A Tribute to Jim Williams

Book Chapters

(1991-2008)

23. The Zoo Circuit

History, Mistakes, and Some Monkeys Design a Circuit

This chapter is dedicated to the memory of Professor Jerrold R. Zacharias, who saved my ass.

A couple of years ago, I was asked to design a circuit for a customer. The requirements were not trivial, and the customer was having difficulty. I worked on this problem for some time and was asked to present my solution in a formal design review at the customer's location.

When I say "formal," I mean it! I came expecting to talk circuits with a few guys over a pizza. Upon arrival, I was taken to a large and very grand room, reminiscent of a movie theater. About 150 engineers were in attendence. There was every audio-visual machine known to humanity at the ready, and I was almost embarrassed to report that I had no slides, overheads, charts, or whatever (although a piece of chalk would be nice). A "senior technical management panel," positioned in a boxed-off section adjacent to the lectern, was to present a prepared list of questions. A video camera duly recorded the proceedings. The whole thing was chaired by somebody who introduced himself as "Dr. So-and-So, senior vice-president of engineering." Everybody in the place talked in whispers and nodded his head a lot. I found myself alternating between intimidation and amusement.

I gave a fairly stiff presentation, clutching my dear little piece of chalk the whole time. Things seemed to go okay, but not great, and then the panel began with their prepared list of questions. The first question went something like, "Can you explain, precisely, where the ideas for this and that piece of the circuit came from? Can you detail what design procedures, programs, and methodologies were helpful?"

I considered various acceptable answers, but decided to simply tell the truth: "Most of the ideas came from history, making mistakes, and the best source of help was some monkeys at the San Francisco Zoo."

You could have heard a pin before it dropped. There was absolute silence for a bit, and then some guy stood up and asked me to elaborate "a little." Everybody cracked up, the mood shifted, and we finally began to really *talk* about the circuit.

This customer originally came to me with a need for a "CMOS voltage-to-frequency converter." The performance requirements were as follows:

Output frequency	0–10 kHz
Input voltage	0–5 V
Linearity	0.04%
Drift	100 ppm/°C
PSRR	100 ppm/V
Temperature range	0°–55°C
Step response	< 5 cycles of output frequency
Output pulse	5 V CMOS-compatible
Power supply	Single 9 V battery (6.5–10 V)
Power consumption	200 μA maximum
Cost	< $6.00/100,000 pieces

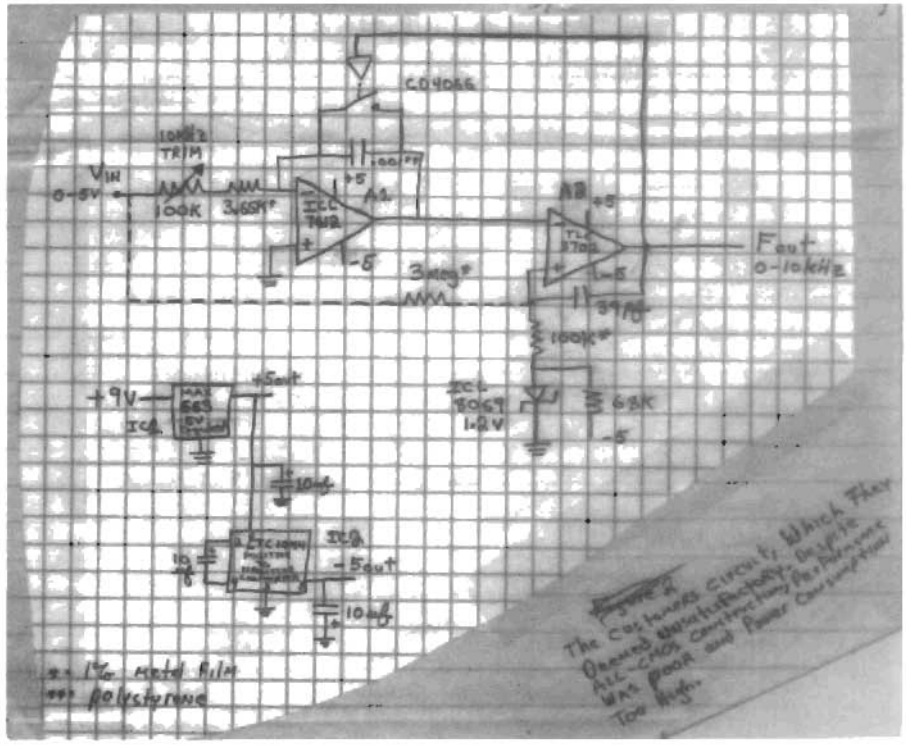

Figure 23-1. The customer's circuit, which was deemed unsatisfactory. Despite all-CMOS construction, performance was poor and power consumption too high.

These people had been working on a design for several months. It functioned, but was described as wholly unsatisfactory. I asked why they needed CMOS and was assured that "the low power requirement is nonnegotiable." Without further comment, I asked them to send me their breadboard. It arrived the next morning, and looked like Figure 23-1.

This is probably the most obvious way to design a V/F converter. The 9 V battery is regulated to 5 V by IC1 and a –5 V rail is derived by IC2. The input voltage causes current flow into A1's summing point. A1 responds by integrating negative, as shown in Figure 23-2, trace A. When A1's output goes low enough, A2 trips high (see trace B in Figure 23-2), turning on the CD4066 switch and resetting the integrator. Local positive feedback around A2 (A2's positive input is trace C) "hangs up" the reset, ensuring a complete integrator discharge. When the positive feedback decays, A1 begins to ramp again. The ramp slope, and hence the repetition frequency, depends upon the input voltage-dependent current into A1's summing point.

As soon as I saw the schematic, I knew I couldn't salvage any portion of this design. A serious drawback to this approach is A1's integrator reset time. This time, "lost" in the integration, results in significant linearity error as the operating frequency approaches it. The circuit's 6 µsec reset (see Figure 23-2, traces A and B) interval introduces a 0.6% error at 1 kHz, rising to 6% at 10 kHz. Also, variations in the reset time contribute additional errors. I added the 3 M resistor (shown in dashed lines) in a half-hearted attempt to improve these figures. This resistor causes A2's trip point to vary slightly with input, partially compensating for the integrator's "lost" reset time. This Band-Aid did improve linearity by more than an order of magnitude, to about 0.4%, but it ain't the way to go.

There are other problems. Quiescent current consumption of this entirely CMOS circuit is 190 µA, rising to a monstrous 700 µA at 10 kHz. Additionally, the polystyrene capacitor's drift alone is –120 ppm/°C, eating up the entire budget. The 1.2

A = 0.5 V/Div.
B = 10 V/Div.
C = 10 V/Div.
Horiz. = 10 μsec/Div.

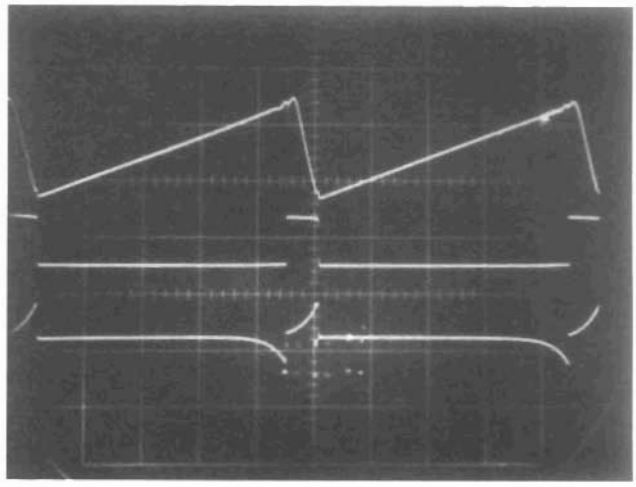

Figure 23-2. Wave forms for Figure 23-1's circuit. Finite reset time prevents good linearity performance.

