.

Exponential V/F Converter for Electronic Music (Continued)

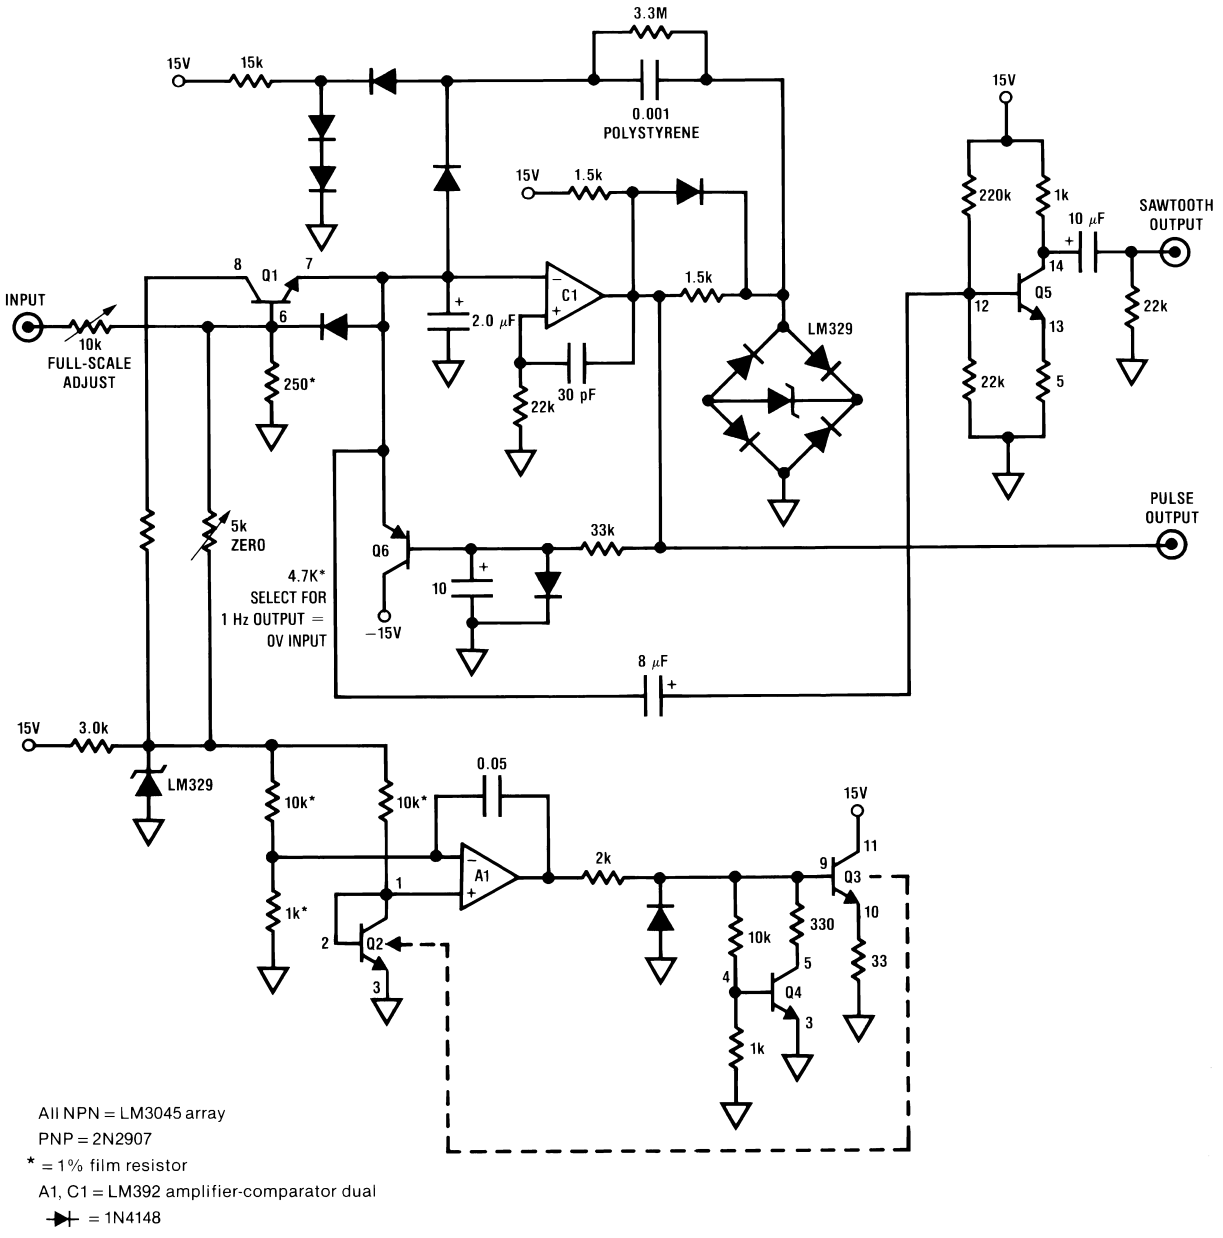

FIGURE 7.

Q1's current output is fed into the summing junction of a charge dispensing I/F converter. C1's output state is used to switch the 0.001 µF capacitor between a reference voltage and C1's "–" input. The reference voltage is furnished by the LM329 zener diode bridge. The comparator's output pulse width is unimportant as long as it permits complete charging and discharging of the capacitor. In operation, C1 drives the 30 pF-22k combination. This RC provides regenerative feedback which reinforces the direction of C1's output. When the 30 pF-22k combination decays, the positive feedback ceases. Thus, any negative going amplifier output will be followed by a positive edge after an amount of time governed by the 30 pF-22k time constant (Waveforms A and B, *Figure 8*). The actual integration capacitor in the circuit is the 2 µF electrolytic. This capacitor is never allowed to charge beyond 10 mV–15 mV because it is constantly being reset by charge dispensed from the switching of the 0.001 µF capacitor (Waveform C, *Figure 8*). Whenever the amplifier's output goes negative, the 0.001 µF capacitor dumps a quantity of charge (Waveform D) into the 2 µF capacitor, forcing it to a lower potential. The amplifier's output going negative also causes a short pulse to be transferred through the 30 pF capacitor to the "+" input. When this negative pulse decays out so that the "+" input is higher than the "–" input, the 0.001 µF capacitor is again able to receive a charge and the entire process repeats. The rate at which this sequence occurs is directly related to the current into C1's summing junction from Q1. Since this current is exponentially related to the

Exponential V/F Converter for Electronic Music (Continued)

circuit's input voltage, the overall I/F transfer function is exponentially related to the input voltage. This circuit can lock-up under several conditions. Any condition which would allow the 2 µF electrolytic to charge beyond 10 mV–20 mV (start-up, overdrive at the input, etc.) will cause the output of the amplifier to go to the negative rail and stay there. The 2N2907A transistor prevents this by pulling the "–" input towards –15V. The 10 µF-33k combination determines when the transistor will come on. When the circuit is running normally, the 2N2907 is biased off and is effectively out of the circuit. To calibrate the circuit, ground the input and adjust the zero potentiometer until oscillations just start. Next, adjust the full-scale potentiometer so that frequency output exactly doubles for each volt of input (e.g., 1V per octave for musical purposes). Repeat these adjustments until both are fixed. C1 provides a pulse output while Q5 AC amplifies the summing junction ramp for a sawtooth output.

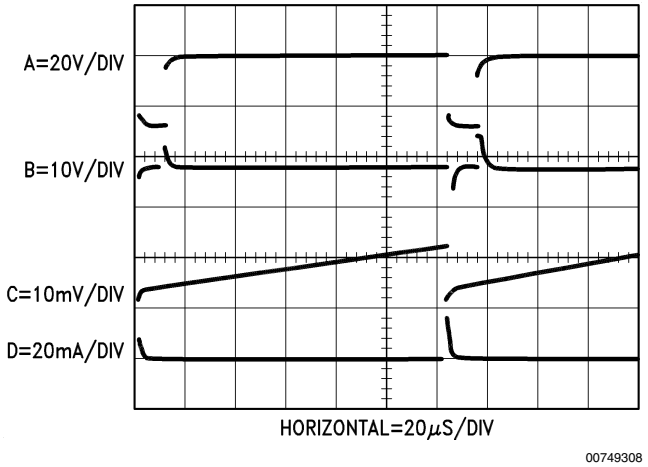

FIGURE 8.

Linearized Platinum RTD Thermometer

In *Figure 9* the LM392 is used to provide gain and linearization for a platinum RTD in a single supply thermometer circuit which measures from 0˚C to 500˚C with ±1˚C accuracy. Q1 functions as a current source which is slaved to the LM103-3.9 reference. The constant current driven platinum sensor yields a voltage drop which is proportionate to temperature. A1 amplifies this signal and provides the circuit output. Normally the slight nonlinear response of the RTD would limit accuracy to about ±3 degrees. C1 compensates for this error by generating a breakpoint change in A1's gain for sensor outputs above 250˚C. When the sensor's output indicates 250˚C, C1's "+" input exceeds the potential at the "–" input and C1's output goes high. This turns on Q2 whose collector resistor shunts A1's 6.19k feedback value, causing a gain change which compensates for the sensor's slight loss of gain from 250˚C to 500˚C. Current through the 220k resistor shifts the offset of A1 so no "hop" occurs at the circuit output when the breakpoint is activated. A precision decade box is used to calibrate this circuit. With the box inserted in place of the sensor, adjust 0˚C for 0.10V output for a value of 1000Ω. Next dial in 2846Ω (500˚C) and adjust the gain trim for an output of 2.60V. Repeat these adjustments until both zero and full-scale are fixed at these points.

Linearized Platinum RTD Thermometer (Continued)

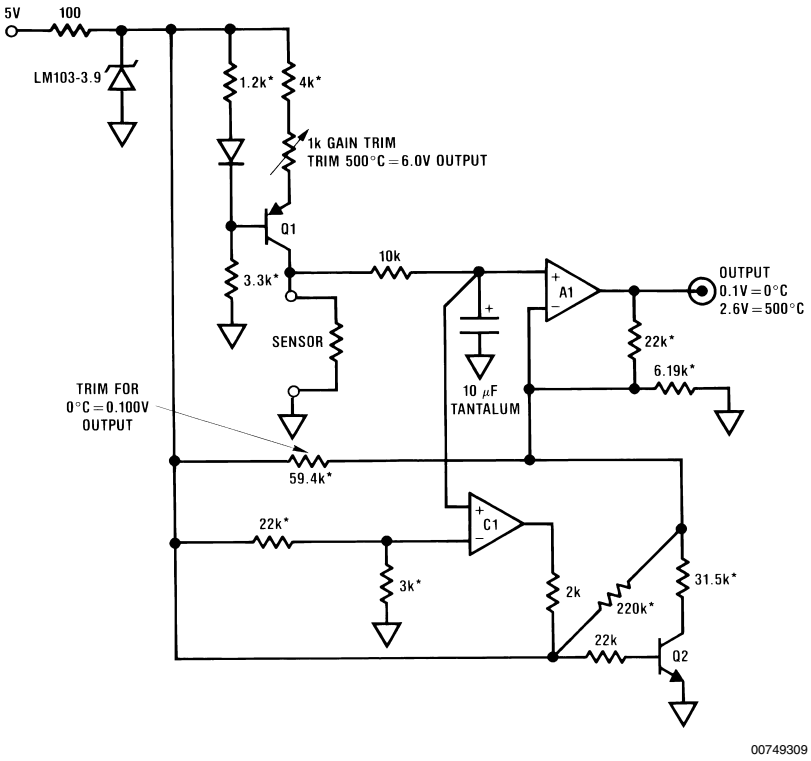

Sensor = Rosemount
118 MF-1000-A
1000Ω at 0˚C

Q1 = 2N2907

Q2 = 2N2222A

A1, C1 = LM392 amplifier-comparator dual

*metal film resistor

FIGURE 9.

Temperature Controller

Figure 10 details the LM392 in a circuit which will temperature-control an oven at 75˚C. This is ideal for most types of quartz crystals. 5V single supply operation allows the circuit to be powered directly from TTL-type rails. A1, operating at a gain of 100, determines the voltage difference between the temperature setpoint and the LM335 temperature sensor, which is located inside the oven. The temperature setpoint is established by the LM103-3.9 reference and the 1k–6.8k divider. A1's output biases C1, which functions as a pulse width modulator and biases Q1 to deliver switched-mode power to the heater. When power is applied, A1's output goes high, causing C1's output to saturate low. Q1 comes on and delivers DC to the heater. When the oven warms to the setpoint, A1's output falls and C1 begins to pulse width modulate the heater in servo control fashion. In practice the LM335 should be in good thermal contact with the heater to prevent servo oscillation.

Temperature Controller (Continued)

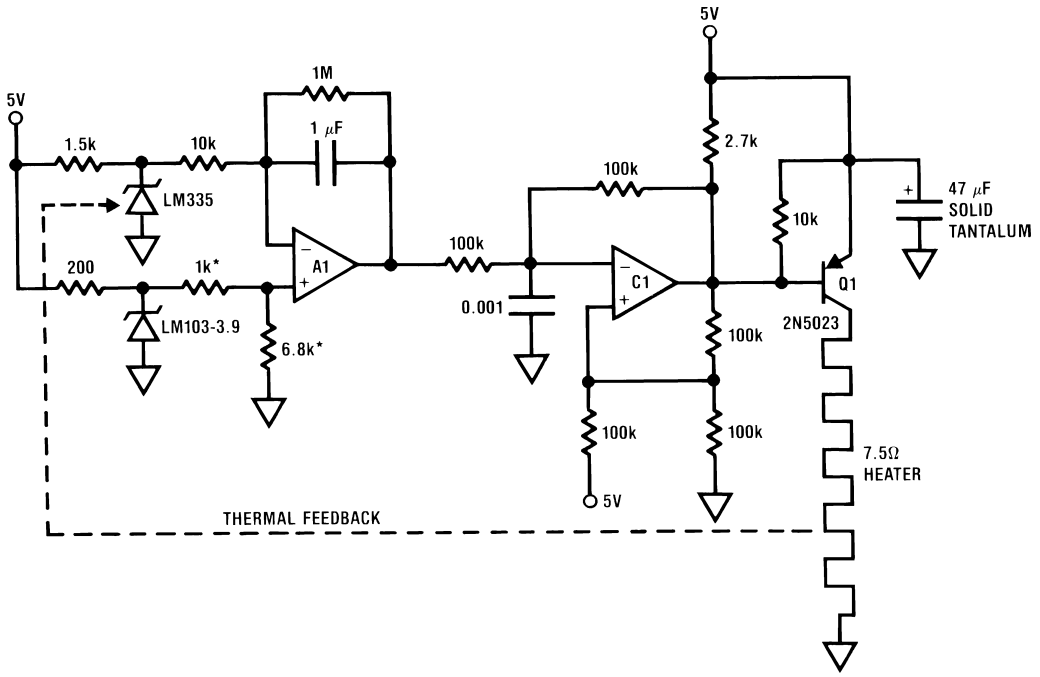

A1, C1 = LM392 amplifier-comparator dual

FIGURE 10.

References

1. *Transducer Interface Handbook,* pp. 220–223; Analog Devices, Inc.

2. "A New Ultra-Linear Voltage-to-Frequency Converter", Pease, R. A.; *1973 NEREM Record* Volume 1, page 167.

National Semiconductor Corporation
Americas
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5639-7560
Fax: 81-3-5639-7507

Simple Circuit Detects Loss of 4-20 mA Signal

National Semiconductor
Application Note 300
Robert A. Pease
May 1982

Four-to-twenty milliampere current loops are commonly used in the process control industry. They take advantage of the fact that a remote amplifier can be powered by the same 4-20 mA current that it controls as its output signal, thus using a single pair of wires for signal and power. Circuits for making 4-20 mA transmitters are found in the LM10, LM163, and LH0045 data sheets.

In general, an expensive isolation amplifier would be required to detect the case of a 4 mA signal falling out of spec (e.g., 3.7 mA) without degrading the isolation of the 4-20 mA current loop.

But this new circuit (*Figure 1*) can detect a loss or degradation of signal below 4 mA, with simplicity and low cost. The LM10 contains a stable reference at pins 1 and 8, 200 mV positive referred to pin 4. As long as the loop current is larger than 4 mA, the I x R drop across the 47.6Ω resistor, R4, is sufficient to pull the LM10's amplifier input (pin 2) below pin 3 and keep its output (pin 6) turned *OFF*.

The 4-20 mA current will flow through the LED in the optoisolator and provide a *LOW* output at pin 5 of the optoisolator.

When the current loop falls below 3.7 mA, the LM10's input at pin 2 will rise and cause the pin 6 output to fall and steal all the current away from the LED in the optoisolator. Pin 5 of the 4N28 will rise to signify a *fault* condition. This *fault* flag will fly for any loop current between 3.7 mA and 0.0 mA (and also in case of reversal or open-circuit). R1 is used to trim the threshold point to the desired value. CR2 is added in series with the LED to make sure it will turn OFF when the LM10's output goes *LOW*. (While the LM10 is guaranteed to saturate to 1.2V, the forward drop of the LED in the 4N28 may be as low as 1.0V, so a diode is added in series with the LED, to insure that it can be shut off.) Note that most operational amplifiers will not respond in a reasonable way if the output pin (6) is connected to the positive supply pin (7), but the LM10 was specifically designed and is specified to perform accurately in this "shunt" mode. (Refer to AN-211 application note, TP-14 technical paper, and the LM10 data sheet.)

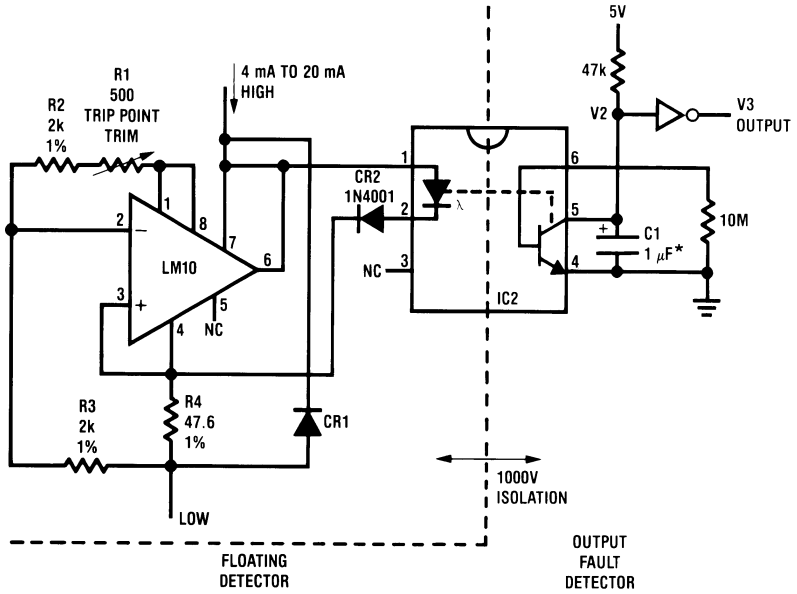

CR1 = 1N4001, optional, in case of signal reversal
LM10 = NSC LM10CLN or LM10CLH amp/reference
IC2 = 4N28 or similar, optoisolator
V2 = Normally low; high signifies fault (I < 3.7 mA)
V3 = Normally high; low signifies fault (I < 3.7 mA) (buffered output)
*C1 = 1 μF optional, to avoid false output when large AC current is superimposed on 4.0 mA. Disconnect this capacitor when using with circuit of *Figure 2*.
▷o = 1/6 MM74C04 or similar, CMOS inverter

00564001

FIGURE 1. Current Loop Fault Detector

While you could manually adjust R1 while observing the status of V3 output, this would be a coarse and awkward trim procedure. *Figure 2* shows an improved test circuit which servos the current through the detector circuit, forcing it to be

at the threshold value. Then that current can be monitored continuously, and the circuit can be trimmed easily. If the current through R107 starts out too small, the output of the 4N28 will be *HIGH* too much of the time, and the op amp output will integrate upwards until the current is at the actual threshold of the detector. The integrator's output will stop at the value where the duty cycle of the 4N28 output is exactly 50%. This occurs when the current through R107 is straddling the threshold value.

The positive feedback via R108 assures that the loop oscillates at approximately 50 cycles per second, with a small, well-controlled sawtooth wave at its output. This mode of operation was chosen to insure that the loop does not oscillate at some high, uncontrolled frequency, as it would be difficult in that case to be sure the duty cycle was exactly 50%. This test circuit is advantageous, because you can measure the trip point directly.

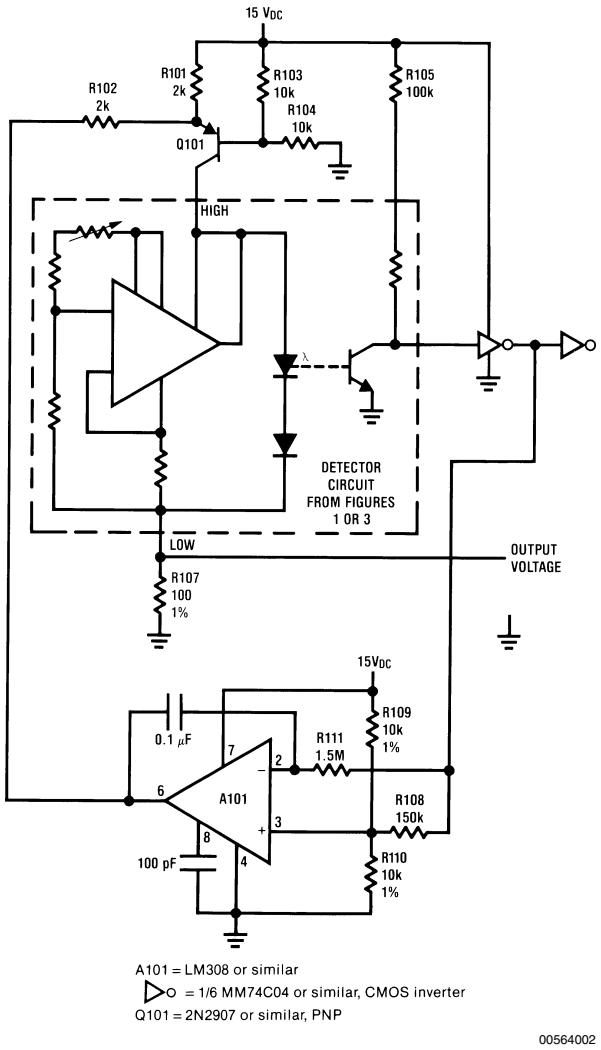

FIGURE 2. Test Circuit for Threshold Detector

The test circuit of *Figure 2* is necessary for trimming the detector in *Figure 3*. This circuit does not have a trim pot, and thus avoids the problem of someone mis-adjusting the circuit after it is once trimmed correctly. It also avoids the compromises between good but expensive trim pots and cheap but unreliable, drifty trim pots. By opening one or more of the links, L1–L4, according to the following procedure, it is easy to trim the threshold level to be within 1% of 3.70 mA (or as desired).

- Observe the DC current through R107 in *Figure 2*
- If $I_{THRESHOLD}$ is larger than 3.950 mA, open link L1;
  —if not, don't
- If $I_{THRESHOLD}$ is larger than 3.830 mA, open link L2;
  —if not, don't
- If $I_{THRESHOLD}$ is larger than 3.760 mA, open link L3;
  —if not, don't
- Then, if $I_{THRESHOLD}$ is larger than 3.720 mA, open link L4;
  —if not, don't

This procedure provides a circuit trimmed to much better than 1% of 3.70 mA, without using any trim pots. Of course, this circuit can be used to detect drop-out of regulation of

other floating signals, while maintaining high isolation from ground, good accuracy, low power dissipation (2 mA x 2.5V typical) and low cost.

Other standard values of current loop are 1 mA–5 mA and 10 mA–50 mA. The version shown in *Figure 4* uses higher resistance values to trip at 0.85 mA. The circuit in *Figure 5* has an additional transistor, to accommodate currents as large as 50 mA without damage or loss of accuracy, and provide an 8.5 mA threshold.

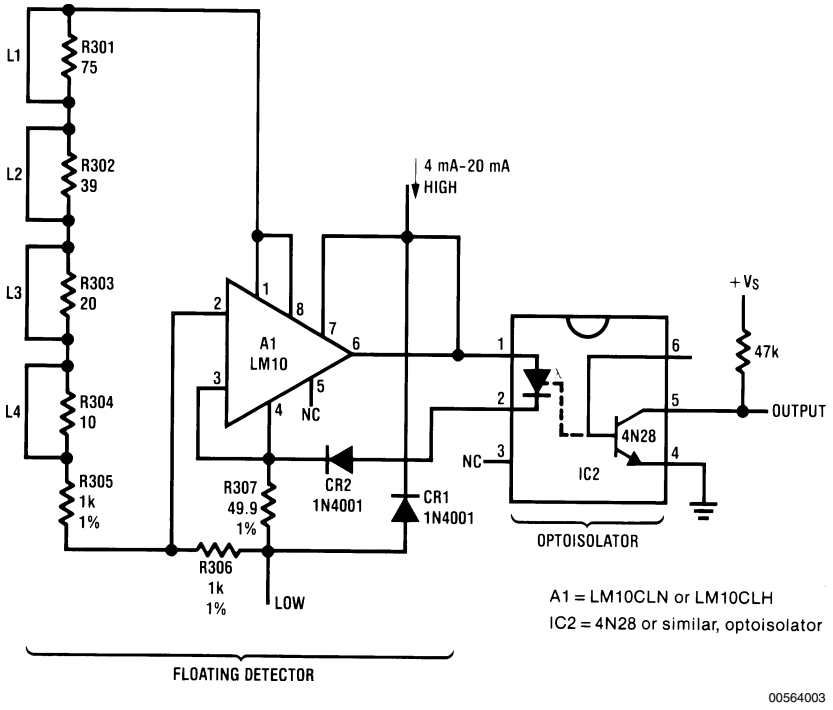

FIGURE 3. Fault Detector with Low-Cost Trim Scheme (To be trimmed in the circuit of *Figure 2*)

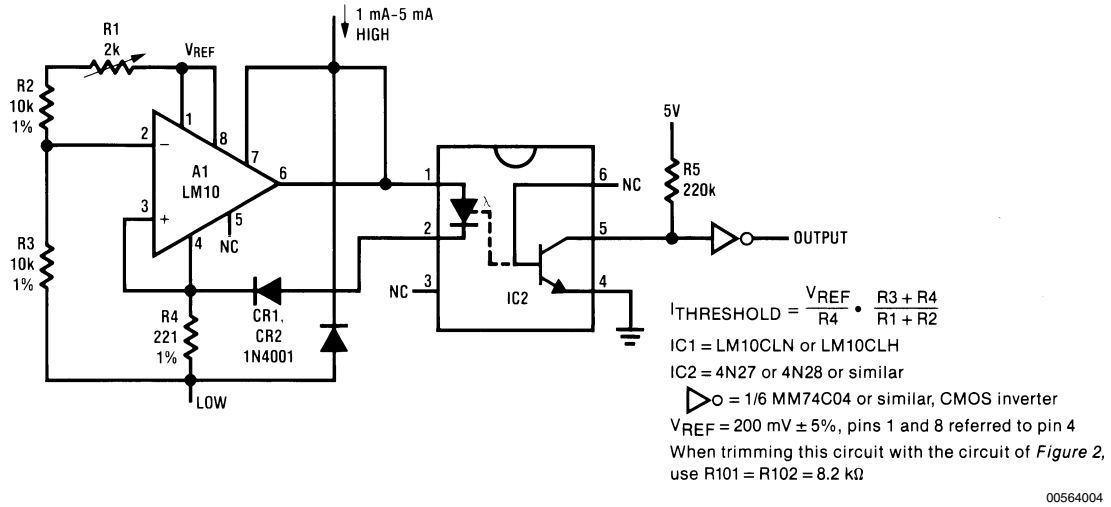

FIGURE 4. Current Loop Fault Detector ($I_{THRESHOLD}$ = 0.85 mA for 1 mA–5 mA Current Loops)

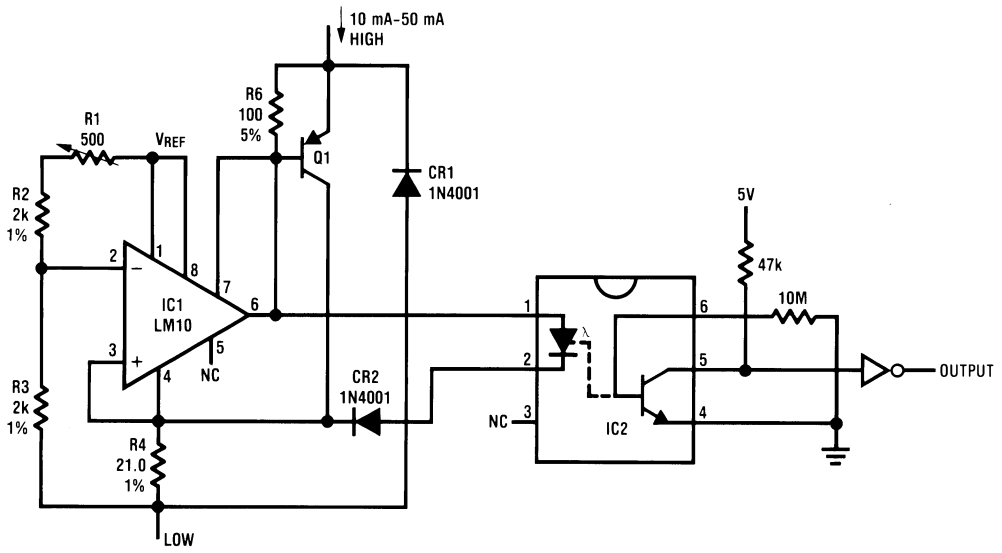

$$-I_{THRESHOLD} = \frac{V_{REF}}{R4} \bullet \frac{R3 + R4}{R1 + R2}$$

IC1 = LM10CLH or LM10CLN op amp and reference

V_{REF} = 200 mV ± 5%, pins 1 and 8 referred to pin 4

IC2 = 4N28 or similar, optoisolator

▷o = 1/6 MM74C04 or similar, CMOS inverter

Q1 = 2N2904 or 2N2907, any silicon PNP

When trimming this circuit with the circuit of *Figure 2*, use R101 = R102 = 820Ω

00564005

FIGURE 5. Current Loop Fault Detector ($I_{THRESHOLD}$ = 8.5 mA, for 10 mA–50 mA Current Loops)

National Semiconductor Corporation
Americas
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5639-7560
Fax: 81-3-5639-7507

LM34/LM35 Precision Monolithic Temperature Sensors

National Semiconductor
Application Note 460
October 1986

Introduction

Most commonly-used electrical temperature sensors are difficult to apply. For example, thermocouples have low output levels and require cold junction compensation. Thermistors are nonlinear. In addition, the outputs of these sensors are not linearly proportional to any temperature scale. Early monolithic sensors, such as the LM3911, LM134 and LM135, overcame many of these difficulties, but their outputs are related to the Kelvin temperature scale rather than the more popular Celsius and Fahrenheit scales. Fortunately, in 1983 two I.C.'s, the LM34 Precision Fahrenheit Temperature Sensor and the LM35 Precision Celsius Temperature Sensor, were introduced. This application note will discuss the LM34, but with the proper scaling factors can easily be adapted to the LM35.

The LM34 has an output of 10 mV/˚F with a typical nonlinearity of only ±0.35˚F over a –50 to +300˚F temperature range, and is accurate to within ±0.4˚F typically at room temperature (77˚F). The LM34's low output impedance and linear output characteristic make interfacing with readout or control circuitry easy. An inherent strength of the LM34 over other currently available temperature sensors is that it is not as susceptible to large errors in its output from low level leakage currents. For instance, many monolithic temperature sensors have an output of only 1 µA/˚K. This leads to a 1˚K error for only 1 µ-Ampere of leakage current. On the other hand, the LM34 may be operated as a current mode device providing 20 µA/˚F of output current. The same 1 µA of leakage current will cause an error in the LM34's output of only 0.05˚F (or 0.03˚K after scaling).

Low cost and high accuracy are maintained by performing trimming and calibration procedures at the wafer level. The device may be operated with either single or dual supplies. With less than 70 µA of current drain, the LM34 has very little self-heating (less than 0.2˚F in still air), and comes in a TO-46 metal can package, a SO-8 small outline package and a TO-92 plastic package.

Forerunners to the LM34

The making of a temperature sensor depends upon exploiting a property of some material which is a changing function of temperature. Preferably this function will be a linear function for the temperature range of interest. The base-emitter voltage (V_{BE}) of a silicon NPN transistor has such a temperature dependence over small ranges of temperature.

Unfortunately, the value of V_{BE} varies over a production range and thus the room temperature calibration error is not specified nor guaranteeable in production. Additionally, the temperature coefficient of about –2 mV/˚C also has a tolerance and spread in production. Furthermore, while the tempo may appear linear over a narrow temperature, there is a definite nonlinearity as large as 3˚C or 4˚C over a full –55˚C to +150˚C temperature range.

Another approach has been developed where the difference in the base-emitter voltage of two transistors operated at different current densities is used as a measure of temperature. It can be shown that when two transistors, Q1 and Q2, are operated at different emitter current densities, the difference in their base-emitter voltages, ΔV_{BE}, is

$$\Delta V_{BE} = V_{BE1} - V_{BE2} = \frac{kT}{q} \ln \frac{(J_{E1})}{J_{E2}} \quad (1)$$

where k is Boltzman's constant, q is the charge on an electron, T is absolute temperature in degrees Kelvin and J_{E1} and J_{E2} are the emitter current densities of Q1 and Q2 respectively. A circuit realizing this function is shown in *Figure 1*.

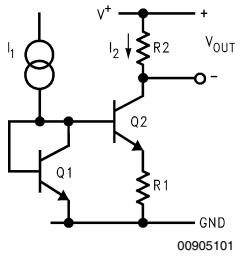

FIGURE 1.

Equation (1) implies that as long as the ratio of I_{E1} to I_{E2} is held constant, then ΔV_{BE} is a linear function of temperature (this is not exactly true over the whole temperature range, but a correction circuit for the nonlinearity of V_{BE1} and V_{BE2} will be discussed later). The linearity of this ΔV_{BE} with temperature is good enough that most of today's monolithic temperature sensors are based upon this principle.

An early monolithic temperature sensor using the above principle is shown in *Figure 2*. This sensor outputs a voltage which is related to the absolute temperature scale by a factor of 10 mV per degree Kelvin and is known as the LM135. The circuit has a ΔV_{BE} of approximately

(0.2 mV/˚K) x (T)

developed across resistor R. The amplifier acts as a servo to enforce this condition. The ΔV_{BE} appearing across resistor R is then multiplied by the resistor string consisting of R and the 26R and 23R resistors for an output voltage of (10 mV/˚K) x (T). The resistor marked 100R is used for offset trimming. This circuit has been very popular, but such Kelvin temperature sensors have the disadvantage of a large constant output voltage of 2.73V which must be subtracted for use as a Celsius-scaled temperature sensor.

Forerunners to the LM34 (Continued)

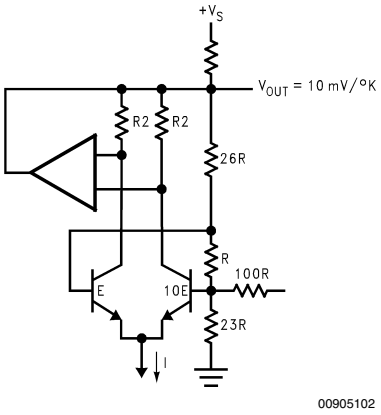

FIGURE 2.

Various sensors have been developed with outputs which are proportional to the Celsius temperature scale, but are rather expensive and difficult to calibrate due to the large number of calibration steps which have to be performed. Gerard C.M. Meijer[4] has developed a circuit which claims to be inherently calibrated if properly trimmed at any one temperature. The basic structure of Meijer's circuit is shown in *Figure 3*. The output current has a temperature coefficient of 1 µA/°C. The circuit works as follows: a current which is proportional to absolute temperature, I_{PTAT}, is generated by a current source. Then a current which is proportional to the V_{BE} drop of transistor Q4 is subtracted from I_{PTAT} to get the output current, I_O. Transistor Q4 is biased by means of a PNP current mirror and transistor Q3, which is used as a feedback amplifier. In Meijer's paper it is claimed that the calibration procedure is straightforward and can be performed at any temperature by trimming resistor R4 to adjust the sensitivity, dI_O/dT, and then trimming a resistor in the PTAT current source to give the correct value of output current for the temperature at which the calibration is being performed.

Meijer's Celsius temperature sensor has problems due to its small output signal (i.e., the output may have errors caused by leakage currents). Another problem is the trim scheme requires the trimming of two resistors to a very high degree of accuracy. To overcome these problems the circuits of *Figure 4* (an LM34 Fahrenheit temperature sensor) and *Figure 5* (an LM35 Celsius temperature sensor) have been developed to have a simpler calibration procedure, an output voltage with a relatively large tempco, and a curvature compensation circuit to account for the non-linear characteristics of V_{BE} versus temperature. Basically, what happens is transistors Q1 and Q2 develop a ΔV_{BE} across resistor R1. This voltage is multiplied across resistor nR1. Thus at the non-inverting input of amplifier A2 is a voltage two diode drops below the voltage across resistor nR1. This voltage is then amplified by amplifier A2 to give an output proportional to whichever temperature scale is desired by a factor of 10 mV per degree.

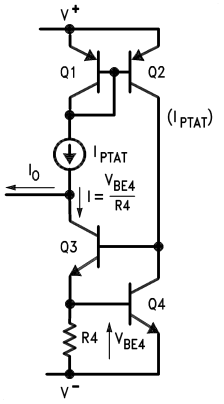

FIGURE 3.

Circuit Operation

Since the two circuits are very similar, only the LM34 Fahrenheit temperature sensor will be discussed in greater detail. The circuit operates as follows:

Transistor Q1 has 10 times the emitter area of transistor Q2, and therefore, one-tenth the current density. From *Figure 4*, it is seen that the difference in the current densities of Q1 and Q2 will develop a voltage which is proportional to absolute temperature across resistor R_1. At 77°F this voltage will be 60 mV. As in the Kelvin temperature sensor, an amplifier, A1, is used to insure that this is the case by servoing the base of transistor Q1 to a voltage level, V_{PTAT}, of ΔV_{BE} x n. The value of n will be trimmed during calibration of the device to give the correct output for any temperature.

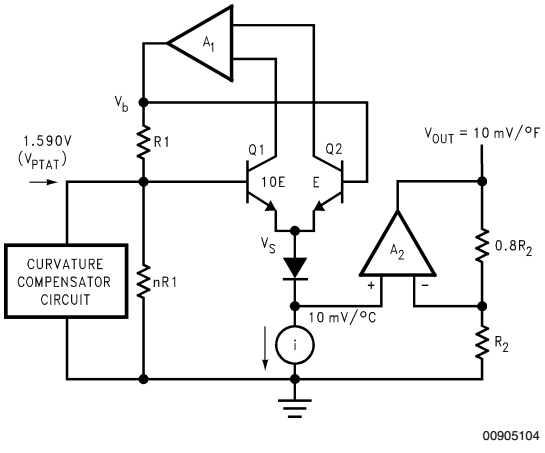

FIGURE 4.

Circuit Operation (Continued)

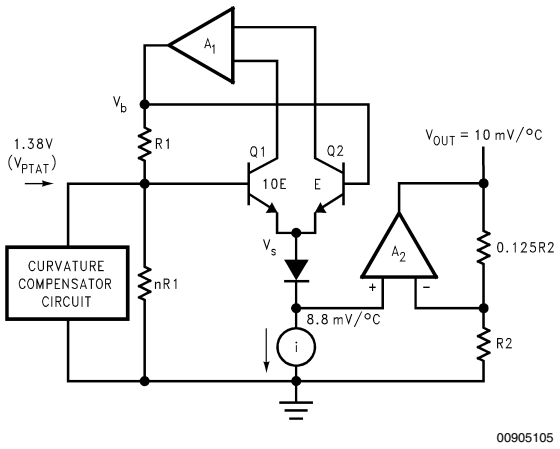

FIGURE 5.

For purposes of discussion, suppose that a value of V_{PTAT} equal to 1.59V will give a correct output of 770 mV at 77˚F. Then n will be equal to $V_{PTAT}/\Delta V_{BE}$ or 1.59V/60 mV = 26.5, and V_{PTAT} will have a temperature coefficient (tempco) of:

$$\frac{nk}{q} \ln \frac{I_1}{I_2} = 5.3 \text{ mV/°C.}$$

Subtracting two diode drops of 581 mV (at 77˚F) with tempcos of −2.35 mV/˚C each, will result in a voltage of 428 mV with a tempco of 10 mV/˚C at the non-inverting input of amplifier A2. As shown, amplifier A2 has a gain of 1.8 which provides the necessary conversion to 770 mV at 77˚F (25˚C). A further example would be if the temperature were 32˚F (0˚C), then the voltage at the input of A2 would be 428 mV−(10 mV/˚C) (25˚C) = 0.178, which would give V_{OUT} = (0.178) (1.8) = 320 mV—the correct value for this temperature.

Easy Calibration Procedure

The circuit may be calibrated at any temperature by adjusting the value of the resistor ratio factor n. Note that the value of n is dependent on the actual value of the voltage drop from the two diodes since n is adjusted to give a correct value of voltage at the output and not to a theoretical value for PTAT. The calibration procedure is easily carried out by opening or shorting the links of a quasi-binary trim network like the one shown in *Figure 6*. The links may be opened to add resistance by blowing an aluminum fuse, or a resistor may be shorted out of the circuit by carrying out a "zener-zap". The analysis in the next section shows that when the circuit is calibrated at a given temperature, then the circuit will be accurate for the full temperature range.

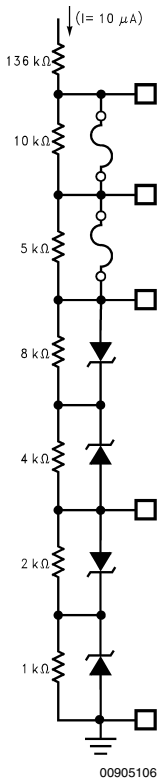

FIGURE 6.

How the Calibration Procedure Works

Widlar[5] has shown that a good approximation for the base-emitter voltage of a transistor is:

$$V_{BE} = V_{G0}\left(1 - \frac{T}{T_0}\right) + V_{beo}\left(\frac{T}{T_0}\right) + \frac{nkT}{q} \ln\left(\frac{T_0}{T}\right) + \frac{kT}{q} \ln \frac{I_C}{I_{co}} \quad (2)$$

where T is the temperature in ˚Kelvin, T_0 is a reference temperature, V_{GO} is the bandgap of silicon, typically 1.22V,

How the Calibration Procedure Works (Continued)

and V_{beo} is the transistor's base-emitter voltage at the reference temperature, T_0. The above equation can be re-written as

$$V_{BE} = (\text{sum of linear temp terms}) + (\text{sum of non-linear temp terms}) \quad (3)$$

where the first two terms of *Equation (1)* are linear and the last two terms are non-linear. The non-linear terms were shown by Widlar to be relatively small and thus will be considered later.

Let us define a base voltage, V_b, which is a linear function of temperature as: $V_b = C_1 \bullet T$. This voltage may be represented by the circuit in *Figure 1*. The emitter voltage is $V_e = V_b - V_{be}$ which becomes:

$$V_e = C_1T - V_{G0}\left(1 - \frac{T}{T_0}\right) - V_{beo}\left(\frac{T}{T_0}\right).$$

If V_e is defined as being equal to C_2 at $T = T_0$, then the above equation may be solved for C_1. Doing so gives:

$$C_1 = \frac{V_{beo} + C_2}{T_0} \quad (4)$$

Using this value for C_1 in the equation for V_e gives:

$$V_e = C_2\frac{T}{T_0} + V_{G0}\left(\frac{T - T_0}{T_0}\right) \quad (5)$$

If V_e is differentiated with respect to temperature, T, *Equation (4)* becomes $dV_e/dT = (C_2 + V_{G0})/T_0$.

