INCREASED RESOLUTION FOR BEAT-PERIOD BASED FREQUENCY STABILITY MEASUREMENTS

By Eric L. Blomberg
Lincoln Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Much of the work of the founders of this meeting has dealt with the characterization of frequency sources by time domain measurements. Typically the sources are offset to provide a low frequency beat signal, statistics of whose period are collected. While the transformation of these time domain statistics is available to the user in manageable form, the actual beat measurement technique has some unexamined limitations. Let us look at the relation of errors in this measurement to their frequency stability interpretation.

The Relation of Beat Frequency and Timing Errors to Measurement Noise Floor

We wish to express the difference between two consecutive frequency measurements in several forms. First we have:

$$\Delta f = f_1 - f_2,$$

since $f = \frac{1}{t}$:

$$\Delta f = \frac{1}{t_1} - \frac{1}{t_2}$$

if $t = t_1$

and $t_2 = t_1 + \Delta t$ with Δt small

then $\Delta f = \frac{1}{t^2}\,\Delta t$

so $$\frac{\Delta f}{f} = \frac{1}{f} \cdot \frac{1}{t^2}\,\Delta t$$

Where t is the period measurement and Δt is the random error in this measurement. We must remember that our measurements are actually averages over the measurement interval; each quantity must be seen as a time average. Hence t represents the nominal beat period, and Δt represents the timing deviation normalized to a single period.

If the beat and averaging periods are both one second, the distinction is hard to see. If we choose a ten-period average of the one second beat period, and further assume that only system timing errors cause the deviations, then the difference is clear. If our timing error for one measurement is one microsecond, then the normalized error over the ten periods is only .1 microseconds, over 100 periods .01 microseconds. For a fixed beat period, the effect of timing errors goes down proportionately to the number of periods averaged.

Suppose that the beat period is changed to ten seconds and the same timing error applies. In this case the Δt term is not decreased, since only one period is averaged. The increased period, however, has a quadratic effect. The $\frac{\Delta f}{f}$ for the first case is:

$$\frac{\Delta f}{f} = \frac{10^{-6} \div 10 \text{ periods}}{(1 \text{ second})^2} = 10^{-7} \cdot \frac{1}{f}$$

for the second:

$$\frac{\Delta f}{f} = \frac{10^{-6} \div 1 \text{ period}}{(10 \text{ seconds})^2} = 10^{-8} \cdot \frac{1}{f}$$

or in general:

$$\frac{\Delta f}{f} = \frac{\Delta t/N}{\tau^2} \cdot \frac{1}{f}$$

Where τ is the beat period

N is the number of contiguous periods measured

Δt is the uncertainty of each measurement

Hence we find different noise floors for different beat periods and the same averaging time. (Measurement interval)

This would lead us to conclude that a longer beat period would yield better results. Unfortunately, for a sine wave beat signal, increasing the beat period proportionately decreases the slew rate which, in turn, proportionately increases timing error for the standard counter configuration. Hence this is a no win game, unless we desensitize our period measurement to slew rate.

The timing errors fall into two categories: those of much higher frequency than the 1 Hz signals of interest, and those of somewhat lower

frequency. The first corresponds to high frequency noise which can be seen simply as thermal noise at the input of the period measuring counter. The effect of this noise is well known and its effect on low slew rate measurements is often specified in commercial counters. Suppose we have generated a 1 Hz beat signal 10 volts in amplitude. Its slew rate at the zero crossings is about 30 volts/ sec. Hence a voltage excursion of 30 microvolts corresponds to an error in the period measurement of 1 microsecond. One commercial counter cites an uncertainty of 100 us for such an input. If two 10 MHz sources are being compared, this implies a noise floor of 10^{-12} for the frequency stability measurement (10 periods measured).

The second source of error corresponds to longer term changes in the operating point of the measurement system. A good example is a time varying input offset voltage of a zero crossing detector or Schmitt trigger. This offset may come from the counter input stage, a preceding amplifier or even the balanced mixer. It is often affected by long term changes in the environment of the measurement system. The change in this offset corresponds to a change in trigger point which, as seen in Fig. 1, represents an over or under estimation of the period. If we could somehow track a feature of the waveform's geometry, such as the peak, rather than a level, we could acquire immunity to such long term drifts.