V reference and the input resistor-trimmer could easily double this figure. There are a host of other problems, but what is really needed is an approach with inherently better linearity and lower power consumption.

There are many ways to convert a voltage to a frequency. The "best" approach in an application varies with desired precision, speed, response time, dynamic range, and other considerations.

Figure 23-3's concept potentially achieves high linearity by enclosing Figure 23-1's integrator in a charge-dispensing loop.

In this approach, C2 charges to $-V_{ref}$ during the integrator's ramping time. When the comparator trips, C2 is discharged into A1's summing point, forcing its output high. After C2's discharge, A1 begins to ramp and the cycle repeats. Because the loop acts to force the average summing currents to zero, the integrator time constant and reset time do not affect frequency. Gain drift terms are V_{ref}, C2, and the input resistor. This approach yields high linearity (typically 0.01%) into the megahertz range.

Figure 23-4 is conceptually similar, except that it uses feedback current instead of charge to maintain the op amp's summing point. Each time the op amp's output trips the comparator, the current sink pulls current from the summing point. Current is pulled from the summing point for the timing reference's duration, forcing the integrator positive. At the end of the current sink's period, the integrators output again heads negative. The frequency of this action is input related.

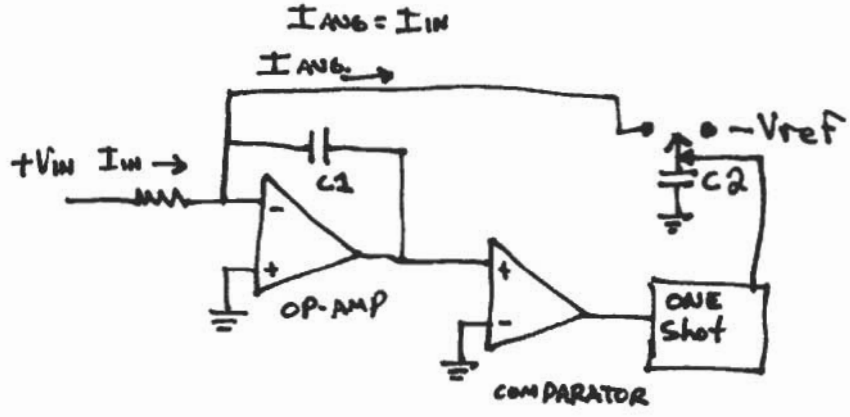

Figure 23-3. Conceptual charge-dispensing type voltage-to-frequency converter.

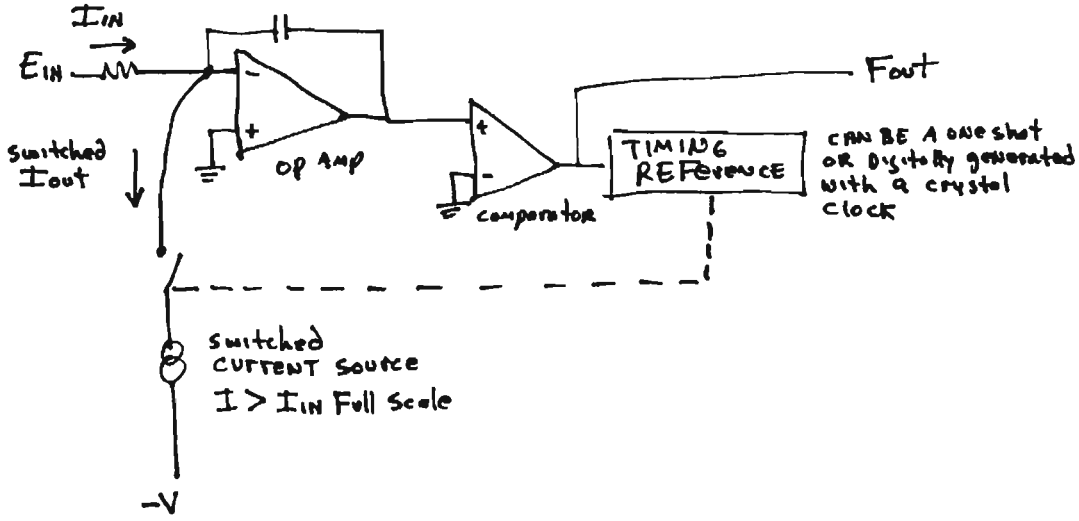

Figure 23-4. Current balance voltage-to-frequency converter.

Figure 23-5 uses DC loop correction. This arrangement offers all the advantages of charge and current balancing except that response time is slower. Additionally, it can achieve exceptionally high linearity (0.001%), output speeds exceeding 100 MHz, and very wide dynamic range (160 dB). The DC amplifier controls a relatively crude V/F converter. This V/F converter is designed for high speed and wide dynamic range at the expense of linearity and thermal stability. The circuit's output switches a charge pump whose output, integrated to DC, is compared to the input voltage.

The DC amplifier forces the V/F converter operating frequency to be a direct function of input voltage. The DC amplifier's frequency compensation capacitor, required because of loop delays, limits response time. Figure 23-6 is similar, except that the charge pump is replaced by digital counters, a quartz time base, and a DAC. Although it is not immediately obvious, this circuit's resolution is not restricted by the DAC's quantizing limitations. The loop forces the DAC's LSB to oscillate around the ideal value. These oscillations are integrated to DC in the loop compensation capacitor. Hence, the circuit will track input shifts much smaller than a DAC LSB. Typically, a 12-bit DAC (4096 steps) will yield one part on 50,000 resolution. Circuit linearity, however, is set by the DAC's specification.

Figure 23-5. Loop-charge pump voltage-to-frequency converter.

If you examine these options, Figure 23-3 looks like the winner for the customer's application. The specifications call for step response inside 5 cycles of output fre-

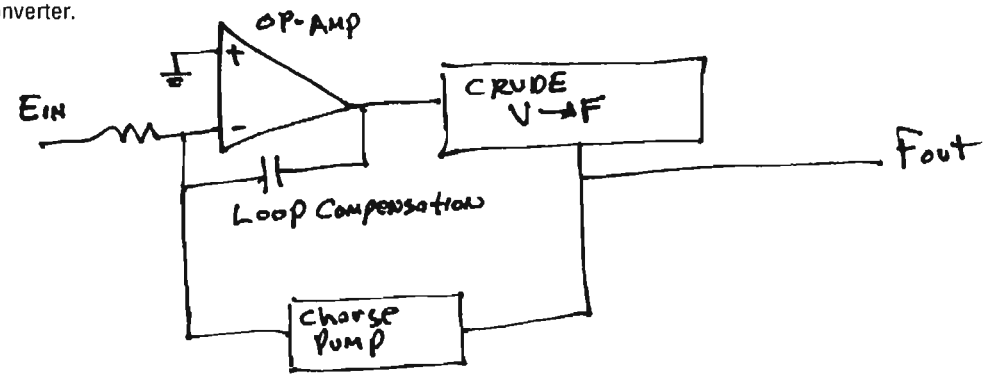

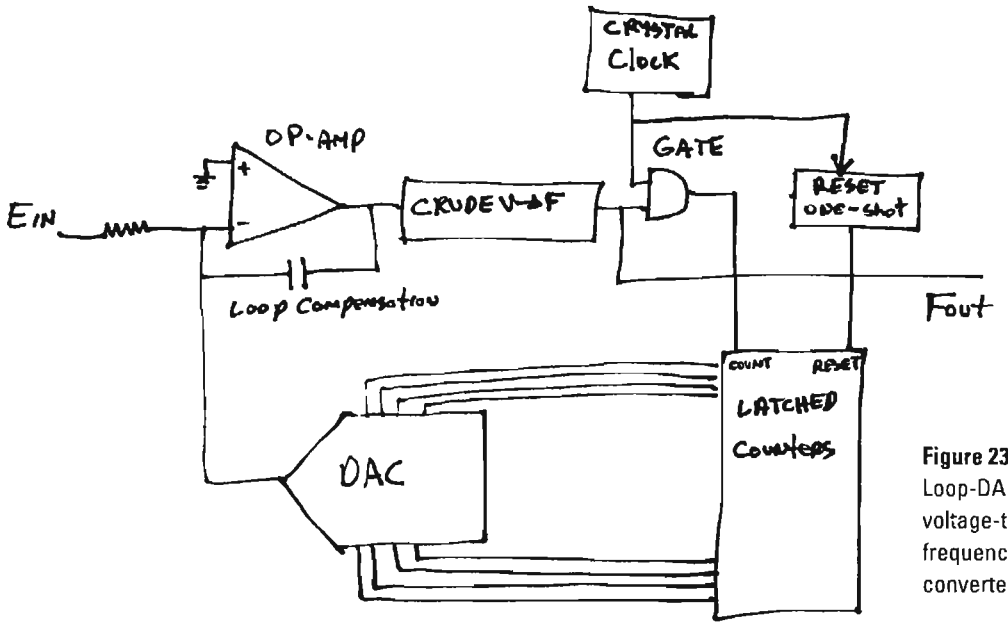

Figure 23-6. Loop-DAC voltage-to-frequency converter.