This equation shows that if V_b is adjusted at T_0 to give $V_e = C_2$, then the rate of change of V_e with respect to temperature will be a constant, independent of the value of V_b, the transistor's beta or V_{be}. To proceed, consider the case where $V_e = C_2 = 0$ at $T_0 = 0^\circ C$. Then

$$\frac{dV_e}{dT} = \frac{V_{G0}}{273.7} = 4.47\ mV/^\circ C$$

Therefore, if V_e is trimmed to be equal to (4.47 mV) T (in °C) for each degree of displacement from 0°C, then the trimming can be done at ambient temperatures.

In practice, the two non-linear terms in *Equation (1)* are found to be quadratic for positive temperatures. Tsividis[6] showed that the bandgap voltage, V_0, is not linear with respect to temperature and causes nonlinear terms which become significant for negative temperatures (below 0°C). The sum of these errors causes an error term which has an approximately square-law characteristic and is thus compensated by the curvature compensation circuit of *Figure 7*.

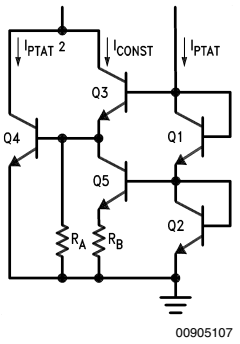

FIGURE 7.

A Unique Compensation Circuit

As mentioned earlier, the base-emitter voltage, V_{BE}, is not a linear function with respect to temperature. In practice, the nonlinearity of this function may be approximated as having a square-law characteristic. Therefore, the inherent non-linearity of the transistor and diode may be corrected by introducing a current with a square-law characteristic into the indicated node of *Figure 4*. Here's how the circuit of *Figure 7* works: transistors Q1 and Q2 are used to establish currents in the other three transistors. The current through Q1 and Q2 is linearly proportional to absolute temperature, I_{PTAT}, as is the current through transistor Q5 and resistor R_B. The current through resistor R_A is a decreasing function of temperature since it is proportional to the V_{BE} of transistor Q4. The emitter current of Q3 is equal to the sum of the current through Q5 and the current through R_A, and thus Q3's collector current is a constant with respect to temperature. The current through transistor Q4, I_{C4}, will be used to compensate for the V_{BE} nonlinearities and is found with the use of the following equation:

$$I_{C4} = I_S\left(e^{\frac{qV_{BE4}}{KT}} - 1\right) \cong I_S\, e^{\frac{qV_{BE4}}{KT}}$$

where $V_{BE4} = V_{BE1} + V_{BE2} - V_{BE3}$.

From the above logarithmic relationship, it is apparent that I_{C4} becomes

$$I_{C4} = \frac{I_{C1}\, I_{C2}}{I_{C3}} = \frac{I_{PTAT}{}^2}{I_{CONST}}.$$

A Unique Compensation Circuit

(Continued)

Thus, a current which has a square law characteristic and is $PTAT^2$, is generated for use as a means of curvature correction.

Processing and Layout

The sensor is constructed using conventional bipolar epitaxial linear processing. SiCr thin-film resistors are used in place of their diffused counterparts as a result of their better tempco matching, an important consideration for resistors which must track over temperature. Such resistors include R1 and nR1 of the bandgap circuit.

Another point of interest in the construction of the device centers around transistors Q1 and Q2 of *Figure 4*. In order for the circuit to retain its accuracy over temperature, the leakage currents of each transistor, which can become quite significant at high temperatures, must be equal so that their effects will cancel one another. If the geometries of the two transistors were equivalent, then their leakage currents would be also, but since Q1 has ten times the emitter area of Q2, the accuracy of the device could suffer. To correct the problem, the circuit is built with Q1 and Q2 each replaced by a transistor group consisting of both Q1 and Q2. These transistor groups have equivalent geometries so that their leakage currents will cancel, but only one transistor of each group, representing Q1 in one group and Q2 in the other pair is used in the temperature sensing circuit. A circuit diagram demonstrating this idea is shown in *Figure 8*.

Using the LM34

The LM34 is a versatile device which may be used for a wide variety of applications, including oven controllers and remote temperature sensing. The device is easy to use (there are only three terminals) and will be within 0.02˚F of a surface to which it is either glued or cemented. The TO-46 package allows the user to solder the sensor to a metal surface, but in doing so, the GND pin will be at the same potential as that metal. For applications where a steady reading is desired despite small changes in temperature, the user can solder the TO-46 package to a thermal mass. Conversely, the thermal time constant may be decreased to speed up response time by soldering the sensor to a small heat fin.

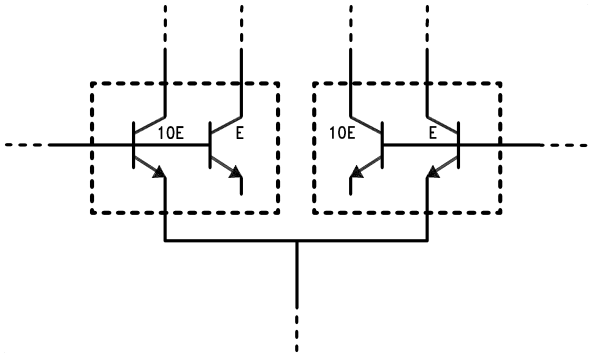

00905108

FIGURE 8.

Fahrenheit Temperature Sensors

As mentioned earlier, the LM34 is easy to use and may be operated with either single or dual supplies. *Figure 9* shows a simple Fahrenheit temperature sensor using a single supply. The output in this configuration is limited to positive temperatures. The sensor can be used with a single supply over the full −50˚F to +300˚F temperature range, as seen in *Figure 10*, simply by adding a resistor from the output pin to ground, connecting two diodes in series between the GND pin and the circuit ground, and taking a differential reading. This allows the LM34 to sink the necessary current required for negative temperatures. If dual supplies are available, the sensor may be used over the full temperature range by merely adding a pull-down resistor from the output to the negative supply as shown in *Figure 11*. The value of this resistor should be $|-V_S|/50\ \mu A$.

For applications where the sensor has to be located quite a distance from the readout circuitry, it is often expensive and inconvenient to use the standard 3-wire connection. To overcome this problem, the LM34 may be connected as a two-wire remote temperature sensor. Two circuits to do this are shown in *Figure 12* and *Figure 13*. When connected as a remote temperature sensor, the LM34 may be thought of as a temperature-dependent current source. In both configurations the current has both a relatively large value,

Fahrenheit Temperature Sensors

(Continued)

$$(20\ \mu A/^{\circ}F) \times (T_A + 3^{\circ}F),$$

and less offset when compared to other sensors. In fact, the current per degree Fahrenheit is large enough to make the output relatively immune to leakage currents in the wiring.

Temperature to Digital Converters

For interfacing with digital systems, the output of the LM34 may be sent through an analog to digital converter (ADC) to provide either serial or parallel data outputs as shown in *Figure 14* and *Figure 15*. Both circuits have a 0 to +128˚F scale. The scales are set by adjusting an external voltage reference to each ADC so that the full 8 bits of resolution will be applied over a reduced analog input range. The serial output ADC uses an LM385 micropower voltage reference diode to set its scale adjust (V_{REF} pin) to 1.28V, while the parallel output ADC uses half of an LM358 low power dual op amp configured as a voltage follower to set its $V_{REF}/2$ pin to 0.64V. Both circuits are operated with standard 5V supplies.

Basic Fahrenheit Temperature Sensor (+5˚ to +300˚F)

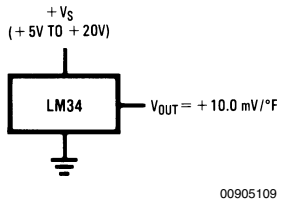

FIGURE 9.

Temperature Sensor, Single Supply, –50˚ to +300˚F

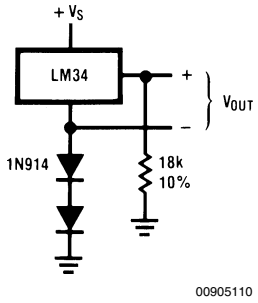

FIGURE 10.

Full-Range Fahrenheit Temperature Sensor

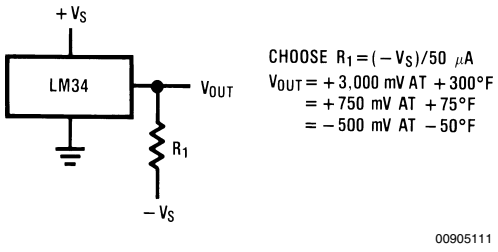

FIGURE 11.

Temperature to Digital Converters (Continued)

Two-Wire Remote Temperature Sensor
(Grounded Sensor)

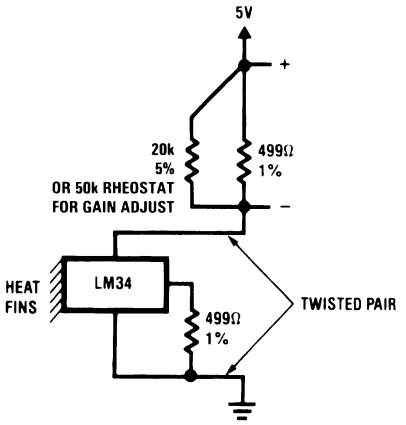

V_{OUT} = 10 mV/°F (T_A + 3°F)
from +3°F to +100°F

FIGURE 12.

Two-Wire Temperature Sensor
(Output Referred to Ground)

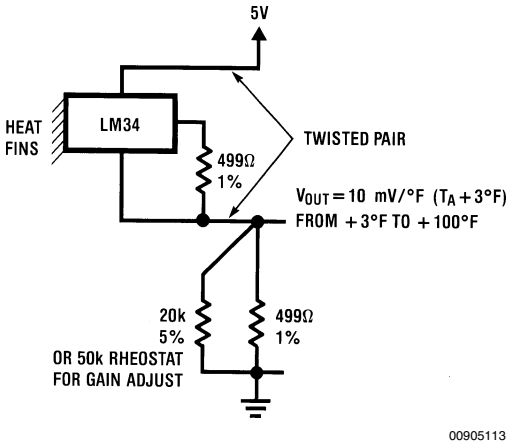

FIGURE 13.

Temperature to Digital Converters (Continued)

Temperature-to-Digital Converter (Serial Output, +128˚F Full Scale)

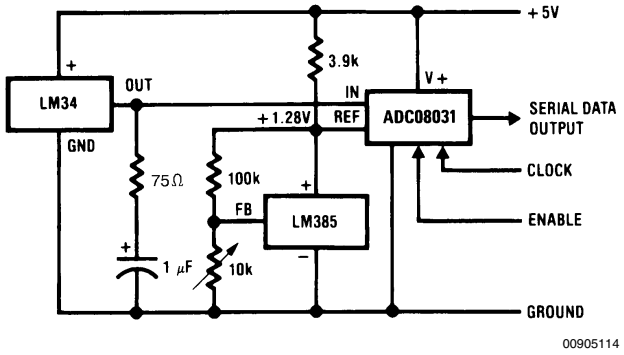

FIGURE 14.

Temperature-to-Digital Converter
(Parallel TRI-STATE Outputs for Standard Data Bus to µP Interface, 128˚F Full Scale)

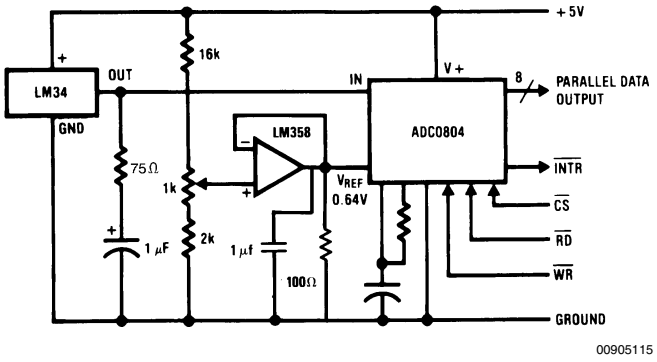

FIGURE 15.

LM34 with Voltage-to-Frequency Converter and Isolated Output
(3˚F to +300˚F; 30 Hz to 3000 Hz)

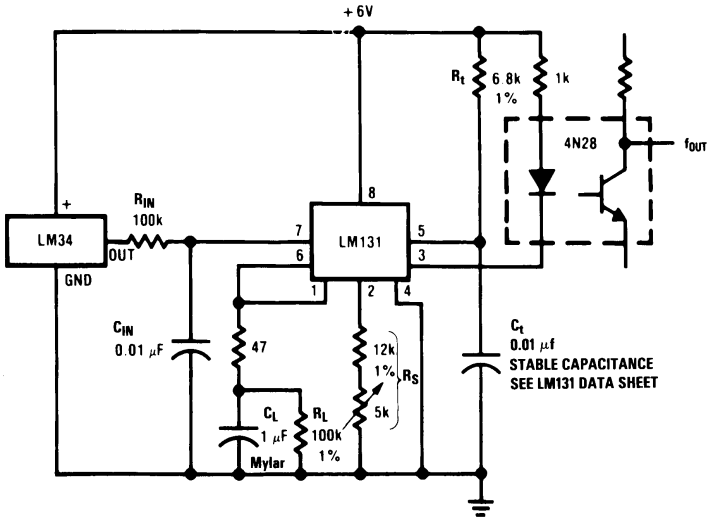

FIGURE 16.

Temperature to Digital Converters (Continued)

Bar-Graph Temperature Display (Dot Mode)

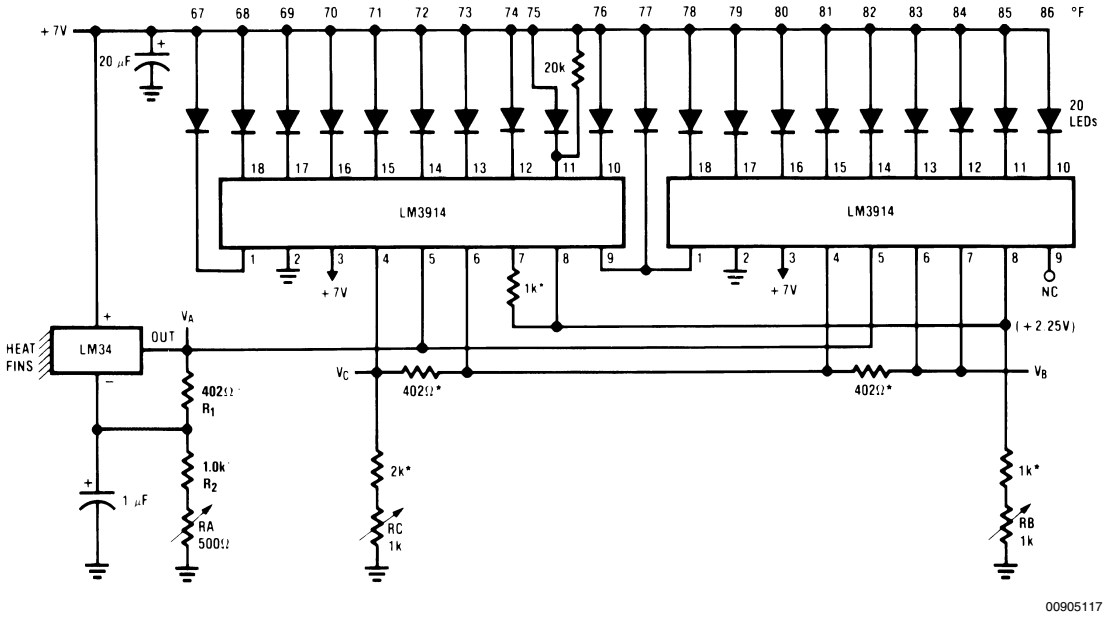

• = 1% or 2% film resistor

— Trim R_B for V_B = 3.525V

— Trim R_C for V_C = 2.725V

— Trim R_A for V_A = 0.085V + 40 mV/˚F x $T_{Ambient}$

— Example: V_A = 3.285V at 80˚F

FIGURE 17.

Temperature-to-Frequency Converter for Remote Sensing

If a frequency proportional to temperature is needed, then the LM34 can be used in conjunction with an LM131 voltage-to-frequency converter to perform the desired conversion from temperature to frequency. A relatively simple circuit which performs over a +3˚F to +300˚F temperature range is shown in *Figure 16*. The output frequency of this circuit can be found from the equation:

$$f_{OUT} = \left(\frac{V_{IN}}{2.09V}\right)\left(\frac{R_S}{R_L}\right)\left(\frac{1}{R_t\, C_t}\right)$$

where resistor R_S is used to adjust the gain of the LM131. If R_S is set to approximately 14.2 kΩ, the output frequency will have a gain of 10 Hz/˚F. Isolation from high common mode levels is provided by channeling the frequency through a photoisolator. This circuit is also useful for sending temperature information across long transmission lines where it can be decoded at the receiving station.

LED Display for Easy Temperature Reading

It is often beneficial to use an array of LED's for displaying temperature. This application may be handled by combining the LM34 with an LM3914 dot/display driver. The temperature may then be displayed as either a bar of illuminated LED's or as a single LED by simply flipping a switch. A wide range of temperatures may be displayed at once by cascading several LM3914's as shown in *Figure 17*.

Without going into how the LM3914 drivers function internally, the values for V_A, V_B, and V_C can be determined as follows:

V_A is the voltage appearing at the output pin of the LM34. It consists of two components, 0.085V and (40 mV/˚F) (T_A).

The first term is due to the LM34's bias current (approximately 70 μA) flowing through the 1 kΩ resistor in series with R_A. The second term is a result of the multiplication of the LM34's output by the resistive string composed of R_1, R_2, and R_A, where R_A is set for a gain factor of 4 (i.e.; 40 mV/˚F).

V_B represents the highest temperature to be displayed and is given by the equation V_B = 0.085V + (40 mV/˚F) (T_{HIGH}). For the circuit in *Figure 17* , V_B = 0.085V + (40 mV/˚F) (86˚F) = 3.525V.

V_C represents the lowest temperature to be displayed minus 1˚F. That is, V_C = 0.085V + (T_{LOW} −1˚F) (40 mV/˚F) which in this case becomes V_C = 0.085V + (67˚F −1˚F) (40 mV/˚F) = 2.725V.

With a few external parts, the circuit can change from dot to bar mode or flash a bar of LED's when the temperature sensed reaches a selected limit (see LM3914 data sheet).

Indoor/Outdoor Thermometer

An indoor/outdoor thermometer capable of displaying temperatures all the way down to –50˚F is shown in *Figure 18*.

Indoor/Outdoor Thermometer

(Continued)

Two sensor outputs are multiplexed through a CD4066 quad bilateral switch and then displayed one at a time on a DVM such as Texmate's PM-35X. The LMC555 timer is run as an astable multivibrator at 0.2 Hz so that each temperature reading will be displayed for approximately 2.5 seconds. The RC filter on the sensors outputs are to compensate for the capacitive loading of the cable. An LMC7660 can be used to provide the negative supply voltage for the circuit.

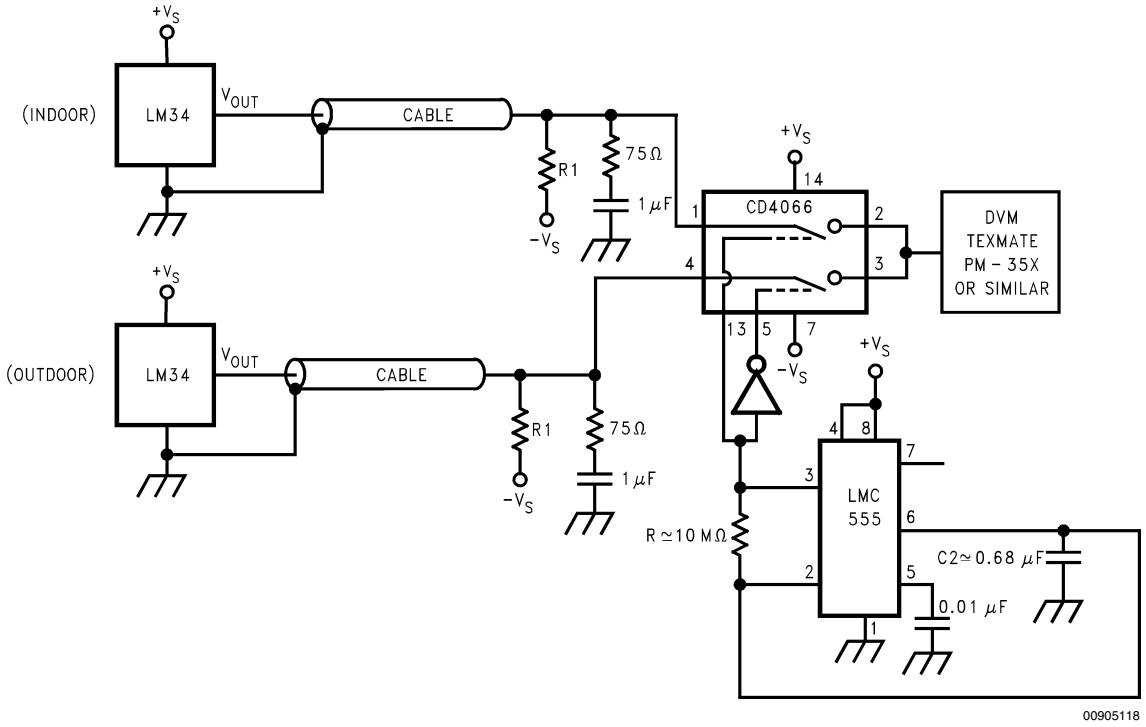

FIGURE 18.

Temperature Controller

A proportional temperature controller can be made with an LM34 and a few additional parts. The complete circuit is shown in *Figure 19*. Here, an LM10 serves as both a temperature setting device and as a driver for the heating unit (an LM395 power transistor). The optional lamp, driven by an LP395 Transistor, is for indicating whether or not power is being applied to the heater.

When a change in temperature is desired, the user merely adjusts a reference setting pot and the circuit will smoothly make the temperature transition with a minimum of overshoot or ringing. The circuit is calibrated by adjusting R2, R3 and C2 for minimum overshoot. Capacitor C2 eliminates DC offset errors. Then R1 and C1 are added to improve loop stability about the set point. For optimum performance, the temperature sensor should be located as close as possible to the heater to minimize the time lag between the heater application and sensing. Long term stability and repeatability are better than 0.5˚F.

Differential Thermometer

The differential thermometer shown in *Figure 20* produces an output voltage which is proportional to the temperature difference between two sensors. This is accomplished by using a difference amplifier to subtract the sensor outputs from one another and then multiplying the difference by a factor of ten to provide a single-ended output of 100 mV per degree of differential temperature.

Differential Thermometer (Continued)

Temperature Controller

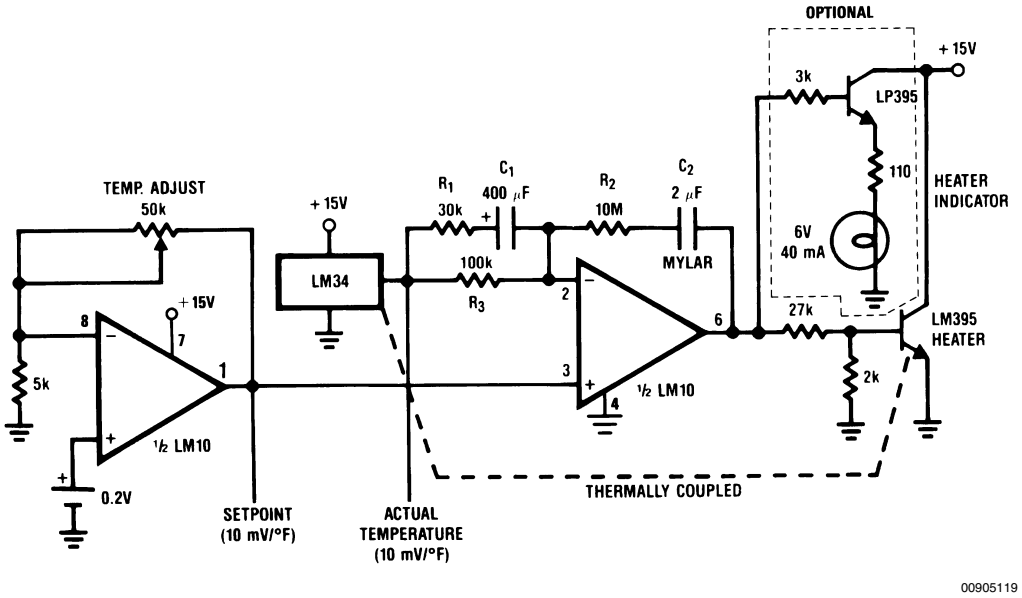

FIGURE 19.

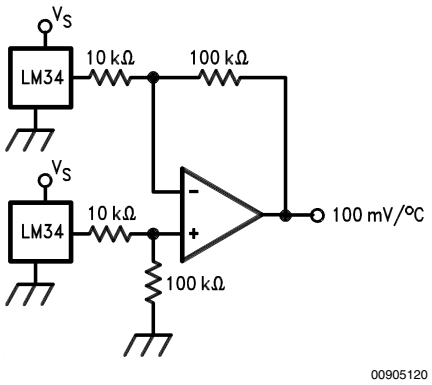

FIGURE 20.

Differential Thermometer (Continued)

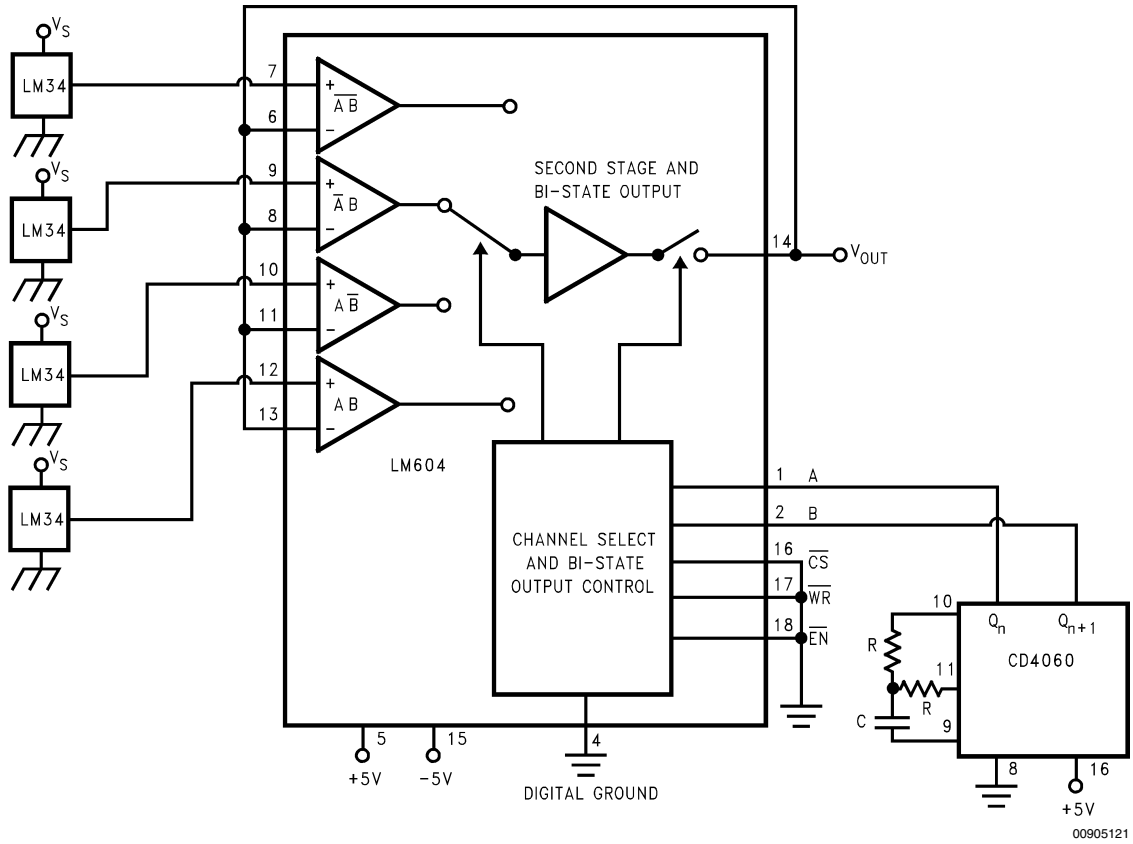

Temperature monitoring

$$\text{rate} = \frac{0.56}{RC}\left(\frac{1}{2^{n-1}}\right)$$

If Q_4 and Q_5 are used with R = 13k and C = 510 pF the rate will be 10 kHz.

FIGURE 21.

Temperature Scanner

In some applications it is important to monitor several temperatures periodically, rather than continuously. The circuit shown in *Figure 21* does this with the aid of an LM604 Mux Amp. Each channel is multiplexed to the output according to the AB channel select. The CD4060 ripple binary counter has an on-board oscillator for continuous updating of the channel selects.

Conclusion

As can be seen, the LM34 and LM35 are easy-to-use temperature sensors with excellent linearity. These sensors can be used with minimal external circuitry for a wide variety of applications and do not require any elaborate scaling schemes nor offset voltage subtraction to reproduce the Fahrenheit and Celsius temperature scales respectively.

References

1. R.A. Pease, “A Fahrenheit Temperature Sensor”, presented at the ISSCC Conference, February 24, 1984.
2. R.A. Pease, “A New Celsius Temperature Sensor”, National Semiconductor Corp., 1983.
3. Robert Dobkin, “Monolithic Temperature Transducer”, in *Dig. Tech. Papers,* Int. Solid State Circuits Conf., 1974, pp. 126, 127, 239, 240.
4. Gerard C.M. Meijer, “An IC Temperature Sensor with an Intrinsic Reference”, *IEEE Journal of Solid State Circuits,* VOL SC-15, June 1980, pp. 370–373.
5. R.J. Widlar, “An Exact Expression for the Thermal Variation of the Emitter — Base Voltage of Bipolar Transistors”, *Proc. IEEE,* January 1967.
6. Y.P. Tsividis, “Accurate Analysis of Temperature Effects in $I_C - V_{BE}$ Characteristics with Application to Bandgap Reference Sources”, *IEEE Journal of Solid-State Circuits,* December 1980, pp. 1076–1084.
7. Michael P. Timko, “A Two-Terminal IC Temperature Transducer”, *IEEE Journal of Solid-State Circuits,* December 1976, pp. 784–788.

LM385 Feedback Provides Regulator Isolation

National Semiconductor
Application Note 715
Robert Pease
Fran Hoffart
November 1990

You can use a conventional 4N27 optocoupler in a feedback arrangement (*Figure 1*) to design a switching regulator with a floating output. The LM3524 switch-mode-regulator IC is configured as a simple flyback power supply with transformer-isolated output. The LM385 acts as a reference and comparison amplifier that satisifes the 4N27's current demands for balancing the dc feedback to pin 2 of the LM3524, thereby closing the loop. The LM385 automatically compensates for any LED degradation or optical-coupling loss.

The 4N27 specs a 0.1 dc to 1.6 dc gain range; the ac gain varies from 0.05 to 1.0. Fortunately, the "Miller" damper from pins 5 to 6 nullifies the effect of the wide gain variance. Moreover, the damper provides excellent loop stability for a wide range of 4N27 optocouplers from several manufacturers. In the example shown, the LM385 provides 0.5 mA to 5 mA to the optocoupler.

If the 5V output available from this circuit does not meet your needs, choose $R_2 = R_1 (V_{OUT} - 1.25)/1.25$. If V_{OUT} is greater than 6V, insert a zener diode between points X and Y, with $V_Z = V_{OUT} - 5V$; this addition prevents the voltage across the LM385 from exceeding its 5.3V max limit. The circuit is suitable for regulated output voltages from 3.2V to 25V.

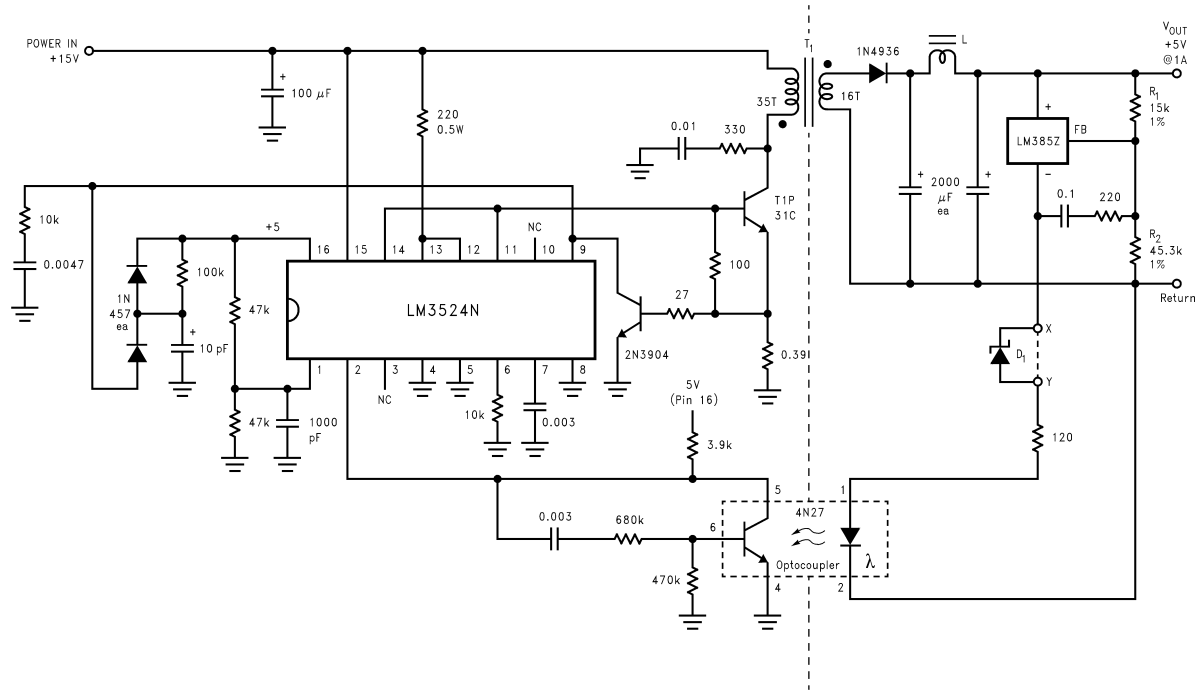

L = 15 turns #22 wire wrapped on ½W resistor body.

D_1 = optional (see text).

FIGURE 1. Optocoupler-based feedback circuit produces galvanically-isolated, regulated voltages. The heavy negative feedback compensates for wide variations in optocoupler gains. The values portrayed in this schematic yield 5V output; by varying R_2, you can obtain voltages from 3.2V to 25V.

The Effect of Heavy Loads on the Accuracy and Linearity of Operational Amplifier Circuits (or, "What's All this Output Impedance Stuff, Anyhow?")

National Semiconductor
Application Note 1485
Bob Pease
May 7, 2008

Introduction

It is well known that the ideal operational amplifier (op amp) should have very high gain, very high bandwidth, very high input impedance, and very low output impedance. [1] **It is possible to get conventional amplifiers with very high gain (120 dB or higher), and very high bandwidth (100 MHz, 1000 MHz, or more). However, most op amps do not have a very low open-loop output impedance (Zout). Only a few are as low as 50 ohms, and can drive a 50-ohm load without any significant degradation of gain (barely 2:1). See** *Figure 1*.

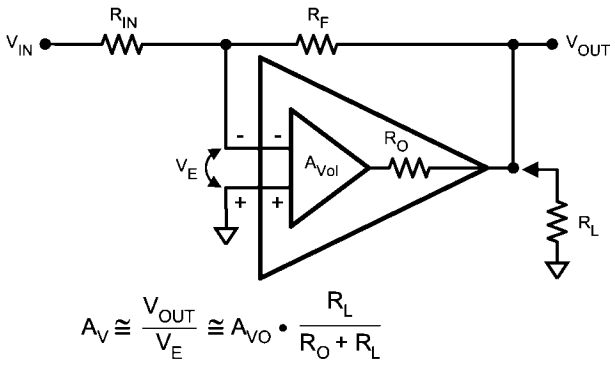

FIGURE 1. Model of Operational Amplifier (Op-Amp) with finite output impedance Ro. If Ro is significant compared to RL, the effective Av (Vout/ VE) will be degraded vs. Avol.

Many op amps designed over the last 50 years have Class B or class A-B emitter-follower output stages, which help provide low output impedance and high efficiency. Many of these use mature bipolar transistor technology, and can operate on ± 15 volts. See *Figure 2*a.

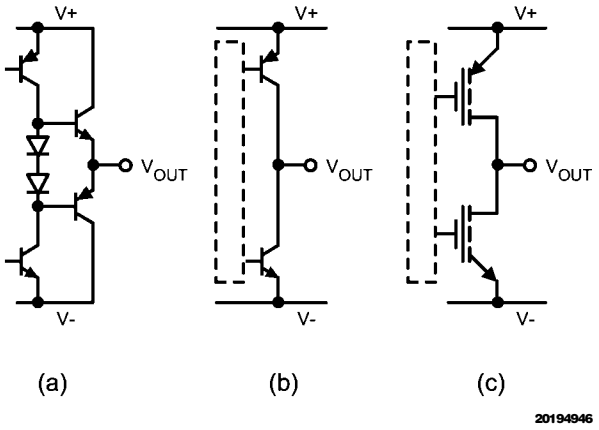

**FIGURE 2. (a) Conventional high-gain Op-Amp with emitter-follower output stage (simplified).
(b) Op-amp with collector-loaded "rail-to-rail" output stage (simplified).
(c) CMOS Op-amp with drain-loaded "rail-to-rail" output stage (simplified)**

It is also known that the closed-loop output impedance of a typical operational amplifier can be MUCH lower than the open-loop output impedance. If an op amp has, for example, an open-loop gain of 10,000, and its open-loop output impedance is 50 ohms, the output impedance after the loop is closed can be as low as 50 milliohms or lower, depending on the application (assuming the amplifier is used at a gain of 5 or lower). So for many applications, at least at low frequencies, it is a fair statement that the closed-loop output impedance can be very low.

However, a new class of amplifiers has been introduced over the last 30 years, which do not have emitter-follower outputs. Why not? Because many of the new amplifiers are designed to operate on low voltages such as ± 5 volts, or ± 2.5 volt, or ± 0.9 volts or sometimes even lower. For maximum signal-to-noise ratio, the output must swing from (nearly) the + rail to (nearly) the - rail. [2]

Obviously any emitter-followers would reduce the output swing by about 0.7 volts in either direction (and even worse at cold temperatures), so amplifier designs that use followers have become obsolete for such low-voltage applications. See *Figure 2*b.

The first "rail-to-rail" output stage was introduced in Bob Widlar's LM10. This was designed and released in 1976, and is still in production. It can operate from ± 20 volts to ± 0.6 volts (or from 40 volts down to 1.2-volts of total power supply) and its output can swing within a few dozens millivolts of the power supply rails. It does not have any output emitter followers. The LM10's output consists of one big NPN output stage to pull the output down, and sink 15 to 20 mA of current, and comparable PNP transistors to source as much as 15 to 20 mA. It has some very complicated bias circuits to make sure these two transistors take turns at driving the load, as required. *Figure 2*b.

More recently, over the last 15 years, there have been dozens of different designs, mostly using CMOS technology, and all have "Drain-loaded" outputs, with N-channel and P-channel FETs which can source and sink many milliamperes, *Figure 2*c. These all have high output impedances. One way to look at it is, that the gain gets lower when you load the output with a heavy load. Another way to look at it is, that the gain RISES when the load gets lighter. See *Figure 3*.

In concept, a Drain-loaded output stage could use negative feedback to an internal stage, so that the gain would not change much as the load gets heavy or light. Practically, it would take a lot of high-value resistors to accomplish this, and such resistors would be very expensive in monolithic IC technology. In practice, the disadvantage that the gain changes as the load is changed, is not serious. This is largely because the gain is very high when the load is heavy, and it just gets higher when the load is lightened.

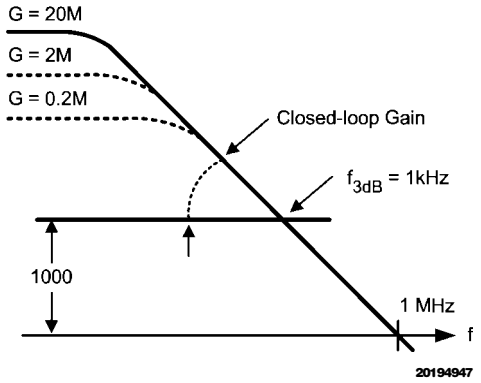

FIGURE 3. Bode plot for high-gain amplifiers with various DC gains, operating at a gain of 1000.

I once heard some engineers argue that there is no advantage when an op amp's gain gets very high, and in concept there may be disadvantages. One argument is that there is no need for any op-amp to have a gain greater than 200,000, because if the gain gets higher, it would have to be tested at very low frequencies, lower than 0.1 Hz. Such testing would take many seconds, and this testing would be quite expensive, and nobody would want to pay for that.

However, this turns out to be untrue, as modern amplifier testing can resolve a "dc gain" as high as 2 million or 20 million, in just a few milliseconds. No 20-second test is required. An operational amplifier with 1 MHz of Gain-Bandwidth Product, operating at a closed-loop gain of 1000, has a closed-loop bandwidth of 1 kHz. Thus its time constant is 160 μseconds, and it can settle in less than 20 tau, or 4 milliseconds, per *Figure 3*.

In its gain test, the output is required to go to its positive rated output, and the input error settles quickly and is then measured, for perhaps 16 milliseconds. The output is then required to go to its negative rated output, the input settles, and then is measured again. The gain depends on the reciprocal of the small difference between those input tests. This is easy to do quickly, even for high gains. It takes less than 1/10 second, not "several" seconds, to test for amplifiers even with a gain of 1 million or 10 million or more. See *Figure 3*.

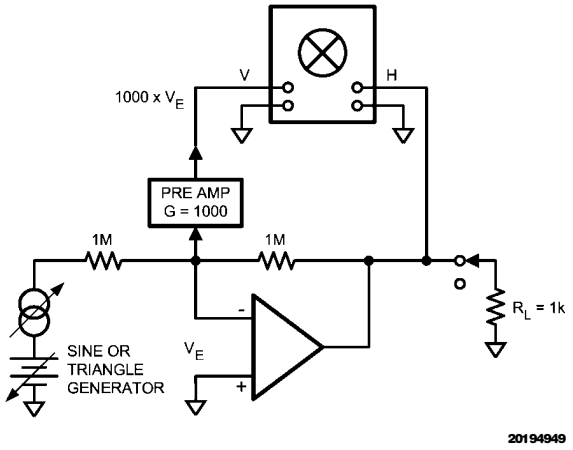

FIGURE 4. Gain test where Av = 1000 x Vout / 1000 x V_E, using separate preamplifier and X-Y oscilloscope.

Another argument is that an amplifier with a gain of 2 million or 20 million, would not be useful except for signals slower than 0.1 Hz. This also turns out to be a misconception. If a modern op-amp is connected for a gain of + 1000.00, and a 1.0 mV dc signal is applied to the input, the output will settle in a few milliseconds, per *Figure 5*. However, an amplifier with a mere gain of 200,000 would settle its output to 995 millivolts. A gain of 2 million would settle to 999.5 millivolts, and an amplifier with a gain of 20 million will settle to 999.95 millivolts - in milliseconds! MUCH better accuracy.