Proposed Circuit

The integrating A/D convertor framework provides three features to achieve this end. It establishes zero-crossings by integration, can cancel DC drifts over intervals longer than one-half the beat period and eliminates the dead-time associated with traditional counter regimes.

Consider the circuit of Fig. 2 with a 1 Hz sine wave input of 10 volts peak to peak. Suppose that the integrator is shorted initially and allowed to operate when the signal level first reaches -.6 volts. When the integrator output returns to zero, it is returned to the shorted condition. If the input is an ideal noiseless sine wave, then the integrator output waveform is symmetric about its peak value, which is analogous to the input's symmetry about its zero crossing. The best estimate of the zero crossing time is a period following the -.6 volt trigger time and equal to one-half the duration of the integrating interval. Clearly there is nothing sacred about the -.6 volt level. It merely provides an arbitrary starting point for the symmetric interval. If the input is corrupted by noise, any components with zero mean and period less than the integrating interval will tend to be rejected. These are precisely the higher frequency components which cause the gate to flap when a conventional zero crossing detector confronts a low slew rate signal. Notice that this technique always estimates the true zero crossing of the input. If there is a residual DC offset to the input, then the trigger point will be skewed relative to the true mean value of

the input wave shape. Nonetheless this solves one of the difficulties of timing the slow rise time signal, even though it exhibits the same vulnerability to time varying offset as conventional means.

It is possible to perform this operation on both rising and falling edges of the input. Taking the midpoint of these two zero crossing estimates has the effect of rejecting DC offset altogether. It might be viewed as a best estimate of the peak of the sine wave input. Whether it actually corresponds to the peak or not is not significant. What this process does in fact accomplish is to establish the timing of a fixed point on the geometry of the waveform, so that our period measurement is actually a measurement of repetition of geometry rather than of some absolute level. While a varying DC offset certainly affects the zero crossings, it would not affect the timing between two consecutive peaks. This technique is vulnerable to changes in offset level which occur between the two integration intervals, but is able to reject such changes outside them. If we are measuring single periods, this method will reject a linear drift occurring over half the cycle. If we are measuring the length of a ten period sequence, it will reject level changes occurring over 9-1/2 of those periods. Since systematic drifts are most likely to occur over long measurement intervals, we find that this technique is extremely effective in removing them. This is not the case with a simple period counter technique, which counts clock pulses over the ten periods and suffers a trigger error equal to the total DC level change times the slew rate of the input.

We have actually formed a bandpass filter for the 1 Hz signal. Since the 10 volt peak-to-peak signal will slew from -.6 volts to +.6 in about 40 milliseconds, we have a high frequency cutoff of 25 Hz. (Actually a $\frac{\sin x}{x}$ response -3 dB at 12 Hz with a first null at 25 Hz.) This is sufficiently higher than the preceding amplifier's bandwidth that the system noise bandwidth is substantially unaffected. The low frequency response rolls off at 6 dB/octave below one half the beat frequency.

Notice that the information content of the 1 Hz signal is not affected. Its information is an FM modulation. Changes in frequency over 1000 second intervals are present as .001 Hz sidebands of the 1 Hz carrier, not as direct .001 Hz signals. Hence we have prevented systematic offsets from irreversibly mixing with the long-term frequency effects we seek to measure.

Construction Details

It was most desirable to implement a synchronous system to preserve resolution and also to minimize the ill effects of attempting to derive logic triggers from slow analog signals. This is, after all, the fundamental problem for which the proposed processor promises solace. The entire operation is that of a four-state machine, whose states may be

identified from the labels given the corresponding wave forms in Fig. 3. Before describing the sequencer in detail, I would first like to explain the overall means of obtaining a period measurement.