quency. This eliminates the circuits in Figures 23-4, 23-5, and 23-6 with their DC amplifiers' response time lag. Figure 23-4 requires a timing reference and a precision switched current source, implying some degree of complexity. In theory, Figure 23-3's approach can meet all the specifications without undue complexity.

This technique is not new. I first saw it back in 1964 in a copy of the *GE Transistor Manual*. T. P. Sylvan used a discrete op amp and a unijunction transistor to form the loop. Hewlett-Packard built rack-mounting V/F converters in the early 1960s which also relied on this approach. In 1972, R.A. Pease developed a commercially produced modular version (Teledyne-Philbrick Model 4701) using a single op amp which routinely achieved 0.01% linearity with commensurate drift performance. Pease's circuit is particularly relevant, and a version of it is shown in Figure 23-7.

Assume C1 sits at a small negative potential. A1's negative input is below its zero-biased positive input, and its output is high. The zener bridge clamps high (at $V_Z + V_{D4} + V_{D2}$) and C2 charges via D6, D7, and D8. The input voltage forces current through R1, and C1 begins to charge positively (trace A, Figure 23-8). When C1 crosses zero volts, A1's output (trace B) goes low and the zener bridge clamps negative, discharging C2 (C2's current is trace C) via the D5–C1 path. The resultant charge removal from C1 causes it to rapidly discharge (trace A). R2–C3 provides positive feedback to A1's positive input (trace D), reinforcing this action and hanging up A1's output long enough for a complete C2 discharge. When the R2–C3 feedback decays, A1's output returns high and the cycle repeats. The frequency of this sequence is directly proportional to the input voltage derived current through R1. Drift terms include R1, C2, and the zener, as well as residual diode mismatches. In theory, all the diode drops cancel and do not contribute toward drift. The R2–C3 "one shot" time constant is not critical, as long as it allows enough time for C2 to completely discharge. Similarly, "integrator" C1's value is unimportant as long as it averages A1's negative input to zero.

Q1 and associated components form a start-up loop. Circuit start-up or input overdrive can cause the circuit's AC-coupled feedback to latch. If this occurs, A1 goes negative and wants to stay there. R3 and C4 slowly charge negative, biasing

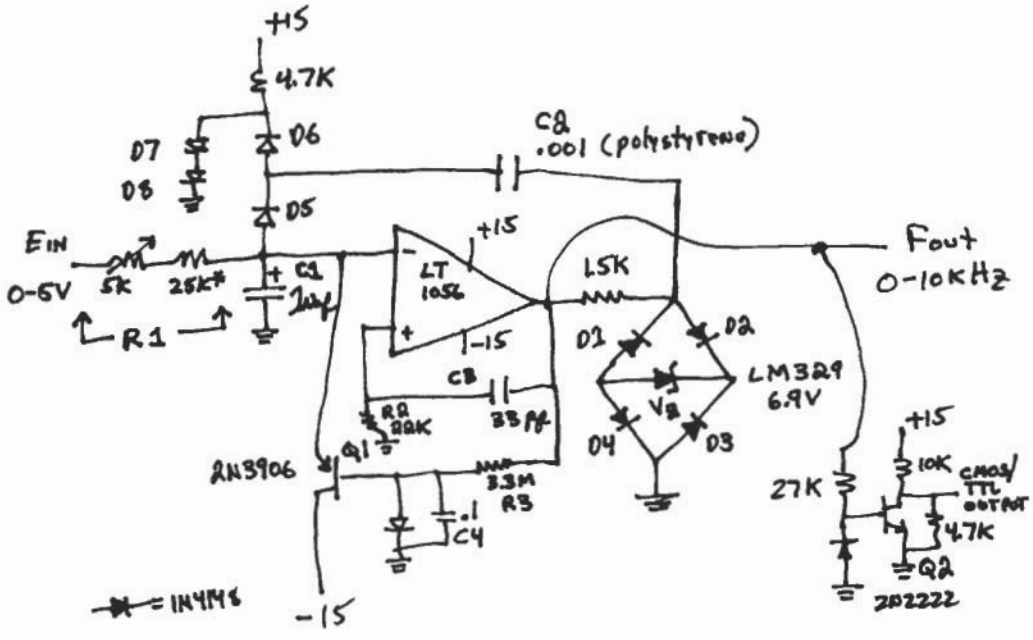

Figure 23-7. A version of Pease's elegant voltage-to-frequency converter circuit.

Q1. Q1 turns on, pulling C1 toward the –15 V rail, initiating normal circuit action. Once the circuit starts, C4 assumes a small positive potential and Q1 goes off. Q2, a simple level shifter, furnishes a logic-compatible output.

Pease's 1972 circuit is a very elegant, practical incarnation of Figure 23-3. With care, it will meet all the customer's requirements except two. It requires a split ±15 V supply, and pulls well over 10 mA. The job now boils down to dealing with these issues.

Figure 23-9 shows my first attempt at adapting Pease's circuit to my customer's needs. Operation is similar to Pease's circuit. When the input current-derived ramp (trace A, Figure 23-10) at C1A's negative input crosses zero, C1A's output (trace B) drops low, pulling charge through C1. This forces the negative input below zero. C2 provides positive feedback (trace D is the positive input), allowing a complete discharge for C1 (C1 current is trace C). When C2 decays, C1A's output goes high, clamping at the level set by D1, D2, and V_{ref}. C1 receives charge, and recycling

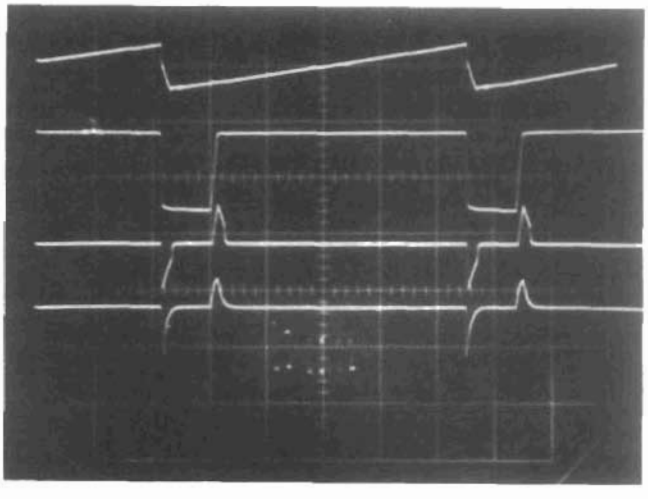

Figure 23-8. Wave forms for the Pease-type voltage-to-frequency converter.

A = 0.02 V/Div.
B = 20 V/Div.
C = 20 mA/Div.
D = 20 V/Div.
Horiz. = 20 μsec/Div.