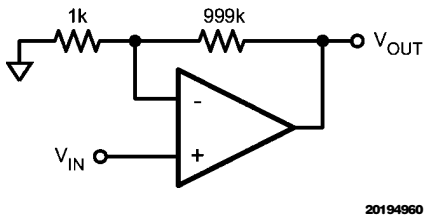

FIGURE 5. High-gain amplifier operating at a gain of 1000. Its precision depends on high Avol (and low Vos).

Furthermore, if you put in 1.0000 millivolt p-p sine waves, at 5 or 10 Hz or 20 Hz, the output amplitudes of those three amplifiers would be, respectively, 995 mv p-p, 999.5 mv p-p, and 999.95 mV, p-p. Even at 10 or 20 Hz, a precision amplifier can provide enhanced accuracy over low-gain amplifiers. The claim that a high open-loop gain at 0.1 or 0.01 Hz is useless, unless your signal is at 0.01 Hz, is just incorrect.

Some other engineers say that an amplifier with high output impedance and good gain (such as 1 million at 1 Hz) can have its dc gain rise to 10 million or more, if the rated load is taken off. The DC gain would rise so high, they claim, that when it starts to roll off, it would roll off too fast, with excessive phase shift, and be unstable. Refer to *Figure 6*. In actuality, all op-amps these days have smooth 6-dB per octave rolloff, all the way back to very low frequencies. Op amps that rolled off at

10 or 12 dB per octave, when the rated load is taken off, have not been seen for over 30 years.

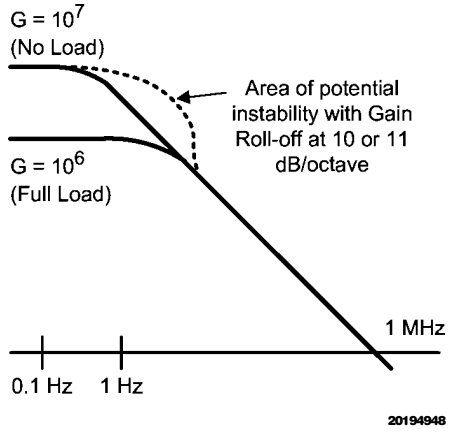

FIGURE 6. High-gain amplifiers with extremely high gain.

So an operational amplifier with very high gain actually does have some good advantages, and not really any disadvantages.

When an op-amp is asked to drive a heavy output current, it can have large errors if it does not have plenty of g_m or transconductance. This is true whether it has output followers and low output impedance, OR if it has high output impedance. So the g_m is very important, and a good amplifier design must have an appropriate amount of g_m - plenty of *mhos* **(milliamperes per millivolt). Many precision amplifiers have many mhos of g_m. As we shall see, precision amplifiers such as the LM627, LMC6022 , and LMP2012 have a g_m of at least 10,000 mhos. Other popular amplifiers have 50 to 500 mhos. Special-purpose amplifiers may have as little as 2 to 20 mhos, which may be adequate for particular needs.**

Instrumentation

Many modern op amps have such high gain that a preamp with a gain such as 1000 is needed, to let you see the gain error. Even then, a time-based scope does not let you resolve the linear and nonlinear components of the gain error. So I used a Tektronix 2465 (analog) oscilloscope in X-Y(cross-plot) mode. One good way to test the amplifier is to connect the Device Under Test (DUT) as a unity-gain inverter as shown in *Figure 4* **and feed the output of the DUT to the scope horizontal display, through a 10X (10 megohm) probe. The output was also fed to one of the vertical channels, so we see the cross-plot of V_{out} versus itself, as a straight line, with a slope of + 45 degrees. Typically, the first amplifiers I tested were high-voltage bipolar amplifiers swinging ± 10 volts, with the scope set at 5 volts per division. The signal source was a Wavetek 193, with adjustable amplitude and variable DC offset. I used the variable offset to adjust the output to swing exactly ± 10 volts, for the high-voltage amplifiers. The output swing was set at ± 4 volt peaks for CMOS amplifiers running on ± 5 volts, and ± 2.0 volts for amplifiers running on ± 2.5 volts, in general.**

The -input voltage (the gain error) is fed to a preamp with a gain of +1000. This was sometimes fed directly to the scope's vertical input (DC coupled) at sensitivities varying from 2000 mV to 5 mV per division, yielding a resolution of 2 mV down to 5 µV per division. By using the cross-plot mode, it was possible to resolve 1 or 2 µV p-p of gain error in the presence of a few microvolts of noise. For amplifiers with large offset voltage, I fed the signal in to the scope's DC input through 11 µF, so that 0.2 Hz signals could be resolved without appreciable phase shift.

The test circuit I actually used was *Figure 7*, **with the amplifier acting as a unity gain inverter for the signals, and acting as a preamp of gain = 1000 for its own error signals. This makes the test set-up easier. The output voltage is plotted in each Test as a straight line at 45° slope, versus the same signal on the horizontal axis. The output voltage is positive and the output is sourcing current on the right side of each Figure, and the output voltage is negative, and sinking current, on the left.**

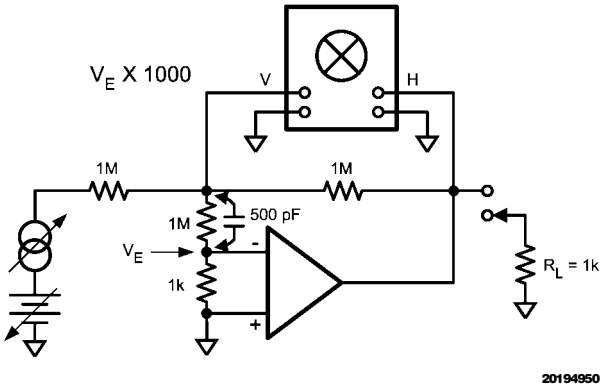

FIGURE 7. Gain test where the DUT acts as its own preamp. Gain = 1001 x (Vvert)/(Vhoriz.)

I generally used a triangle wave for almost all tests. This gave better resolution of p-p errors for the gain test, and it allowed me to run at moderate frequencies (1 to 10 to 80 Hz) and not get the DC gain error signal confounded by the ac gain error. Refer to *Figure 31*, **the plot of Test A11. Even though the gain at 8 Hz on this amplifier was just 2,500,000, I was easily able to resolve the 2.5 µV of gain error, which is completely independent from the AC gain error. The AC gain error (due to finite gain-bandwith product) causes the upper and lower traces to separate by 8 µV, yet we can still see the "DC" gain error of 2.5 µV (gain of 8 million). The gain error is the SLOPE of either the upper or the lower trace, as the output ramps back and forth. This gives much better resolution than a sine wave, and is easier to instrument at a higher sweep rate.**

For example... measuring the dc gain of the LMP2012 with Test F01A would require operating at sine frequencies below 0.1 Hz; but by using a 2 Hz triangle wave, we could see that the DC slope would be less than 1 µV at 0.01Hz, by "subtracting" the opening between the upper and lower traces.

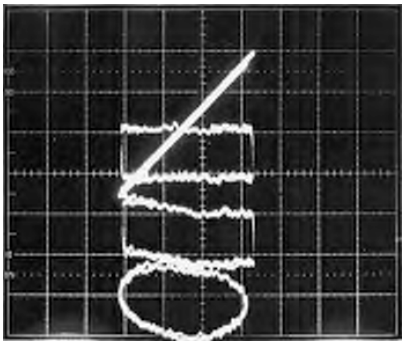

20194907

Test D07B, LMC6022 F = 2 Hz.
Vs = ± 5 Vdc; Vout = ± 4 volts peak, Iout = ± 4 mA peak.
Upper Trace: Gain Error, No Load, 4μV p-p at 20μV/div.
Middle Trace: Gain Error, Full Load, 7μV p-p at 20μV/div. (TRIANGLE)
Lower Trace: Gain Error, Full Load, 7μV p-p at 20μV/div. (SINE)

FIGURE 8.

Also, when we start seeing nonlinearity, we can easily resolve what is nonlinear, because the error correlates with the location on the X-Y plot. In *Figure 8*, **we see the curves taken from Test D07B. This is an example of an amplifier, the LMC6022, with distortion down near 1/2 ppm (+/- 2 μV). When we use triangle waves it is easy to see this distortion, per the middle trace. If we relied on sine waves, it would be hard to resolve this amount of distortion, per the lowest curve.**

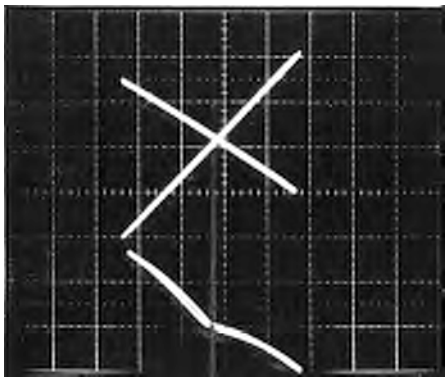

20194921

Test A01, LM709, Curve of Gain Error, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 480μV p-p at 200μV/div.
Lower Trace: Gain Error, Full Load, 540μV p-p at 200μV/div.

FIGURE 9.

Bipolar Amplifiers -- And Funny Errors

I started by measuring the old LM709, one of the first monolithic op-amps, almost 40 years old. This was a good test. The gain error at 1 megohm load was 480 microvolts, so the Av was 42,000 at 10 Hz. This was safely better than the 25,000 published spec of the device. I then applied the 1 kilohm load.

Most of these op amps were rated to drive a 2 kilohm load, but I put on a 1k load, to see what was going to happen. It made errors about twice as big as they would have been with a 2k load, which was slightly unfair, but helped the resolution of the errors, which were often pretty small. (On rare occasions, I could tell that a 1k load was unfair, so I would re-test at a 2k load, to see what was really going on at rated load.)

In the case of the LM709, (*Figure 9***) a 1k load actually caused the overall gain slope to degrade by about 60 μV. This corresponds to an output impedance of about 120 ohms, not too bad. However, there was definitely some non-linear error - about 80 μV p-p. Where did that came from? This nonlinear error seems to be bigger than the linear error caused by Zout.**

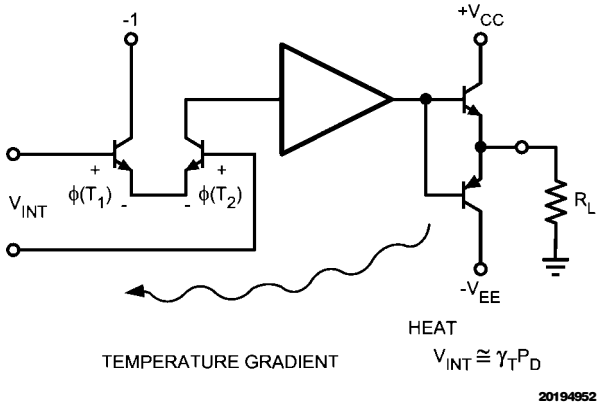

20194952

FIGURE 10. AN-A shows that the heat from an amplifier's output transistors can flow past its input transistors.

This has been thoroughly analyzed in a 1975 technical article, known as NSC Application Note A (AN-A) by James Solomon. This App Note analyzes the circuit and layout of a monolithic amplifier, where an output stage drives a heavy load. One or the other of the output transistors gets warm to the extent of 25 or 50 milliwatts, and sends thermal gradients across the IC chip. See *Figure 12*. **For a complete overview, refer to AN-A at http://www.national.com/an/AN/AN-A.pdf . But in their simplest form, the drawings from AN-A are included here.** *Figure 11* **and** *Figure 13* **show that a mere 49 milliwatts can cause a 40 milli-degree C temperature gradient, between the input transistors of the op-amp, located 10 milli-inches apart. If the input transistors were laid out transverse to the heat gradient, they would be heated to the same extent, and the thermal error referred to input could have been quite small. The LM709's input transistors are Q1 and Q2, per** *Figure 14*. **They were located along the gradient, 10 milli-inches apart, (***Figure 15***) and did a very good job of detecting the thermal gradient. Every competitor who copied Widlar's LM709 was afraid to change anything, for fear of making something worse! Even after the LM709 became obsolescent, other amplifiers' layouts still did not do a very good job of rejecting the thermal gradients, for many years.**

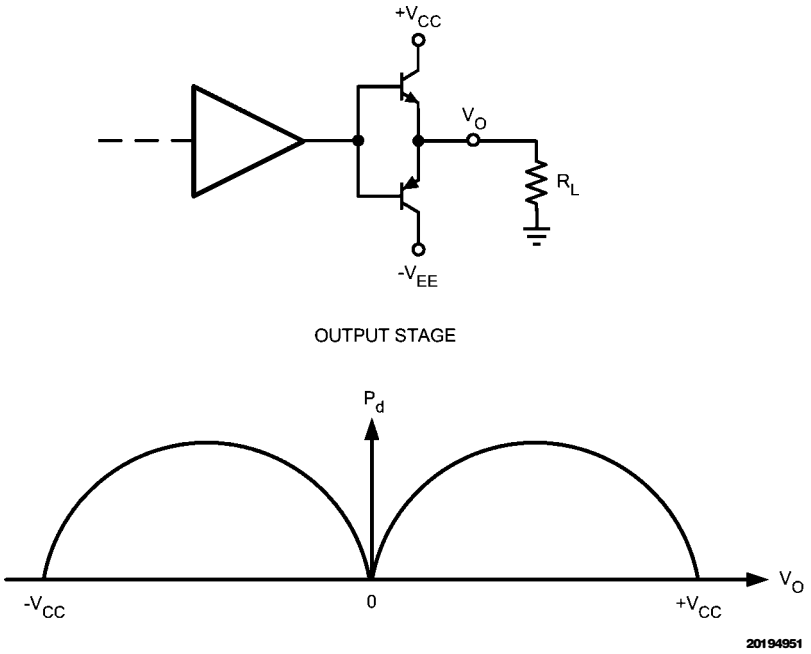

FIGURE 11. AN-A shows the shape of the input error caused by output heat flow, when cross-plotted vs. Vout on an X-Y scope.

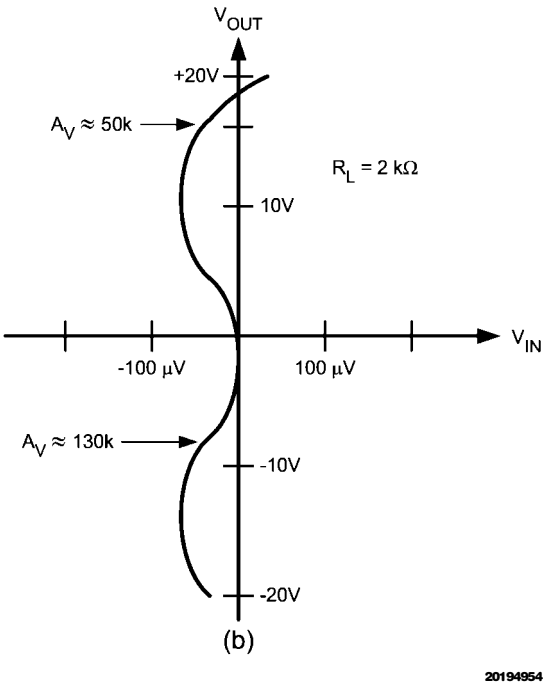

FIGURE 12. AN-A shows the shape of the input error caused by output heat flow, when cross-plotted vs. Vout on an X-Y scope.

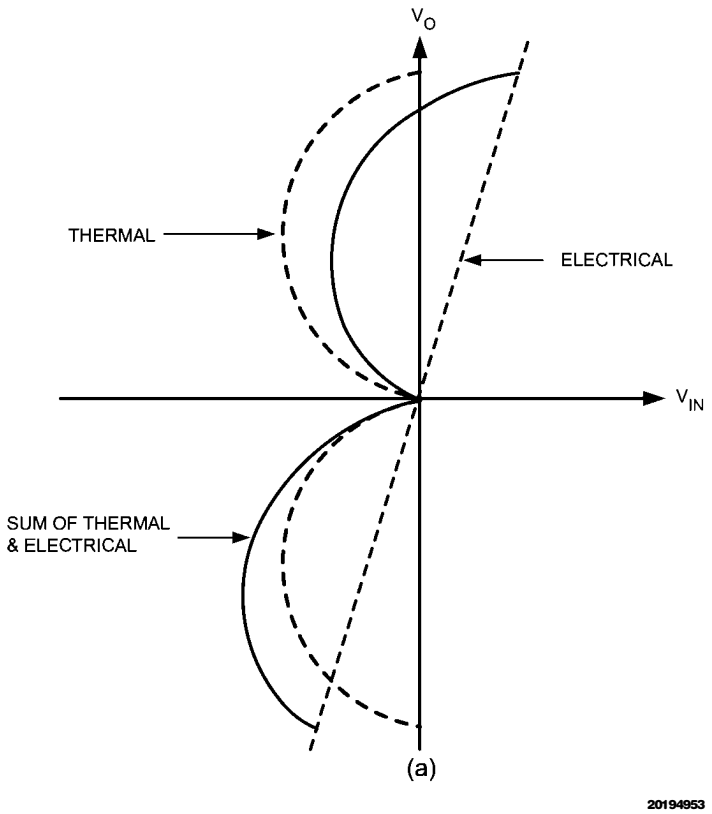

FIGURE 13. AN-A shows that the shape of the error voltage (gain error) can be a summation of electrical and thermal errors. Compare to Figure 9.

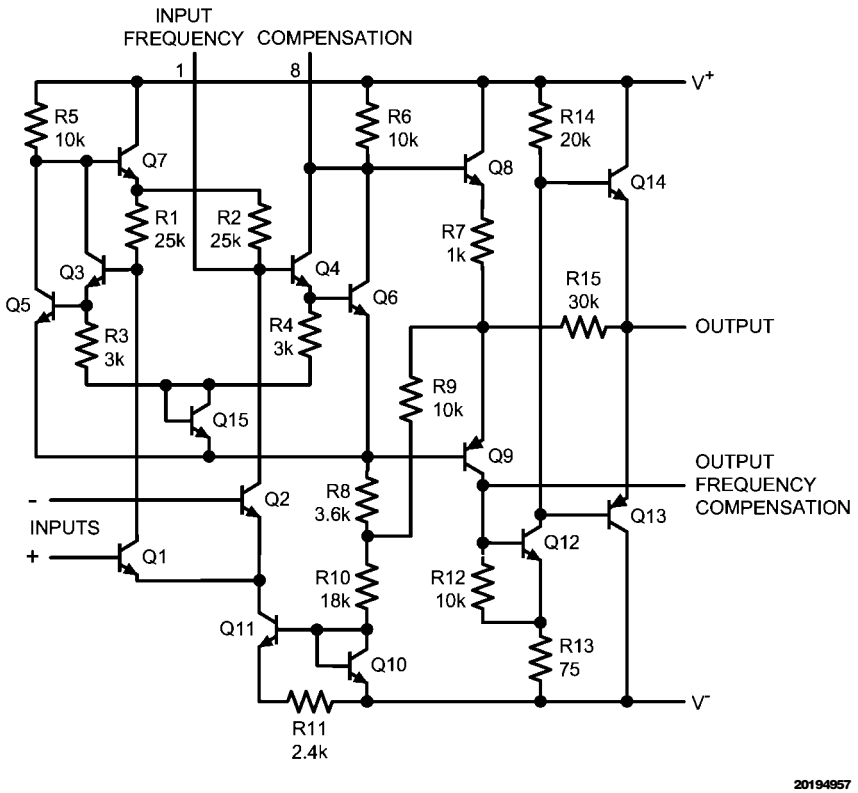

FIGURE 14. Schematic Diagram of the LM709.

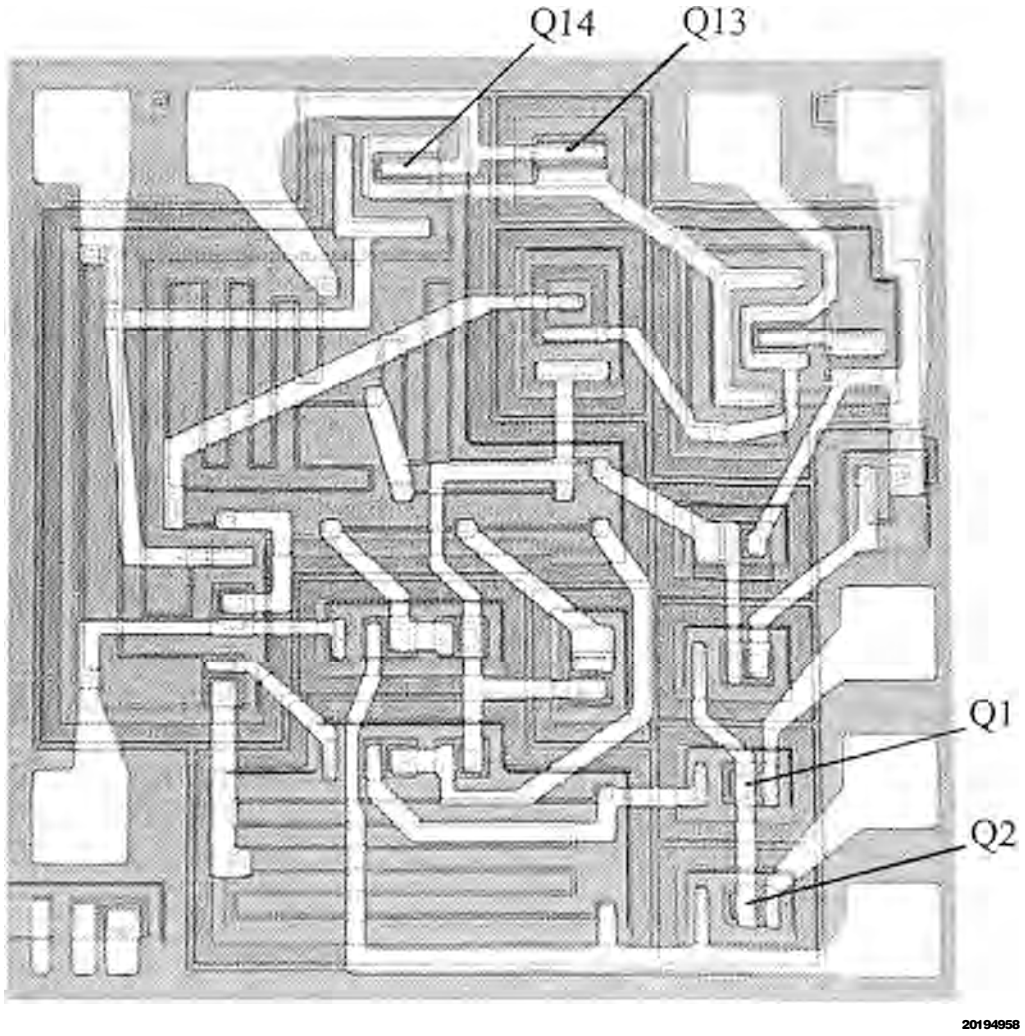

FIGURE 15. Layout of LM709 Die. The spacing from Q1 to Q13 or Q14 is 56 milli-inches, and to Q2 is 10 mils.

Eventually, newer amplifiers took advantage of symmetry and common-centroid layouts (See at "What's All This Common-Centroid Stuff, Anyhow?") http://www.elecdesign.com/Articles/Index.cfm?ArticleID=6121 [3] to reject thermal gradients. Most of the CMOS amplifiers we will study, below, do not have any appreciable thermal errors, because the CMOS amplifiers were carefully laid out with good layouts to reject thermal gradients. These were accomplished mostly with the use of symmetry, and not with the use of computers. That is because computers are not generally suitable for analyzing the heat flow among the millions of points inside a silicon die, not to mention the thousands of points in time, when a thermal transient occurs. Also, if an amplifier is modelled in SPICE, the SPICE models of most transistors do not allow the transistors to be at different temperatures. New and improved transistor models do now (2001 to 2006) have the ability to analyze temperature differences, but these models are bulky and slow and not highly successful. Symmetry generally works much better.

Other than that, the LM709's gain error was quite adequate for most applications. And if the 709 was run with a load of not such a heavy resistance as 1k or 2k, but 4 k ohms total, its linearity would be as low as 2 ppm. So even the oldest amplifier designs are not too bad.

The next example of a good amplifier with imperfect thermal layout is the familiar LM301A, per *Figure 16*. Its no-load gain was measured at 280,000. But at full load, its non-linear error is also about 80 μV p-p. This, too, gives acceptable over-all performance. *Figure 17* shows an LM301A's thermal errors at various frequencies. The errors at 2 Hz are as expected. When the frequency increases to 20 Hz, the thermal errors are decreasing rapidly. At 200 Hz, they have shrunk to a low level, so the distortion is much less. This is a characteristic of thermal errors, that they decrease rapidly at high frequencies.

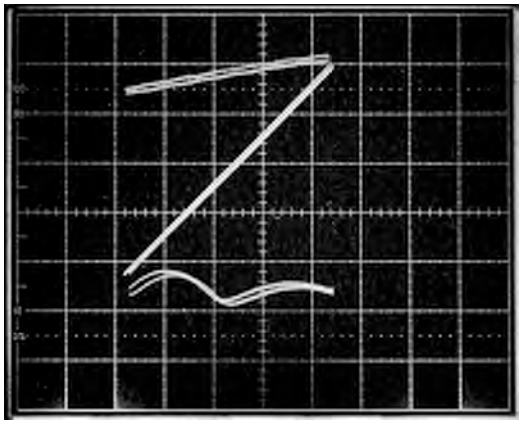

20194961

Test A02, LM301A, Curve of Gain Error, F = 5 Hz
Vs = +/- 15 Vdc; Vout = +/- 10 volts peak, I out = +/-10 mA peak.
Upper Trace: Gain Error, No Load, +75 μV p-p at 100 μV/div.
Lower Trace: Gain Error, Full Load, 70 μV p-p at 100 μV/div.

FIGURE 16.

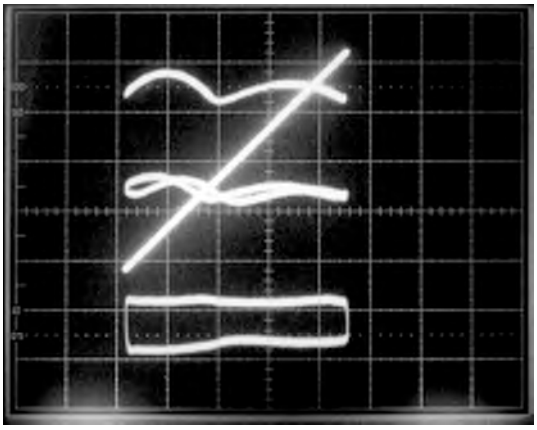

20194962

Test A02B, LM301A, Curve of Gain Error.
Vs = +/- 15 Vdc; Vout = +/- 10 volts peak, I out = +/-10 mA peak.
Upper Trace: Gain Error, Full Load, 60μVp-p at 100μV/div., F = 2 Hz
Middle Trace: Gain Error, Full Load, 60μVp-p at 100μV/div., F = 20 Hz
Lower Trace: Gain Error, Full Load, 25μVp-p at 100μV/div., F = 200 Hz.

FIGURE 17.

Group A: High-Voltage Amplifiers

Now we will go through a big list of operational amplifiers that run on ± 15 volts, and are designed with mostly bipolar transistors. Many of these older amplifiers have imperfect thermal errors, but there are some exceptions.

Test A01 is the LM709, as we have already discussed, included here just for comparison with the other amplifiers in this group.

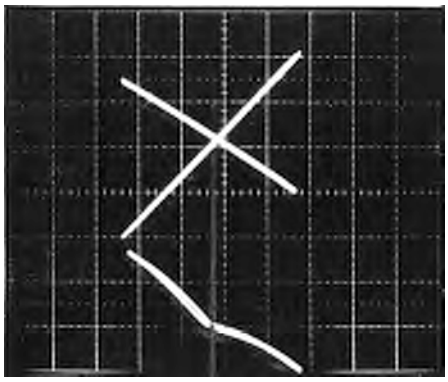

20194921

Test A01, LM709, Curve of Gain Error, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 480µV p-p at 200µV/div.
Lower Trace: Gain Error, Full Load, 540µV p-p at 200µV/div.

FIGURE 18.

Test A02 is the LM301, included again, for comparison.

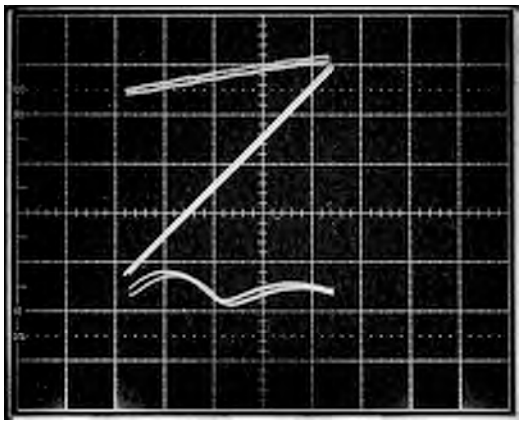

20194961

Test A02, LM101A, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, +28µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 30µV p-p at 50µV/div.

FIGURE 19.

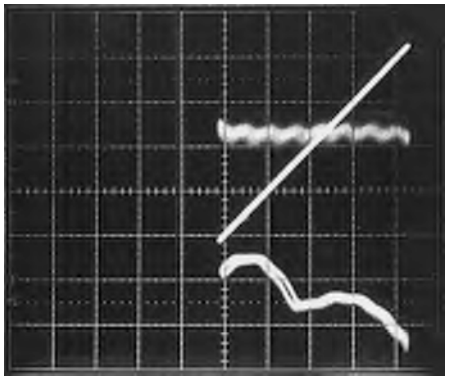

20194923

Test A02B, LM301, F = 4 Hz
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA Peak.
Upper Trace: Gain Error, No Load, 6µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 90µV p-p at 50µV/div.

FIGURE 20.

Test A03 is an LM741. It, too has significant thermal errors. Note that the left-side hump is larger than the right-hand hump, indicating that the output transistor that sinks current has more thermal effect than the one that sources.

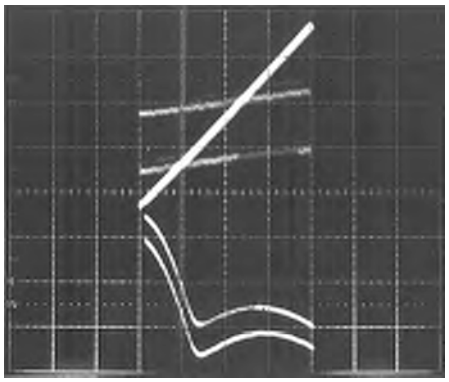

20194924

Test A03, LM741, F = 2 Hz
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA Peak.
Upper Trace: Gain Error, No Load, +9µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 120µV p-p at 50µV/div.

FIGURE 21.

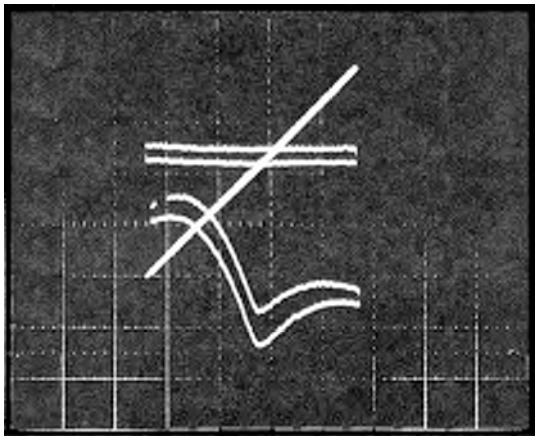

20194925

Test A03B, LM741, F = 1 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 20μV p-p at 50μV/div.
Lower Trace: Gain Error, Full Load, 120μV p-p at 50μV/div.

FIGURE 22.

Test A04 **is the old LM725, not based on the Fairchild μA725. This amplifier had much lower thermal errors than the amplifiers we have seen so far, reflecting an improved layout. This was a 3-stage amplifier with much higher dc gain, about 2 million at no load, and 1.8 million even at full load. However, this design had a large die, was expensive, was hard to provide with damping components, and was never popular. But it did have improved linearity and low thermal errors.**

NOTE **that the frequency response caused the p-p dynamic error to be about 5μV p-p larger at the right-hand side, than at the left. This is because the LM725 was damped largely by diode capacitances, rather than by discrete damping capacitors. The capacitance was larger when the output voltage was positive. It is also noticeable to see the little blip as the output has a bit of cross-over distortion at 0 volts. Still, we are only seeing these tiny errors (with a resolution of about 2μV) because this amplifier's gain and noise are better than most of the previous amplifiers. At moderate loads such as 4 k ohms, it was capable of about 0.1 ppm linearity.**

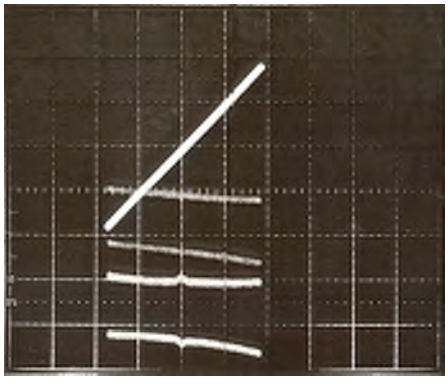

20194926

Test A04, LM725*, F = 75 Hz.
V_S = ±15 Vdc; Vout = ±9 volts peak, Iout = ±9 mA peak.
Upper Trace: Gain Error, No Load, 10μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 11μV p-p at 20μV/div.

FIGURE 23.

Test A05 **is the old LM308. Its thermal errors hump up on one side, and down on the other side, indicating that the thermal errors couple into the input stage differently for outputs sourcing vs. sinking currents. This, too, is down near 1 or 2 ppm of error. However, the LM308 was only rated for a 2 mA load, and this unit was run at just 5 mA, as it could not drive a 1k resistive load.**

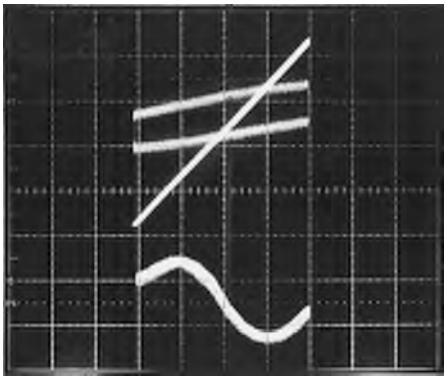

20194927

Test A05, LM308, F = 100 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±5 mA peak. (RL = 2 k)
Upper Trace: Gain Error, No Load, 40μV p-p at 50μV/div.
Lower Trace: Gain Error, Full Load, 80μV p-p at 50μV/div. (F = 2 Hz)

FIGURE 24.

Test A05B **is an older LM308. We don't know how old -- perhaps 25 or 30 years -- but this shows that the chip layout was quite different, with a distinctly different thermal signature, compared to Test A05. The total thermal error is not much better than A05, but it sure is different! Robert Widlar made many experiments of different layouts. Most amplifier designers made one layout, but Widlar knew that it was important to try different layouts, as layout is such an important factor in amplifier performance.**

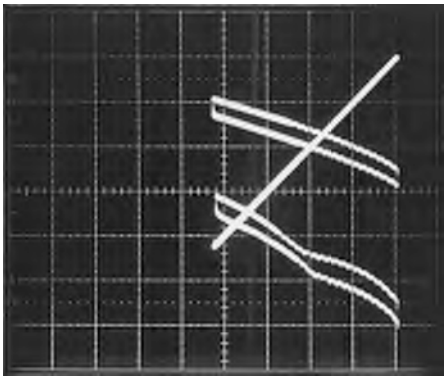

20194928

Test A05B, LM308, OLD, F = 4 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 160μV p-p at 100μV/div.
Lower Trace: Gain Error, Full Load, 260μV p-p at 100μV/div.

FIGURE 25.

Test A06 is the old LM318. This is not a perfect design, and not a very good thermal layout, but it was very fast, and ran rather rich, and hot, and its mediocre thermal errors are acceptable compared to general-purpose amplifiers.

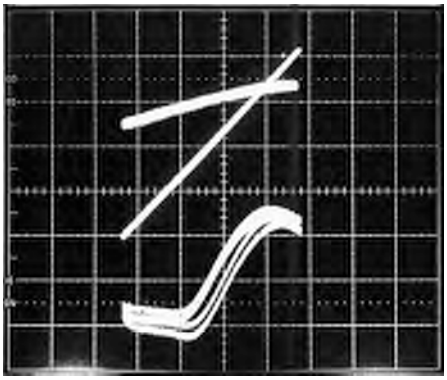

20194929

Test A06, LM318, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 40μV p-p at 50μV/div.
Lower Trace: Gain Error, Full Load, + 125μV p-p at 50μV/div.

FIGURE 26.

Test A07 is an NSC OP-07. Its thermal errors are not appreciably better than normal.

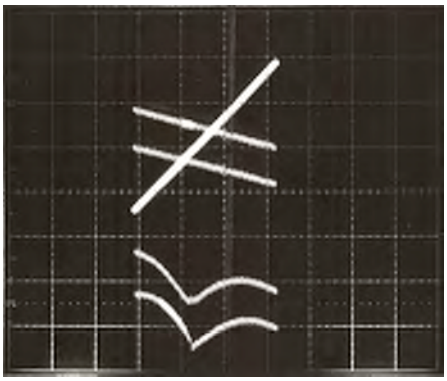

20194930

Test A07, OP-07*, F = 1.2 Hz.
V_S = ±15 Vdc; Vout = ±7.5 volts peak, Iout = ±7.50 mA peak.
Upper Trace: Gain Error, No Load, 18μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 26μV p-p at 20μV/div.

FIGURE 27.

Test A08 is the LF411, with BiFET (TM) input FETs. It used a very complicated layout, that did not work especially well, in terms of gain or thermals. No better than average.

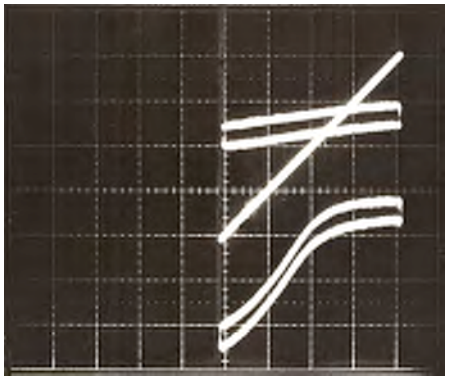

20194931

Test A08, LF411, F = 6 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 24μV p-p at 50μV/div.
Lower Trace: Gain Error, Full Load, + 140μV p-p at 50μV/div.

FIGURE 28.

Test A09 was the older LF356 BiFET amplifier. It had a unique and proprietary output stage, that worked just so-so. It did provide adequate output impedance at 4 MHz, so it was a little faster than most of the general-purpose amplifiers. But its nonlinearity was only average.

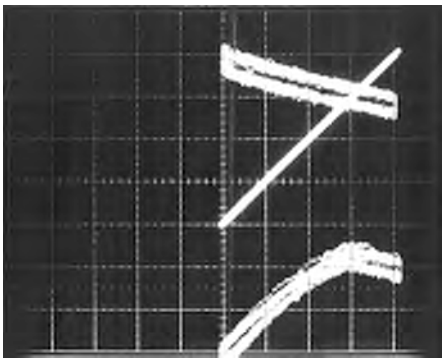

20194932

Test A09, LF356, F = 6Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 55μV p-p at 50μV/div.
Lower Trace: Gain Error, Full Load, + 130μV p-p at 50μV/div.

FIGURE 29.

Test A10 is the LM607. At one time it was the ~ best op amp in the world, but it was discontinued as no customers ever found out about it. Its non-linearity is down at the 0.2 ppm level. The distortion does not look so good, only because the gain is turned up twice as high as ~ any previous amplifier. I used to think the LM607 had a perfect design and layout, but it does seem to show a couple microvolts of thermal error, mostly on the positive side, when sourcing current. This could easily lead to a nonlinearity of 0.15 ppm.

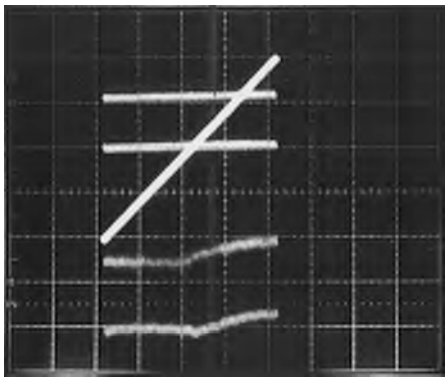

20194933

Test A10, LM607*, F = 3 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 1µV p-p at 10µV/div.
Lower Trace: Gain Error, Full Load, + 5µV p-p at 10µV/div.

FIGURE 30.

Test A11, the LM627, was a similar design to the LM607, but the layout must have gotten lucky, and the thermal errors are down below 1 microvolt, even at the heavy load. I must admit, I am not sure why the gain tends to go from (+ 10 million) at no load, to (+ 4 million) at full load. Adding a heavy load does not usually cause the gain to go more (positive). This amplifier also was not well promoted, was not well known, and was discontinued.

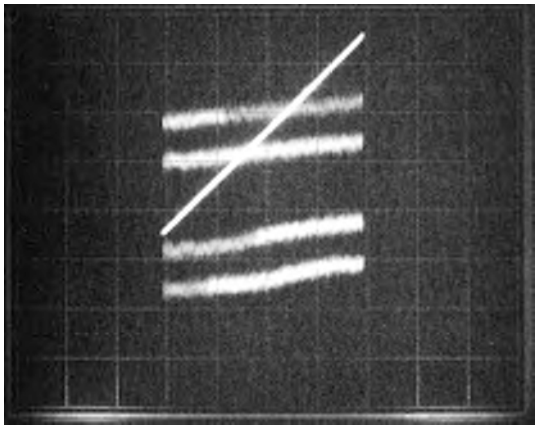

20194970

Test A11, LM627*, F = 8 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 2µV p-p at 10µV/div.
Lower Trace: Gain Error, Full Load, +5µV p-p at 10µV/div.

FIGURE 31.

Test A12 is the LM10. As mentioned earlier, this is the first amplifier with a "rail-to-rail" output. This amplifier met many dc characteristics with miraculous accuracy, but the ac linearity was NOT quite as good as you would expect from Widlar. Later, Widlar's LM12 showed that he could do excellent accuracy for dynamic errors and linearity, but the LM10 was primarily a DC amplifier. Its errors look "pretty bad", but actually its non-linearity was no worse than general-purpose amplifiers -- barely 1 or 2 ppm. Its cross-over distortion was NOT very good, even at 1 Hz, and at higher frequencies, it is not good at all. The LM10 was NOT a good, linear audio amplifier.

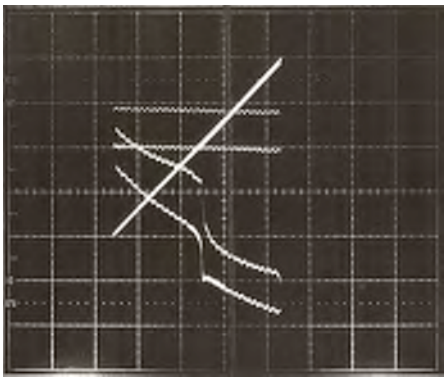

20194935

Test A12, LM10, F = 1 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 4µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 170µV p-p at 50µV/div.

FIGURE 32.

Test A13 shows an LM10 running slightly faster, at 10 Hz. If you look at the lower trace, done with a sine wave, it looks very distorted and confusing , and it is hard to see what is going on. The upper trace shows the error using a triangle wave. This looks just like a speeded up version of the curve at Test A12. This is one of the major reasons we prefer using triangle waves, rather than sines -- so we can see and understand what is going on.