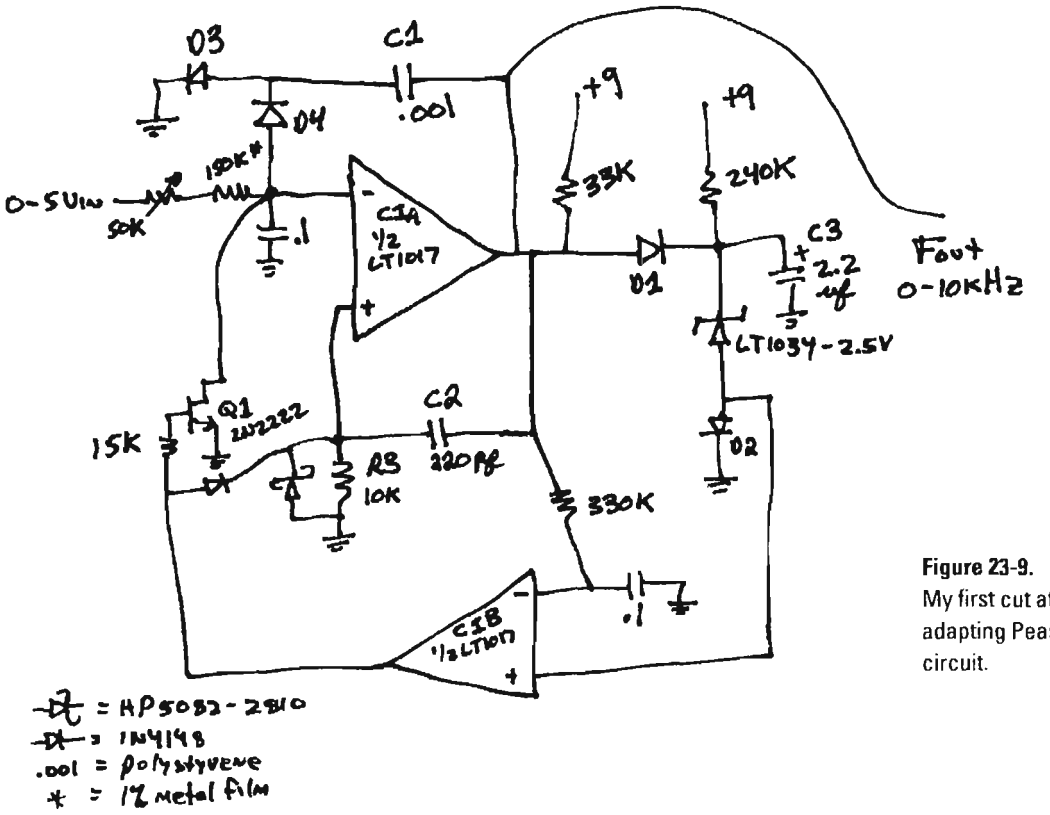

Figure 23-9. My first cut at adapting Pease's circuit.

occurs when C1A's negative input again arrives at zero. The frequency of this action is related to the input voltage. Diodes D3 and D4 provide steering and are temperature compensated by D1 and D2. C1A's sink saturation voltage is uncompensated but small. (These temperature coefficient assumptions are first order and will require more care later.) Although the LT1017 and LT1034 have low operating currents, this circuit pulls almost 400 µA. The AC current paths include C1's charge-discharge cycle, and C2's branch. The DC path through D2 and V_{ref} is particularly costly. C1's charging must occur quickly enough for 10 kHz operation, meaning the clamp seen by C1A's output must have low impedance at this frequency. C3 helps, but significant current still must come from somewhere to keep impedance low. C1A's current-limited output (≈30 µA source) cannot do the job unaided, and the resistor from the supply is required. Even if C1A could supply the necessary current, V_{ref}'s settling time would be an issue. Dropping C1's value will reduce impedance requirements proportionally and would seem to solve the problem. Unfortunately, such reduction magnifies the effects of stray capacitance at the D3–D4 junction. It also mandates increasing R_{in}'s value to keep scale factor constant. This lowers operating currents at C1A's negative input, making bias current and offset more significant error sources.

C1B, Q1, and associated components form a start-up loop which operates in similar fashion to the one in Pease's circuit (Figure 23-7).

Figure 23-11 shows an initial attempt at dealing with these issues. This scheme is similar to Figure 23-9, except that Q1 and Q2 appear. V_{ref} receives switched bias via Q1, instead of being on all the time. Q2 provides the sink path for C1. These transistors invert C1A's output, so its input pin assignments are exchanged. R1 provides a light current from the supply, improving reference settling time. This

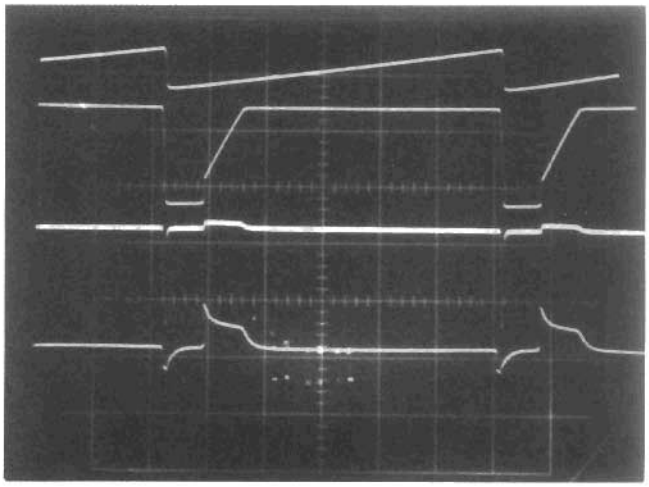

Figure 23-10. Wave forms for the circuit in Figure 23-9.

A = 50 mV/Div.
B = 2 V/Div.
C = 2 mA/DIv.
D = 1 V/Div.
Horiz. = 20 μsec/Div.

arrangement decreases supply current to about 300 μA, a significant improvement. Several problems do exist, however. Q1's switched operation is really effective only at higher frequencies. In the lower ranges, C1A's output is low most of the time, biasing Q1 on and wasting power. Additionally, when C1A's output switches, Q1 and Q2 simultaneously conduct during the transition, effectively shunting R2 across the supply. Finally, the base currents of both transistors flow to ground and are lost. Figure 23-12 shows the wave form traces for this circuit. The basic temperature compensation is as before, except that Q2's saturation term replaces the comparator's. This temperature compensation scheme looks okay, but we're still hand waving.

Figure 23-11. The second try. Q1 and Q2 switch the reference, saving some power.

Figure 23-13 is better. Q1 is gone, Q2 remains, but Q3, Q4, and Q5 have been

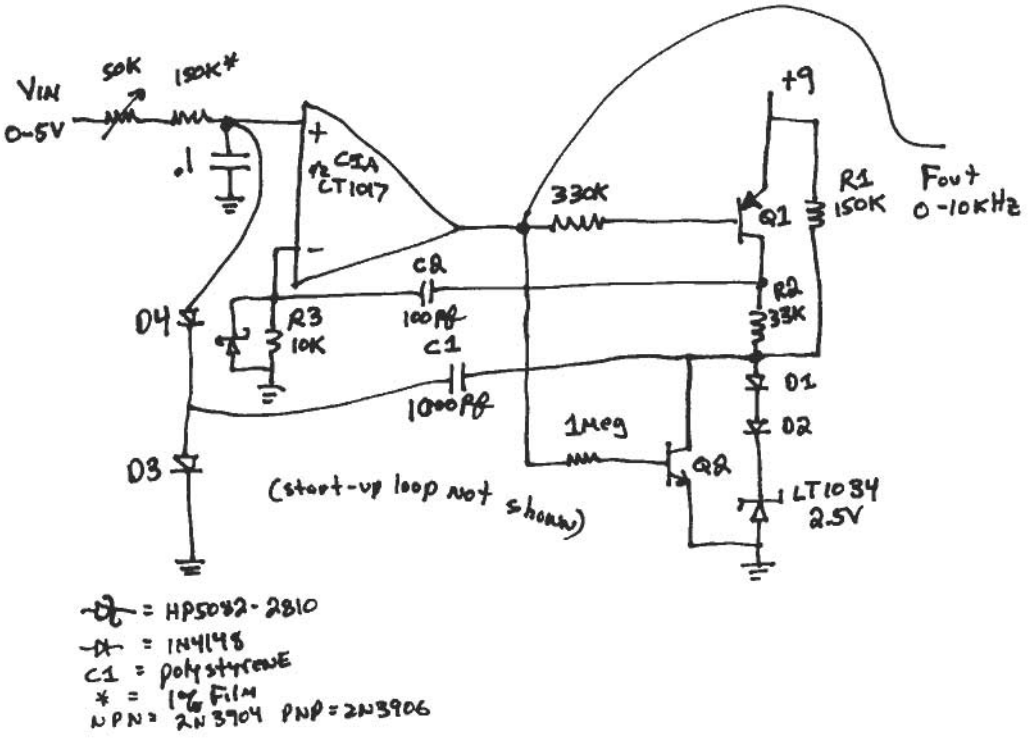

A = 5 V/Div.
B = 5 V/Div.
C = 2 V/Div.
D = 100 µA/Div.
Horiz. = 10 µsec/Div.

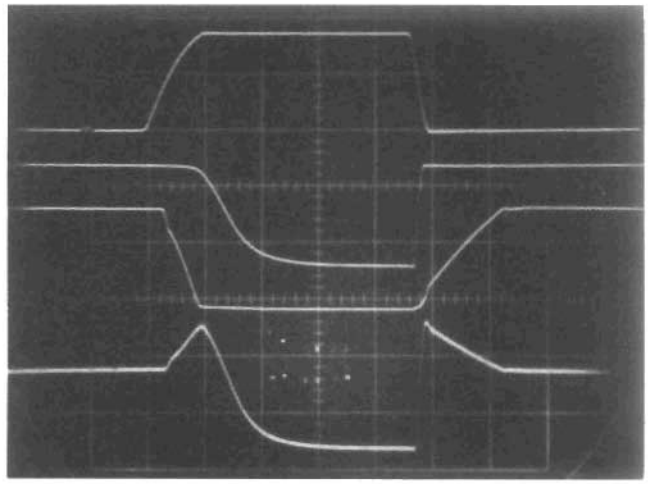

Figure 23-12. Figure 11's wave forms. Traces A, B, C, and D are C1A output, Q1 collector, Q2 collector, and R2 current, respectively. Q1–Q2 simultaneous conduction problem is evident in trace D.

added. V_{ref} and its associated diodes are biased from R1. Q3, an emitter-follower, is used to source current to C1. Q4 temperature compensates Q3's V_{be}, and Q5 switches Q3.

This method has some distinct advantages. The V_{ref} string can operate at greatly reduced current because of Q3's current gain. Also, Figure 23-11's simultaneous conduction problem is largely alleviated because Q5 and Q2 are switched at the same voltage threshold out of C1A. Q3's base and emitter currents are delivered to C1. Q5's currents are wasted, although they are much smaller than Q3's. Q2's small base current is also lost. The values for C2 and R3 have been changed. The time constant is the same, but some current reduction occurs due to R3's increase.

Operating wave forms are shown in Figure 23-14, and include C1's output (trace

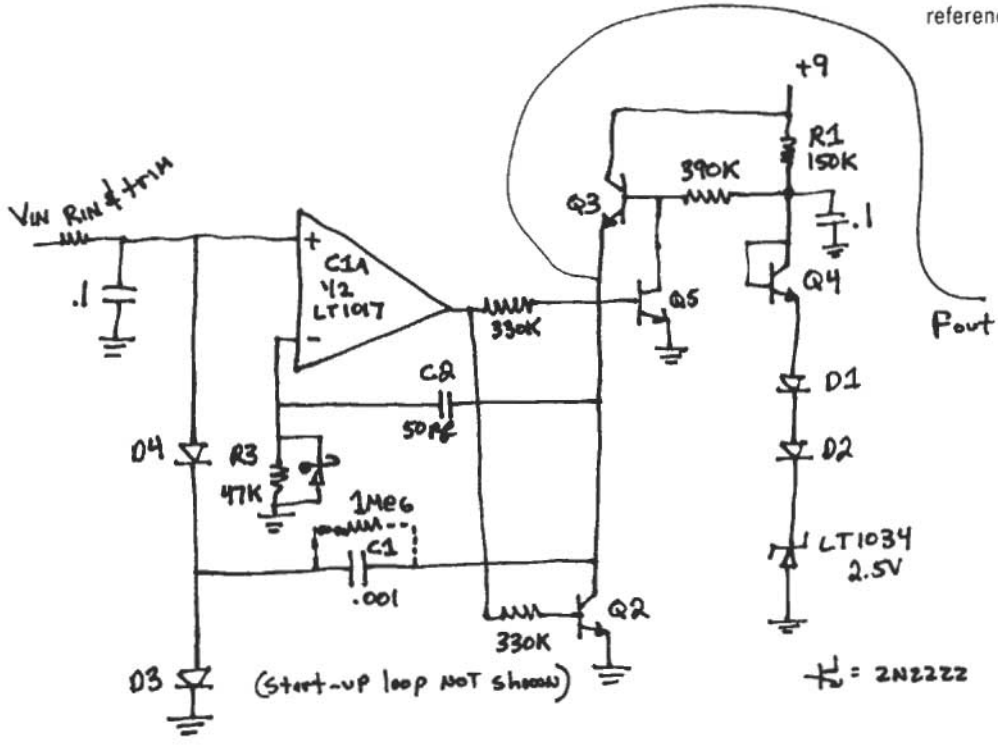

Figure 23-13. A better scheme for switching the reference.

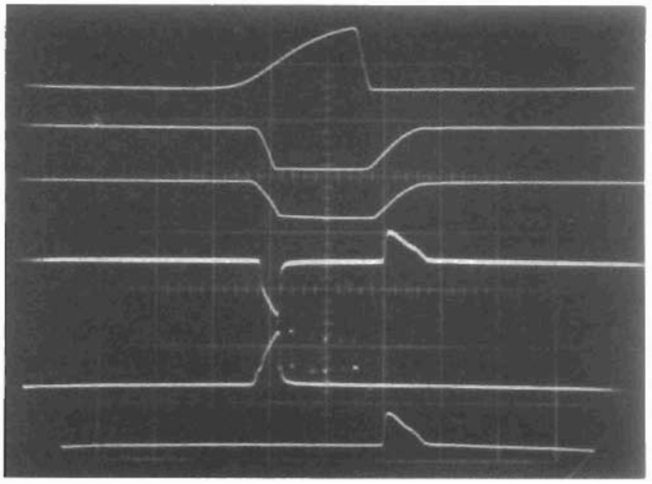

Figure 23-14. Figure 23-13's operation. Traces D, E, and F reveal no simultaneous conduction problems.

A), Q5's collector (trace B), Q2's collector (trace C), Q2's collector current (trace D), C1's current (trace E), and Q3's emitter current (trace F). Note that the current steering is clean, with no simultaneous conduction problems.

This circuit's 200 µA power consumption was low enough to make other specifications worth checking. Linearity came in at 0.05%, and dropped to 0.02% when I added a 1 M resistor (dashed lines) across C1. The D4–Q2 path cannot *fully* switch C1 because of junction drop limitations. The resistor squeezes the last little bit of charge out of C1, completing the discharge and improving linearity.

Power supply rejection ratio (PSRR) was not good enough. Supply shifts show up as current changes through R1. The LT1034 is relatively insensitive to this, but the Q4, D1, D2 trio shift value. As such, I measured 0.1%/V PSRR. R1 really needs to be a current source, or some compensation mechanism must be used.