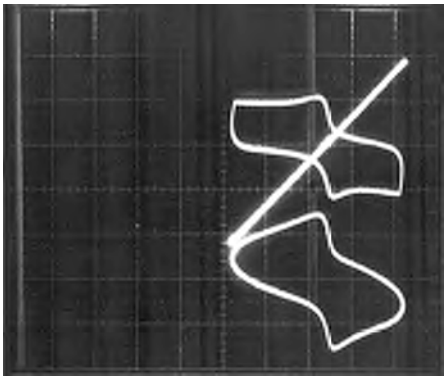

20194936

Test A13, LM10, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 230µV p-p at 100µV/div., TRIANGLE wave.
Lower Trace: Gain Error, Full Load, 290µV p-p at 100µV/div., SINE wave.

FIGURE 33.

Test A14 shows an LM307, a version of the LM301 with a 30 pF compensation capacitor built in. This re-layout caused somewhat different thermal errors. The distortion is about typical for general-purpose amplifiers.

This completes the study of single high-voltage amplifiers.

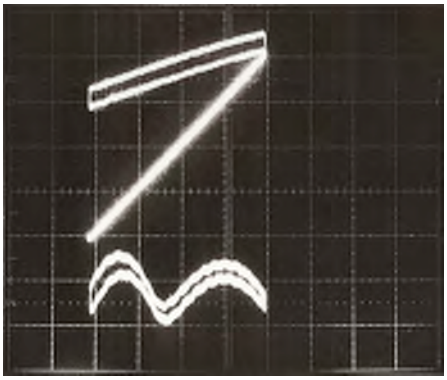

20194937

Test A14, LM307J, F = 2 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 55µV p-p at 50µV/div.
Lower Trace: Gain Error, Full Load, 60µV p-p at 50µV/div.

FIGURE 34.

Section B, High-Voltage (±15V) DUAL Amplifiers

Test B01, the LM358, is the dual version of the LM324. No study of amplifiers would be complete without a mention of the pioneering LM324/LM358. This is the first amplifier whose honest gain is so non-linear. That is because the output stage has a Darlington to source the output current, but only one vertical PNP to drive the sinking current. So it really is deficient in gain, for negative currents. The DC distortion is STILL at the 1.5 ppm level. But the thermal errors are negligible.

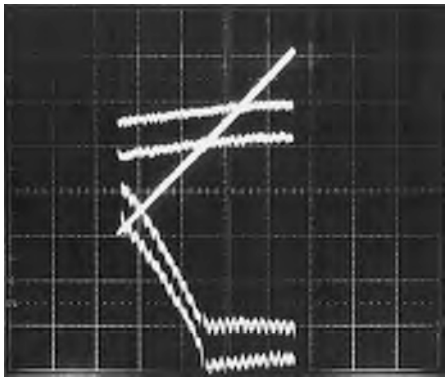

20194938

Test B01, LM358, F = 1.5 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 10µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 65µV p-p at 20µV/div.

FIGURE 35.

LM324s really are used for audio amplifiers and preamps. But who would use an amplifier with poor linearity like that for an audio amplifier? It's easy: the output of the amplifier gets a pre-load or pull-down resistor, such as 5 k from the output to the - supply, so the output voltage can swing up and down a couple volts, but the output current is only sourcing. This provides very adequate linearity for small signals. The LM324 or LM358 can only swing a couple volts at 10 kHz, but that is adequate for preamps.

Test B02 is an LF412, a dual version of the LF411. Despite strenuous efforts to make a good layout, its thermal errors are only a little better than average (about 1/4 ppm).

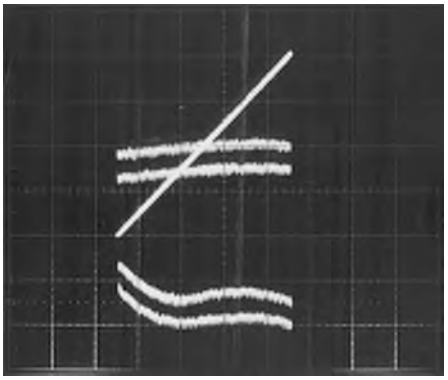

20194939

Test B02, LF412, F = 4 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 5µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 15µV p-p at 20µV/div.

FIGURE 36.

Test B03 is the LF442, a low-power version of the LF412. It was not rated to drive more than 2 mA, and driving 5 mA did cause poor gain, hundreds of microvolts of gain error, and not very linear. When driving light loads, less than 1 mA, the LF442 was a good general-purpose amplifier.

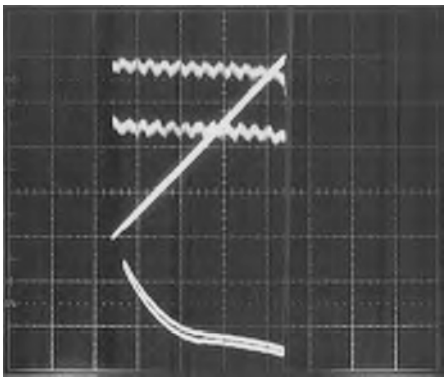

20194940

Test B03, LF442, F = 2 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 6µV p-p at 20µV/div.
Lower Trace: Gain Error, Full Load, 380µV p-p at 200µV/div.

FIGURE 37.

Test B04 is the LM833, an amplifier optimized for audio applications. It has reasonably good linearity, under rated conditions but is not able to drive more than the over-load of −8 mA without some distortion.

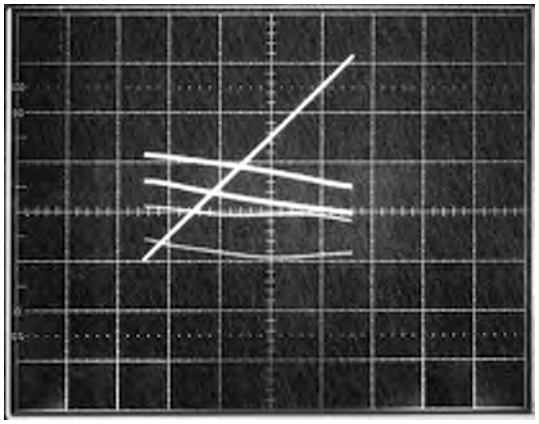

20194963

Test B04, LM833, F = 10 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±8 mA peak.
Upper Trace: Gain Error, No Load, 6μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 380μV p-p at 200μV/div. (Rl = 2k)

FIGURE 38.

Test B05 is another general-purpose amplifier, the LM1458, basically, a dual LM741. Its errors are only a little worse than typical.

Note that the humps are upside down, compared to most of the other amplifiers. This just means the heat-sensing inputs are arranged to detect the thermal gradients in the reverse sense.

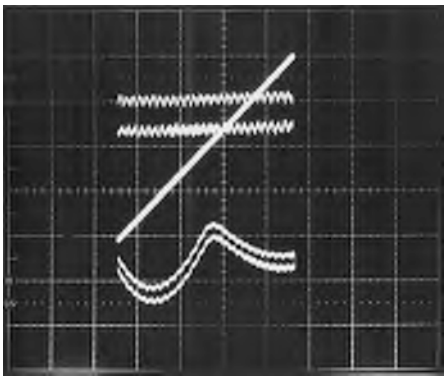

20194941

Test B05, LM1458, F = 1.1 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, + 3μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 70μV p-p at 50μV/div.

FIGURE 39.

Test B06 is an LM6182, and its gain errors are quite large - the voltage gain is just 2,500, and the gain error degrades 1 millivolt with the 1k load. Its gm is only 20 mhos. Who would be interested in an amplifier with such mediocre gain? It's not even as good gain as an old LM709!

The answer is, the LM6182 is quite fast. Its distortion at DC is not great, but the distortion holds low even up to 10 MHz (-50 dBc). So while we would not say it is a good general-purpose amplifier, it actually is a fairly popular amplifier for high-speed applications. This is one of the first current-mode amplifiers we have seen.

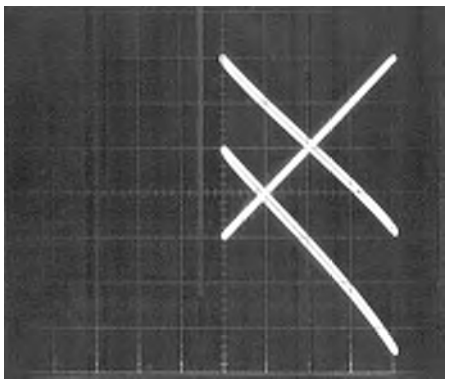

20194942

Test B06, LM6182, F = 500 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 7.8 mV p-p at 2 mV/div.
Lower Trace: Gain Error, Full Load, 8.8 mV p-p at 2 mV/div.

FIGURE 40.

Test B07 is the LM6142, a rail-to-rail amplifier. We don't expect its gain to not change with load - and its gain DOES change with load. But its voltage gain falls from just 3 million to 1/4 million. Its nonlinearity is still about average, with a 1k load. Note that its cross-over distortion is MUCH improved over the LM10 (Test A12). This amplifier, running on less than 0.7 mA per channel, has a 17 MHz gain-bandwidth product, much improved over the slow LM10.

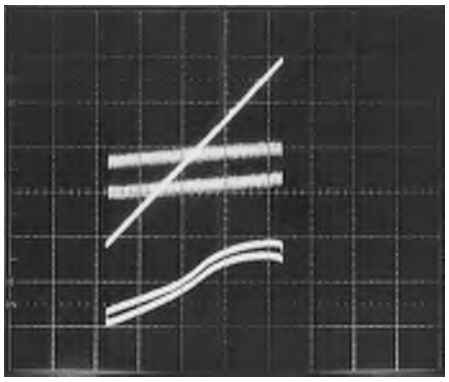

20194943

Test B07, LM6142, F = 20 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±5 mA peak.
Upper Trace: Gain Error, No Load, + 7μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 80μV p-p at 50μV/div. (RL = 2k)

FIGURE 41.

Test B08 is the LM6152, a faster 75 MHz amplifier, which also is a "rail-to-rail" Test. Its nonlinearity at 1k load (middle trace) is mediocre, but at its rated 2k load (lower trace) its linearity is well below 1 ppm.

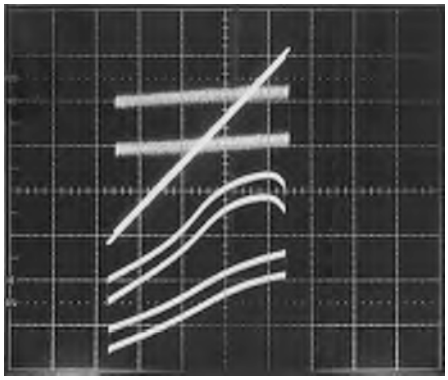

20194944

Test B08, LM6152, F = 100 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±5 mA peak.
Upper Trace: Gain Error, No Load, −7μV p-p at 20μV/div.
Middle Trace: Gain Error, ±10 mA Load, 120μV p-p at 50μV/div.
Lower Trace: Gain Error, ±5 mA Load, 76μV p-p at 50μV/div.

FIGURE 42.

Test B09 is the LM8262, another fast amplifier. Its gain is high at no-load, but the gain falls to 2700 at the 1k load. The crossover distortion is not very good, either. But it is fast. Also, it is tolerant of capacitive loads.

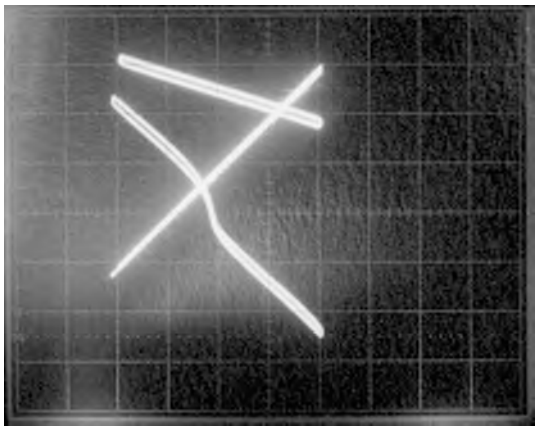

20194964

Test B09, LM8262 F = 200Hz.
V_S = ±11 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, 1.3μV p-p at 1mV/div.
Lower Trace: Gain Error, Full Load, 4.8mV p-p at 1mV/div.

FIGURE 43.

In Test B10, we have "saved the best for last". This precision amplifier, the LME49720, (also known as an LM4562) not only tests good, but it sounds good. The distortion is not only down somewhere below 0.15 ppm at 25 Hz, but it keeps improving at frequencies up to 1 kHz. It was designed as a precision audio amplifier, but is well suited for many other precision op-amp functions, with the best, lowest distortion in the industry. As you can plainly see, the thermals found in most other bipolar transistor op-amps have been banished by excellent layout. Distortion as low as -159 dB has been observed as an inverter, even driving a 2 kilohm load. For a study of how to test an op-amp with such low distortion at 1 kHz, refer to AN-1671.

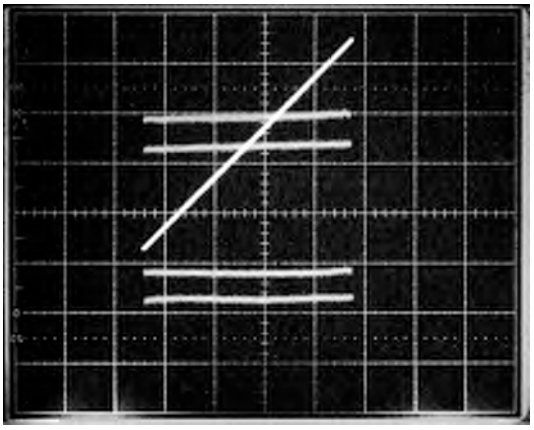

20194965

Test B10, LM4562 (also known as LME49720) F = 25 Hz.
V_S = ±15 Vdc; Vout = ±10 volts peak, Iout = ±10 mA peak.
Upper Trace: Gain Error, No Load, +1.5μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 1.5μV p-p at 10μV/div.

FIGURE 44.

Group C: Single CMOS Op-Amps

I did not include or test any of these; I tested the more popular dual amplifiers.

Group D: Dual CMOS Op-Amps

These are all rated to run on ±7.5 volts. I operated them on ±5.0 volts, and I required them to drive a 1 kilohm load to ±4 volts.

Test D01 is the basic old LMC662, a dual version of the LMC660, NSC's first CMOS amplifier. Its gain error looks quite non-linear; however, it is really not bad. The peak error is 27µV p-p, and the p-p nonlinear error is about 13µV p-p. If tested with a 4k load, it would have a nonlinearity of better than 1 ppm (as a unity-gain inverter, for example). The designer, Dennis Monticelli, pointed out that this amplifier design has 3 honest gain stages for sinking current (left side of the trace) but 4 stages of gain for sourcing current (right side of the trace). Since gain for sourcing current is usually considered more important, he let the design go as "plenty good enough". I tend to agree that a linearity of 1 ppm is "plenty good enough" for any general-purpose amplifier.

Here is an amplifier where the output impedance really is high. When the load is lightened from 1k to 2k, 4k, 8k, etc., the gain keeps going up. How high does it go? It's almost impossible to resolve how high the gain goes, or how high the output impedance is. The gain goes up by AT LEAST a factor of 30, and quite possibly 60 or more. So the output impedance goes up to at least 30k, and maybe 100 or 200k. Is the exact number important? Is it important if the gain goes up to 4 million, or 8 million? In theory, it is fun to imagine that a gain of 4 million is not quite as good as 8 million. Or that if the gain goes up to 8 million, then the low-frequency gain roll-off starts falling from the DC gain of 8 million at 0.1 Hz. But as you can see, these amplifiers are well-behaved, and the loop is obviously stable for all conditions. If you only looked at the left-hand side (Vout = negative) where the output is sinking load current, the gain may be finite, but this amplifier is very well behaved. Likewise on the right-hand side, it is a very high-gain amplifier -- and very well behaved. If the amplifier runs anywhere in the middle, or on either end -- the amplifier is STILL very well-behaved. It just has a small bit of nonlinearity. We don't usually think of 1 ppm as a significant amount of nonlinearity -- but that is the only thing wrong with this amplifier! We are discussing this at great length, primarily because it shows that very high gain, whether at no load or at heavy load, does not cause any problems. Also because several other CMOS amplifiers have very similar characteristics.

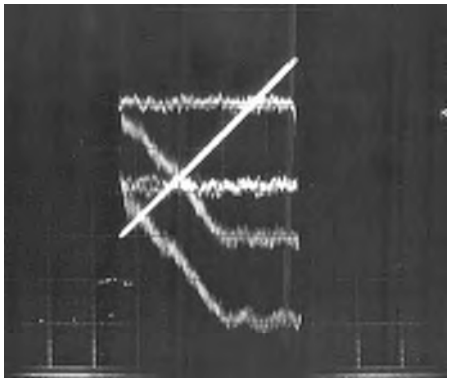

20194901

Test D01, LMC662, F = 6 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 1µV p-p at 10µV/div.
Lower Trace: Gain Error, Full Load, 27µV p-p at 10µV/div.

FIGURE 45.

The traces on D01B are for the same amplifier. In the topmost trace, a 500 pF filter capacitor is added across the 1 megohm gain-setting resistor in *Figure 7*, to cut the noise a little. The middle trace shows how noisy this set-up was, when I neglected to ground the operator's body while pushing the shutter button! The standard noise was on the lower trace. Note that even though these traces seem noisy, the noise is barely 3 or 4 µV p-p, and the gain errors as large as 1 or 2µV can be resolved, nicely.

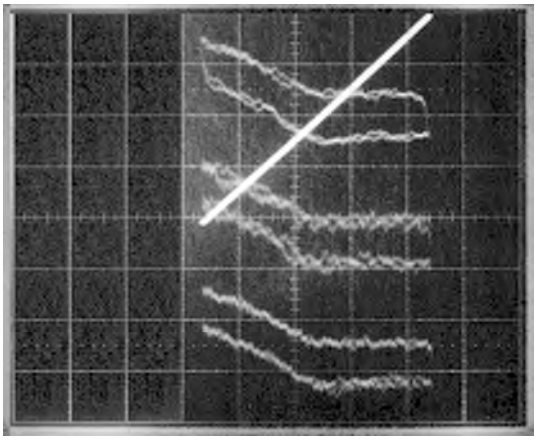

20194902

Test D01B, LMC662, F = 6 Hz.
Upper Trace: Gain Error, Full Load, 27µV p-p at 20µV/div., C= 500pF.
Middle Trace: Gain Error, Full Load, 27µV p-p at 20µV/div., with 60 Hz Ambient Noise..
Lower Trace: Gain Error, Full Load, 27µV p-p at 20µV/div., Normal Test.

FIGURE 46.

Test D02 is a LMC6492, a standard CMOS amplifier similar to the LMC6482, with rail-to-rail inputs and output, rated from -40 to + 125 degrees C.

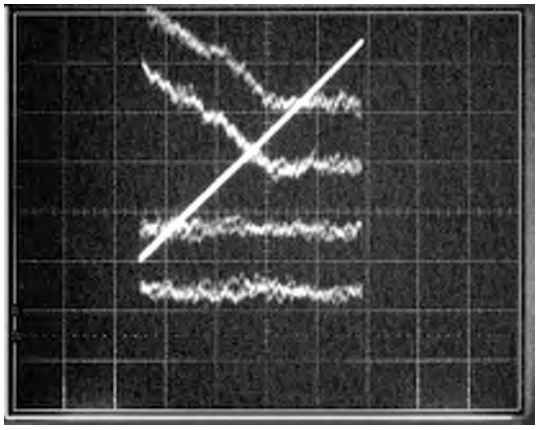

20194971

Test D02, LMC6492, F = 6 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 1μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 22μV p-p at 10μV/div.

FIGURE 47.

Test D03 is a standard CMOS amplifier, the LMC6482, similar to LMC6492, rated from -40 to +85 degrees C. Its gain curves are - typical.

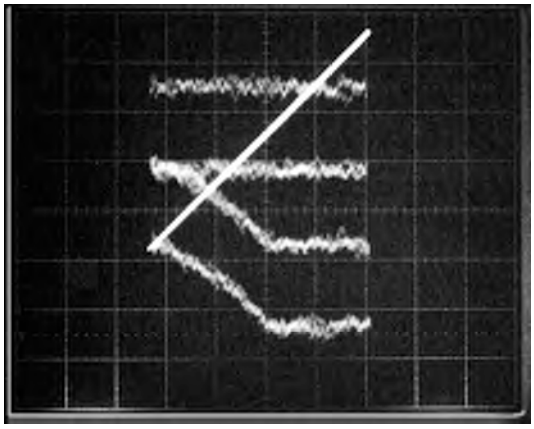

20194972

Test D03, LMC6482, F = 6 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 1μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 18μV p-p at 10μV/div.

FIGURE 48.

Test D05A is a micropower amplifier, the LMC6572, drawing just 40 μA of current. Even though it is running very lean, internally, it can drive a ± 4 mA load with a gain over 1 million, and a nonlinearity better than 0.2 ppm. It is characterized down to 2.7 volts of power supply.

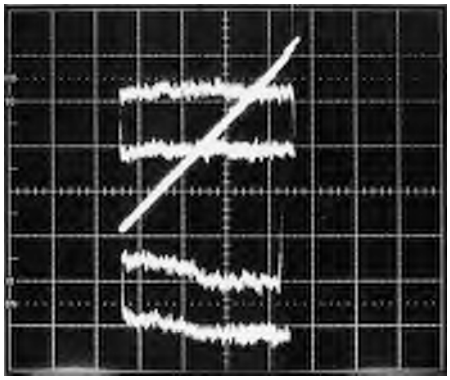

20194904

Test D05A, LMC6572, F = 0.8 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 2μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 5μV p-p at 10μV/div.

FIGURE 49.

Test D06A is an LMC6042, another micropower amplifier, running on just 10 μA. Its gain and linearity are about as good as the previous example, with a gain over 1 million and gain linearity below 0.2 ppm. It is only rated to run from + 15 volts down to + 5 volts of total power supply.
Test D06B is another LMC6042.

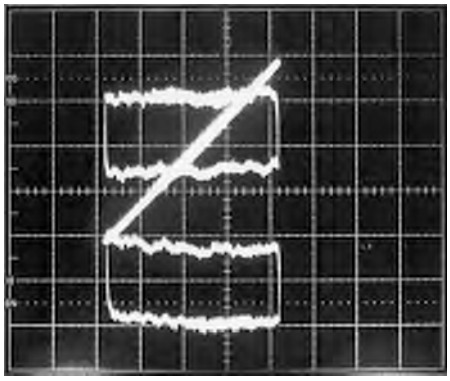

20194905

Test D06A, LMC6042, F = 0.6 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 2μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 6μV p-p at 20μV/div.

FIGURE 50.

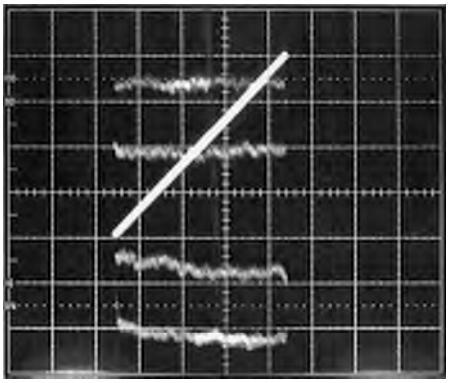

20194906

Test D06B, LMC6042, F = 0.6 Hz
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 3μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 6μV p-p at 20μV/div.

FIGURE 51.

Test D07B **is another low-power amplifier, requiring less than 100 μA per channel. Its nonlinearity is down below 0.3 ppm. As noted earlier, our testing with triangle waves help us resolve non-linearities below 1 ppm. If we were testing with sine waves, as in the lower trace, it would be hard to resolve these small sub-ppm errors.**

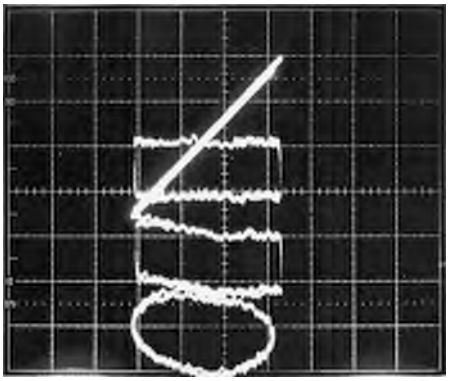

20194907

Test D07B, LMC6022, F = 2 Hz.
V_S = ±5 Vdc; Vout = ±4 volts peak, Iout = ±4 mA peak.
Upper Trace: Gain Error, No Load, 4μV p-p at 20μV/div.
Middle Trace: Gain Error, Full Load, 7μV p-p at 20μV/div. (TRIANGLE)
Lower Trace: Gain Error, Full Load, 7μV p-p at 20μV/div. (SINE)

FIGURE 52.

Test D08 **is the LMC6062, a precision amplifier with Vos as good as 350 μV, max. Its linearity is down near 0.2 ppm.**

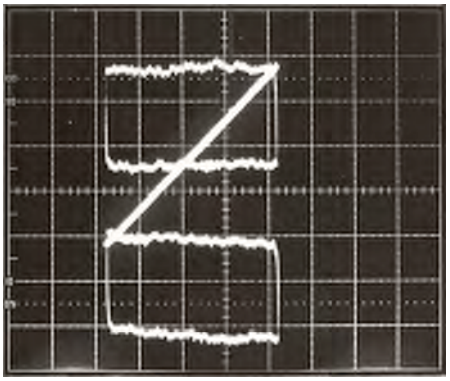

20194908

Test D08, LMC6062, F = 0.6 Hz.
V_S = ±5 Vdc; Vout = ±3.5, −4.5 volts peak, Iout = 8 mA p-p.
Upper Trace: Gain Error, No Load, 3μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 6μV p-p at 20μV/div.

FIGURE 53.

Group E: Low-Voltage Single Amplifiers (± 2.5-volt Supplies)

Test E01 is the LMV715, a low-voltage amplifier. Its linearity is as good as 1.5 ppm, 3μV p-p at the input compared to 4 volts p-p of output swing.Of course, at lighter loads, the linearity would improve.

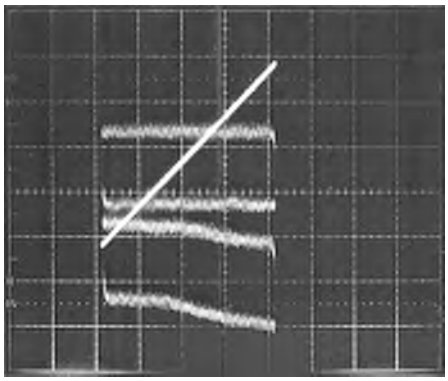

20194910

Test E01, LMV715, F = 26 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 2μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 12μV p-p at 20μV/div.

FIGURE 54.

Test E02 is an LMV751. This amplifier has poorer gain for positive swings (sourcing current). The no-load gain curve (lower trace) is obviously well under 1 μV p-p. The linearity with a 4k load would be slightly better than 1 ppm, even though the gain error looks pretty bad! The LMV751 has very low noise, about 6.5 nV per square-root Hertz.

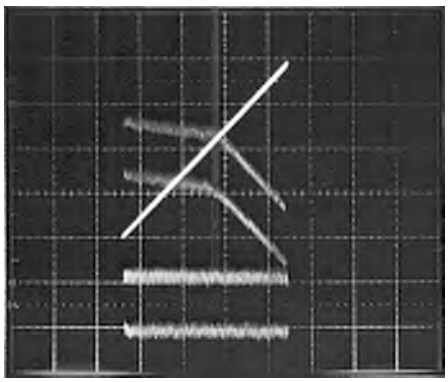

20194911

Test E02, LMV751, F = 12 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, Full Load, 11μV p-p at 5μV/div.
Lower Trace: Gain Error, No Load, 1μV p-p at 5μV/div.

FIGURE 55.

Test E03 is an LMV771, with mediocre gain in both directions! It looks awful - yet the nonlinearity with a 4k load would be still be better than 1/2 ppm.

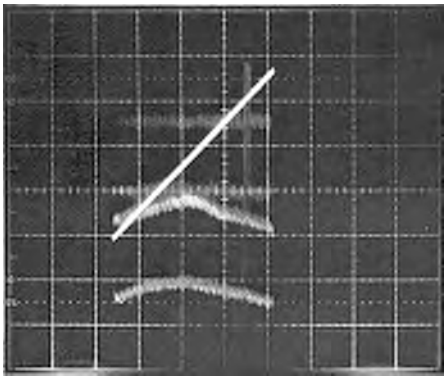

20194912

Test E03, LMV771, F = ~6 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 1μV p-p at 5μV/div.
Lower Trace: Gain Error, Full Load, 3μV p-p at 5μV/div.

FIGURE 56.

Test E04 is an LMV301 (bipolar, not CMOS) with very high gain and linearity better than 1/2 ppm.

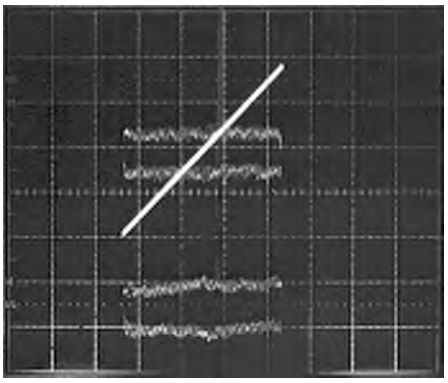

20194913

Test E04, LMV301, F = 12 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 2μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 4μV p-p at 20μV/div.

FIGURE 57.

Group F: Low-Voltage Duals (± 2.5-volt Supplies)

Test F01A is the LMP2012, a chopper-stabilized amplifier with gain well over 2 million. The linearity seems to be better than 1/4 ppm. The offset voltage is typically below 4μV.

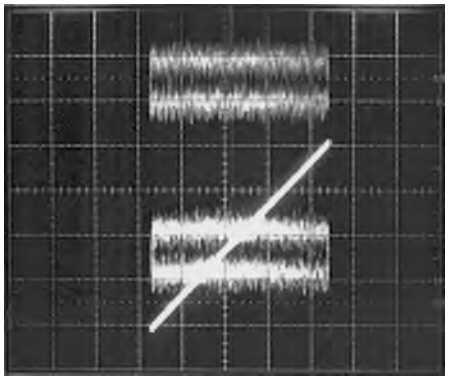

20194914

Test F01A, LMP2012, Side A, F = 2 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 1μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 2μV p-p at 10μV/div.

FIGURE 58.

Trace F01D is an LMP2012 with the 500 pF filter capacitor added, to help resolve the signals down in the noise; linearity is still below 1/4 ppm.

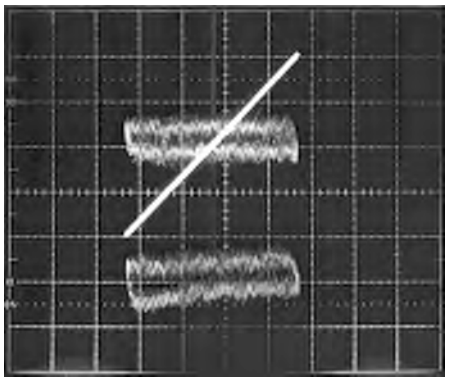

20194915

Test F01D, LMP2012, Side B, F = ~2 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 1μV p-p at 10μV/div.
Lower Trace: Gain Error, Full Load, 3μV p-p at 10μV/div.

FIGURE 59.

Test F02 is an LMV932, with 1/4 ppm, most of which is its cross-over distortion.

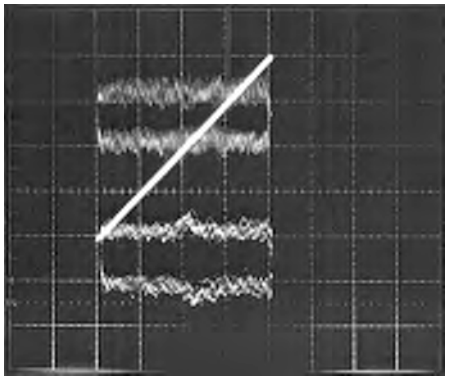

20194916

Test F02, LMV932, F = 12 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, No Load, 3μV p-p at 20μV/div.
Lower Trace: Gain Error, Full Load, 8μV p-p at 20μV/div.

FIGURE 60.

Test F03

The low-voltage LMV358 does not have the exact same shape of nonlinearity as the LM358 (see test B01, Figure 35) but a somewhat different shape. Its gain is OK, but its non-linearity when driving a 4 kilohm load is about 6 ppm. This is noticeably inferior to many other modern op-amps -- but yet, when do you measure an amplifier with linearity worse than 3 ppm, or complain about it? As with the LM358, the LMV358 can provide excellent linearity if the output has a pre-load (pull-down or pull-up resistor) connected.

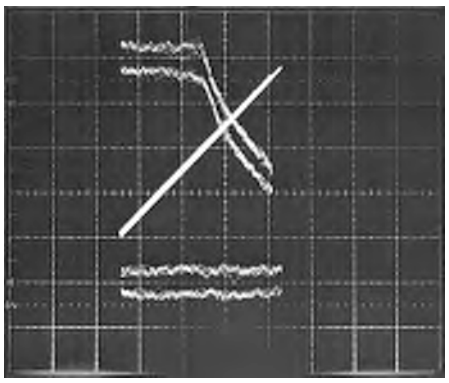

20194917

Test F03, LMV358, F = 20 Hz.
V_S = ±2.5 Vdc; Vout = ±2 volts peak, Iout = ±2 mA peak.
Upper Trace: Gain Error, Full Load, 150μV p-p at 50μV/div.
Lower Trace: Gain Error, No Load, 5μV p-p at 20μV/div.

FIGURE 61.

Test X06 is a very low-voltage amplifier, running on ±0.45 volts, with gain error below 7µV p-p, and linearity near 2 ppm.

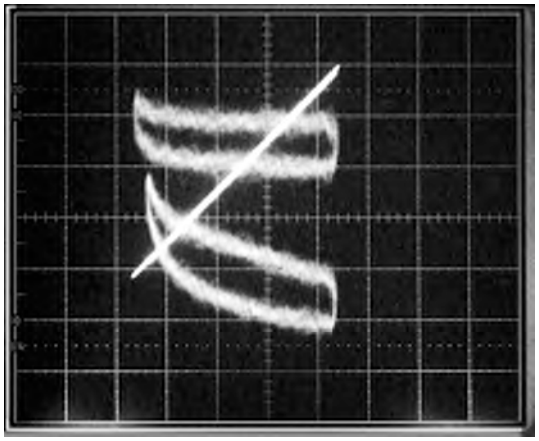

20194973

Test X06, LMV751 F = 75 Hz.
V_S = ±0.5 Vdc; Vout = ±0.4 volts peak, Iout = ±0.4 mA Peak.
Upper Trace: Gain Error, No Load, 7µV p-p at 10µV/div.
Lower Trace: Gain Error, Full Load, 20µV p-p at 10µV/div.

FIGURE 62.

Conclusions

There are many interesting things to learn about an amplifier's gain, not just one number on a datasheet. Not all amplifiers are the same - or even SIMILAR!! Amplifiers with output followers are not simple to analyze, when thermal errors can cause bigger errors than the gain error. CMOS amplifiers with high output impedance, would seem to have a major source of error at heavy loads, but in actuality, good amplifiers can drive loads with accuracy and linearity much better than 1 ppm. A high output impedance can allow the gain to go extremely high at light loads, and this may be useful in precision applications.

Design Engineers have many things to think about. The gain for positive outputs versus negative outputs may be important for precision amplifiers. Thermal problems may also have to be studied, in areas where computers are not helpful.

Mask Designers have to be concerned with precise placement of critical components. They have to make sure they are given complete instructions on placement and matching.

Applications Engineers have to measure and characterize the new amplifiers, to make sure the characteristics are as good as expected. The data sheet may need to be revised, to show good or bad features of an amplifier's gain.

The Customer does not have to worry so much about the internal design of the amplifier, but he/she may have to be concerned, for critical applications, about some of these features of amplifiers.

Philosophical Insights

Many engineers have opinions or preconceptions that operational amplifiers made with bipolar transistors have better, higher voltage gain than CMOS amplifiers. Many people have a sense that bipolar op-amps are more linear than CMOS amplifiers. We have showed that this is not exactly true. There are many amplifiers of each Type that are very good -- with linearity better than 0.3 parts per million. Some are barely as good as 2 parts per million -- but at light loads, they can be used with excellent accuracy and linearity. And of course, many applications do not require linearity better than 1 ppm!

Amplifiers are not simple. Silicon is not simple. Understanding circuits is not simple, but it is possible.

Appendix A: List of Amplifiers with Low and Lower Distortion

The testing of amplifiers in this Applications Note was done on amplifiers that were mostly rated with a 2 kilohm load. I ran most of the tests with a heavy load of 1 kilohm, to make sure I had enough nonlinearity to see a signal.

For this Appendix, the engineering was done for a 10k/10k unity-gain inverter, with a 6.67k load, making a virtual 4 kilohm total load, so the nonlinearity would be done with a moderate load (half the current of the rated 2k load, not double the current). The nonlinearity was sort of interpolated as 1/4 of the nonlinearity with a 1k load. As you will see, many of the amplifiers have surprisingly good linearity, even though the curves with RL = 1k looked pretty bad. They are listed in order of improving linearity. All data are approximate, and typical. No data are guaranteed. Availability of old amplifier types denoted by * is not guaranteed, and are very unlikely.

Example: An LM709, per the data shown on Test A01, has a 100µV p-p nonlinear error at its summing point, driving a 1 k load. That is the total p-p deviation from the best-fit straight line. When it is driving 4k of total load, the error would be 25µV p-p, referred to input. The 709's error will be decreased quite strictly by this factor of 4, because it is a thermal error, which heats the input transistors in a highly predictable way.

A unity gain inverter runs at a Noise Gain of 2, so its output would have 50µV p-p. Its output swing is 20 volts p-p. Therefore we will call the distortion, 2.5 ppm, as it is 2.5 ppm of the total output swing. All other amplifiers get the same conversion done for them. It is true that SOME amplifiers will not improve by this transformation, by the exact factor of 4, but it is still approximately correct. The computations were done in terms of p-p errors, as RMS computations would probably not be applicable for such nonlinear signals. If you wanted an LM709 to have better linearity than 2.5 ppm, you could run it with a lighter load, or, choose a better amplifier. Or get a helper amplifier to put out most of the load current.

Amplifiers with Bipolar Transistors and with ± 10-Volt output swing (supplies = ± 15 volts)

Type	Test	Nonlinearity
LM8262	(B10B)	12 ppm
LF442	(B03)	8 ppm (light load)
LM6182	(B06)	6 ppm
LM709*	(A01)	2.5 ppm
LM318	(A06)	2.1 ppm
LM741	(A03)	2 ppm
LM301A	(A02)	1.5 ppm
LM10	(A12)	1.5 ppm
LF411	(A08)	1.4 ppm
LM308	(A05)	1.3 ppm
LM1458	(B05)	1.3 ppm
LM307J	(A14)	1.3 ppm
LF356	(A09)	1.2 ppm
LM358N	(B01)	1.0 ppm
LM6142	(B07)	0.5 ppm
OP-07*	(A07)	0.4 ppm
LM833N	(B04)	0.4 ppm
LF412N	(B02)	0.3 ppm
LM6152	(B08)	0.3 ppm
LM607*	(A10)	0.12 ppm
LM725*	(A04)	0.10 ppm
LM627*	(A11)	0.04 ppm
LM4562	(B10)	0.025 ppm

* Amplifiers denoted by * are obsolete and are no longer available from NSC.

CMOS AMPLIFIERS with ~ Rail-to-Rail outputs and with ± 4-Volt output swing (supplies = ± 5 volts)

Type	Test	Nonlinearity
LMC662	(D01)	1.4 ppm
LMC6482	(D03)	1.1 ppm
LMC6492	(D02)	1.1 ppm
LMC6022	(D07B)	0.3 ppm
LMC6042	(D06A)	0.3 ppm
LMC6062	(D08)	0.2 ppm
LMC6572	(D05A)	0.2 ppm

Low Voltage Amplifiers with ~ Rail-to-Rail outputs and with ± 2-Volt output swing (supplies = ± 2.5 volts)

Type	Test	Nonlinearity
LMV358	(F03)	6 ppm
LMV715	(E01)	1 ppm
LMV771	(E04)	0.6 ppm
LMV751	(E02)	0.4 ppm
LMV771	(E03)	0.4 ppm
LMV932	(F02)	0.3 ppm
LMP2012	(F01A)	0.2 ppm

Very Low Voltage Amplifier with ~ Rail-to-Rail outputs and with ± 0.4Volt output swing (supplies = ± 0.5 volts)

Type	Test	Nonlinearity
LMV751	(X06)	5 ppm

Footnotes

1. Ideal amplifiers are characterized in T. Frederiksen's book, "Intuitive IC Opamps", NSC 1984, p. 23.

2. Some wise engineers have pointed out that even a "rail-to-rail" output stage can not literally swing all the way to the rail, even driving as light a load as a megohm, or even 10 megohms. There are practical reasons why an amplifier can not drive a 1 or 10 μA load much closer than 10 or 20 mV to either power supply rail: if they tried to run with such a starved bias, the output loops would go out of control. For loads as heavy as 100 μA, 20 to 50 mV is a practical overhead or "drop-out" limitation. For 1 or 2 mA, the drop-out is in the vicinity of 100 to 200 mV. For typical real data, refer to the specific amplifier's data sheet. The typical curves of "Output Characteristics, Current Sourcing" and "Output Characteristics, Current Sinking" will show what you can expect to get, for this dropout. It may not be terribly small, but at moderate loads, it is a lot better than the 600 or 700 mv of the best amplifiers with emitter followers.

3. "What's All This Common-Centroid Stuff, Anyhow?" http://www.elecdesign.com/Articles/Index.cfm?ArticleID=6121 R. A. Pease, Electronic Design, October 1, 1996.

4. "What's All This Output Impedance Stuff, Anyhow?" R. A. Pease, Electronic Design.

5. Appendix A., List of Operational amplifiers with low distortion at dc and low frequencies

Notes

For more National Semiconductor product information and proven design tools, visit the following Web sites at:

Products		Design Support	
Amplifiers	www.national.com/amplifiers	WEBENCH	www.national.com/webench
Audio	www.national.com/audio	Analog University	www.national.com/AU
Clock Conditioners	www.national.com/timing	App Notes	www.national.com/appnotes
Data Converters	www.national.com/adc	Distributors	www.national.com/contacts
Displays	www.national.com/displays	Green Compliance	www.national.com/quality/green
Ethernet	www.national.com/ethernet	Packaging	www.national.com/packaging
Interface	www.national.com/interface	Quality and Reliability	www.national.com/quality
LVDS	www.national.com/lvds	Reference Designs	www.national.com/refdesigns
Power Management	www.national.com/power	Feedback	www.national.com/feedback
Switching Regulators	www.national.com/switchers		
LDOs	www.national.com/ldo		
LED Lighting	www.national.com/led		
PowerWise	www.national.com/powerwise		
Serial Digital Interface (SDI)	www.national.com/sdi		
Temperature Sensors	www.national.com/tempsensors		
Wireless (PLL/VCO)	www.national.com/wireless		

National Semiconductor Americas Technical Support Center
Email: support@ nsc.com
Tel: 1-800-272-9959

National Semiconductor Europe Technical Support Center
Email: europe.support@ nsc.com
German Tel: +49 (0) 180 5010 771
English Tel: +44 (0) 870 850 4288

National Semiconductor Asia Pacific Technical Support Center
Email: ap.support@ nsc.com

National Semiconductor Japan Technical Support Center
Email: jpn.feedback@ nsc.com

A Comprehensive Study of the Howland Current Pump

National Semiconductor
Application Note 1515
Robert A. Pease
January 29, 2008

A Comprehensive Study

It is well known to analog experts that you can use the positive and negative inputs of an operational amplifier to make a high-impedance current source (current pump) using a conventional operational amplifier (op amp). This basic circuit can put out both + and - output current (or zero current) into various loads. The theory is simple. But the practical problems involved are not so simple or obvious.