Temperature compensation was next. Now it was time to stop hand waving and take a hard look. Q4 supposedly compensates Q3, with D1 and D2 opposing D3 and D4. Unfortunately, these devices operate under different dynamic and DC conditions, making precise cancellation difficult. In practice, R1's value should be established to source the current through Q4–D1–D2, which provides optimum circuit temperature coefficient. Assuming perfect cancellation, and no LT1034 or input resistor drift, we still must deal with Q2's V_{ce} saturation term. At 100 mV saturation, Q2 will drift about +0.3%/°C (see the Motorola 2N2222 data sheet), causing about a –300 µV/°C shift in the voltage C1 discharges toward. This works out to about –100 ppm/°C (C1 charges to 3 V) temperature coefficient, which will force a similar *positive* shift in output frequency. C1, a polystyrene type, drifts about –120 ppm/°C, contributing further overall positive temperature coefficient (as C1, or the voltage it charges to, gets smaller, the circuit must oscillate faster to keep the summing point at zero). So the best case is about 220 ppm/°C, and reality dictates that all the other junctions won't match precisely. Temperature testing confirmed all this. Initially, the breadboard showed about 275 ppm/°C, and, by varying R1, bottomed out at about 200 ppm/°C. This certainly wasn't production-worthy engineering but pointed the way toward a solution.

How could I reduce the temperature coefficient and fix the PSRR? Additionally, power consumption was still marginal, although linearity was close. Replacing R1 with a current source offered hope for PSRR, but reliable temperature compensation and lower power needed another approach. I pined for inspiration but got nothing. I was stuck.

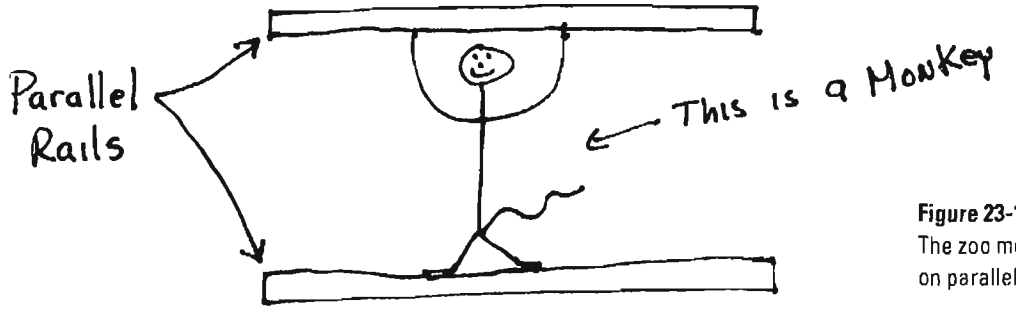

Figure 23-15. The zoo monkey on parallel rails.

Something that *had* inspired me for a couple of months was a physician I'd been seeing. We really had a good time together—a couple of playful kids. There was much dimension to this woman, and I really enjoyed just how relaxed I felt being with her. Things were going quite nicely, and I sometimes allowed myself the luxury of wondering what would become of us.

One weekday afternoon, we played hookey and went to the San Francisco Zoo. The weather was gorgeous, no crowds, and the Alfa ran great. (On our second date it threw a fan belt.) We saw bears, elephants, tigers, birds, and ate lots of junk food. The lions got fed; they were *loud* and *hungry*. Strolling around, eating cheeseburgers, and doing just fine, we came to the monkeys.

These guys are actors; they love an audience. There was the usual array of grinning, simian catcalls, cheeping, squawking, lots of jungle bar performances, wondrous feats of balance, and other such theatrics. One character particularly caught my eye. He did a little routine between two parallel rails. First, he hung by his hands as shown in figure 23-15.

Then, very quickly, he flipped over, simultaneously rotating, so he ended up inverted (see Figure 23-16).

He did this over and over at great speed; it was his act. Standing there, watching the little fellow do his inverting routine between the rails, I saw my circuit problems simply melt. I felt very lucky. I had a good lady, and a good circuit too.

If you look inside a CMOS logic inverter, the output stage looks like Figure 23-17.

The MOS output transistors connect the output terminal to the supply or ground rail. The input circuitry is arranged so only one transistor is on at a time; simultaneous conduction cannot occur. Typically, channel-on resistance is 100–200 Ω. There are no junction effects; the transistor channels are purely ohmic. The device's input pin appears almost purely capacitive, drawing only picoamperes of bias current.

Figure 23-18 shows what happens when the CMOS inverter is dropped into the gizzard of Figure 23-13's circuit. C1 is charged and discharged via the CMOS inverter's ohmic output transistors. Q3 now drives the inverter's supply pin, and Q2 goes away. Along with Q2's departure goes its 100 ppm/°C temperature coefficient

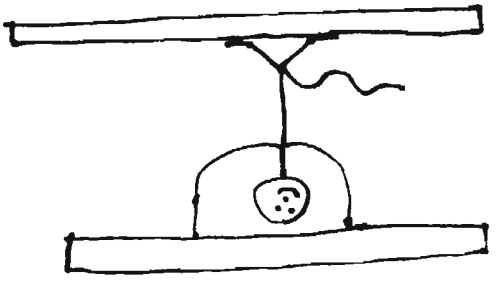

Figure 23-16. The zoo monkey on parallel rails, inverted.

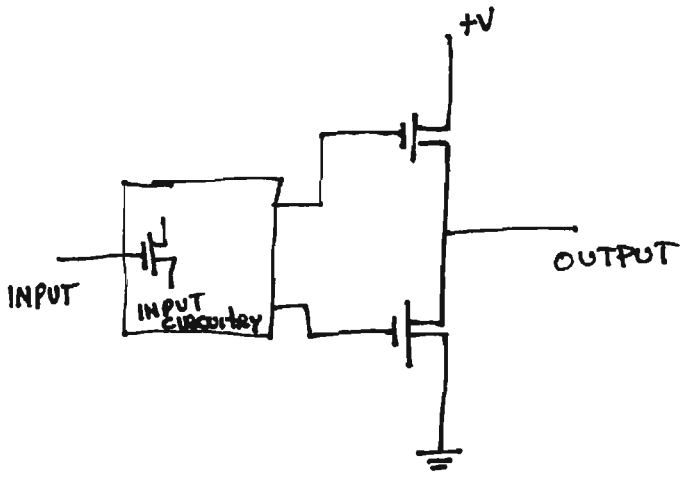

Figure 23-17. Conceptual CMOS inverter.

error. Also. Q2's base current is eliminated, along with Q5's base and collector current losses.

This scheme promises both lower temperature drift and lower power. Assuming ideal junction compensation, the remaining uncompensated drift terms are C1's −120 ppm temperature coefficient and the input resistor. Unfortunately, this configuration does nothing to fix the PSRR problem. The only realistic fix for that is to replace R1 with a current source. The current source doesn't have to be very stable but must run with only 2 V of headroom because the circuit has to work down to 6.5 V. The simplest alternative is the monolithic LM134. This three-terminal, resistor-programmable device will function with only 800 mV across it, although it does have a 0.33%/°C temperature coefficient. This temperature coefficient seemed small enough to avoid causing any trouble. The LT1034 shouldn't care, but what about D1, D2, and Q4? When I calculated the effect of current-source shift with temperature on these devices, I realized I had just inherited the world. It came out

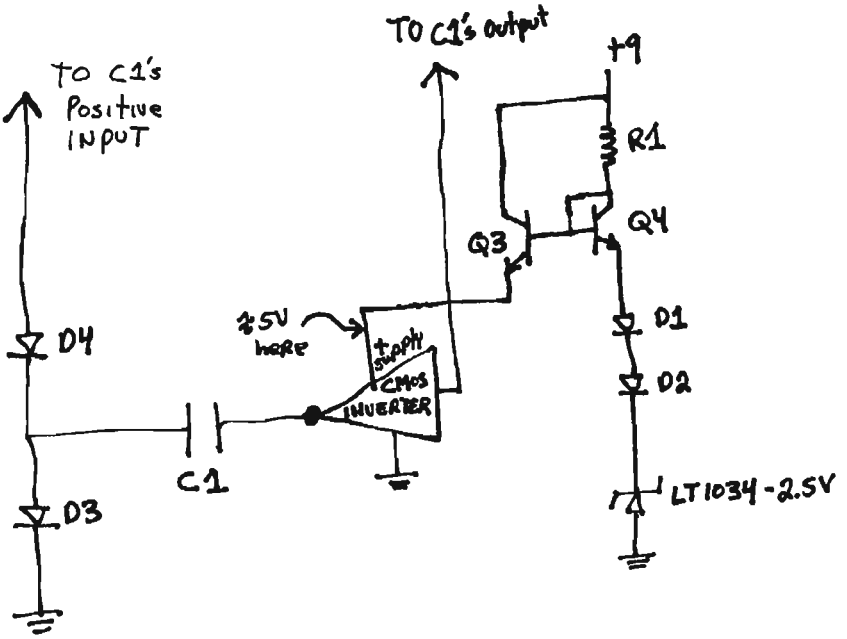

Figure 23-18. Adding the CMOS inverter to the circuit in Figure 23-13.

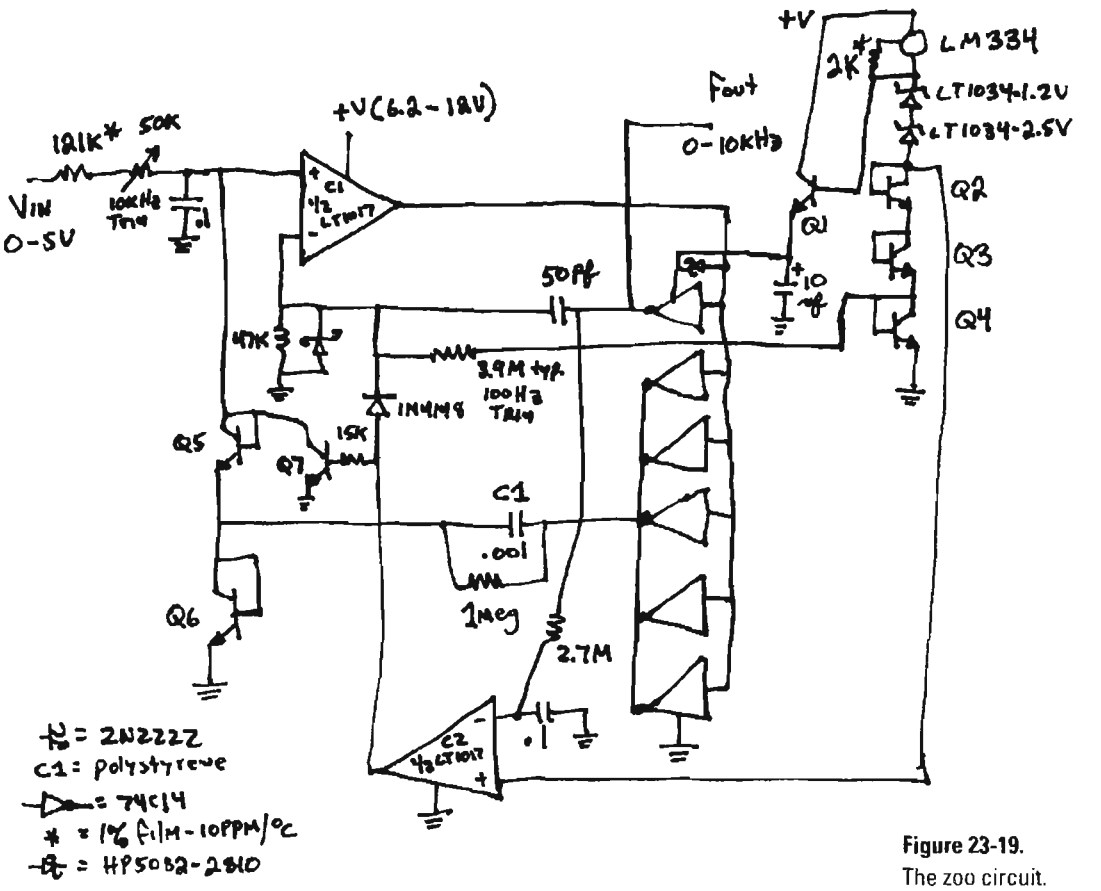

Figure 23-19.
The zoo circuit.

positive 180 ppm/°C! This tends to cancel the capacitor's –120 ppm/°C term. Additionally, increasing the LT1034's reference voltage by about 50% would pull the compensation down to +120 ppm/°C, further reducing drift. This also aids overall temperature coefficient by making the residual junction mismatches a smaller percentage of the total reference voltage. The current source's low headroom capability allows this, while maintaining operation down to $V_{supply} = 6.2$ V. The sole uncompensated term is the input resistor, which can be specified for low temperature drift.

Figure 23-19 is the final circuit. It meets or exceeds every customer specification.

A 0–5 V input produces a 0–10 kHz output, with a linearity of 0.02%. Gain drift is 40 ppm/°C, and PSRR is inside 40 ppm/V. Maximum current consumption is 145 µA, descending to 80 µA for $V_{in} = 0$. Other specifications appear in Table 2's summary. Much of this circuit should be, by now, familiar. Some changes have occurred, but nothing too drastic. The diodes have been replaced with transistors for lower leakage and more consistant matching. Also, paralleling the CMOS inverters provides lower resistance switching. The start-up loop has also been modified.

To maintain perspective, it's useful to review circuit operation. Assume C1's positive input is slightly below its negative input (C2's output is low). The input voltage causes a positive-going ramp at C1's positive input (trace A, Figure 23-20). C1's output is low, biasing the CMOS inverter outputs high. This allows current to flow from Q1's emitter, through the inverter supply pin to the 0.001 µF capacitor. The 10 µF capacitor provides high-frequency bypass, maintaining a low impedance

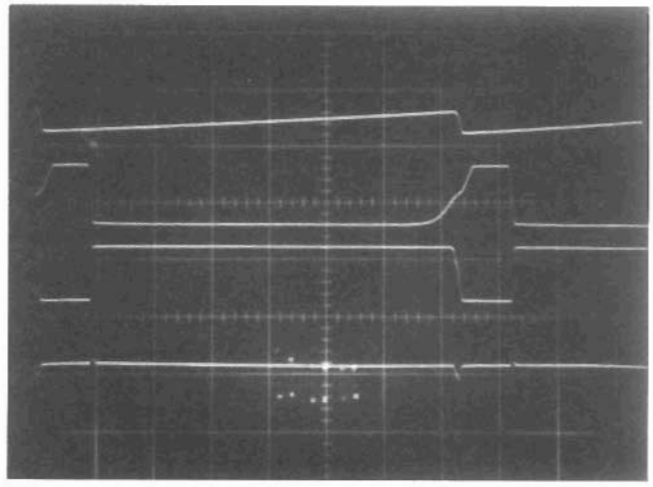

Figure 23-20. Figure 23-19's wave forms.

at Q1's emitter. Diode connected Q6 provides a path to ground. The voltage that the 0.001 µF unit charges to is a function of Q1's emitter potential and Q6's drop. When the ramp at C1's positive input goes high enough, C1's output goes high (trace B) and the inverters switch low (trace C). The Schottky clamp prevents CMOS inverter input overdrive. This action pulls current from C1's positive input capacitor via the Q5–0.001 µF route (trace D). This current removal resets C1's positive input ramp to a potential slightly below ground, forcing C1's output to go low. The 50 pF capacitor connected to the circuit output furnishes AC positive feedback, ensuring that C1's output remains positive long enough for a complete discharge of the 0.001 µF capacitor. As in Figure 23-13, the 1 MΩ resistor completes C1's discharge.