There are two basic circuits -- the Basic Howland Current Pump, *Figure 1*, and the "Improved" Howland Current Pump. The Basic circuit does good service for simple applications, but if its weaknesses are unacceptable, the "Improved" circuit may do much better for critical tasks. See *Figure 5*.

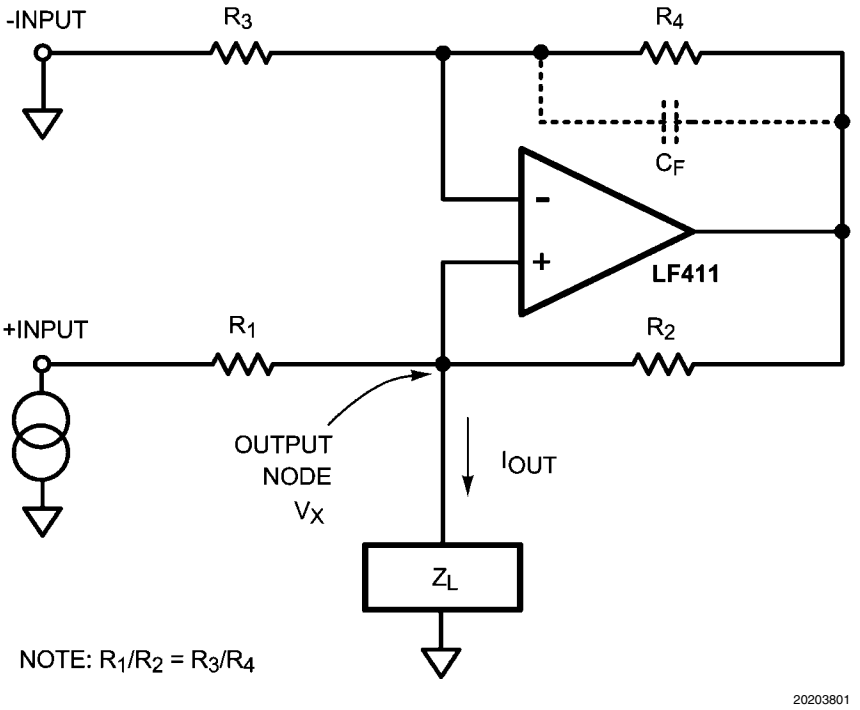

FIGURE 1. The Basic Howland Current Pump

Applications for the Howland Current Pump

Sometimes a unidirectional current source (or sink) is just right. It is easy to make them with high output impedance and wide range, using an op-amp and some Darlington-connected transistors. But sometimes you need a current pump that can put out a current in either direction – or even AC currents.

The Howland current pump is usually excellent for that. Current sources are often used for testing other devices. They can be used to force currents into sensors or other materials. They can be used in experiments, or in production test. They can bias up diodes or transistors, or set test conditions. When you need them, they are useful — even if you only need them once or twice a year.

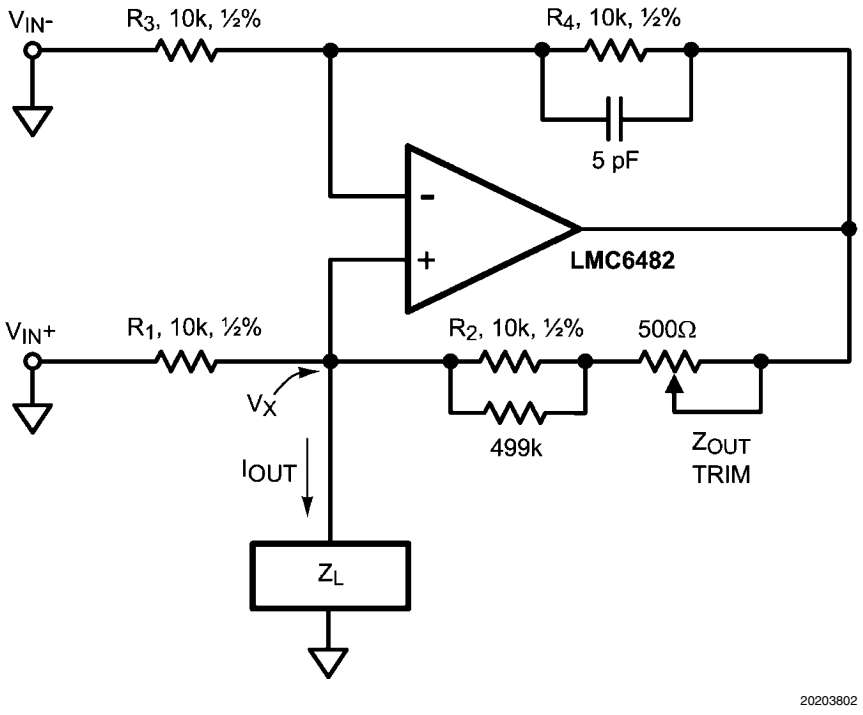

FIGURE 2. Basic Howland Current Pump with Trim of Z_{OUT}

The "Basic Howland Current Pump" was invented by Prof. Bradford Howland of MIT, about 1962, and the invention was disclosed to his colleague George A. Philbrick (the analog computer pioneer who was head of Philbrick Researches, Boston MA at that time). This circuit was not patented. The Howland Current Pump was first published in the January 1964 "Lightning Empiricist", Volume 12, Number 1. It is Figure 5A on page 7 of an article by D. H. Sheingold, "Impedance & Admittance Transformations using Operational Amplifiers". This can be found at http://www.philbrickarchive.org/1964-1_v12_no1_the_lightning_empiricist.htm. It was also included in the Philbrick Researches Applications Manual, in 1965. Its elegance arises because the feedback from the output to both the + and - inputs is at equal strength -- the ratios of R1/R2 and R3/R4 are the same. While it is possible to analyze this circuit mathematically, it is easiest to just analyze it by inspection:

If the "output" node Vx -- which is the + input of the op amp -- is grounded, it is easy to see that the "gain" is 1/R1, that is, the output current per change of the input voltage is equal to 1/R1. So you don't need a fancy set of equations for that. The resistors R2, R3, and R4 have no effect when the output is grounded, and only the + input voltage is active.

When you move the - input upward, the gain to the grounded output node is – R4/R3 x 1/R2. Since the ratio of the resistors is defined to be R1/R2 = R3/R4, then that gain is also equal to - 1/R1. That is easy to remember! Note that the gain is reversed for the - input.

Thus it is easy to see that if both Vin+ and Vin- are moved together, then there is no change of Iout. When Vin + rises, the "gain" to the output node is "1/R1". Then it follows that the gain for the - input is also "1/R1" but with a negative sign. So this current pump can accept positive or negative inputs. It has true differential inputs. Now all that we need to show, is that the output impedance is high, so that the gain is correct for all output voltages and impedances, and for all inputs.

It is easy to see that the output impedance is very high, using this analysis: If both signal inputs are grounded, and if the "output" node Vx is lifted up, somebody has to drive the resistance "R1". But as the op-amp's + input is lifted, the - input must also rise up, and the output also rises, providing just enough current through R2 to cancel the current flowing through R1, thus making the output impedance very high indeed. The principle of linear superposition says that no matter what is Vin+ or Vin-, and no matter what is Zload, and no matter what is Vout (within the limitations that you shouldn't ask the op-amp output to put out more than it can do, in voltage or current), the Iout will be (Vin+ - Vin-) x 1/R1. If you like to see a lot of fancy equations, see at Appendix A.

Most applications notes just indicate the circuit and the ratio, that R1/R2 must be equal to R3/R4. However they do not indicate how important it is to have precise matched or trimmed resistors. If all 4 resistors were 10 k ohms with a 1% tolerance, the worst-case output impedance might be as bad as 250 k ohms -- and it might be plus 250k, or it might be minus 250k! For some applications, this might be acceptable, but for full precision, you might want to use precision resistors such as 0.1% or even 0.01%. These are not inexpensive! But it may be preferable to use precision resistors rather than to use a trim pot, which has to be trimmed (and which may get mistrimmed).

Note that if you use adjacent resistors from a tape of 1% resistors, the odds are that they will match better than 1/2%. But that is not guaranteed!

Figure 2, *Figure 3*, and *Figure 4* show ways to use a trim-pot to make the output impedance very high. Typically, using 1/2% resistors and one trim pot, you can trim the output impedance to be 5 ppm of I (full scale) per volt.

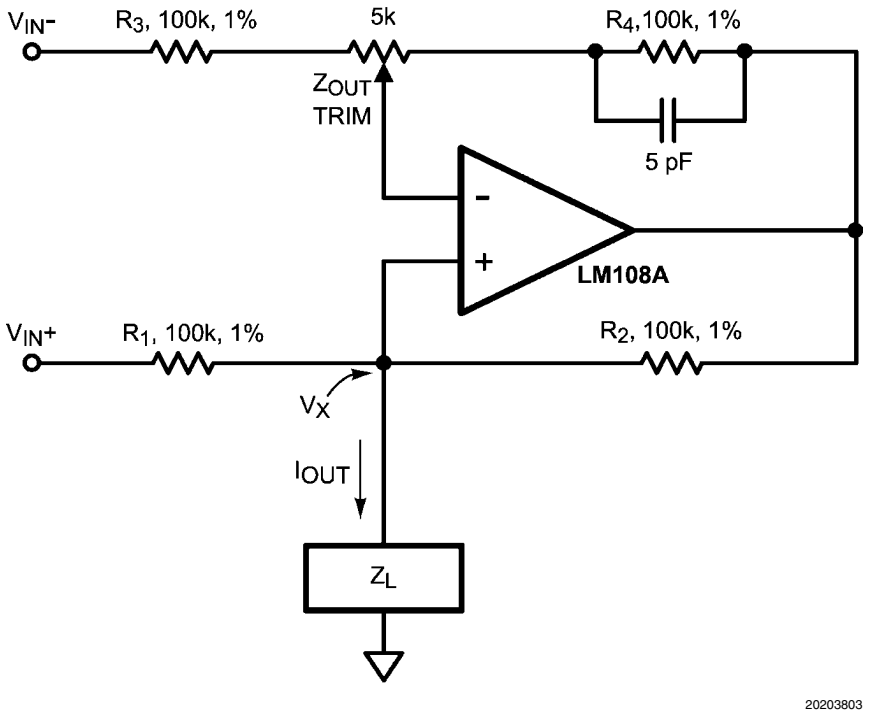

FIGURE 3. Basic Howland Current Pump with Trim of Z_{OUT}

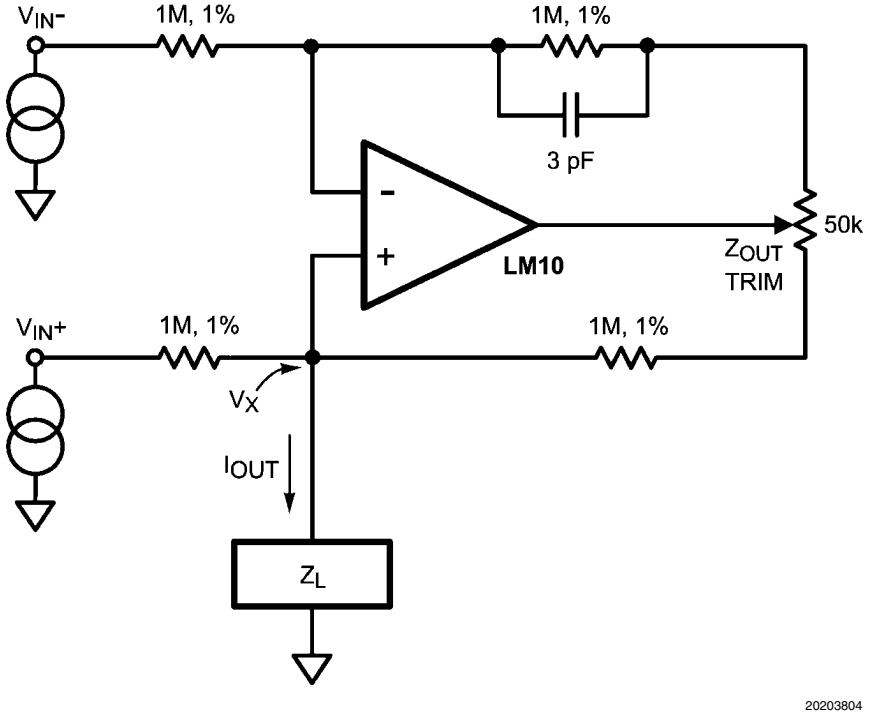

FIGURE 4. Basic Howland Current Pump with Trim of Z_{OUT}

However the resistor tolerance is not the only thing that needs to be trimmed out. The CMRR of the amplifier needs to be accommodated. Fortunately, an amplifier CMRR of 60 dB would cause the output impedance to degrade only to 10 megohms, not even as bad as 0.1% resistors would cause, in the example above. And many amplifiers have CMRR better than 80 dB. However, the CMRR of an op-amp is not always linear -- it may be curvy or it may be otherwise non-

linear. Some amplifiers that have the advantage of rail-to-rail inputs may have a nonlinear Vos which may jump a millivolt or more as the CM signal gets within a couple volts of the + rail. Amplifiers with bipolar inputs often do have this kind of nonlinearity. Amplifiers such as LM6142 and LM6152 have nonlinearities of this type. See Appendix B. Some CMOS amplifiers such as the LMC6482, LMC6462, etc. (See Appendix C) have a fairly linear curve of Vos, with no jumps, due to proprietary input process and circuit design.

One of the weaknesses of the Basic Howland Current Pump is its output capability. Its output node does not normally swing very close to the rail. For example, the basic 10k/10k/10k/10k scheme can only swing its output node to + or - 5 or 6 volts , with ± 15-volt supplies. If the output node rises a lot, the op-amp's output would have to rise about twice as high. When that is no longer possible, the "Improved Howland" should be considered.

If you kept the gain resistor R1 as 10k, and change R2 and R4 to 1k, you could make a 10k/1k, 10k/1k circuit, that would let the output node rise to 10 volts with a good amplifier. However, this is a little less accurate, with more offset and noise.

Another weakness of the Basic circuit is the inefficiency. If you want to have a gain of (1/100 ohms), with R1 at 100 ohms, the amplifier has to put out a lot of drive, if the load voltage swings a lot. If the load is a low voltage, such as a diode, that may not be so bad. If the load only rises a half volt, only a few mA will be wasted. But if it had to rise 5 volts, that is a lot of power wasted!

If you had to drive a heavy load, you do not have to have equal resistances at R1 and R3. You could have 100 ohms, 100 ohms for R1/R2, and 10k/10k for R3/R4. Then if you want to drive the - input, the input impedance will not be very heavy. However, when you have this imbalance, you must be careful that the amplifier's Ib does not cause a big error, which may be significant if a bipolar input op amp is used.

To avoid these weaknesses of the Basic Howland, the "Improved" Howland generally does solve many of these problems, very well. See *Figure 5*.

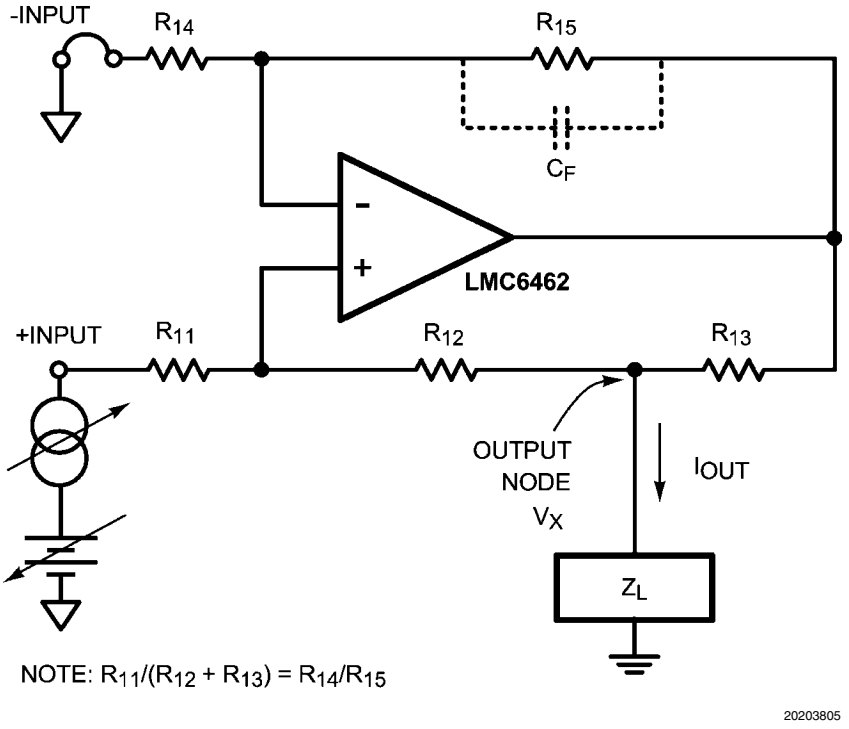

FIGURE 5. The "Improved Howland" Current Pump

The "Improved Howland" Current Pump

In this case, you still have to trim the R's to get good CMRR and high output impedance. But the gain is set by R13, modified by the ratio of R14/R15 (which is typically 1/1). Consequently you can use low values for R13, and keep all the other resistors high in value, such as 100k or 1 Megohm.

In the "Improved" Howland, note that it is not just the ratio of R11/R12 that must match R14/R15; it is the ratio R11 / (R12 + R13) that must be equal to R14/R15. If you do the intuitive analysis as mentioned above, you can see that if R14 = R15, R12 will normally be (R11 - R13). Conversely, you could make R11 a little higher, to get the gain to balance out. You could put a 2k pot in series with R11. This "improved" circuit can now force many milliamperes (or as low as microamperes, if you want) into voltages as large as 10 volts, with good efficiency. See *Figure 6*.

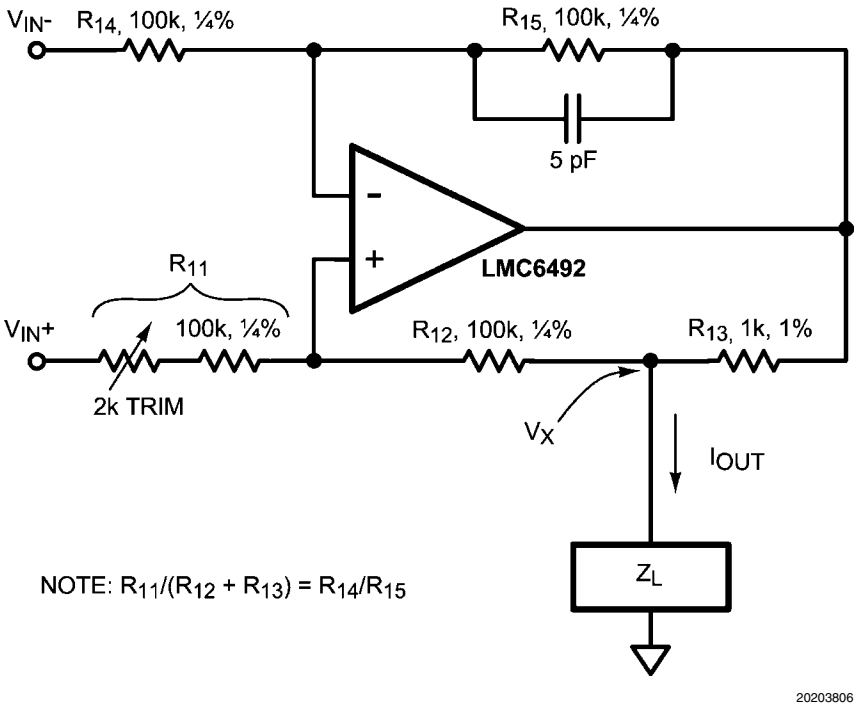

FIGURE 6. The "Improved Howland" with Trim for Z_{OUT}

Dynamics

Most engineers know (if you remind them) that it's a good idea to add a feedback capacitor across the feedback resistor of an inverting amplifier. The Howland Current Pump does like a little bit of feedback capacitor there, across R4 (or R15). A small feedback cap of 3 to 5 to 10 pF is almost always a good idea. If you are putting in a really slow current, and if the rate of change of the output voltage at the output node is not high, you could make the Cf equal to 100 or 1000 pf, to cut down the bandwidth and the noise.

Most engineers have not analyzed the dynamics of this circuit. According to the detailed analysis (at Appendix D) the "output capacitance", as seen at the output node can be as large as 80 pF, for an ordinary 1 MHz op amp. However, there are many fast amplifiers available these days, so it is usually easy to select one with a lot more bandwidth than that, if you need it. But you have to remember to design for that.

The equation for the output capacitance of the Improved Howland is derived in the latter part of Appendix E. This may be slightly better than for the standard Howland.

Choice of Amplifiers

Almost any op-amp can be used in a Howland current pump. However, if you need a wide output voltage range, a high-voltage amplifier, running on ±15 volts (or more) may be needed. Conversely, if you only need a small Vout range, a low-voltage CMOS amplifier may work just fine. As with any amplifier application, choosing the amplifier may take some engineering, to choose the right type. For high impedance applications (resistors higher than 0.1 Megohm), FET inputs may be a good choice. If you have one left-over section of LM324, it can even do an adequate job, for resistors below 100k. A list of amplifiers with Bipolar inputs (and generally wider signal ranges) is found in Appendix B. A list of CMOS amplifiers with very high Zin (but smaller output range) is in Appendix C.

If you needed a ±6 volt output swing, you might not need a ±15-volt op-amp. The "Improved Howland" may be able to swing that far, using a good CMOS op-amp running on ± 7 volts, such as LMC6482. The "Improved" Howland is much more effective in terms of output swing.

Current-feedback amplifiers are not normally good choices for the Howland current pump, as they mostly work best at low impedances. And their CMRR is rarely as good as 58 dB. But if you need a blindingly fast current pump at fairly low impedance levels (100 ohms to 2k), current-feedback amplifiers can do a good job. Be aware that they may need a good bit of trimming to counteract their poor CMRR.

Special Applications

The Howland Integrator

One of the obscure applications for the Howland Current pump is the "Howland Integrator", shown in *Figure 7*. This is sometimes called a "DeBoo Integrator". If a capacitor is used as the load, the Amplifier's Vout can be easily seen to be: Vout = 2 x 1/RC x $\int$ Vin/dt . (This assumes that R1/R2 = R3/R4 = 1.) Of course, this integrator has to have some means to reset it, just as every integrator does. However, it is fairly easy to reset the integrator with a single FET or switch to ground, as shown. The Howland Integrator only works using its + input, as the - input must be grounded. It has a positive gain, as opposed to the conventional inverting integrator.

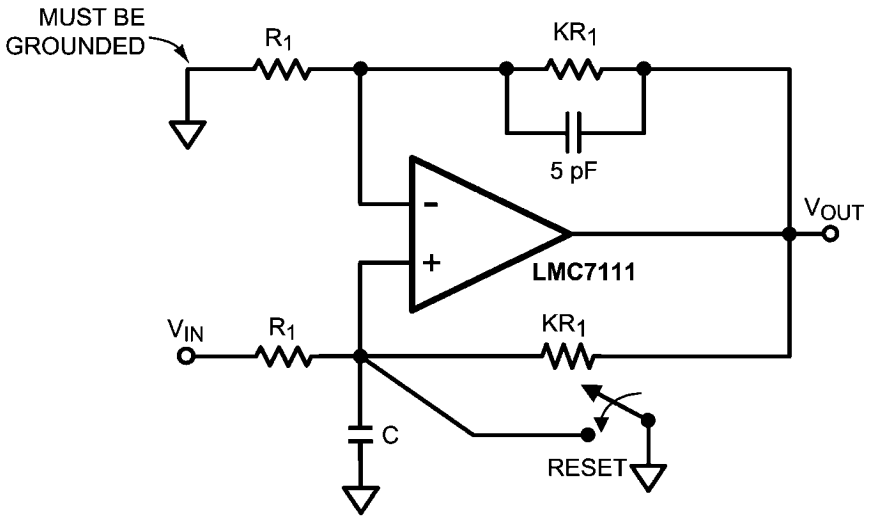

$$V_{OUT} = (1 + K) \frac{1}{(R_1 C)} \int V_{IN}\, dt$$

$$= \left(\frac{2}{R_1 C} \right) \int V_{IN}\, dt, \text{ if } K = 1$$

20203807

FIGURE 7. The Howland Integrator

This is much easier to use than the rarely-used Positive Integrator (see *Figure 8*), which would need TWO FET switches to reset it. That integrator is rarely used, for obvious reasons, but it can be used in loops which inherently provide some feedback to bring the output back to a low level, and to keep it zeroed. For example: a servo integrator, that will pull an error back to zero, can work well. This "positive integrator" actually is a differential integrator, with positive and negative signal input gains. It can be used with either input active, or BOTH.

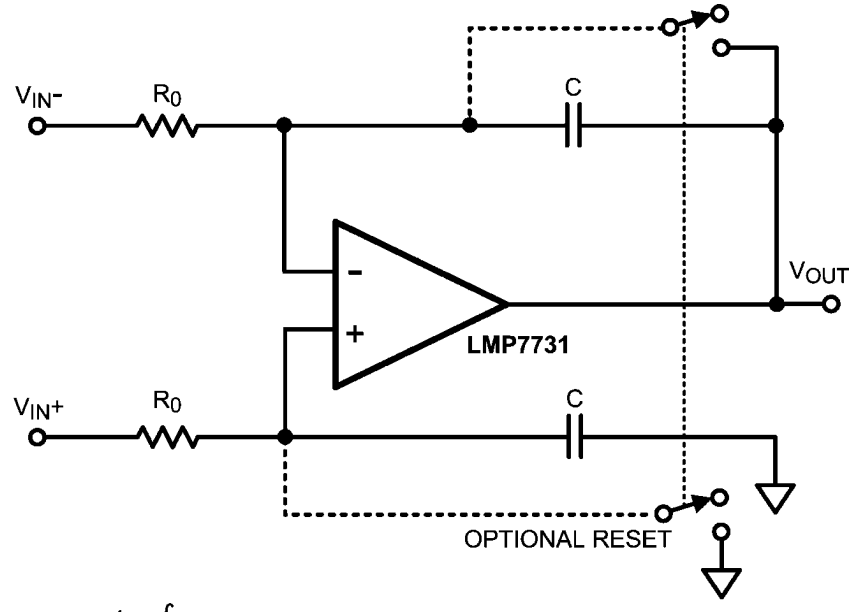

$$V_{OUT} = \frac{1}{R_0 C} \int (+V_{IN}+ - V_{IN}-)\, dt$$

20203808

FIGURE 8. The Positive or Differential Integrator (not a Howland Circuit)

Multi-Range Current Pump

If you want to use a Current Pump with various different ranges (such as connecting in various Gain resistors of 1 ohm, 10 ohms, ... , 1k, 10k, ..., 1 M, etc., etc....) it is possible to add one precision op amp to allow you to change ranges without affecting the other resistors. A precision FET-input op amp with good CMRR can be used as the unity-gain follower, as shown in *Figure 9*. The resistors should be trimmed to take into account, (and trim out the effect of) the CMRR of BOTH operational amplifiers. A typical trim scheme is shown.

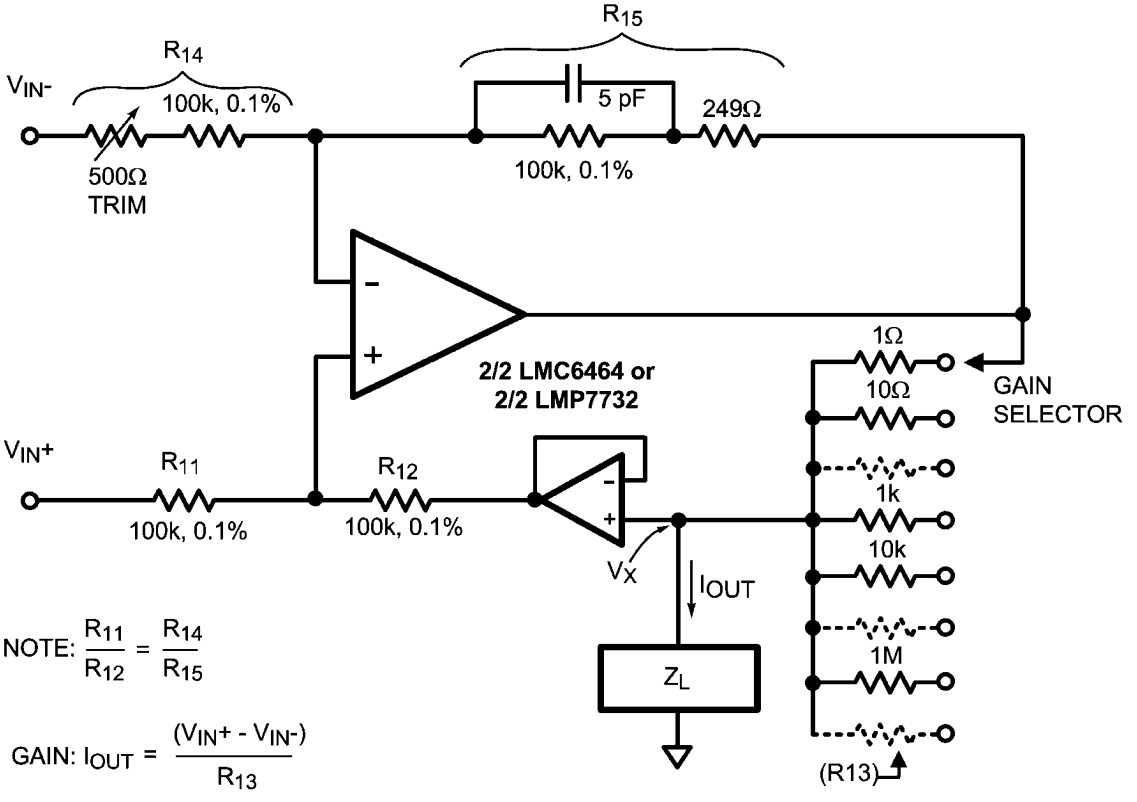

FIGURE 9. Multi-Range Current Pump

Appendix A. Output Impedance as a Function of Trimming

If a basic Howland Current pump has 4 equal resistors, it is easy to see that the tolerance of any one resistor can cause the output impedance of the circuit to be as bad as (R1 x 1/ tolerance). If the tolerance is 1% and you are using 10k resistors, the output impedance could be as high as 1 megohm. In fact, the output impedance could be as high as " plus 1 megohm" or it could be a negative output impedance of MINUS 1 megohm.

Most engineers are not very familiar with the concept of "negative resistance". In a case like this, if the Howland Current Pump is connected to a small capacitor, a conventional "positive" output resistance would mean the output would gradually droop down toward ground (assuming small Vos of the op amp). If the output impedance were very high, the Vout would just stay constant. If the output impedance were negative, the output voltage would rise up faster and faster! Or it would descend NEGATIVE, faster and faster!

If R1 is too low, the output impedance will be "positive". If R4 is also too low, that helps make the output impedance lower, and positive. If R2 and R3 are too high, they all add up to the worst case, where Zout would be 1/4 megohm.

If R1 and R4 were HIGH in tolerance, and R2 and R3 were too LOW, the output impedance could be as poor as NEGATIVE 250k. Both of these cases would be usually unacceptable, for any precision application, because the circuit is normally capable of being 200x higher than that, in its output impedance.

Some of the possible solutions to this tolerance problem are

1. Buy resistors with tighter tolerances - or
2. Sort and match them by pairs. This is cheap and simple if you only need a small number of well-matched resistors, and you don't want to go out and buy special parts. Just match the R's you have.
3. Add a trim-pot in series with one of the resistors, and add a compensating resistor of about half that size, on the opposite side, to let you trim the ratio up or down a little. Of course, if any pot can be adjusted, it can also be misadjusted.... Per *Applications for the Howland Current Pump*, *Figure 3* or *Figure 4*.
4. Buy matched sets of tightly-matched resistors. These can be purchased in sets of 4 for a couple dollars, in an SOIC package, such as four 10k resistors matched to 0.1% or even to 0.01% of ratio. See *Figure 10*. Resistors such as Caddock T914's (in SIP packages) can be found at: http://www.caddock.com/Online_catalog/Mrktg_Lit/TypeT912_T914.pdf These are available with matching

down to 0.01%. Resistors such as Vishay Beyschlag ACAS 0612's are at www.vishay.com The catalog lists these as good as 0.1% matching, in Surface Mount packages. Bourns has thin-film networks in DIPS (4100T Series) and in surface mount packages (4400T Series) at: http://www.bourns.com. Their tolerance is better than 0.1% but the matching specs are not listed.

5. Use the techniques of NSC's Linear Brief LB-46 to make snip-trimmed resistors, to avoid the problems with pots. See *Figure 11*.

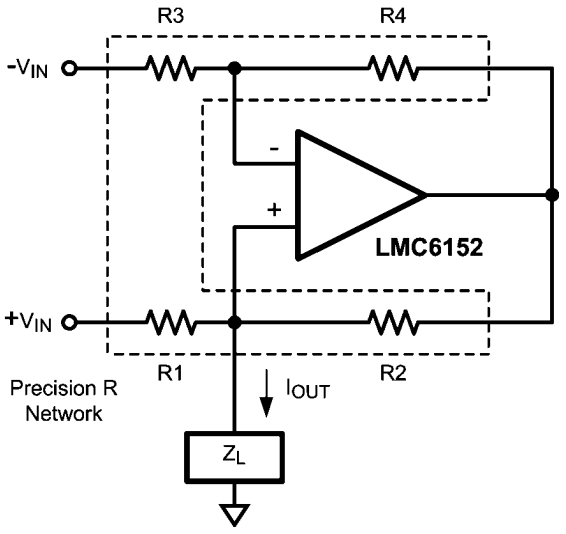

FIGURE 10. Basic Howland Current Pump with Precision Resistor Network

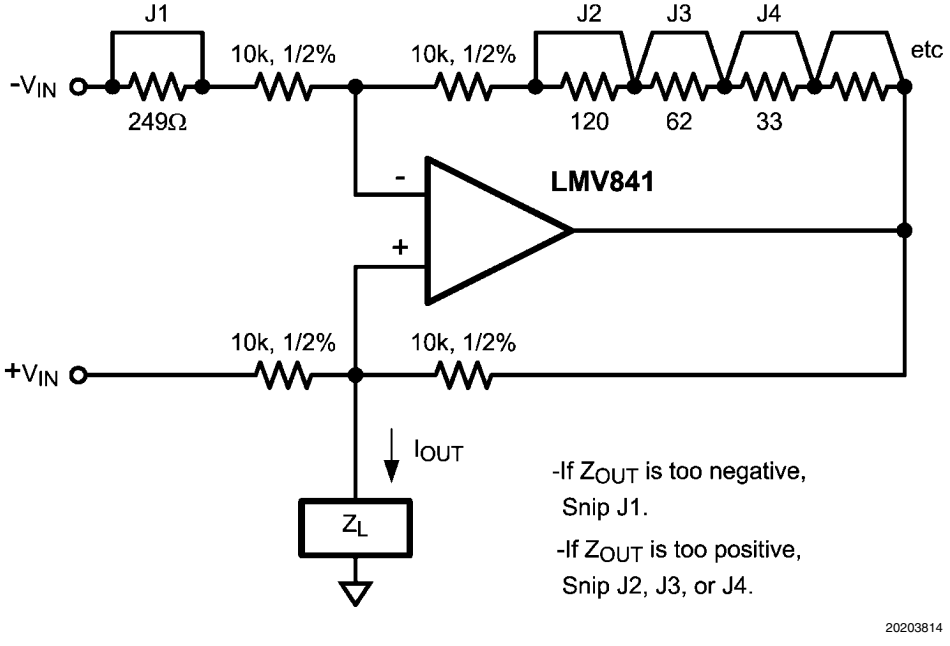

FIGURE 11. Trimming Z_{OUT} without a Trim-Pot

Appendix B

List of NSC amplifiers with rail-to-rail inputs using bipolar transistors (NPN and PNP):

- LMV931, LMV932 (dual), LMV934 (quad) - down to 1.8-volt supplies.
- LMV7301 (2.2 to 30 V supplies)
- LM6132 (dual), LM6134 (quad) - 10 MHz per 0.36 mA
- LM6142 (dual), LM6144 (quad) - 17 MHz
- LM6152 (dual), LM6154 (quad) - 75 MHz.
- LMV981 (single), LMV982 (dual)
- LMH6645, 6646, 6647.
- LM8261, LM8262 (dual)
- LMP7731 (single).
- LMP7732 (dual)
- and many more.

Note, Rail-to-Rail Input Common-Mode Range is not usually needed for a Howland Current Pump, but may be advantageous for the "Improved Howland".

To find the full range of these amplifiers with Rail-to-Rail common-mode range, go to NSC's home page at http://www.national.com and find the Selection Guide at: "Select" "Amplifiers" "Rail-to-Rail Input" and select "Maximum input current" which leads to a selection chart and look for bias currents larger than 10 nA.

Appendix C

List of NSC amplifiers with rail-to-rail inputs using CMOS inputs:

- LMC6492 (dual), LMC6494 (quad) - standard
- LMC6482 (dual), LMC6484 (quad) - precision
- LMC6462 (dual), LMC6464 (quad) - precision low power
- LMC7111 - single, low power
- LMC8101 - 2.7 V
- LMP710 - 5MHz
- LMV7701
- LMV712 (dual) - with shutdown
- LMV841
- LMP7704

and several more...

Note, Rail-to-Rail Input Common-Mode Range is not usually needed for a Howland Current Pump, but may be advantageous for the "Improved Howland".

To find the full range of these amplifiers with Rail-to-Rail common-mode range, go to NSC's home page at http://www.national.com and find the Selection Guide at: "Select" "Amplifiers" "Rail-to-Rail Input" and select "Maximum input current" which leads to a selection chart, and look for bias currents smaller than 10 nA.

Appendix D

Output Capacitance of the Basic Howland Current Pump (*Figure 1*).

An operational amplifier can move its output quickly, only if there is a significant Vin transient or Error Voltage (Vε) applied across the inputs. Typically, this can be millivolts, or dozens of millivolts, according to the need for fast output speed.

The rate of change of Vout, dVout/dt, is equal to: dVout/dt = - (2πfh) x Vin where (fh) is the Gain Bandwidth Product. Thus, Vε = dVout/dt x 1/(2 π fh)

The best way to analyze this current pump is to apply a long, slow ramp at Vx. The dVx/dt is the same as the dV(sum)/dt, and dVout/dt = (R4 + R3)/R3 x dVx/dt.

To move dVout/dt at this rate, there must be an error voltage at the - input (V(sum)), even when Vx = 0 and moving. When we apply a long ramp from a negative voltage and Vx just passes 0 volts, with a rate of change equal to dVx/dt, this error voltage Vε will be generated:

When the momentary voltage at the applied voltage Vx is: Vx = 0 + dVx/dt, V(sum) = dVx/dt - Vε,

Then Vout has a rate of change, (dVx/dt) (R4 + R3)/R3 and a momentary offset of:

- Vε x (R4 + R3)/R3 = -1/(2 π fh) x dVx/dt x ((R4 + R3)/R3)^2.

The current i1 through R1 = 0, since Vx = 0. The current through R2 is:

i2 = Vout/R2 = - dVx/dt x 1/(2 π fh) x 1/R2 x ((R4 + R3)/R3)^2.

This current acts as a capacitive current, as it is a direct function of dVx/dt. This virtual capacitance is:

C = 1/(2 π fh) x 1/R2 x ((R4 + R3)/R3)^2.

For a 1 MHz op amp, in a typical application when R1 = R2 = R3 = R4 = 10k, this capacitance will be 200 pf / π, or about 64 pF, (in addition to the actual capacitance at the + input of the operational amplifier). This capacitance is inversely proportional to the gain bandwidth product fh. It gets smaller as a faster amplifier is employed. This capacitance will also be inversely proportional to R2, so for 1 kilohm, it would be 636 pF. For 100 kilohms, it would be just 6.4 pF. This capacitance will also be modified by the ratio of (R4 + R3)/R3, if that is not 2. The apparent capacitance may be different if a large feedback capacitance is connected, across R4. A small feedback capacitance of 3 to 10 pF across R4 is normally a good idea, even if fast signals are not contemplated.

With modern fast op-amps, this capacitance may or may not be a significant factor, but it should be taken into account, depending on the application.

Appendix E. Output Capacitance of the "Improved" Howland Current Pump

The analysis of this circuit is similar to the analysis of the basic Howland Current Pump of *Figure 1*.

An operational amplifier can move its output quickly, only if there is a significant Vin transient or Error Voltage (Vε) applied across the inputs. Typically, this can be millivolts, or dozens of millivolts, according to the need for fast output speed.

The rate of change of Vout, dVout/dt, is equal to: dVout/dt = - (2 π fh) x Ve where (fh) is the Gain Bandwidth Product. Thus Vε = dVout/dt x 1/(2 π fh)

The best way to analyze this is to apply a long, slow ramp at Vx. To move dVout/dt at a quick rate, there must be an error voltage at the - input (V(sum)), even when Vx = 0 and moving. When we apply a long ramp, starting from a negative voltage in a positive direction, and when it exactly passes 0 volts, with Vx = dVx/dt, this error voltage Vε will be generated:

When the momentary voltage at the applied voltage Vx is: Vx = 0 + dVx/dt,

The rate of change at V+ is R11/(R11 + R12) x dVx/dt, and the rate of change at Vsum is the same dV/dt, plus an offset: Vsum = R11/(R11 + R12) x dVx/dt + Vε

The rate of change of Vout is the rate of change at V(sum), magnified by (R14 + R15)/R14: (R14 + R15)/R14 x dV+/dt = R11/(R11 + R12)x (R14 + R15)/(R14) x dVx/dt - Vε x (R14 + R15)/R14:

Since dVout/dt = 2 (π) fh x Vε, then Ve = dVx/dt x 1/(2 π fh) x R11/(R11 + R12) (R14 + R15)/R14 ,

and the momentary value of Vout = (R14 + R15) / R14 x Vε =

= -1/(2πfh) x dVx/dt x R11/(R11+R12) x[(R14+R15)/R14]^2 .

When the momentary voltage at the applied voltage Vx is: Vx = 0 + dVx/dt,

the current i12 through R12 = 0. The current through R13 is: i13 = Vout/R13 = - 1/(2 π fh) x dVx/dt x R11/(R11 + R12) x [(R14 + R15)/(R14)]^2 x (1/R13).

This current acts as a capacitive current, as it is related only to dVx/dt. This current is equivalent to Cx x dVx/dt This virtual capacitance is: Cx = 1/(2 π fh) x (1/R13) x R11/(R11 + R12) x [(R14 + R15)/R14]^2.

For a 1 MHz op amp, in a typical application when R11 = 11k, and R12 = R14 = R15 = 10k, and R13 is relatively small compared to R12, such as R13 = R12/10 = 1k, this capacitance will be about 334 pF.

If R11 = 110k and R12 = R14 = R15 = 100k, and R13 = R12/10 = 10k, the Cx will be 34 pF. The capacitance depends inversely on the value of R13, and is also inversely proportional to the op amp's gain bandwidth product fh. It gets smaller as the amplifier is faster. This capacitance will also be inversely proportional to R13, so for 1 kilohm, it would be 334 pF. For 100 kilohms, it would be just 3.4 pF. This capacitance will also be modified by the ratio of (R15 + R14)/R14, if that ratio is not 2. The capacitance does not depend on R11, R12, R14, or

R15, but only on their ratios. The apparent capacitance may be different if a large feedback capacitance is connected, across R15. A small feedback capacitance of 3 to 10 pF is normally a good idea, even if fast signals are not contemplated.

With modern fast op-amps, this capacitance may or may not be a significant factor, but it should be taken into account, depending on the application.

Appendix F. Testing and Trimming of Z_{OUT}

Techniques for the testing of output impedance of a current pump are not well known, nor published, so they will be presented here for the first time. In concept, you could make dc measurements of output current at different levels of Vout, but that is not effective. Visual testing on an oscilloscope is better.