The Schottky diode prevents C1's input from being driven outside its negative common-mode limit. When the 50 pF unit's feedback decays, C1 again switches low and the entire cycle repeats. The oscillation frequency depends directly on the input voltage–derived current.

Q1's emitter voltage must be carefully controlled to get low drift. Q3 and Q4 temperature compensate Q5 and Q6 while Q2 compensates Q1's V_{be}. The two LT1034s are the actual voltage reference and the LM334 current source provides excellent supply immunity (better than 40 ppm/V PSRR) and also aids circuit temperature coefficient. It does this by utilizing the LM334's 0.3%/°C temperature coefficient to slightly temperature modulate the voltage drop in the Q2–Q4 trio. This correction's sign and magnitude directly oppose that of the –120 ppm/°C 0.001 µF polystyrene capacitor, aiding overall circuit stability.

The Q1 emitter-follower delivers charge to the 0.001 µF capacitor efficiently. Both base and collector current end up in the capacitor. The paralleled CMOS inverters provide low loss SPDT reference switching without significant drive losses. Additionally, the inverter specified is a Schmitt input type, minimizing power loss due to C1's relatively slow rising edges. The 0.001 µF capacitor, as small as accuracy permits, draws only small transient currents during its charge and discharge cycles. The 50 pF–47 K positive feedback combination draws insignificantly small switching currents. Figure 23-21, a plot of supply current versus operating frequency, reflects the low power design. At zero frequency, the LT1017's quiescent current and the 35 µA reference stack bias accounts for all current drain. There are no other paths for loss. As frequency scales up, the charge–discharge cycle of the 0.001 µF capacitor introduces the 7 µA/kHz increase shown. A smaller value

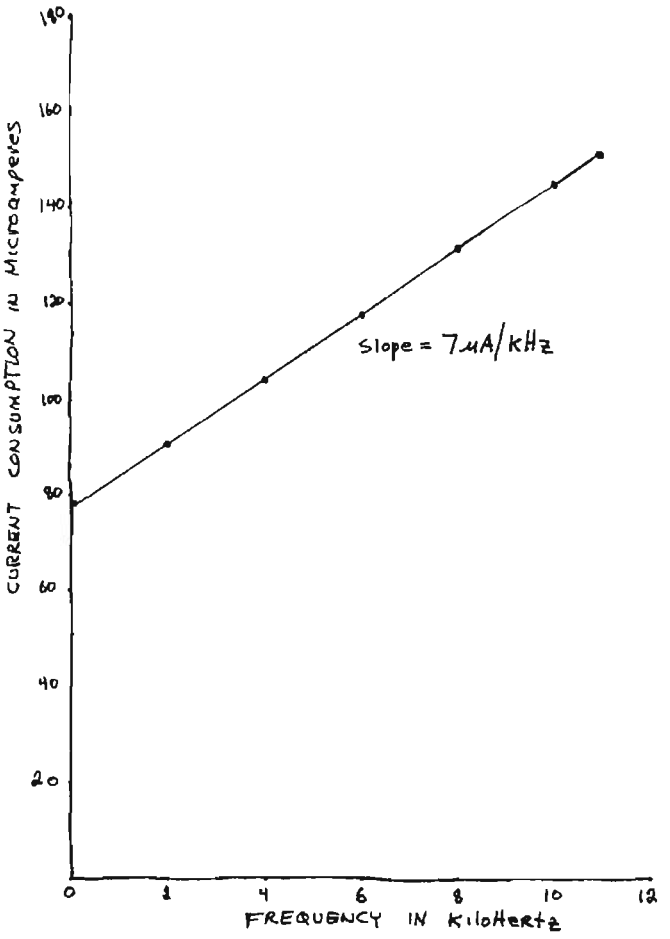

Figure 23-21. Current consumption versus frequency for Figure 23-19.

capacitor would cut power, but the effects of stray capacitance, charge imbalance in the 74C14, and LT1017 bias currents would introduce accuracy errors. For example, if C1 is reduced to 100 pf (along with other appropriate changes), the circuit consumes only 90 μA at 10 kHz, but linearity degrades to .05%.

Circuit start-up or overdrive can cause the circuit's AC-coupled feedback to latch. If this occurs, C1's output goes high. C2, detecting this via the inverters and the 2.7 M–0.1 μF lag, also goes high. This lifts C1's negative input and grounds the positive input with Q7, initiating normal circuit action.

Because the charge pump is directly coupled to C1's output, response is fast. Figure 23-22 shows the output (trace B) settling within one cycle for a fast input step (trace A).

To calibrate this circuit, apply 50 mV and select the value at C1's input for a 100 Hz output. Then, apply 5 V and trim the input potentiometer for a 10 kHz output.

Here's what the customer ended up getting:

Summary: Voltage-to-Frequency Converter	
Output frequency	0–10 kHz
Input voltage	0–5 V
Linearity	0.02%
Drift	40 ppm/°C

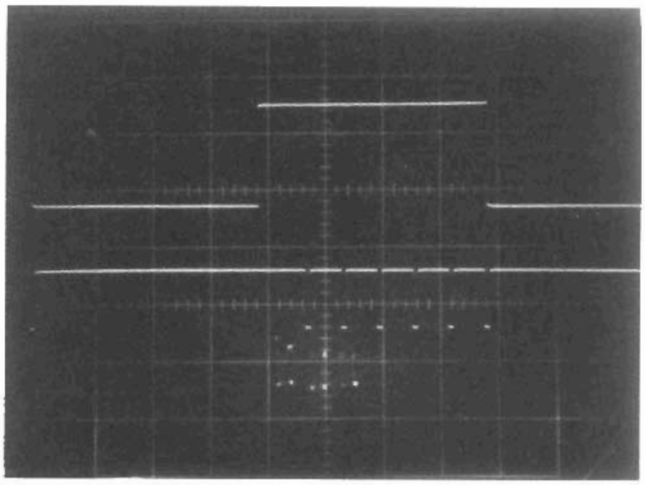

A = 2V/DIV
B = 5V/DIV
Horiz. = 200 μsec/DIV

Figure 23-22. Figure 23-18's step response.

PSRR	40 ppm/V
Temperature range	0–70° C
Step response	1 cycle of output frequency
Output pulse	5 V CMOS-compatible
Power supply	Single 9 V battery (6.2–12 V)
Power consumption	145 μA maximum, 80 μA quiescent
Cost	< $6.00/100,000 pieces

The zoo circuit made my customer happy, even if it is almost entirely bipolar. The inverter is the only piece of CMOS in the thing. I'm fairly certain the customer wouldn't mind if I had used 12AX7s[1] as long as it met specifications. It runs well in production, and they make lots of them, which makes my boss and the stockholders happy.

This circuit has received some amount of attention in the technical community. I am aware of some spectacularly complex mathematical descriptions of it, along with some arcane explanations of its behavior. Similarly, it has been shown that the circuit could have only been arrived at with the aid of a computer. Given this undue credit, the least I could do is come clean about the circuit's humble origins.

I hope it was as much fun to read about the circuit as it was to build it.

References

1. "Voltage to Frequency Converter," *General Electric Transistor Manual*, page 346, General Electric Company, Syracuse, New York, 1964.
2. R.A. Pease, "A New Ultra-Linear Voltage-to-Frequency Converter," 1973 *NEREM Record*, Vol. I, page 167.
3. R.A. Pease, assignee to Teledyne, "Amplitude to Frequency Converter," U.S. patent 3,746,968, filed September, 1972.
4. J. Williams, "Micropower Circuits for Signal Conditioning," 10 kHz Voltage-to-Frequency Converter, pp. 10-13, Linear Technology Corporation, Application Note 23, 1987.
5. J. Williams, "Designs for High Performance Voltage-to-Frequency Converters," Linear Technology Corporation, Application Note 14, 1986.

1. For those tender of years, 12AX7s are thermionically activated FETs, descended from Lee DeForest.