In concept, you can apply a sine wave through a transformer to the Vout node of a current pump. Read the current that flows out of the bottom of the transformer, at the input of an oscilloscope, which is very close to ground. See *Figure 12*. This does seem to work, but to be sure it is working well, you have to be sure that the transformer does not have any leakage or capacitive strays that can corrupt the reading. That is hard to prove, so it is not a recommended technique. A floating (battery-powered) sine or triangle-wave generator could be used, but that is not necessary.

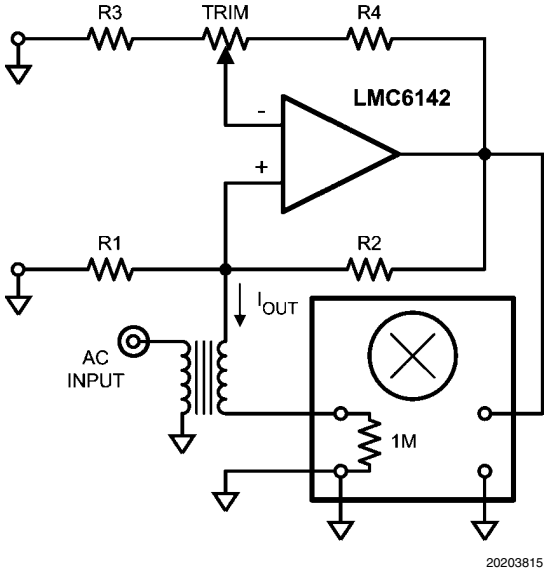

FIGURE 12. Transformer-Coupled Test for Z_{OUT}

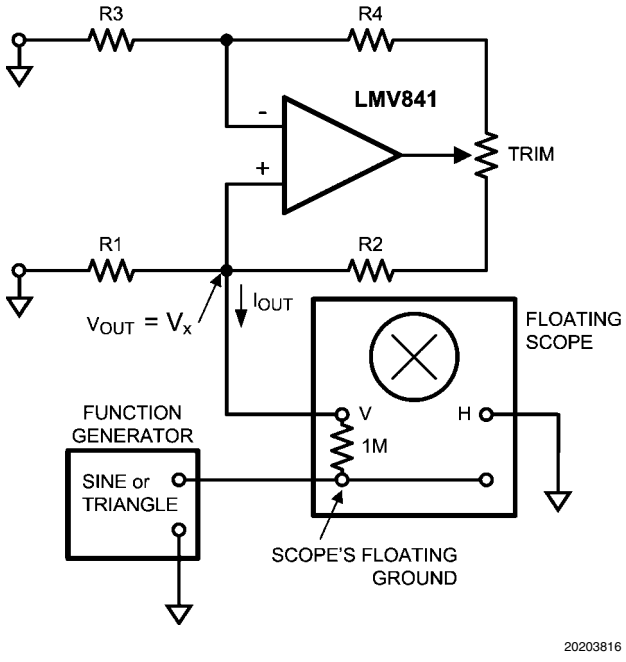

FIGURE 13. Good Test for Z_{OUT} Using Floating Scope

Refer to *Figure 13*. An ordinary sine or triangle wave function generator can be used to drive the Vout node. A floating scope is inserted between the function generator and the Vout. Any current flowing into or out of Vout, can be seen on the scope's vertical display. When this is cross-plotted against the main Vout, by connecting ground to the horizontal (X) input, it is easy to see the slope of Iout versus Vout. Then it is easy to adjust the trim pot as in *Figure 2*, *Figure 3*, or *Figure 4*, to trim for very high Zout. It is easy to get the curve just about flat, for the range of interest to you. Of course, if you want a very wide voltage range, you may find some curvature or nonlinearity. Most op-amps are not perfect for linearity of Common Mode Range, but some are better than others.

A scope with its 1 megohm of Zin can be used to easily resolve an output impedance better than 100 or 1000 megohms. If you have very low resistor values such as 1k or 100 ohms in your current pump, you can put a comparable low-value resistor across the scope's input. If you needed even higher resolution, a floating preamp could be added.

R. A. Pease. October 2007

Gain and Linearity Testing for Precision Operational Amplifiers

National Semiconductor
Application Note 1671
Robert A. Pease
May 21, 2008

Introduction

Many operational amplifiers are very linear — as a unity-gain inverter, they often have linearity in the range of 5 to 1 to 0.5 ppm. However, testing for gain linearity is not a standard or well-known test. Test procedures will be shown here, to permit any user to resolve the non-linearity errors.

Why Linearity?

It is possible to use op-amps for precision applications, even if they are not more linear than 5 ppm. Most people cannot hear an amplifier's distortion below 1 ppm, even if you cascade 10 of such stages. However, for precision instrumentation, or for high-fidelity audio, it is usually a good idea to select an op-amp with good linearity. Refer to AN-1485 which shows examples of good, moderate, and poor linearity. It shows Bipolar and also CMOS amplifiers. Contrary to many opinions, CMOS amplifiers with Drain-Loaded outputs do not necessarily have inferior linearity versus old bipolar-transistor amplifiers.

Linearity is often specified in dB below the signal level, (such as -110dB), or in percent (such as 0.0003%). However, most engineers are not familiar with nonlinearity below -120dB, so in this Linear Brief we will usually express the linearity in parts per million, (3 ppm or 0.3 ppm) compared to the output signal.

Linearity is often measured at low frequencies, because the testing is easier there. Linearity can be measured for an inverter or for a follower. In concept, one could subtract the input from the output, and these deviations could be plotted as a function of V_{OUT}. However, it is nearly impossible to do this with full resolution. So a real-time subtraction technique is employed. Test results can be seen in just a few seconds.

Conversely, a Distortion Analyzer such as an Audio Precision System Two Cascade SYS-2522 could be used to provide a direct reading of the distortion. However, even this procedure has limitations, when the distortion gets down below 3 ppm.

Test Fixturing and Test Procedure

The circuit of *Figure 1* 1 is a classic test fixture that has been used for over 40 years. Rin and Rf are used to force the input voltage (typically 20 volts p-p) to cause an output voltage of 20 volts p-p, as observed on an oscilloscope. The signal source should have an offset adjust control, which the user adjusts so the output is actually swinging ±10 volts with no offset. The Vos of the op-amp is thus effectively trimmed out. Meanwhile the R3 and R4 are used to magnify the error voltage (Vε) by a factor of ~1000, the ratio of R3 to R4. So if the amplifier's gain is 1 million, and its error voltage is 20µVp-p, the voltage at the oscilloscope's vertical input will be 20mVp-p. A good op-amp with a gain of 1 million can easily fulfill both gain requirements. This arrangement will avoid the need for a low-noise preamp, which is sure to add some extra noise. See *Figure 1*.

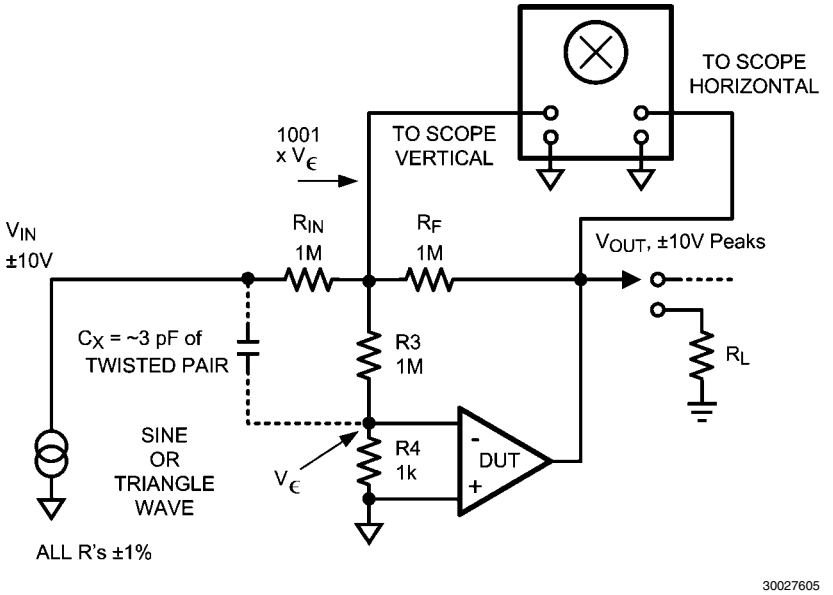

FIGURE 1. Gain Test
where A_V = 1001 ($V_{HORIZ.\ P\text{-}P}$ / $V_{VERT.\ P\text{-}P}$)

Crossplot Display

In concept you could look at the voltage "1000 Vε" with any kind of meter or scope. But to see the linear and non-linear components, it is best to use a scope in X-Y mode (Crossplot). Then any bumps or curves will correlate with the related level on the output swing. The slopes and bumps and curves can be resolved despite noise that is bigger than these signals! The eye can correlate out errors like noise, and offset, and linear errors, to let you see the non-linear errors.

It is also advantageous for many tests, to use a triangle wave input, rather than a sine wave. For a more complete discussion of why triangle waves are superior, see AN-1485. A frequency around 3 to 30Hz is usually suitable for a fast amplifier. See *Figure 2*.

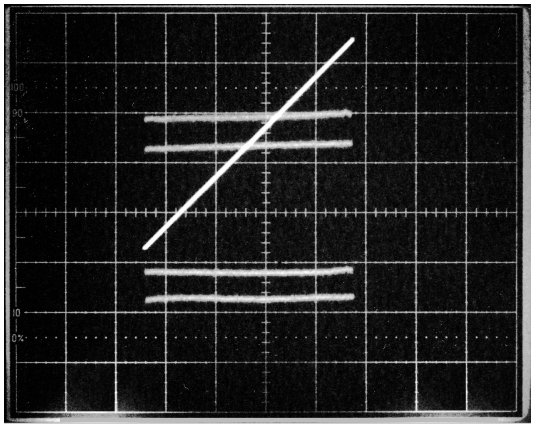

30027607

FIGURE 2. Test B10, LM4562 (also known as LME49720), F = 25Hz
V_S = ±15V_{DC}; V_{OUT} = ±10 volts peak, I_{OUT} = ±10mA peak
Upper Trace: Gain Error, No Load, +1.5μV_{P-P} at 10μV/div
Lower Trace: Gain Error, Full Load, 1.5μV_{P-P} at 10μV/div

Audio Distortion

The LM4562 is a good fast 54MHz amplifier, (GBW Product). But at 1kHz with 20 volts p-p of output signal, the summing-point error (Vε) will be as big as ±118μV. The signal at the vertical scope input will be ±118mV peak, or 236mV p-p. When you turn down the scope's sensitivity to keep this signal on-scale, you cannot resolve 1μV any more - you can barely see 10μV. So how can we see 1μV of error at 1kHz?

The solution is to connect a small amount of twisted pair, to be used as an adjustable capacitor, Cx, as shown in *Figure 1*. Use Teflon-insulated wire. When you wind (or un-wind) the twisted pair -- to cancel out most of the dynamic errors of the op-amp -- the signal becomes just a few mVp-p, and you can resolve the equivalent 1 μV of nonlinearity, even in the presence of some noise. The amount of non-linearity on this amplifier seems to be down near 0.08 ppm, referred to input. That is about -159dB, referred to the 7-volt rms signal.

This circuit, using 1M,1M, 1k, 1M is a little noisy - which is not a surprise, as the noise of a 1k resistor is larger than the noise of the LM4562. To get a better resolution of the distortion, you can use new, lower R values in place of the 1M/1M resistors. I just went in and slapped in a 20k across the 1 meg's, and 20 ohms across the 1k. I chose this scale factor because it was easy for me to add a150pF variable air capacitor across the "3pF" capacitor. This worked very well. I could see some quadrature error, even with a triangle wave, and with a sine wave, I could see a small amount of distortion down below 0.08 ppm This seems to be second harmonic, at 2kHz, at about -159db down. However, we have not yet given up on the possibility of further improvements.

Distortion as a Non-Inverting Amplifier (Unity-Gain Follower)

Some of the early tests on the LM4562 showed it could get down to 0.3 ppm of nonlinearity. This was true, because at voltages larger than 3 volts (rms) of input swing, non-linearity started to be degraded, due to imperfect (nonlinear) CMRR. This was when applied as a unity gain follower as in *Figure 3*. If you are using an amplifier as a unity-gain follower, the LM4562 does have some nonlinearity, but this can be improved. Write to request advice. But the circuit of *Figure 3* does provide so much gain (100) that you can see the signal, and the noise and the distortion. This gain is the ratio of (R5 + R6 + R7)/R6.

Analysis of Results

Figure 2 shows that the summing-point error has about 1/2 μVp-p of nonlinearity with a 1 Megohm load (upper trace.) With a 1k load (lower trace), it is closer to 1 μVp-p deviation from the best straight line. (1 mVp-p at the scope vertical input). The gain is therefore up near 10 million to 40 million, and the non-linearity is down near 1/20 part per million, also. This is one of the better amplifiers we can test, the LM4562. Other amplifiers with comparable non-linearity are the LME49710 and LME49740. If amplifier's linearity is any better than 1/20 ppm, the nonlinearity will be below the noise, and may require extreme techniques, such as adding a Spectrum Analyzer or Distortion Analyzer after the circuit of *Figure 1*.

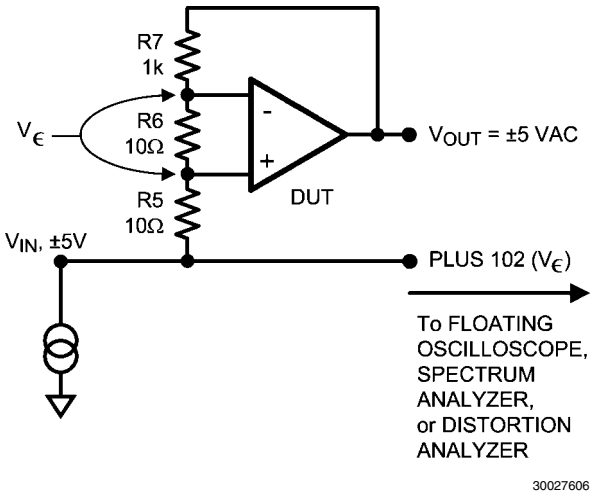

FIGURE 3. Test Circuit for Gain and Distortion of Amplifier at G = +1 (Including the nonlinearity of the CMRR)

Conclusion

The LM4562 is one of the best linear amplifiers in the world. It has some of the lowest distortion of any op-amp or any linear amplifier in the world, better than 0.2 PPM from 0.01 to 2kHz. This measurement technique is running out of capabilities to see much lower than 1/10 ppm. But it can be used with a Spectrum Analyzer to resolve a little lower. Unlike other amplifiers with good numbers and good plots, but which may not sound so great, the LM4562 also is said to "sound good", according to people who have good audio acuity (who have "good ears"). This is probably related to its good dynamic stability and clean slew rate. The electronic testing is consistent with these observations.

Wide-Range Current-to-Frequency Converters

National Semiconductor
Application Note 240
Robert A. Pease
May 1980

Does an analog-to-digital converter cost you a lot if you need many bits of accuracy and dynamic range? Absolute accuracy better than 0.1% is likely to be expensive. But a capability for wide dynamic range can be quite inexpensive. Voltage-to-frequency (V-to-F) converters are becoming popular as a low-cost form of A-to-D conversion because they can handle a wide dynamic range of signals with good accuracy.

Most voltage-to-frequency (V-to-F) converters actually operate with an input current which is proportional to the voltage input:

$$I_{IN} = \frac{V_{IN}}{R_{IN}}$$

(*Figure 1*). This current is integrated by an op amp, and a charge dispenser acts as the feedback path, to balance out the average input current. When an amount of charge Q=I•T (or Q=C•V) per cycle is dispensed by the circuit, then the frequency will be:

$$f = \left(\frac{V_{IN} - V_{OS}}{R_{IN}} + I_b\right) \times \frac{1}{Q}.$$

When V_{IN} is large:

$$f \cong \frac{V_{IN}}{R_{IN}} \times \frac{1}{Q}.$$

.

When V_{IN} covers a wide dynamic range, the V_{OS} and I_b of the op amp must be considered, as they greatly affect the usable accuracy when the input signal is very small. For example, when the full-scale input is 10V, a signal which is 100 dB below full-scale will be only 100 µV. If the op amp has an offset drift of ± 100 µV, (whether caused by time or temperature), that would cause a ±100% error at this signal level. However, a current-to-frequency converter can easily cover a 120 dB range because the voltage offset problem is not significant when the input signal is actually a current source. Let's study the architecture and design of a current-to-frequency converter, to see where we can take advantage of this.

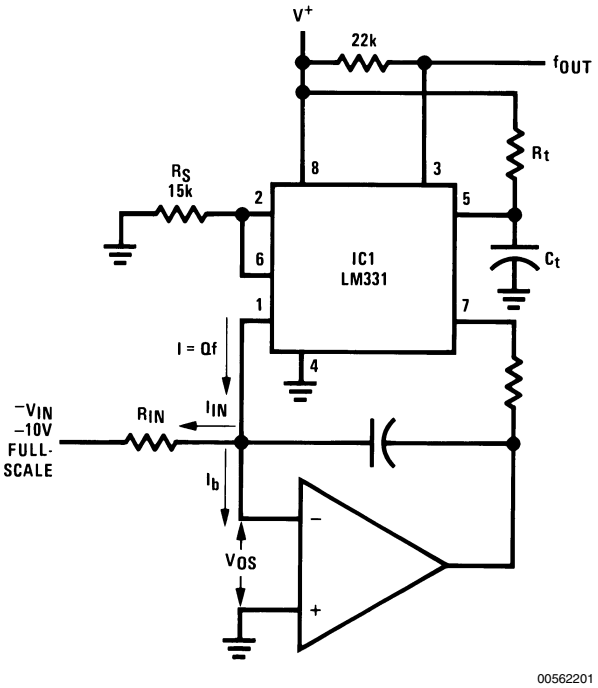

FIGURE 1. Typical Voltage-to-Frequency Converter

When the input signal is a current, the use of a low-voltage-drift op amp becomes of no advantage, and low bias current is the prime specification. A low-cost BI-FET™ op amp such as the LF351A has I_b <100 pA, and temperature coefficient of I_b less than 10 pA/˚C, at room temperature. In a typical circuit such as *Figure 2*, the leakage of the

BI-FET™ is a trademark of National Semiconductor Corp.

charge dispenser is important, too. The LM331 is only specified at 10 nA max at room temperature, because that is the smallest current which can be measured economically on high-speed test equipment. The leakage of the LM331's current-source output at pin 1 is usually 2 pA to 4 pA, and is always less than the 100 pA mentioned above, at 25˚C.

The feedback capacitor C_F should be of a low-leakage type, such as polypropylene or polystyrene. (At any temperature above 35˚C, mylar's leakage may be excessive.) Also, low-leakage diodes are recommended to protect the circuit's input from any possible fault conditions at the input. (A 1N914 may leak 100 pA even with only 1 millivolt across it, and is unsuitable.)

After trimming this circuit for a low offset when I_{IN} is 1 nA, the circuit will operate with an input range of 120 dB, from 200 µA to 100 pA, and an accuracy or linearity error well below (0.02% of the signal plus 0.0001% of full-scale).

The zero-offset drift will be below 5 or 10 pA/˚C, so when the input is 100 dB down from full-scale, the zero drift will be less than 2% of signal, for a ±5˚C temperature range. Another way of indicating this performance is to realize that when the input is 1/1000 of full-scale, zero drift will be less than 1% of that small signal, for a 0˚C to 70˚C temperature range.

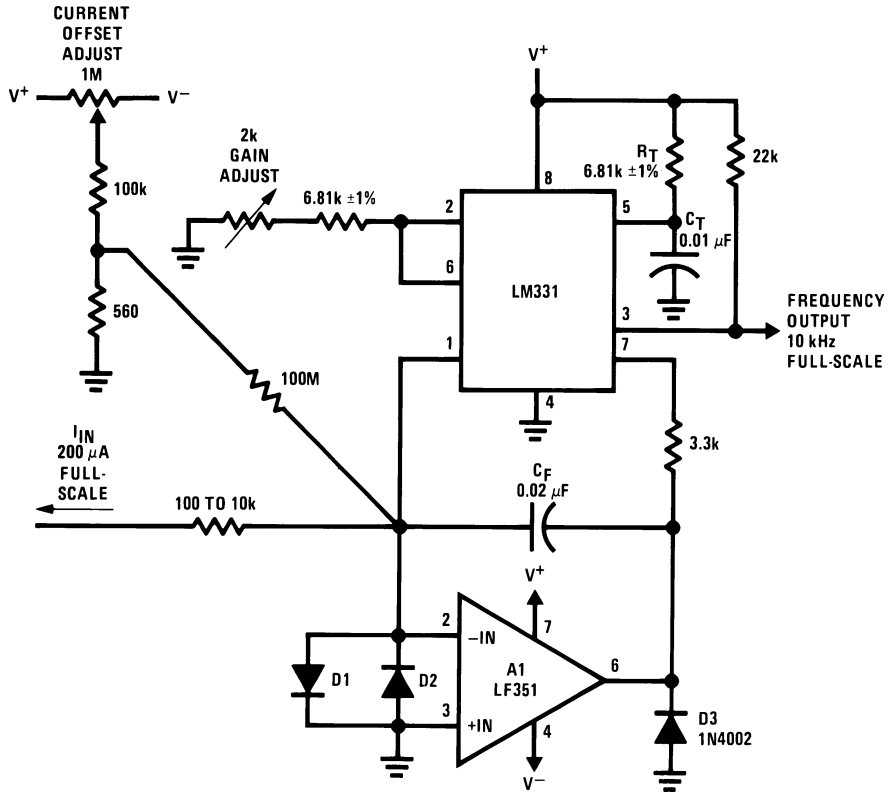

D1, D2=1N457, 1N484, or similar low-leakage planar diode

FIGURE 2. Practical Wide-Range Current-to-Frequency Converter

What if this isn't good enough? You *could* get a better op amp. For example, an LH0022C has 10 pA max I_b. But it is silly to pay for such a good op amp, with low V offset errors, when only a low input current specification is needed. The circuit of *Figure 3* shows the simple scheme of using FET followers ahead of a conventional op amp. An LF351 type is suitable because it is a cheap, quick amplifier, well suited for this work. The 2N5909s have a maximum I_b of 1.0 pA, and at room temperature it will drift only 0.1 pA/˚C. Typical drift is 0.02 pA/˚C.

The voltage offset adjust pot is used to bring the summing point within a millivolt of ground. With an input signal big enough to cause f_{OUT}=1 second per cycle, trim the V offset adjust pot so that closing the *test* switch makes no effect on the output frequency (or, output period). Then adjust the input current offset pot, to get f_{OUT}=1/1000 of full-scale when I_{IN} is 1/1000 of full-scale. When I_{IN} covers the 140 dB range, from 200 µA to 20 pA, the output will be stable, with very good zero offset stability, for a limited temperature range around room temperature. Note these precautions and special procedures:

1. Run the LM331 on 5V to 6V to keep leakage down and to cut the dissipation and temperature rise, too.
2. Run the FETs with a 6V drain supply.
3. Guard all summing point wiring away from all othervoltages.

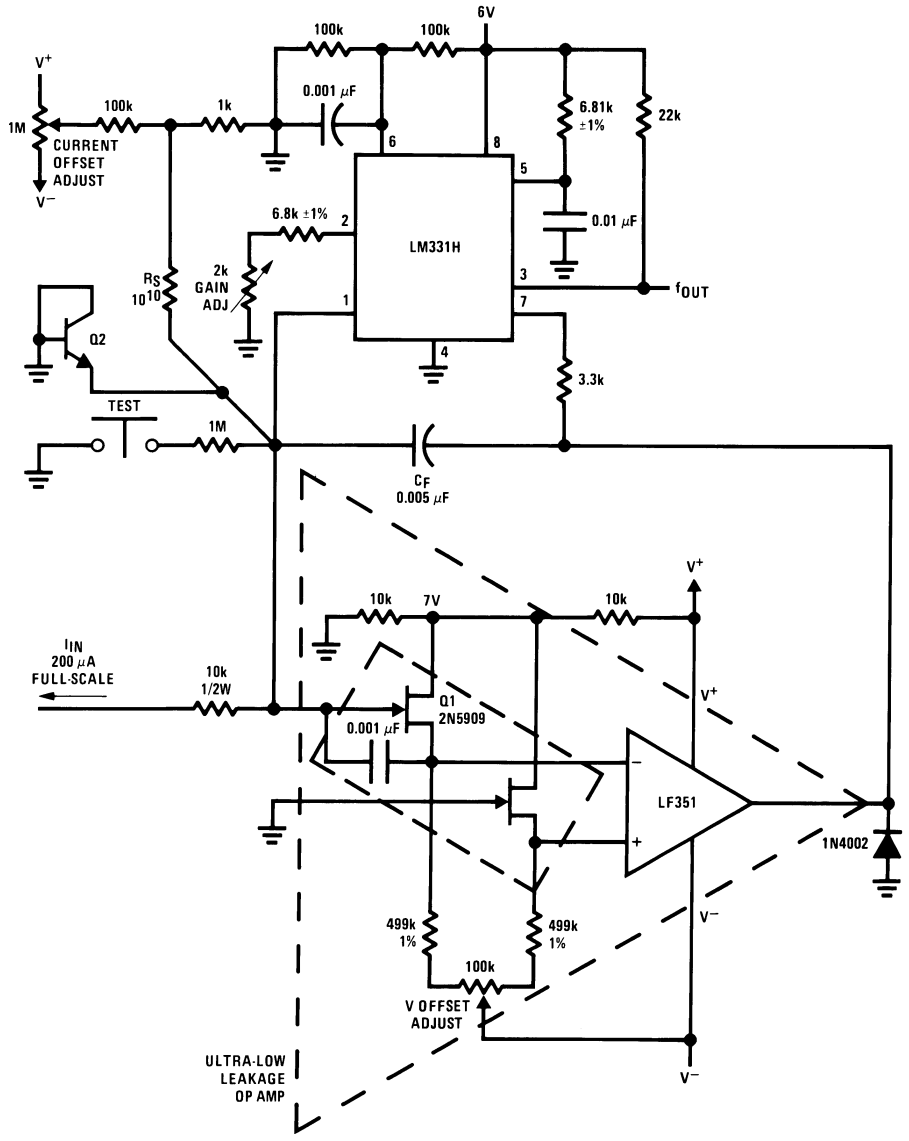

Q1 - 2N5909 or similar
1G<1 pA
Q2 - 2N930 or 2N3565

FIGURE 3. Very-Wide-Range Current-to-Frequency Converter

An alternate approach, shown in *Figure 4*, uses an LM11C as the input pre-amplifier. The LM11C has much better voltage drift than any of the other amplifiers shown here (normally less than 2 μV/˚C) and excellent current drift, less than 1 pA/˚C by itself, and typically 0.2 pA/˚C when trimmed with the 2N3904 bias current compensation circuit as shown. Of course, the LM331's leakage of 1 pA/˚C will still double every 10˚C, so that having an amplifier with excellent I_b characteristics does not solve the whole problem, when trying to get good accuracy with a 100 pA signal. For that job, even the leakage of the LM331 must be guarded out!

What if even lower ranges of input current must be accepted? While it might be possible to use a current-to-voltage converter ahead of a V-to-F converter, offset voltage drifts would hurt dynamic range badly. Response and zero-drift of such an I-V will be disappointing. Also, it is not feasible to starve the LM331 to an arbitrary extent.

For example, while its I_{OUT} (full-scale) of 280 μA DC can be cut to 10 μA or 28 μA, it cannot be cut to 1 μA or 2.8 μA with good accuracy at 10 kHz, because the internal switches in the integrated circuit will not operate with best speed and precision at such low currents.

Instead, the output current from pin 1 of the LM331 can be fed through a current attenuator circuit, as shown in *Figure 5*. The LM334 (temperature-to-current converter IC) causes −120 mV bias to appear at the base of Q2. When a current flows out of pin 1 of the LM331, 1/100 of the current will flow out of Q1's collector, and the rest will go out of Q2's collector.

As the LM334's current is linearly proportional to Kelvin temperature, the –120 mV at Q2's base will change linearly with temperature so that the Q1/Q2 current divider stays at 1:100, invariant of temperature, according to the equation:

$$i_1/i_2 = e^{\frac{q(Vb1 - Vb2)}{kT}}$$

This current attenuator will work stably and accurately, even at high speeds, such as for 4 µS current pulses. Thus, the output of Q1 is a charge pump which puts out only 10 picocoulombs per pulse, with surprisingly good accuracy. Note also that the LM331's leakage is substantially attenuated also, by a factor of 100 or more, so that source of error virtually disappears. When Q1 is off, it is really *OFF*, and its leakage is typically 0.01 pA if the summing point is within a millivolt or two of ground.

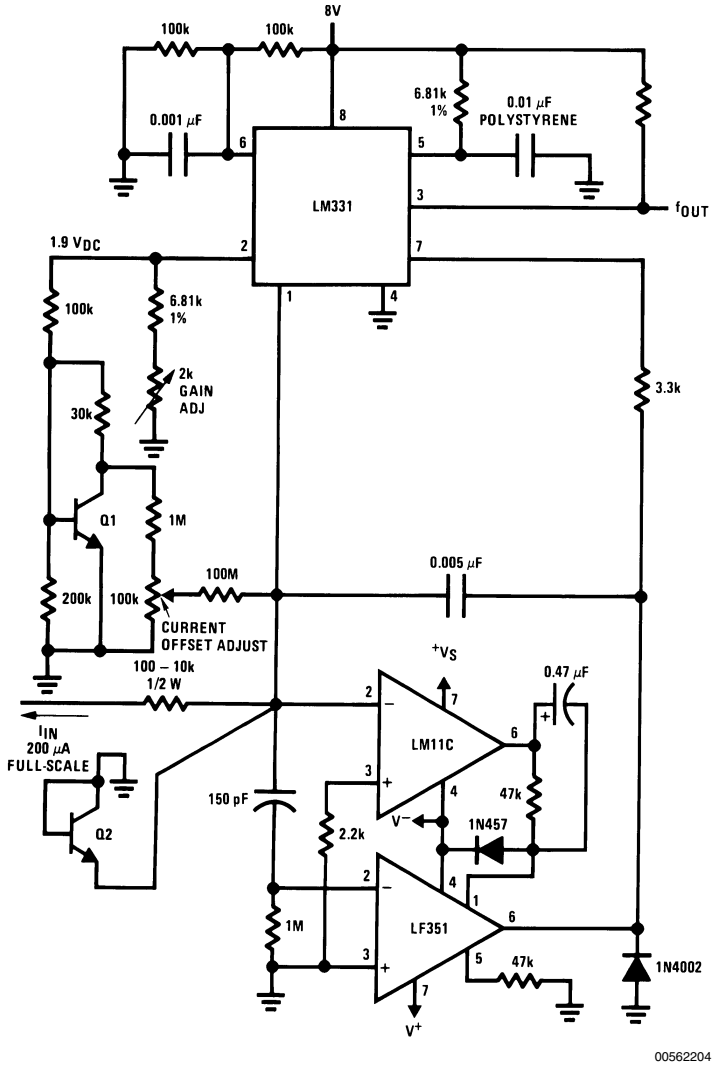

Q1, 2N3904 or any silicon NPN

Q2, 2N930 or 2N3565

FIGURE 4. Very-Wide-Range I-to-F Converter with Low Voltage Drift

To do justice to this low leakage of the VFC, the op amp should be made with MOSFETs for Q3 and Q4, such as the Intersil 3N165 or 3N190 dual MOSFET (with no gate-protection diodes). When MOSFETs have relatively poor offset voltage, offset voltage drift, and voltage noise, this circuit does not care much about these characteristics, but instead takes advantage of the MOSFET's superior current leakage and current drift.

Now, with an input current of 1 µA, the full-scale output frequency will be 100 kHz. At a 1 nA input, the output frequency will be 100 Hz. And, when the input current is 1 pA, the output frequency will drop to 1 cycle per 10 seconds or 100 mHz. When the input current drops to zero, frequencies as small as 500 µHz have been observed, at 25˚C and also as warm as 35˚C. Here is a wide-range data converter whose zero drift is *well* below 1 ppm per 10˚C! (Rather more

like 0.001 ppm per°C.) The usable dynamic range is better than 140 dB, with excellent accuracy at inputs between 100% and 1% and 0.01% and 0.0001% of full-scale.

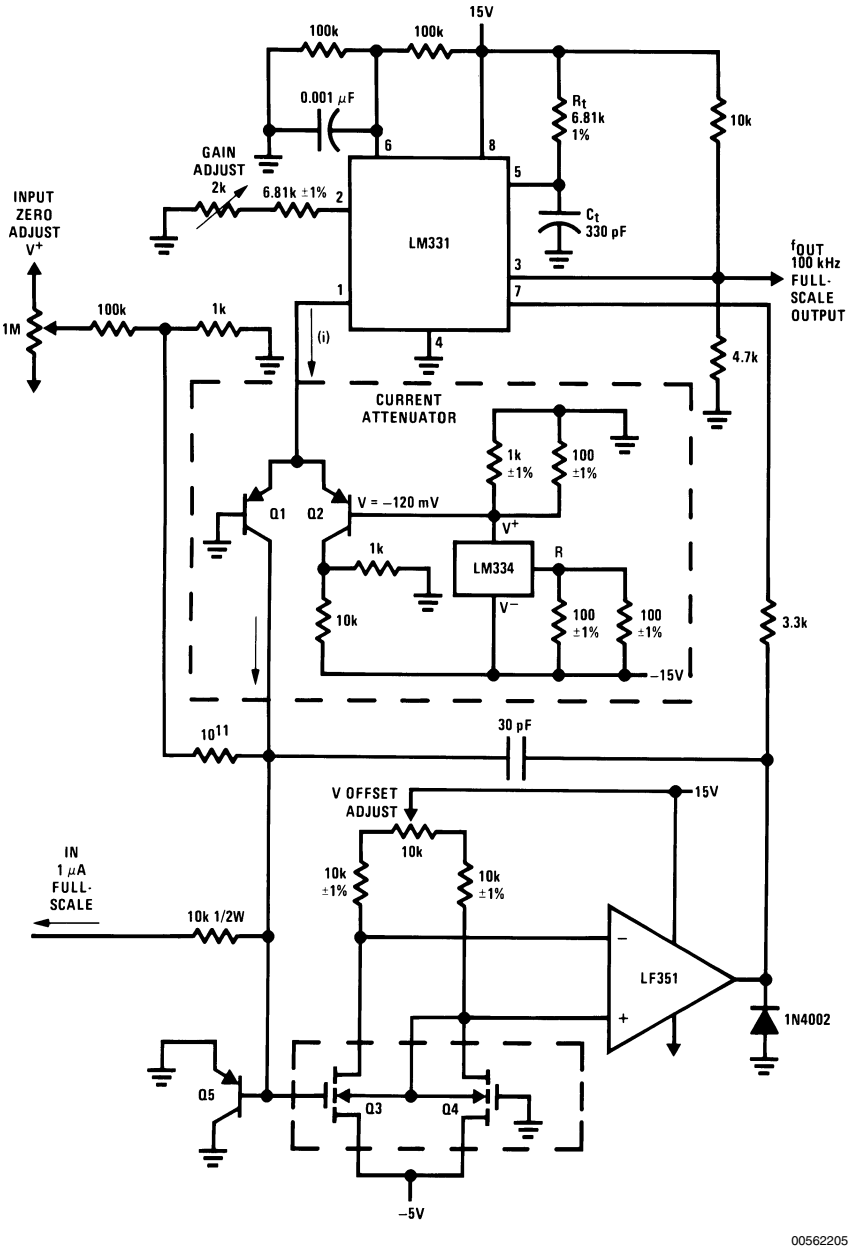

Q1, Q2, Q5 - 2N3906, 2N4250 or similar

Q3, Q4 - 3N165, 3N190 or similar. See text

Keep Q1, Q2 and LM334 at the same temperature

FIGURE 5. Picoampere-to-Frequency Converters

If a positive signal is of interest, the LM331 can be applied with a current reflector as in *Figure 6*. This current reflector has high output impedance, and low leakage. Its output can go directly to the summing point, or via a current attenuator made with NPN transistors, similar to the PNP circuit of *Figure 5*. This circuit has been observed to cover a wide (130 dB) range, with 0.1% of signal accuracy.

What is the significance of this wide-range current-to-frequency converter? In many industrial systems the question of using an inexpensive 8-bit converter instead of an expensive 12-bit data converter is a battle which is decided everyday. But if the signal source is actually a current source, then you can use a V-to-F converter to make a cheap 14-bit converter or an inexpensive converter with 18 bits of dynamic range. The choice is yours.

Why use an I-to-F converter?

- It is a natural form of A-to-D conversion.
- It naturally facilitates integration, as well.
- There are many signals in the world, such as photospectrometer currents, which like to be digitized and integrated as a standard part of the analysis of the data.
- Similarly: photocurrents, dosimeters, ionization currents, are examples of currents which beg to be integrated in a current-to-frequency meter.
- Other signal sources which provide output currents are:
  - —Phototransistors
  - —Photo diodes
  - —Photoresistors (with a fixed voltage bias)
  - —Photomultiplier tubes
  - —Some temperature sensors
  - —Some IC signal conditioners

Why have a fast frequency out?

- A 100 kHz output full-scale frequency instead of 10 kHz means that you have 10 times the resolution of the signal. For example, when I_{IN} is 0.01% of full-scale, the f will be 10 Hz. If you integrate or count that frequency for just 10 seconds, you can resolve the signal to within 1% – a factor of 10 better than if the full-scale frequency were slower.

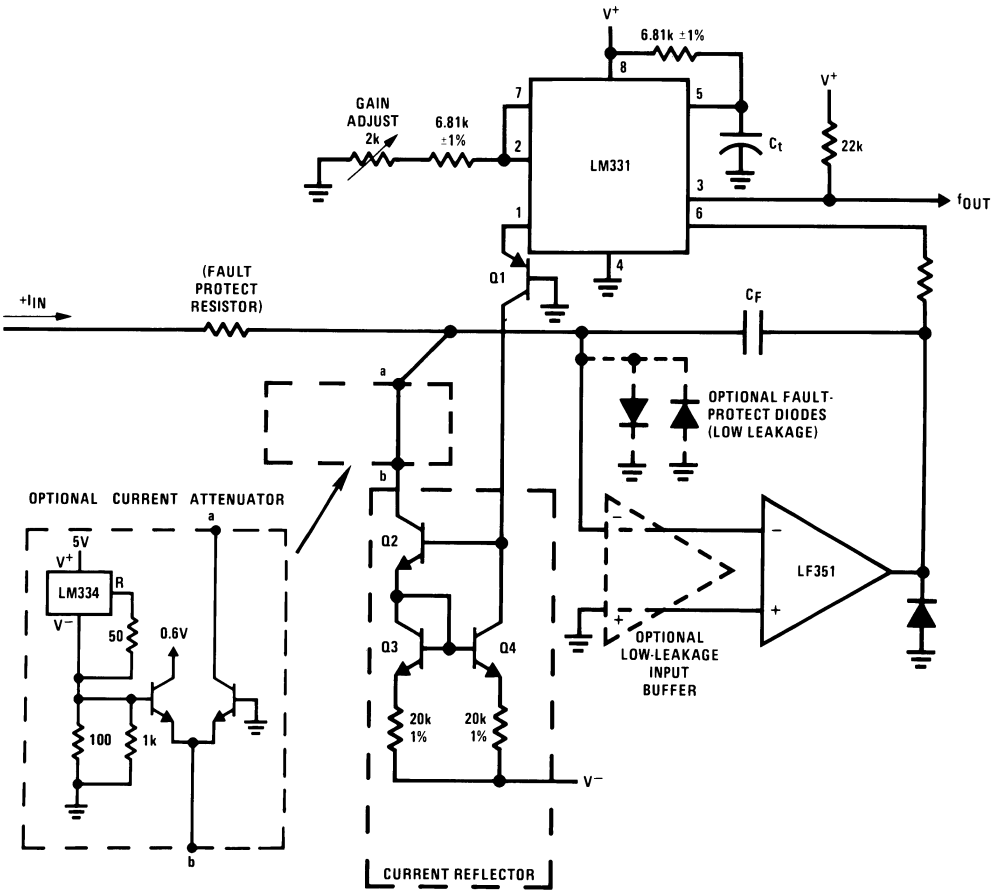

Q1 - 2N4250 or 2N3906

Q2, Q3, Q4 - 2N3904 or 2N3565

FIGURE 6. Current-to-Frequency Converter For Positive Signals

V/F Converter ICs Handle Frequency-to-Voltage Needs

National Semiconductor
Application Note C
Robert A. Pease
August 1980

Simplify your F/V converter designs with versatile V/F ICs. Starting with a basic converter circuit, you can modify it to meet almost any application requirement. You can spare yourself some hard labor when designing frequency-to-voltage (F/V) converters by using a voltage-to-frequency IC in your designs. These ICs form the basis of a series of accurate, yet economical, F/V converters suiting a variety of applications.

Figure 1 shows an LM331 IC (or LM131 for the military temperature range) in a basic F/V converter configuration (sometimes termed a stand-alone converter because it requires no op amps or other active devices other than the IC). (Comparable V/F ICs, such as RM4151, can take advantage of this and other circuits described in this article, although they might not always be pin-for-pin compatible).

This circuit accepts a pulse-train or square wave input amplitude of 3V or greater. The 470 pF coupling capacitor suits negative-going input pulses between 80 µs and 1.5 µs, as well as accommodating square waves or positive-going pulses (so long as the interval between pulses is at least 10 µs).

IC Handles the Hard Part

The LM331 detects an input-signal change by sensing when pin 6 goes negative relative to the threshold voltage at pin 7, which is nominally biased 2V lower than the supply voltage. When a signal change occurs, the LM331's input comparator sets an internal latch and initiates a timing cycle. During this cycle, a current equal to V_{REF}/RS flows out of pin 1 for a time $t = 1.1\ R_tC$. The 1 µF capacitor filters this pulsating current from pin 1, and the current's average value flows through load resistor R_L. As a result, for a 10 kHz input, the circuit outputs 10 V_{DC} across R_L with good (0.06% typical) nonlinearity.

Two problems remain, however: the output at V1 includes about 13 mVp-p ripple, and it also lags 0.1 second behind an input frequency step change, settling to 0.1% of full-scale in about 0.6 second. This ripple and slow response represent an inherent tradeoff that applies to almost every F/V converter.

The Art of Compromise

Increasing the filter capacitor's value reduces ripple but also increases response time. Conversely, lowering the filter capacitor's value improves response time at the expense of larger ripple. In some cases, adding an active filter results in faster response and less ripple for high input frequencies.

Although the circuit specifies a 15V power supply, you can use any regulated supply between 4 V_{DC} and 40 V_{DC}. The output voltage can extend to within 3 V_{DC} of the supply voltage, so choose R_L to maintain that output range.

Adding a 220 kΩ/0.1 µF postfilter to the circuit slows the response slightly, but it also reduces ripple to less than 1 mVp-p for frequencies from 200 Hz to 10 kHz. The reduction in ripple achieved by adding this passive filter, while not as good as that obtainable using an active filter, could suffice in some applications.

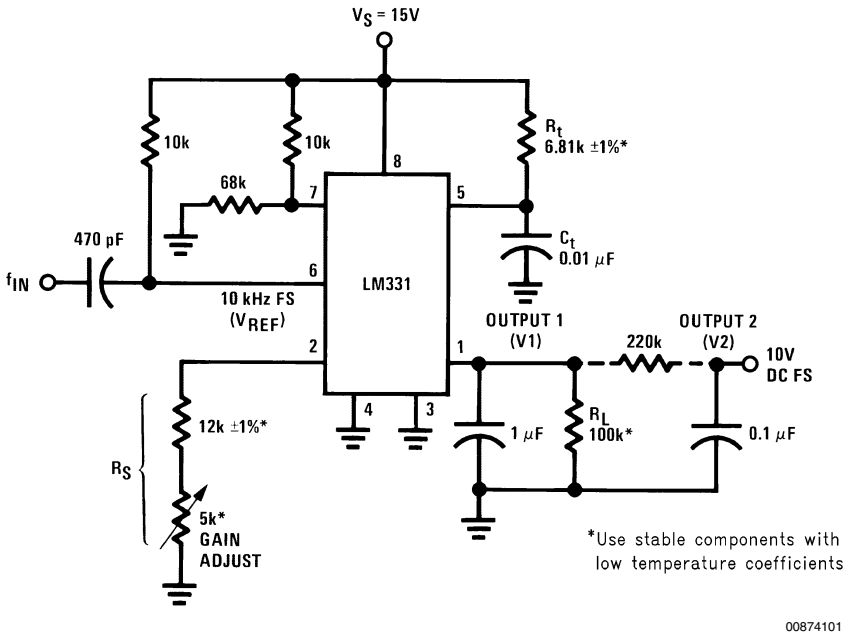

FIGURE 1. A Simple Stand-Alone F/V Converter Forms the Basis for Many Other Converter-Circuit Configurations

Improving the Basic Circuit

Further modifications and additions to the basic F/V converter shown in *Figure 1* can adapt it to specific performance requirements. *Figure 2* shows one such modification, which improves the converter's nonlinearity to 0.006% typical.

Reconsideration of the basic stand-alone converter shows why its nonlinearity falls short of this improved version's. At low input frequencies, the current source feeding pin 1 in the LM331 is turned off most of the time. As the input frequency increases, however, the current source stays on more of the time, and its own impedance attenuates the output signal for an increasing fraction of each cycle time. This disproportionate attenuation at higher frequencies causes a parabolic change in full-scale gain rather than the desired linear one.

In the improved circuit, on the other hand, the PNP transistor acts as a cascade, so the output impedance at pin 1 sees a constant voltage that won't modulate the gain. Also, with an alpha ranging between 0.998 and 0.990, the transistor exhibits a temperature coefficient of between 10 ppm/°C and 40 ppm/°C—a fairly minor effect. Thus, this circuit's nonlinearity does not exceed 0.01% maximum for the 10V output range shown and is normally not worse than 0.01% for any supply voltage between 4V and 40V.

Add an Output Buffer

The circuit in *Figure 3* adds an output buffer (unity-gain follower) to the basic single-supply F/V converter. Either an LM324 or LM358 op amp functions well in a single-supply circuit because these devices' common-mode ranges extend down to ground. But if a negative supply is available, you can use any op amp; types such as the LF351B or LM308A, which have low input currents, provide the best accuracy.

The output buffer in *Figure 3* also acts as an active filter, furnishing a 2-pole response from a single op amp. This filter provides the general response

$$V_{OUT}/I_{OUT} = R_L/(1 + K1p + K2p^2).$$

(p is the differential operator d/dt.) As shown, R_L controls the filter's DC gain. The high frequency response rolls off at 12 dB/octave. Near the circuit's natural resonant frequency, you can choose the damping to give a little overshoot—or none, as desired.

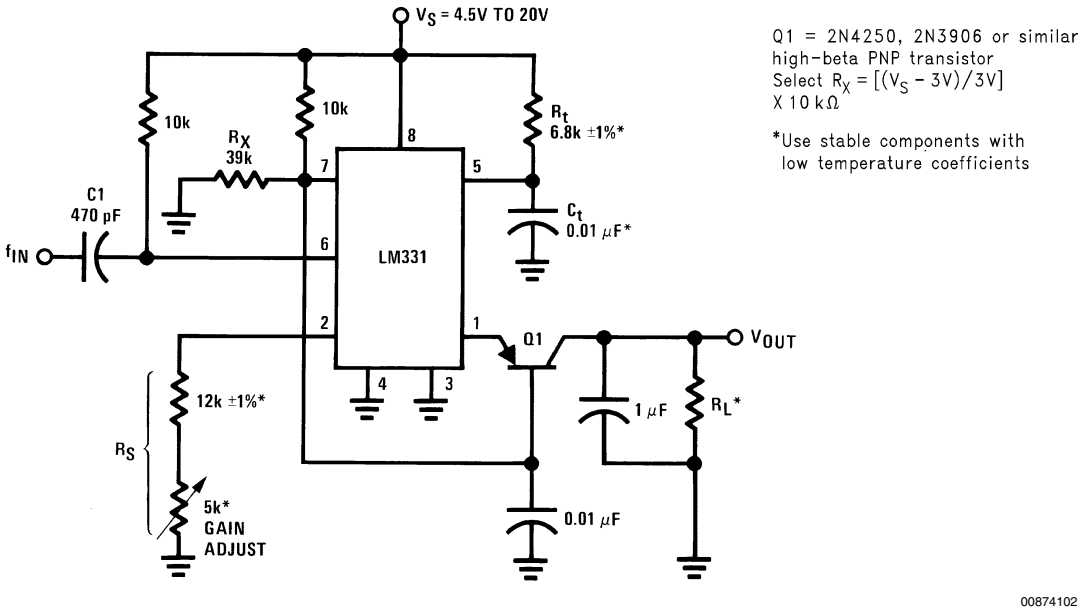

FIGURE 2. Adding a Cascade Transistor to the LM331's Output Improves Nonlinearity to 0.006%

Add an Output Buffer (Continued)

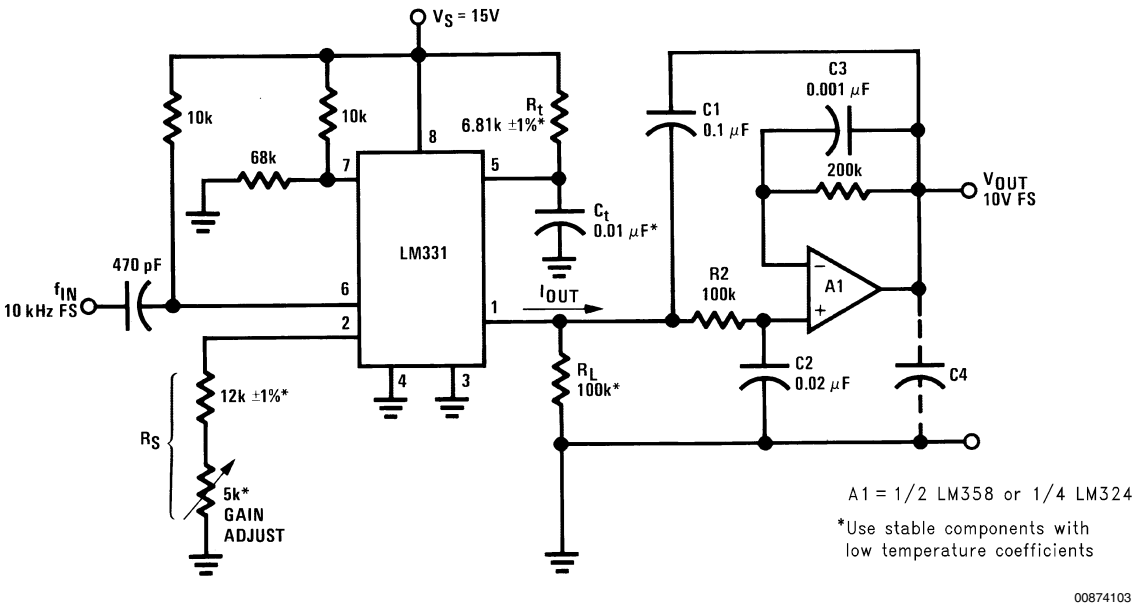

FIGURE 3. The Op Amp on This F/V Converter's Output Acts as a Buffer as Well as a 2-Pole Filter

Dealing with F/V Converter Ripple

Voltage ripple on the output of F/V converters can present a problem, and the chart shown in *Figure 4* indicates exactly how big a problem it is. A simple, slow, RC filter exhibits low ripple at all frequencies. Two-pole filters offer the lowest ripple at high frequencies and provide a 30-times-faster step response than RC devices.

To reduce a circuit's ripple at moderate frequencies, however, you can cascade a second active-filter stage on the F/V converter's output. That circuit's response also appears in *Figure 4* and shows a significant improvement in low-ripple bandwidth over the single-active-filter configuration, with only a 30% degradation of step response.

Figure 5 and *Figure 6* show filter circuits suitable for cascading. The inverting filter in *Figure 5* requires closely matched resistors with a low TC over their temperature range for best accuracy. For lowest DC error, choose R5 = R2 + (R_{IN}|R_F). This circuit's response is

$-V_{OUT}/V_{IN} = n/(1 + (R_F + R2 + nR2)C4p + R_FR2C3C4p^2)$.

where n = DC gain. If $R_{IN} = R_F$ and n = 1,

$-V_{OUT}/V_{IN} = 1/(1 + (R_F + 2R2)C4p + R_FR2C3C4p^2)$.

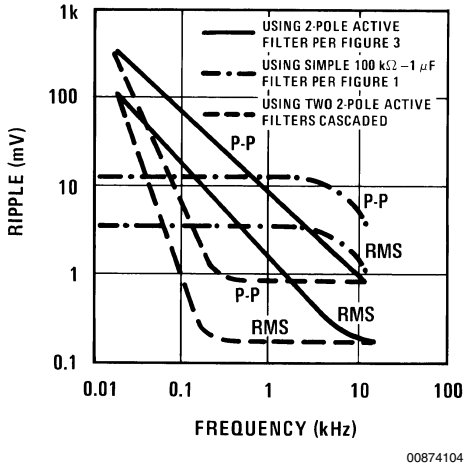

FIGURE 4. Output-Ripple Performance of Several Different F/V Converter Configurations Illustrates the Effect of Voltage Ripple

Dealing with F/V Converter Ripple

(Continued)

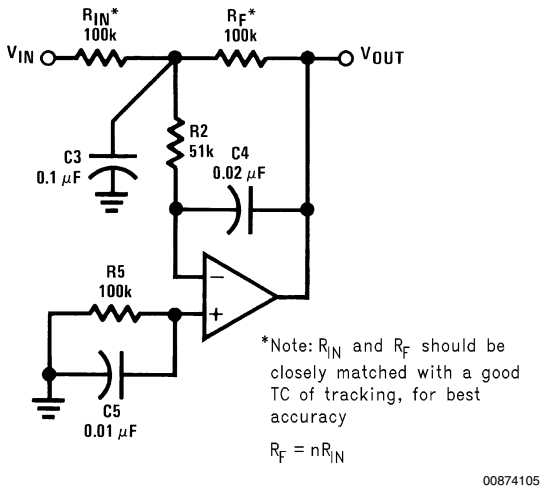

FIGURE 5. You Can Cascade This 2-Pole Inverting Filter onto an F/V Converter's Output

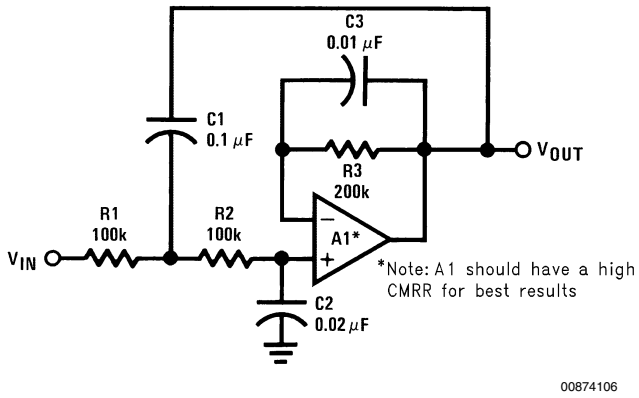

FIGURE 6. This 2-Pole Noninverting Filter Suits Cascade Requirements on F/V Converter Outputs

The circuit shown in *Figure 6* does not require precision passive components, but for best accuracy, choosing an A1 with a high CMRR is critical. An LM308A op amp's 96 dB minimum CMRR suits this circuit well, but an LM358B's 85 dB typical figure also proves adequate for many applications. Circuit response is

$V_{OUT}/V_{IN} = 1/(1 + (R1 + R2)\ C2p + R1R2C1C2p^2)$.

For best results, choose R3 = R1 + R2.

Components Determine Response

The specific response of the circuit in *Figure 3* is

$V_{OUT}/I_{OUT} = R_L/(1 + (R_L + R2)C2p + R_LR2C1C2C2p^2)$.

Making C2 relatively large eliminates overshoot and sine peaking. Alternatively, making C2 a suitable fraction of C1 (as is done in *Figure 3*) produces both a sine response with 0 dB to 1 dB of peaking and a quick real-time response having only 10% to 30% overshoot for a step response. By maintaining *Figure 3*'s ratio of C1:C2 and R2:R_L, you can adapt its 2-pole filter to a wide frequency range without tedious computations.

This filter settles to within 1% of a 5V step's final value in about 20 ms. By contrast, the circuit with the simple RC filter shown in *Figure 1* takes about 900 ms to achieve the same response, yet offers no less ripple than *Figure 3*'s op amp approach.

As for the other component in the 2-pole filter, any capacitance between 100 pF and 0.05 μF suits C3 because it serves only as a bypass for the 200 kΩ resistor. C4 helps reduce output ripple in single positive power-supply systems when VOUT approaches so close to ground that the op amp's output impedance suffers. In this circuit, using a tantalum capacitor of between 0.1 μF and 2.2 μF for C4 usually helps keep the filter's output much quieter without degrading the op amp's stability.

Avoid Low-Leakage Limitations

Note that in most ordinary applications, this 2-pole filter performs as well with 0.1 μF and 0.02 μF capacitors as the passive filter in *Figure 1* does with 1 μF. Thus, if you require a 100 Hz F/V converter, the circuit in *Figure 3* furnishes good filtering with C1 = 10 μF and C2 = 2 μF, and eliminates the 100 μF low-leakage capacitor needed in a passive filter.

Note also that because C1 always has zero DC voltage across it, you can use a tantalum or aluminum electrolytic capacitor for C1 with no leakage-related problems; C2, however, must be a low-leakage type. At room temperature, typical 1 μF tantalum components allow only a few nanoamperes of leakage, but leakage this low usually cannot be guaranteed.

Compensating for Temperature Coefficients

F/V converters often encounter temperature-related problems usually resulting from the temperature coefficients of passive components. Following some simple design and manufacturing guidelines can help immunize your circuits against loss of accuracy when the temperature changes.

Capacitors fabricated from Teflon or polystyrene usually exhibit a TC of −110 ±30 ppm/°C. When you use such a component for the timing capacitor in an F/V converter (such as C_t in the *figure*) the circuit's output voltage — or the gain in terms of volts per kilohertz — also exhibits a −110 ppm/°C TC.

But the resistor-diode network (R_X, D1, D2) connected from pin 2 to ground in the *figure* can cancel the effect of the timing capacitor's large TC. When R_X = 240 kΩ, the current flowing through pin 1 will then have an overall TC of 110 ppm/°C, effectively canceling a polystyrene timing capacitor's TC to a first approximation. Thus, you needn't find a zero-TC capacitor for C_t, so long as its temperature coefficient is stable and well established. As an additional advantage, the resistor-diode network nearly compensates to zero the TC of the rest of the circuit.

Bake it for a While

After the circuit has been built and checked out at room temperature, a brief oven test will indicate the sign and the size of the TC for the complete F/V converter. Then you can add resistance in series with R_X, or add conductance in parallel with it, to greatly diminish the TC previously observed and yield a complete circuit with a lower TC than you could obtain simply by buying low TC parts.

For example, if the circuit increases its full-scale output by 0.1% per 30˚C (33 ppm/˚C) during the oven test, adding 120 kΩ in series with R_X = 240 kΩ cancels the temperature-caused deviation. Or, if the full-scale output decreases by −0.04% per 20˚C (−20 ppm/˚C), just add 1.2 MΩ in parallel with R_X.

Note that to allow trimming in both directions, you must start with a finite *fixed* TC (such as the −110 ppm/˚C of C_t), which then nominally cancels out by the addition of a finite *adjustable* TC. Only by using this procedure can you compensate for whatever polarity of TC is found by the oven test.

You can utilize this technique to obtain TCs as low as 20 ppm/˚C, or perhaps even 10 ppm/˚C, if you take a few passes to zero-in on the best value for R_X. For optimum results, consider the following guidelines:

- Use a good capacitor for C_t; the cheapest polystyrene capacitors can shift value by 0.05% or more per temperature cycle. In that case, you would not be able to distinguish the actual temperature sensitivity from the hysteresis, and you would also never achieve a stable circuit.
- After soldering, bake or temperature-cycle the circuit (at a temperature not exceeding 75˚C in the case of polystyrene) for a few hours to stabilize all components and to relieve the strains of soldering.
- Do not rush the trimming. Recheck the room temperature value before *and* after you take the high temperature data to ensure a reasonably low hysteresis per cycle.
- Do not expect a perfect TC at −25˚C if you trim for ±5 ppm/˚C at temperatures from +25˚C to 60˚C. None of the components in the *figure's* circuit offer linearity much better than 5 ppm/˚C or 10 ppm/˚C cold, if trimmed for a zero TC at warm temperatures. Even so, using these techniques you can obtain a data converter with better than 0.02% accuracy and 0.003% linearity, for a ±20˚C range around room temperature.
- Start out the trimming with R_X installed and its value near the design-center value (e.g., 240 kΩ or 270 kΩ), so you will be reasonably close to zero TC; you will usually find the process slower if you start without any resistor, because the trimming converges more slowly.
- If you change R_X from 240 kΩ to 220 kΩ, do not pull out the 240 kΩ part and put in a new 220 kΩ resistor—you will get much more consistent results by adding a 2.4 MΩ resistor in parallel. The same admonition holds true for adding resistance in series with R_X.
- Use reasonably stable components. If you use an LM331A (±50 ppm/˚C maximum) and RN55D film resistors (each ±100 ppm/˚C) for R_L, R_t and R_S, you probably won't be able to trim out the resulting ±350 ppm/˚C worst-case TC. Resistors with a TC specification of 25 ppm/˚C usually work well. Finally, use the same resistor value (e.g., 12.1 kΩ ±1%) for both R_S and R_t; when these resistors come from the same manufacturer's batch, their TC tracking will usually rate at better than 20 ppm/˚C.

Whenever an op amp is used as a buffer (as in *Figure 3*), its offset voltage and current (±7.5 mV maximum and ±100 nA, respectively, for most inexpensive devices) can cause a ±17.5 mV worst-case output offset. If both plus and minus supplies are available, however, you can easily provide a symmetrical offset adjustment. With only one supply, you can add a small positive current to each op amp input and also trim one of the inputs.

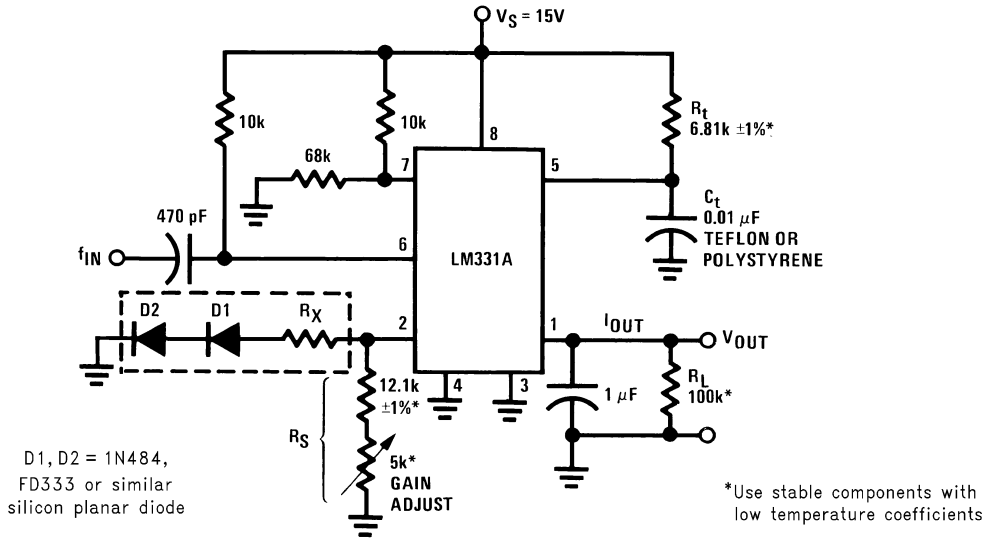

00874107

Two Diodes and a Resistor Help Decrease an F/V Converter's Temperature Coefficient

Need a Negative Output?

If your F/V converter application requires a negative output voltage, the circuit shown in *Figure 7* provides a solution with excellent linearity (±0.003% typical, ±0.01% maximum). And because pin 1 of the LM331 always remains at 0 V_{DC}, this

Need a Negative Output? (Continued)

circuit needs no cascade transistor. (Note, however, that while the circuit's nonlinearity error is negligible, its ripple is not.)

The circuit in *Figure 7* offers a significant advantage over some other designs because the offset adjust voltage derives from the stable 1.9 V_{DC} reference voltage at pin 2 of the LM331; thus any supply voltage shifts cause no output shifts. The offset pot can have any value between 200 kΩ and 2 MΩ.

An optional bypass capacitor (C2) connected from the op amp's positive input to ground prevents output noise arising from stray noise pickup at that point; the capacitance value is not critical.

A Familiar Response

The circuit in *Figure 7* exhibits the same 2-pole response—with heavy output ripple attenuation—as the noninverting filter in *Figure 3*. Specifically,

$$V_{OUT}/I_{OUT} = R_F/(1 + (R4 + R_F)C4p + R4R_FC3C4p^2).$$

Here also, R5 = R4 + R_F = 200 kΩ provides the best bias current compensation.

The LM331 can handle frequencies up to 100 kHz by utilizing smaller-value capacitors as shown in *Figure 8*. This circuit increases the current at pin 2 to facilitate high-speed switching, but, despite these speed-ups, the LM331's 500 ppm/°C TC at 100 kHz causes problems because of switching speed shifts resulting from temperature changes.

To compensate for the device's positive TC, the LM334 temperature sensor feeds pin 2 a current that decreases linearly with temperature and provides a low overall temperature coefficient. An Ry value of 30 kΩ provides first-order compensation, but you can trim it higher or lower if you need more precise TC correction.

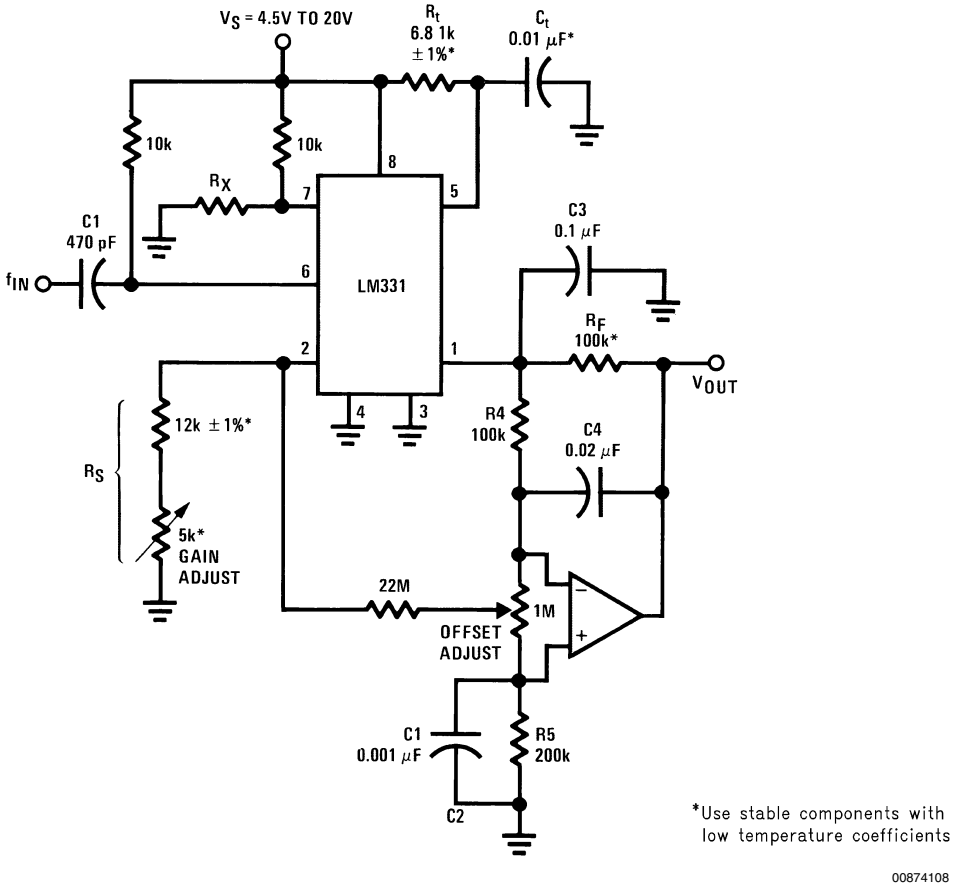

FIGURE 7. In This F/V Circuit, the Output-Buffer Op Amp Derives its Offset Voltage from the Precision Voltage Source at Pin 2 of the LM331

A Familiar Response (Continued)

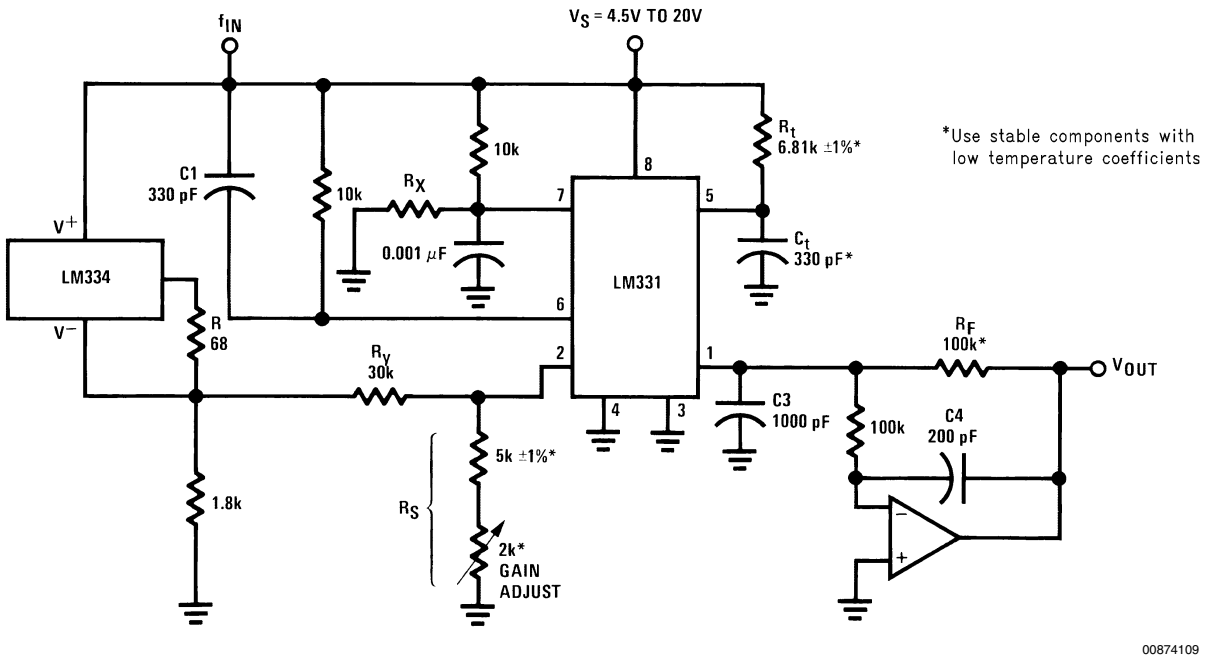

FIGURE 8. An LM334 Temperature Sensor Compensates for the F/V Circuit's Temperature Coefficient

Detect Frequencies Accurately

Using an F/V converter combined with a comparator as a frequency detector is an obvious application for these devices. But when the F/V converter is utilized in this way, its output ripple hampers accurate frequency detection, and the slow filter frequency response causes delays.

If a quick response is not important, though, you can effectively utilize an LM331-based F/V converter to feed one or more comparators, as shown in *Figure 9*. For an input frequency drop from 1.1 kHz to 0.5 kHz, the converter's output responds within about 20 ms. When the input falls from 9 kHz to 0.9 kHz, however, the output responds only after a 600 ms lag, so utilize this circuit only in applications that can tolerate F/V circuits' inherent delays and ripples.

Detect Frequencies Accurately (Continued)

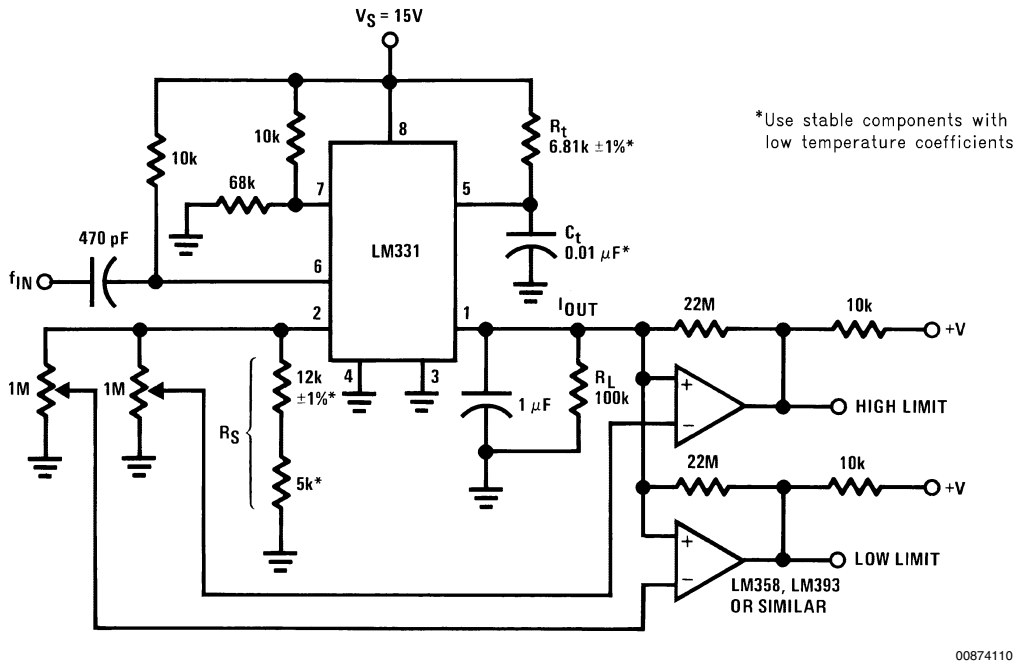

FIGURE 9. Combining a V/F IC with Two Comparators Produces a Slow-Response Frequency Detector

Author's Biography

Bob Pease is a staff scientist in the Advanced Linear Integrated Circuit Group at National Semiconductor Corp., Santa Clara, CA. Holder of four patents, he earned a BSEE from MIT. Bob lists tracking abandoned railroad roadbeds and designing V/F converters as hobbies.

National Semiconductor Corporation
Americas
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5639-7560
Fax: 81-3-5639-7507

Get Fast Stable Response From Improved Unity-Gain Followers

National Semiconductor
Linear Brief 42
Robert A. Pease
August 1978

In many applications, a unity-gain follower (e.g. any operational amplifier with tight feedback to the inverting input) may oscillate or exhibit bad ringing when required to drive heavy load capacitance. For example, the LM110 follower will normally drive a 50 pF load capacitor, but will not drive 500 pF, because the open-loop outout impedance is lagged by such a large capacitive load. The frequency at which this lag occurs is comparable to the gain-bandwidth product of the amplifier, and when the phase margin is decreased to zero, oscillation occurs.

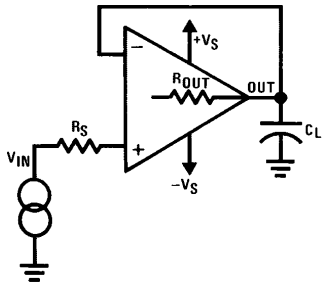

TL/H/8491–1

FIGURE 1. Unity-Gain Follower Attempting to Drive Capacitive Load

While the solution to this problem is not widely known, an analysis of the general problem shown in *Figure 2* can lead to a useful approach. It is generally known that increasing the noise gain of an op amp's feedback network will improve tolerance of capacitive load. In *Figure 2*, adding a resistor R2 $\cong R_F/10$ will do this. (A moderate capacitor C2 is usually inserted in series with R2, to prevent the DC noise gain from increasing also—to avoid degrading DC offset, drift and inaccuracy.) If the op amp has a 1 MHz gain bandwidth product, and R1 = R_F, the closed-loop frequency response will be ½ MHz. Adding R2 = $R_F/10$ will drop the closed-loop frequency response to 90 kHz, where the amplifier can usually tolerate a much larger C_L;

$$\text{Noise Gain} = \frac{R_F}{R1} + \frac{R_F}{R2} + 1 \text{ (AC)}$$

$$\text{Noise Gain} = \frac{R_F}{R1} + 1 \text{ (DC)}$$

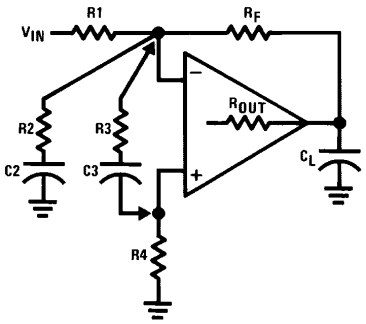

TL/H/8491–2

FIGURE 2. Stabilizing an Operational Amplifier for Capacitive Load

A similar result will occur if you install R3 and C3, instead of R2. Now the (AC) noise gain will be:

$$1 + \frac{R4}{R3} + \frac{R_F}{R3} + \left(\frac{R_F}{R1}\right)\left(\frac{R3 + R4}{R3}\right)$$

As a simplification, if R1 is an open circuit, the AC noise gain will be: (R4/R3 + R_F/R3 + 1). Now it can be seen that noise gain can be raised by having a low value of R3 and a high value of R4 or R_F (or both).

In particular, where R_F is required to be 0Ω, as in a follower, the noise gain can be raised by adding a large R4 and a small R5, as shown in *Figure 3*. If R_S is low, the AC noise gain will be R4/R5 + 1. (If R_S is large and constant, R4 may be unnecessary, and the noise gain would then be R_S/R5 + 1.) For LM110/LM310's R4 = 10 kΩ is recommended and when R5 = 3.3 kΩ, C5 = 200 pF, the LM110 will stably drive C_L up to 600 pF.

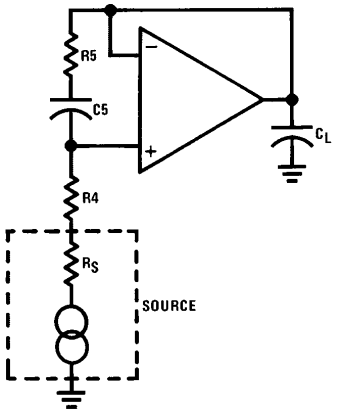

TL/H/8491–3

FIGURE 3. Stabilizing a Unity-Gain Follower for Capacitive Load

RRD-B30M115/Printed in U. S. A.

Another application of this technique is for making a fast follower with a high slew rate. An LF356 is specified as a follower, but an LF357 must be applied at an "$A_V = 5$" minimum, because it has been "decompensated" with a smaller internal capacitor. Most people do not realize how easy it is to apply an LF357 as a follower. In *Figure 4*, an LF357 will have fast, stable response just like an LF356 does, when R_S is $< 1\ k\Omega$, but it will have a $50V/\mu s$ slew rate (typical) vs. $12V/\mu s$ for an LF356.

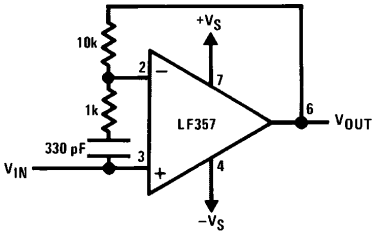

TL/H/8491–4

FIGURE 4. Unity-Gain Follower With Fast Slew Rate

Similarly, an LM348 is a fast decompensated quad op amp. Its bipolar input stage has a finite bias current, 200 nA max. For best results, the resistance which makes up the noise gain should be put equally in the **plus** and **minus** input circuits, as shown in *Figure 5*. The LM349 can slew at $2V/\mu s$ typical, and is much faster for handling audio signals without distortion than the LM348 (which at $0.5V/\mu s$ is only as fast as an ordinary LM741). The same approach can be used for an LM101 with a 5 pF damping capacitor. While these circuits give faster slewing, the bandwidth may degrade if the source impedance R_S increases. Also, when the AC noise gain is raised, the AC noise will also be increased. While most modern op amps have low noise, a noise gain of 10 may make a significant increase in output noise, which the user should check to insure it is not objectionable.

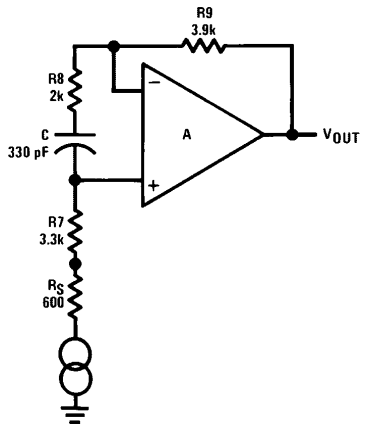

TL/H/8491–5

FIGURE 5. Application of Fast Follower With Balanced Resistors, R9 = R7 + R_S, A = 1/4 LM349 (or LM101 with 5 pF Capacitor)

If the series capacitor is much larger than necessary, noise will be increased more than necessary. In general, choose the C5 for *Figure 3*, (e.g.) per these guidelines: (where f_V = unity-gain bandwidth of op amp)

$$C5\ \text{Min} = \frac{4 \bullet \left(1 + \frac{R4}{R5}\right)}{2\pi R5 \bullet f_V} = \frac{R4 + R5}{\frac{\pi}{2} \bullet f_V \bullet (R5)^2}$$

For best results, choose the design center value of C5 to be 2 or 3 times C5 min.

National Semiconductor Corporation
1111 West Bardin Road
Arlington, TX 76017
Tel: 1(800) 272-9959
Fax: 1(800) 737-7018

National Semiconductor Europe
Fax: (+49) 0-180-530 85 86
Email: cnjwge@tevm2.nsc.com
Deutsch Tel: (+49) 0-180-530 85 85
English Tel: (+49) 0-180-532 78 32
Français Tel: (+49) 0-180-532 93 58
Italiano Tel: (+49) 0-180-534 16 80

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductor Japan Ltd.
Tel: 81-043-299-2309
Fax: 81-043-299-2408

Get More Power Out of Dual or Quad Op-Amps

National Semiconductor
Linear Brief 44
Bob Pease
April 1979

Although simple brute-force paralleling of op-amps is a bad scheme for driving heavy loads, here is a good scheme for dual op-amps. It is fairly efficient, and will not overheat if the load is disconnected. It is *not* useful for driving active loads or nonlinear loads, however.

In *Figure 1*, an LF353N mini-DIP can drive a 600Ω load to ±9V typical (±6V min guaranteed) and will have only a 47˚C temperature rise above free air. If the load R is removed, the chip temperature will rise to +50˚C above free air. Note that A2's task is to drive half of the load. A1 could be applied as a unity-gain follower *or* inverter, or as a high-gain or low-gain amplifier, integrator, etc.

While *Figure 1* is suitable for sharing a load between 2 amplifiers, it is not suitable for 4 or more amplifiers, because the circuit would tend to go out of control and overheat if the load is disconnected.

Instead, *Figure 2* is generally recommended, as it is capable of driving large output currents into resistive, reactive, nonlinear, passive, or active loads. It is easily expandable to use as many as 2 or 4 or 8 or 20 or more op-amps, for driving heavier loads.

It operates, of course, on the principle that every op-amp has to put out the same current as A1, whether that current is plus, minus, or zero. Thus if the load is removed, all amplifiers will be unloaded together. A quad op-amp can drive 600Ω to ±11 or 12 volts. Two quads can put out ±40 mA, but they get only a little warm. A series R-C damper of 15Ω in series with 0.047 µF is useful to prevent oscillations (although LM324's do not seem to need any R-C damper).

Of course, there is no requirement for the main amplifier to run only as a unity-gain amplifier. In the example shown in *Figure 3*, A1 amplifies a signal with a gain of +10. A2 helps it drive the load. Then A3 operates as a unity-gain inverter to provide V2 = −V1, and A4 helps it drive the load. This circuit can drive a floating 2000Ω load to ±20V, accurately, using a slow LM324 or a quick LF347.

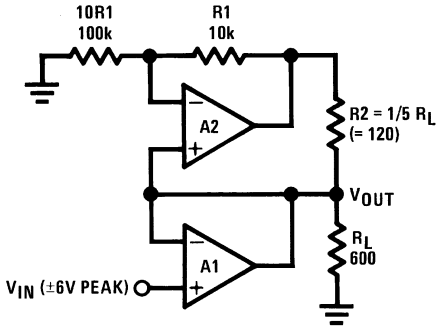

A1, A2 = 1/2 LM747 or 1/2 LF353 or any op-amp.

FIGURE 1. A1 and A2 Share the Load

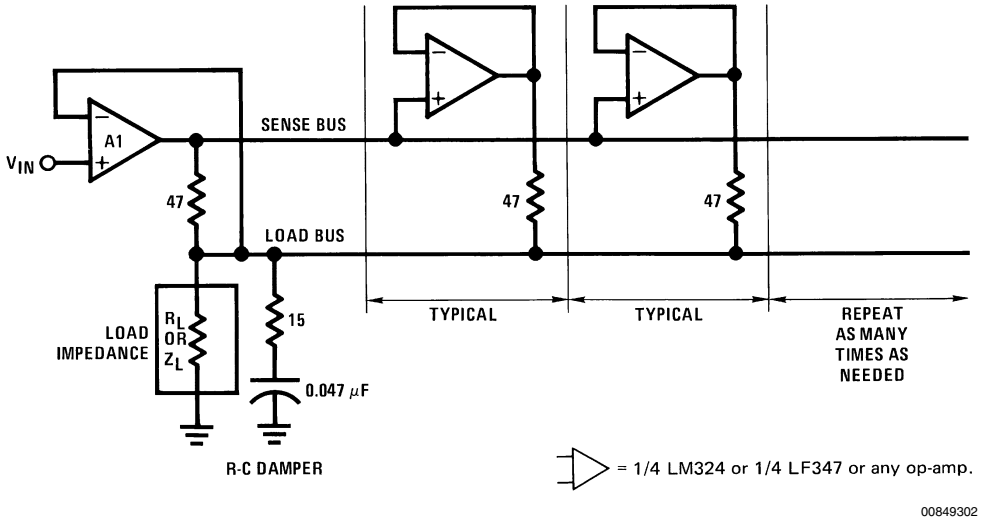

FIGURE 2. Improved Load-Sharing Circuit

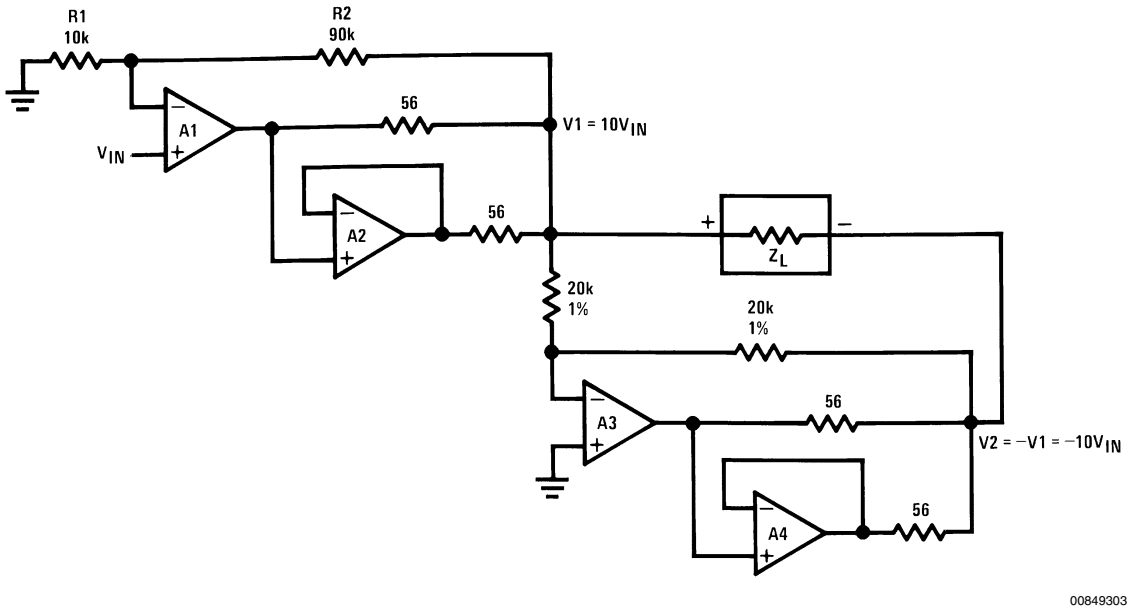

FIGURE 3. Typical Application of Load-Sharing

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT AND GENERAL COUNSEL OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.

National Semiconductor Corporation
Americas
Email: support@nsc.com

www.national.com

National Semiconductor Europe
Fax: +49 (0) 180-530 85 86
Email: europe.support@nsc.com
Deutsch Tel: +49 (0) 69 9508 6208
English Tel: +44 (0) 870 24 0 2171
Français Tel: +33 (0) 1 41 91 8790

National Semiconductor Asia Pacific Customer Response Group
Tel: 65-2544466
Fax: 65-2504466
Email: ap.support@nsc.com

National Semiconductor Japan Ltd.
Tel: 81-3-5639-7560
Fax: 81-3-5639-7507

A New Production Technique for Trimming Voltage Regulators

National Semiconductor
Linear Brief 46
Robert A. Pease
July 1979

Three-terminal adjustable voltage regulators such as the LM317 and LM337 are becoming popular for making regulated supplies in instruments and various other OEM applications. Because the regulated output voltage is easily programmed by two resistors, the designer can choose any voltage in a wide range such as 1.2V to 37V. In a typical example (*Figure 1*) the output voltage will be:

$$V_{OUT} = V_{REF}\left(\frac{R2}{R1} + 1\right) + R2 \bullet I_{ADJ}$$

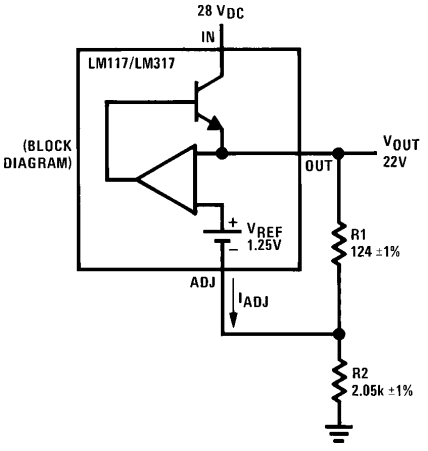

TL/H/8495–1

FIGURE 1. Basic Regulator

In many applications, when R1 and R2 are inexpensive ±1% film resistors, and the room temperature accuracy of the LM117 is better than ±3%, the overall accuracy of ±5% will be acceptable. In other cases, a tighter tolerance such as ±1% is required. Then a standard technique is to make up part of R2 with a small trim pot, as in *Figure 2*. The effective range of R2 is 2.07k ±10%, which is adequate to bring V_{OUT} to exactly 22.0V. (Note that a 200Ω rheostat in series with 1.96 kΩ ±1% would not necessarily give a ±5% trim range, because the end resistance and wiper resistance could be as high as 10Ω or 20Ω; and the maximum value of an inexpensive 10% or 20% tolerance trimmer might be as low as 180Ω or 160Ω.)

In some designs, the engineering policy may frown on the use of such trim pots, for one or more of the following reasons:

— Good trim pots are more expensive.
— Inexpensive trim pots may be drifty or unreliable.
— Any trim pot which can be adjusted can be *mis*adjusted, sooner or later.

To get a tighter accuracy on a regulated supply, while avoiding these disadvantages of trim pots, consider the scheme in *Figure 3*.

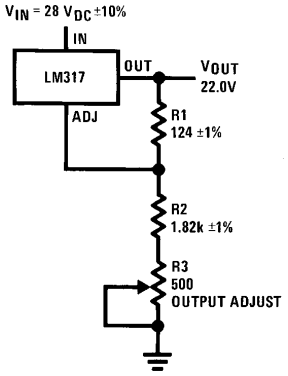

TL/H/8495–2

FIGURE 2. Regulator with Small Adjustment Range

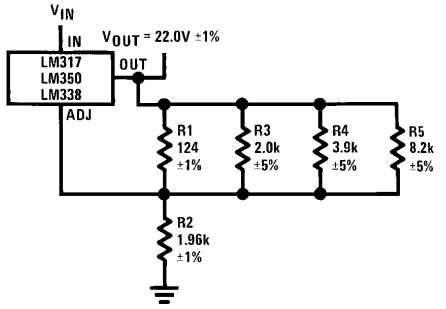

TL/H/8495–3

FIGURE 3. Regulator with Trimmable Output Voltage

When first tested, V_{OUT} will tend to be 4% to 6% higher than the 22.0V target. Then, while monitoring V_{OUT}, snip out R3, R4, and/or R5 as appropriate to bring V_{OUT} closer to 22.0V. This procedure will bring the tolerance inside ±1%:

— If V_{OUT} is 23.08V or higher, cut out R3 (if lower, don't cut it out).
— Then if V_{OUT} is 22.47V or higher, cut out R4 (if lower, don't).
— Then if V_{OUT} is 22.16V or higher, cut out R5 (if lower, don't).

The entire production distribution will be brought inside 22.0V ±1%, with a cost of 3 inexpensive carbon resistors, much lower than the cost of any pot. After the circuit is properly trimmed, it is relatively immune to being misadjusted by a screwdriver. Of course, the resistors' carcasses must be properly removed and disposed of, for full reliability to be maintained.

An alternate scheme shown in *Figure 4* has R6, R7, and R8 all shorted out initially with a stitch or jumper of wire. The

RRD-B30M115/Printed in U. S. A.

trim procedure is to open up a link to bring a resistor into effect. The advantage of this circuit is that V_{OUT} starts out *lower* than the target value, and never exceeds that voltage during trimming. In this scheme, note that a total "pot resistance" of 215Ω is plenty for a 10% trim span, because the *minimum* resistance is always below 1Ω, and the maximum resistance is always more than 200Ω—it can cover a much wider range than a 200Ω pot.

The circuit of *Figure 5* shows a combination of these trims which provides a new advantage, if a ±2% max tolerance is adequate. You may snip out R4, or link L1, or both, to accommodate the worst case tolerance, but in most cases, the output will be within spec without doing any trim work at all. This takes advantage of the fact that most ±1% resistors are well within ±⅓%, and most LM337's output voltage tolerances are between −½% and +1½%, to cut the average trim labor to a minimum. Note that L1 could be made up of a 2.7Ω ±10% resistor which may be easier to handle than a piece of wire.

In theory, a 10% total tolerance can be reduced by a factor of $(2^n - 1)$ when n binary-weighted trims are used. In practice, the factor would be $(1.8^n - 1)$ if ±10% trim resistors are used, or $(1.9^n - 1)$ if ±5% resistors are used. For n = 2, a 10% tolerance can be cut to 3.8% p-p or ±1.9%. For n = 3, the spread will be 1.7% p-p or ±0.85%, and most units will be inside ±0.5%, perfectly adequate for many regulator applications.

National Semiconductor manufactures several families of adjustable regulators including LM117, LM150, LM138, LM117HV, LM137, and LM137HV, with output capabilities from 0.5A to 5A and from 1.2V to 57V. For complete specifications and characteristics, refer to the appropriate data sheet or the 1982 Linear Databook.

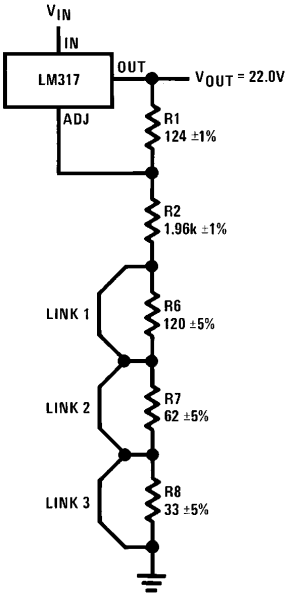

TL/H/8495–4

If V_{OUT} is lower than 20.90V, snip link 1 (if not, don't).
Then if V_{OUT} is lower than 21.55V, snip link 2 (if not, don't).
Then if V_{OUT} is lower than 21.82V, snip link 3 (if not, don't).

FIGURE 4. Alternate Trim Scheme

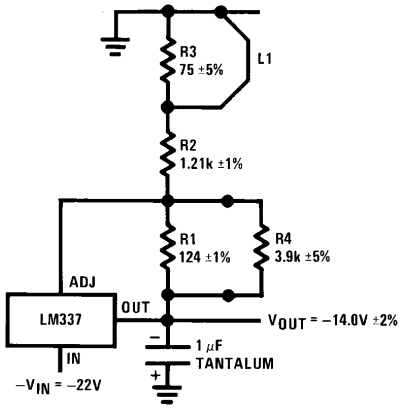

TL/H/8495–5

If $|V_{OUT}|$ is smaller than 13.75V, snip L1 and it will get bigger by 6%.
Then if $|V_{OUT}|$ is bigger than 14.20V, snip R4 and it will get smaller by 3%.

FIGURE 5. Circuit Which Usually Needs No Trim to Get V_{OUT} Within ±2% Tolerance

National Semiconductor Corporation
1111 West Bardin Road
Arlington, TX 76017
Tel: 1(800) 272-9959
Fax: 1(800) 737-7018

National Semiconductor Europe
Fax: (+49) 0-180-530 85 86
Email: cnjwge@tevm2.nsc.com
Deutsch Tel: (+49) 0-180-530 85 85
English Tel: (+49) 0-180-532 78 32
Français Tel: (+49) 0-180-532 93 58
Italiano Tel: (+49) 0-180-534 16 80

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductor Japan Ltd.
Tel: 81-043-299-2309
Fax: 81-043-299-2408

A Low-Noise Precision Op Amp

National Semiconductor
Linear Brief 52
Robert A. Pease
December 1980

It is well known that the voltage noise of an operational amplifier can be decreased by increasing the emitter current of the input stage. The signal-to-noise ratio will be improved by the increase of bias, until the base current noise begins to dominate. The optimum is found at:

$$I_{e(optimum)} = \frac{KT}{q} \frac{\sqrt{h_{FE}}}{r_s}$$

where r_s is the output resistance of the signal source. For example, in the circuit of *Figure 1*, when r_s = 1 kΩ and h_{FE} = 500, the I_e optimum is about 500 µA or 560 µA. However, at this rich current level, the DC base current will cause a significant voltage error in the base resistance, and even after cancellation, the DC drift will be significantly bigger than when I_e is smaller. In this example, I_b = 1 µA, so I_b x r_s = 1 mV. Even if the I_b and r_s are well matched at each input, it is not reasonable to expect the I_b x r_s to track better than 5 or 10 µV/˚C versus temperature.

A new amplifier, shown in *Figure 2*, operates one transistor pair at a rich current, for low noise, and a second pair at a much leaner current, for low base current. Although this looks like the familiar Darlington connection, capacitors are added so that the noise will be very low, and the DC drift is very good, too. In the example of *Figure 2*, Q2 runs at I_e = 500 µA and has very low noise. Each half of Q1 is operated at 11 µA = I_e. It will have a low base current (20 nA to 40 nA typical), and the offset current of the composite op amp, $I_{b1}-I_{b2}$, will be very small, 1 nA or 2 nA. Thus, errors caused by bias current and offset current drift vs. temperature can be quite small, less than 0.1 µV/˚C at r_s = 1000Ω.

The noise of Q1A and Q1B would normally be quite significant, about 6 $nV/\sqrt{Hz}$, but the 10 µF capacitors completely filter out the noise. At all frequencies above 10 Hz, Q2A and Q2B act as the input transistors, while Q1A and Q1B merely buffer the lowest frequency and DC signals.

For audio frequencies (20 Hz to 20 kHz) the voltage noise of this amplifier is predicted to be 1.4 $nV/\sqrt{Hz}$, which is quite small compared to the Johnson noise of the 1 kΩ source, 4.0 $nV/\sqrt{Hz}$. A noise figure of 0.7 dB is thus predicted, and has been measured and confirmed. Note that for best DC balance R6 = 976Ω is added into the feedback path, so that the total impedance seen by the op amp at its negative input is 1 kΩ. But the 976Ω is heavily bypassed, and the total Johnson noise contributed by the feedback network is below ½ $nV/\sqrt{Hz}$.

To achieve lowest drift, below 0.1 µV/˚C, R1 and R2 should, of course, be chosen to have good tracking tempco, below 5 ppm/˚C, and so should R3 and R4. When this is done, the drift referred to input will be well below 0.5 µV/˚C, and this has been confirmed, in the range +10˚C to +50˚C.

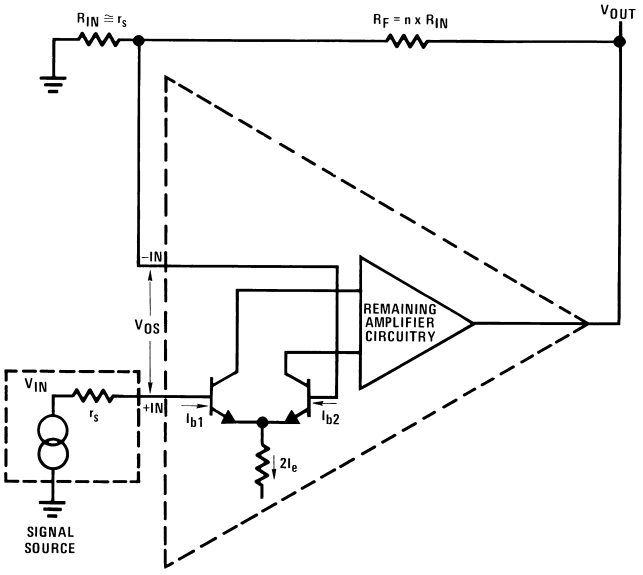

$$V_{OUT} \simeq (n+1)\, V_{IN} + V_{OS} \times (n+1) + (I_{b2} - I_{b1}) \times r_s \times (n+1) + V_{noise} \times (n+1) + i_{noise} \times (r_s + R_{IN}) \times (n+1)$$

FIGURE 1. Conventional Low-Noise Operational Amplifier

Overall, we have designed a low-noise op amp which can rival the noise of the best audio amplifiers, and at the same time exhibits drift characteristics of the best low-drift amplifiers. The amplifier has been used as a precision pre-amp (gain = 1000), and also as the output amplifier for a 20-bit DAC, where low drift and low noise are both important.

To optimize the circuit for other r_s levels, the emitter current for Q2 should be proportional to $1/\sqrt{r_s}$. The emitter current of Q1A should be about ten times the base current of Q2A. The base current of the output op amp should be no more than 1/1000 of the emitter current of Q2. The values of R1 and R2 should be the same as R7.

Various formulae for noise:

Voltage noise of a transistor,

$$\text{per } \sqrt{\text{Hz}},\ e_n = KT\sqrt{\frac{2}{qI_C}}$$

Current noise of a transistor,

$$\text{per } \sqrt{\text{Hz}},\ i_n = \sqrt{\frac{2qI_C}{h_{FE}}}$$

Voltage noise of a resistor, per $\sqrt{\text{Hz}}$, $e_n = \sqrt{4\ KTR_s}$

For a more complete analysis of low-noise amplifiers, see AN-222, “Super Matched Bipolar Transistor Pair Sets New Standards for Drift and Noise”, Carl T. Nelson.

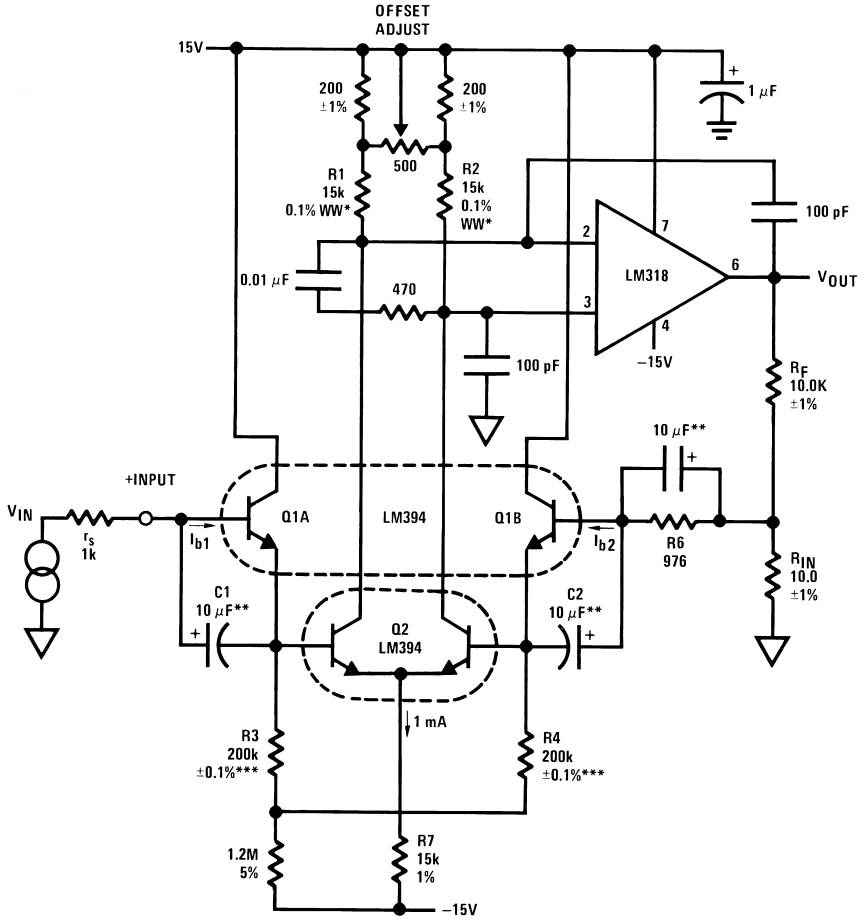

*Tracking TC < 5 ppm/°C

**Solid tantalum

***Tracking TC < 5 ppm/°C, Beckman 694-3-R100K-D or similar

FIGURE 2. New Low-Noise Precision Operational Amplifier as Gain-of-1000 Pre-Amp

Reducing DC Errors in Op Amps

National Semiconductor
Technical Paper 15
December 1980

Robert J. Widlar
Apartado Postal 541
Puerto Vallarta, Jalisco
Mexico

Abstract. *An IC op amp design that reduces bias currents below 100 pA over a −55°C to +125°C temperature range is discussed. Super-gain bipolar transistors with on-wafer trimming are used, providing low offset voltage and drift. The key to low bias current is the control of high temperature leakage currents along with the development of reasonably accurate nanoampere current sources with low parasitic capacitance.*

introduction

A bipolar replacement for the LM108 [1] drastically reduces offset voltage, bias current and temperature drift. This design, the LM11, does not depend on new technology. Instead, the improvements result from a better understanding of transistor behavior, new circuit techniques and the application of proven offset trimming methods. Table I summarizes the results obtained. The combination of low offset voltage and low bias current is unique to IC op amps, while the performance at elevated temperatures represents an advance in the state of the art.

TABLE I. Input error terms of the LM11 show an improvement over FET op amps even at room temperature. There is little degradation in performance from −55°C to 125°C. Other important specifications are somewhat better than LM108A.

Parameter	$T_j = 25°C$		$-55°C \le T_j \le +125°C$	Units
	Typ	Max	Max	
Input Offset Voltage	0.1	0.3	0.6	mV
Input Offset Current	0.5	10	30	pA
Input Bias Current	25	50	150	pA
Offset Voltage Drift	1		3	μV/°C
Offset Current Drift	20			fA/°C
Bias Current Drift	0.5		0.5	pA/°C

junction FETs

At first glance, field effect transistors seem to be the ideal input stage for an op amp, mainly because they have a low gate current, independent of their operating current. Practically, they do provide an attractive combination of performance characteristics in a relatively simple design. But there are serious shortcomings.

For one, FETs do not match as well as bipolar devices: the offset voltage is at least an order of magnitude worse. Laser trimming can compensate for this to some extent. But with FETs, low offset voltage does not guarantee low drift, as it does with bipolars. FETs are also sensitive to mechanical strains and subject to offset shifts during assembly or with temperature cycling.

Typically, long term stability is about 100 μV/year, although this can go to 1 mV/year with no prior warning in early life. This contrasts to a 10 μV/year long term stability for bipolar pairs.

Lastly, although the input current of FETs is low at room temperature, it doubles for every 10°C increase. This, coupled with high offset voltage drift, makes FETs much less attractive as operating temperature is increased.

MOS FETs

Field effect transistors, with a metal gate and oxide insulation, give the ultimate in low input current. Practically, this advantage disappears when diodes are included to protect the gate from static charges encountered in normal handling. Further, the offset voltage problems of JFETs go double for MOS FETs. They are also subject to offset shifts due to contamination.

Interesting designs are on the horizon for various chopper-stabilized complementary MOS ICs. These solve most offset voltage problems, but not that of input leakage current. Even at moderate temperatures, this input current will seriously degrade the low offset voltage and drift even with relatively low source resistances. Chopper-stabilized amplifiers have added problems with overload recovery and noise, especially with high source impedances. These problems have limited solutions, but chopper stabilization is not usually suitable for general purpose applications.

bipolar op amps

Offset voltage, its drift or long term stability has not been a serious problem with bipolar-input op amps. Such techniques as cross-coupling or zener-zap trimming have reduced offset voltage to 25 μV in production. The real problem has been bias current. The LM108, introduced in 1968, has represented the state of the art in low bias currents for standard bipolar devices. At 3 nA, maximum over temperature, the bias current is lower than FETs above 85°C.

A Darlington version of the LM108, the LM216, provided bias currents in the 50 pA range; but this design was seriously marred by high offset voltage, drift, excessive low frequency noise and anomalous leakage currents at higher temperatures.

Improvements in this design were thwarted by the inability to provide nanoampere bleed currents to stabilize the Darlington input and the erroneous belief that uncontrollable surface states created the anomalous leakage.

RRD-B30M115/Printed in U. S. A.

a new design

With bipolar transistors, there is a tradeoff between current gain and breakdown voltage. Super-gain transistors are devices that have been diffused for maximum current gain at the expense of breakdown voltage (which is typically a couple volts for a current gain of 5000). These low voltage transistors can be operated in a cascode connection with standard transistors to give a composite device with both high gain and breakdown voltage.

Figure 1 shows a modified Darlington input stage for a super-gain op amp. Common base standard transistors (Q5 and Q6, drawn with a wider base) are bootstrapped to the super-gain input transistors so that the latter are operated at near zero collector base voltage. In addition to permitting the use of super-gain inputs, this connection also isolates the input transistors from common-mode variations, increasing common-mode rejection.

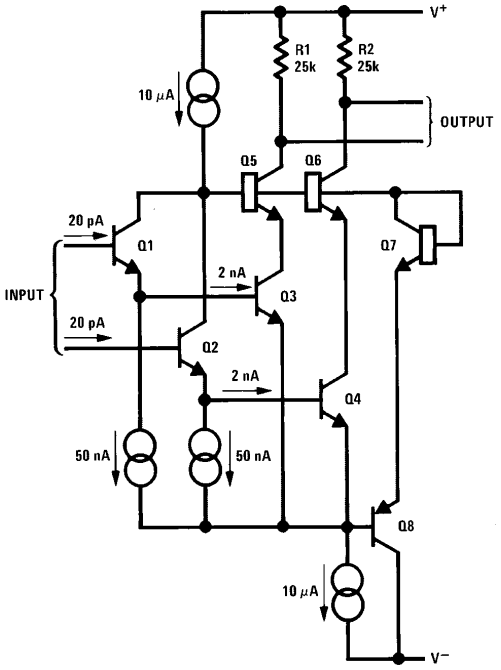

TL/H/8722–1

Figure 1. Bootstrapped input stage using super-gain transistors in modified-Darlington connection. The objectionable characteristics of the Darlington are virtually eliminated by operating the input transistors at a much larger current than the base current of the transistors they are driving.

The usual problems with the Darlington connection are avoided by providing a bleed current that operates the input transistors, Q1 and Q2, at a current much higher than the base current of the transistors they are driving, Q3 and Q4. This is necessary because the base currents are not that well matched, especially over temperature, and have excess low frequency noise.

a nanoampere current source

A circuit that generates the 50 nA bleed current is shown in *Figure 2*. A super-gain transistor operated in the forward mode is used to bias a standard transistor in the reverse mode. The reverse connection is used because the capacitance of an ordinary collector tub would reduce the common-mode slew rate from 2 V/μs to 0.02 V/μs.

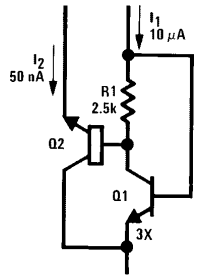

$$I_2 = \frac{h_{fe2}A_2I_1}{h_{fe1}A_1} \exp \frac{qI_1R_1}{kT}$$

TL/H/8722–2

Figure 2. Forming a nanoampere current source with low parasitic capacitance. Design takes advantage of predictable V_{BE} difference between standard and super-gain transistors and fact that V_{BE} of a transistor is the same when operated in forward or reverse mode.

At first look, this biasing scheme would seem to be subject to a number of process variations. This is not so. For one, the V_{BE} of a transistor depends on the base Gummel number ($Q_B/\overline{\mu}_B$), the number of majority carriers per unit area divided by their effective mobility. Since the Gummel number and the effective area are unchanged when the collector and emitter are interchanged, the V_{BE} will be the same in either connection, provided that base recombination is not excessive. In standard IC transistors, reverse h_{fe} is about 30, indicating that recombination is not a significant factor. Measured reverse h_{fe} is much lower, but this is the result of a parasitic PNP that does not affect V_{BE} or α_E, the common base current gain.

The bleed current depends also on the ratio of super-gain to standard transistor h_{fe}, as indicated by the equation in *Figure 2*. Intuition suggests that super-gain h_{fe} will increase much faster than standard transistor h_{fe} with increasing emitter diffusion time, giving lower bleed current with higher super-gain h_{fe}. However, measurements with variances of standard LM108 processing indicate that the bleed current remains within 25% of design center.

As shown in *Figure 2*, higher current ratios can be obtained by increasing the area of Q1 relative to Q2 or by including R1. The equation in *Figure 2* assumes that I_1 varies as absolute temperature. If the voltage drop across R1 is equal to kT/q, changes in the V_{BE} of Q1 with small changes in I_1 will be cancelled by changes in the voltage drop across R1. This makes input bias current essentially unaffected by variations in supply or common-mode voltage as long as I_1 is reasonably well controlled.

leakage currents

The input leakage currents of bipolar op amps can be kept under control because small geometry devices are satisfactory and because the collector-base junction can be operated at an arbitrarily low voltage if bootstrapping is used.

Simple theory predicts that bulk leakage saturates for reverse biases above 2kT/q. But generation in the depletion zone dominates below 125°C. Because the depletion width varies with reverse bias, so does leakage. The characteristics of a high quality junction plotted in *Figure 3* show that leakage current can be reduced with lower bias.

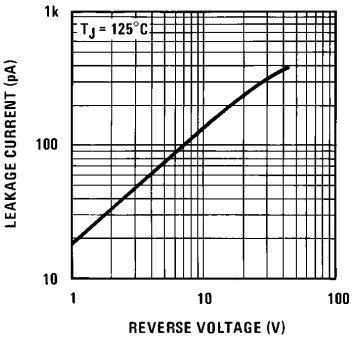

TL/H/8722–3

Figure 3. Voltage sensitivity of collector base leakage indicates that generation in the depletion zone dominates even at 125°C.

When more than one junction is involved, minimum leakage is not necessarily obtained for zero bias. This is illustrated in *Figure 4*, a plot of I_{CBO} for a junction isolated NPN transistor. A parasitic PNP is formed between the base and the isolation as diagrammed in the inset. Zero leakage is obtained when V_{CB} is set so that the PNP diffusion current equals I_{CBO} of the NPN.

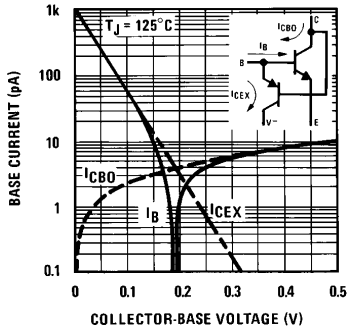

TL/H/8722–4

Figure 4. Plot above explains "anomalous" leakage of NPN transistors in ICs. As collector base bias is reduced, base current reverses then increases exponentially. This excess current is the forward diffusion current of parasitic PNP to substrate (see inset).

input protection

The input clamps perform a dual function. Most important, they protect the emitter base junction of the input transistors from damage by in-circuit overloads or static charges in handling. Secondly, they limit the voltage change across junction capacitances on low current nodes under transient conditions. This minimizes recovery delays.

The clamp circuitry is shown in *Figure 5*. Emitters are added on the input transistors and cross-coupled to limit the differential input voltage. Another transistor, Q5, has been added to limit voltage on the input transistors if the inputs are driven below V^-.

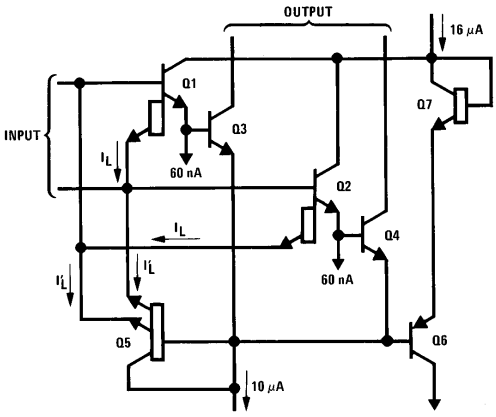

TL/H/8722–5

Figure 5. Separate clamps are used for differential and common-mode overloads. Leakage currents, I_{CES} of forward and reverse connected transistors, cancel.

The differential clamp transitors do contribute to input current because $V_{CB} > 0$, so collector current is not zero for $V_{BE} \cong 0$ ($I_{CES} \cong 100$ pA at 125°C). The common-mode input clamp, Q5, is also operated at $V_{BE} = 0$ and $V_{CB} > 0$, although in the inverted mode. The resulting error is diffusion current, dependent only on the characteristic V_{BE} of the transistors. Thus, the current contributed by the differential clamp transistors is cancelled, within a couple percent, by that from the common-mode clamp.

bias current

Figure 6 shows some results of the design approach described here. A room temperature bias current of 25 pA is obtained, and this his held to 60 pA over a −55°C to 125°C temperature range. The figure also shows the results of

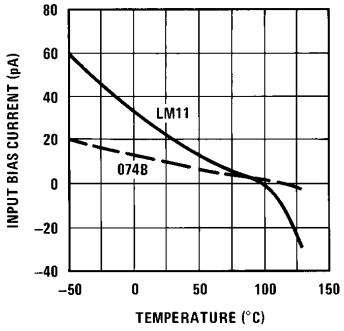

TL/H/8722–6

Figure 6. Input bias current of the LM11 remains low over military temperature range. Improvements in development give even better results (074B). Offset current is usually below 1 pA.

some improvements in development that have reduced bias current to 20 pA over the full operating temperature range.

Figure 6 shows that bias current is very nearly a linear function of temperature, at least from −55°C to +100°C. This, coupled with the fact that bias current is virtually unaffected by changes in common-mode or supply voltage, suggests that bias current compensation can be provided for critical applications. An appropriate circuit is shown in *Figure 7*. Details are given in reference [2], but properly set up it should be possible to hold bias currents to less than 20 pA over a −55°C to +100°C temperature range or 5 pA over a 15°C to 55°C range with a simple room temperature adjustment.

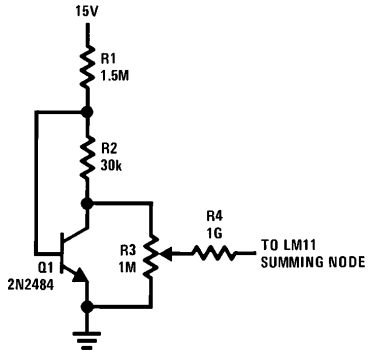

TL/H/8722–7

Figure 7. Bias current of LM11 varies linearly with temperature so it can be effectively compensated with this circuit. Bias currents less than 5 pA over 15°C to 55°C range or 20 pA over −55°C to +100°C are practical.

noise

The broadband noise of a bipolar transistor is given by

$$e_n = kT\sqrt{2\Delta f/qI_c}. \quad (1)$$

Therefore, operating the input transistors at low collector current does increase noise. Because the noise of most op amps is greater than the theoretical noise voltage of the input transistors, the noise increase from low current input buffers is not as great as might be expected. In addition,

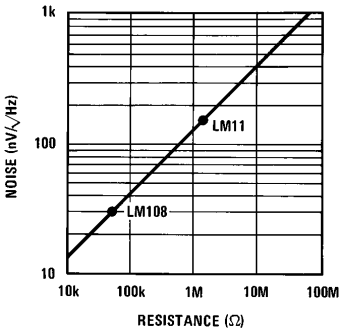

TL/H/8722–8

Figure 8. Increased noise of LM11 is consequence of low collector current in input transistors. But in high impedance applications, op amp noise is masked by the thermal noise of source resistance given above.

when operating from higher source resistances, op amp noise is obscured by resistor noise, as shown in *Figure 8*.

Low frequency noise is not as easily accounted for as broadband noise, but lower operating currents increase noise in much the same fashion. The low frequency noise of the LM11, shown in *Figure 9*, is a bit less than FETs but greater than that of the LM108 when it is operated from source resistances less than 500 kΩ.

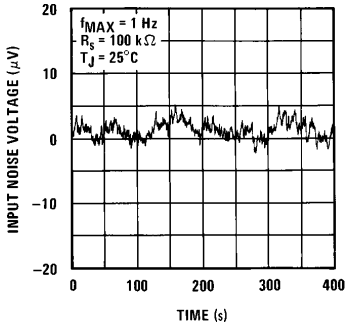

TL/H/8722–9

Figure 9. Low frequency noise of LM11 is high compared to other bipolar devices but somewhat less than FETs. It is equal to LM108 operating from 500 kΩ source resistances.

complete circuit

A schematic diagram of an IC op amp using the techniques described is shown in *Figure 10*. Other than the input stage, the circuitry is much like the LM112, a compensated version of the LM108 that includes offset balancing.

One significant change has been the inclusion of wafer level trimming for offset voltage. This is done using zener-zap trimming across portions of the input stage collector load resistors, R4 and R5. This kind of zener is simply the emitter base junction of an NPN transistor. When pulsed with a large reverse current at wafer sort, the junction is destroyed by the formation of a low resistance filament between the emitter and base contact beneath the protective oxide. This shorts out a portion of the collector load resistor. The process is repeated on binary weighted segments until the offset voltage has been minimized.

Offset voltage of the LM11 is conservatively specified at 300 μV. Although low enough for most applications, offset voltage trimming is provided for fine adjustment. Balance range is determined by the resistance of the balance potentiometer, varying from ±5 mV at 100 kΩ to ±400 μV at 1 kΩ. Incidentally, when nulling offset voltages of 300 μV, the thermal matching of balance-pot resistance to the internal resistors is not a significant factor.

The actual balancing is done on the emitters of lateral PNP transistors, Q9 and Q10, that imbalance the collector loads of the input stage. This particular arrangement was used so that no damage would result from accidental connection of the balance pins to voltages outside either supply. Not obvious is that a balance pin voltage 15V more negative than V^+ can effectively short these PNP transistors with a parallel P-channel MOS transistor, forcing the output to one limit or another.

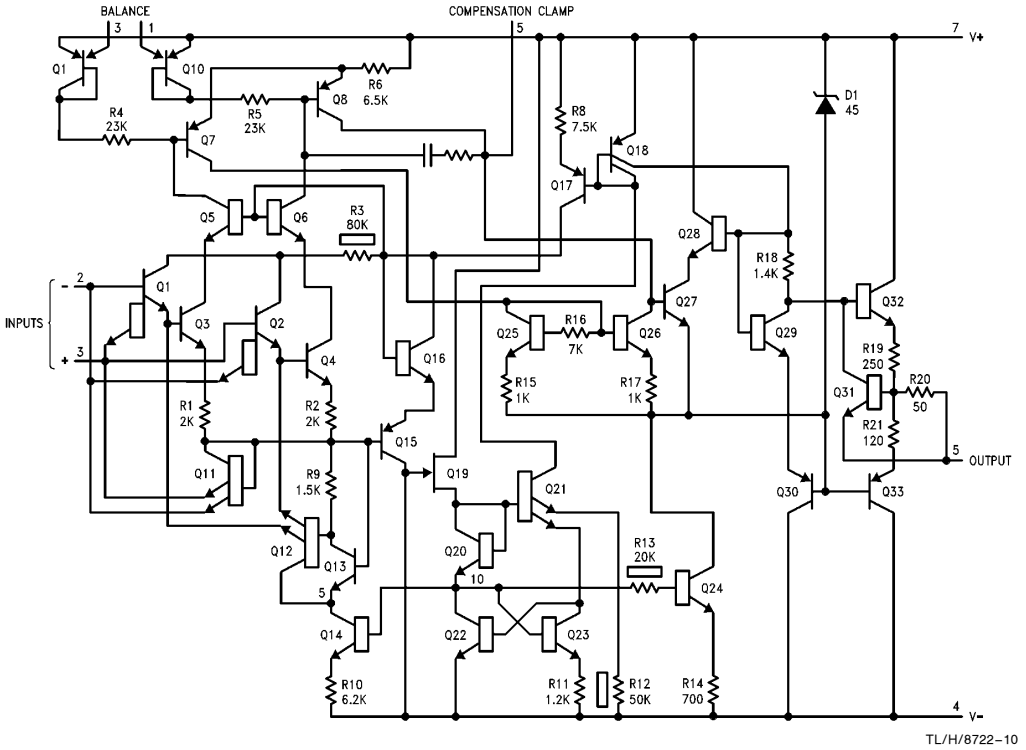

TL/H/8722–10

Figure 10. Complete schematic of the LM11. Except for the input stage, circuit is much like the LM112, a compensated version of the LM108 that includes offset balancing.

Although the LM11 is specified to a lower voltage than the LM108, the minimum common-mode voltage is a diode drop further from V^- because the bleed current generator, Q12 and Q13, has been added.

Proceeding from the input stage, the second stage amplifier is a differential pair of lateral PNPs, Q7 and Q8. These feed a current mirror, Q25 and Q26, which drive a super-gain follower, Q27. The collector base voltage of Q26 is kept near zero by including Q28. The current mirror is bootstrapped to the output so that second stage gain error depends only on how well Q7 and Q8 match with changes in output voltage. This gives a gain of 120 dB in a two stage amplifier. Frequency compensation is provided by MOS capacitor C1.

The output stage is a complementary class-B design with current limiting. Biasing has been altered so that the guaranteed output current is twice the LM108. A zener diode, D1, limits output voltage swing to prevent stressing the MOS capacitor to the point of catastrophic failure in the event of gross supply transients.

The main bias current generator design (Q20–Q23) is due to Dobkin [3]. It is powered by Q19, a collector FET. The circuit is auto-compensated so that output current of Q14 and Q21 varies as absolute temperature and changes by less than 1% for a 100:1 shift in Q19 current.

speed

With a unity gain bandwidth of 500 kHz and a 0.3 V/μs slew rate the LM11 is not fast. But it is no slower than might be expected for a supply current of only 300 μA.

If the precision of the LM11 is required along with greater speed, the circuit in *Figure 11* might be used. Here, the LM11 senses input voltage and makes appropriate adjustments to the balance terminals of a fast FET amplifier. The main signal path is through the fast amplifier.

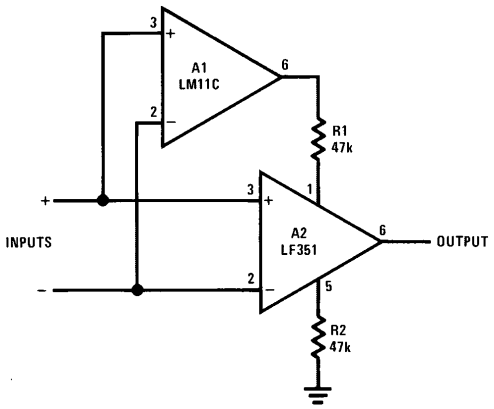

TL/H/8722–11

Figure 11. The LM11 can zero offset of fast FET op amp in either inverting or non-inverting configurations. Speed is that of fast amplifier. FET amplifier can be capacitively coupled to critical input to eliminate its leakage current.

Surprisingly, this connection will work even as a voltage follower. The common-mode slew recovery of the LM11 is about 10 μs to 1 mV, even for 30V excursions. This was accomplished by minimizing or bootstrapping stray capacitances and providing clamping to limit the voltage excursion across the strays.

When bias current is an important consideration, it will be advisable to ac couple the FET op amp to the critical input. Reference [2] discusses this and other practical aspects of fast operation with the LM11.

conclusions

A new IC op amp has been described that can not only increase the performance of existing equipment but also creates new design possibilities. Op amp error has been reduced to the point where other problems can dominate. Many of the practical difficulties encountered in high impedance circuitry are discussed in reference [4] along with solutions. A number of tested designs using these techniques are given in reference [2].

The LM11 is not the result of any breakthrough in processing technology. It is simply a modification of ICs that have been in volume production for over 10 years. The improvements have resulted primarily from an understanding of strange behavior observed on the earlier ICs and taking advantage of certain inherent characteristics of bipolar transistors that were not fully appreciated.

As users of the LM11 may have discovered, the offset voltage and bias current specifications are quite conservative. It seems possible to offer 50 μV offset voltage and perhaps 1 μV/°C drift even on low cost parts. Taking full advantage of 5 pA bias current would require guarded 10-pin TO-5 packages or 14-pin DIP packages. Further, the feasibility of reducing low frequency noise to 2 μV and 0.1 pA, peak to peak, has been demonstrated on prototype parts.

acknowledgement

The author would like to acknowledge the contributions of Dennis Foltz for solving the rather formidable production test problems of the LM11.

references

[1] R.J. Widlar, ''IC op amp beats FETs on input current'', National Semiconductor AN-29, December 1969.

[2] R.J. Widlar, R. Pease and M. Yamatake, ''Applying a new precision op amp'', National Semiconductor AN-242, April 1980.

[3] R. Dobkin, U.S. patent no. 3930172.

[4] R.J. Widlar, ''Working with high impedance op amps'', National Semiconductor AN-241, February 1980.

National Semiconductor Corporation
1111 West Bardin Road
Arlington, TX 76017
Tel: 1(800) 272-9959
Fax: 1(800) 737-7018

National Semiconductor Europe
Fax: (+49) 0-180-530 85 86
Email: cnjwge@tevm2.nsc.com
Deutsch Tel: (+49) 0-180-530 85 85
English Tel: (+49) 0-180-532 78 32
Français Tel: (+49) 0-180-532 93 58
Italiano Tel: (+49) 0-180-534 16 80

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semiconductor Japan Ltd.
Tel: 81-043-299-2309
Fax: 81-043-299-2408