Consistent with my goal of a zero "dead time" system and my goals of simplicity and low cost, an interpolation technique is used. A counter with a 1 kHz clock (my "coarse" register) is used in precisely the traditional way to measure the period between -.6 volt initial trigger points (the point which initiates the entire double integration cycle). Its slow clock rate permits it to be strobed and reset on the fly, so that no period goes unmeasured. This counter is possessed of all the flaws previously described. Another register, the "fine," is controlled by the integrator, and contains, after both integration cycles are done, the timing from the somewhat arbitrary, but nonetheless synchronous, -.6 volt point to the midpoint of the positive half cycle of the wave. As mentioned earlier, this point is referred to the geometry of the waveform and is DC level independent. The best estimate for a given period is found by adjusting the coarse timing by the interpolations before and after it.

$$T = t_{c_n} + t_{f_n} - t_{f_{n-1}}$$

Where t_{c_n} is the contents of the coarse register (its measurement of the n^{th} period) and t_{f_n} is the interpolation time for the end of the coarsely determined n^{th} period to the next cycle's midpoint. Since the output of the system is alternating readouts of t_f and t_c, it is a small matter to arrive at a running readout of actual periods. Let us now consider the operation of the integrator.

State A of Fig. 3 is the primordial state, the pre-interpolation mode. The coarse counter is making its measurement and the integrator is held in an auto-zero mode. This has an advantage over simply shorting the integrator in that it nulls out the integrator's offset voltage as well. Since the integrator never works for more than one-half second after an auto-zero, we can expect that long-term systematic changes in the integrator itself are well below the total system noise floor.

State B is the estimation of the positive zero crossing time. It is entered when the input first exceeds -.6 volt, synchronously with the 1 kHz coarse clock. This prevents a gross 1 millisecond quantization error. At this moment the contents of the coarse counter are strobed into the output latches and the coarse counter is reset to zero, long before its next 1 kHz clock edge. Thus the entry point of State B has an uncertainty range of ± 1 millisecond. At a 30 volt per second slope of the input, this is only a voltage band of ± 30 millivolts about -.6. The absolute time is unimportant as all residues are absorbed into the

interpolation measurement. When the integrator output is driven back to zero, the symmetric interval about the true zero crossing is over. This event is determined by a comparator which senses integrator output. Notice that the comparator is not included in the auto-zero loop. Its trigger point is actually biased .1 volts below zero. This prevents the oscillation which usually accompanies attempted linear operation of a comparator. This also provides a well defined comparator output. It will always be low during auto-zero and during positive integrator output. Thus it goes high only after the integrator has in fact crossed zero at the end of B. Since the integrator is slewing approximately as fast as the input here, there is a delay of 5 milliseconds until the comparator trigger point is reached. The comparator has a window of about 2 millivolts around its nominal trigger point within which noise will cause multiple transitions. Since the 5 millisecond delay is uniform from measurement to measurement and is only affected by changes in the slope of the input, and since the noise at the integrator output is quite small, we lose nothing to take the first comparator transition as the integrator output zero crossing. The net delay is subtracted out in the period determination algorithm and is uniform from period to period.

Thus the comparator causes the system to enter State C, which is almost identical to A. The integrator is auto-zeroed, and the coarse counter is still running, but the duration of State C is being timed in the fine counter.

State D is the exact analogue of State B, only it functions on the falling slope edge of the input. It is triggered by the input signal reaching +.6 volts (with no synchronization required this time). This state defines a time interval symmetric about the zero crossing and passes to State A when this is done, with the same convention on the comparator as State B.

To construct an interpolation interval from the durations of States B, C, and D is a simple matter. Suppose we assign the entry into State B as our time origin. This is also the timing edge of the coarse register, so it is the logical place to begin an interpolation. Clearly the time to the best estimate of the zero crossing (rising edge) is half the duration of State B. Similarly, the time to the best estimate of the zero crossing of the falling edge is the sum of the durations of B and C and half of D. Simple algebra yields a time to the midpoint of the zero crossings and shows that a counter clocked with relative rates of 3:2:1 during B, C, and D respectively will yield this same interval. (See Appendix 1). Specifically, rates of 15, 10, and 5 MHz will allow interpolation to 100 nanoseconds. CMOS clocking rates limited me to one quarter of this resolution. (See Fig. 4)

Thus we see how the interpolation is actually generated for the 1 Hz signal with a measurement of 1 second. Multiple period measurements

required adding a period counter to the trigger of State B. This way an interpolation will be performed initially and on the final period of the number selected, and not for intervening periods.

Preliminary Results

Time permitted only a crude evaluation of this technique, but the outcome shows the technique to be worth investigating further. A 1 Hz test signal was derived from a 5 MHz quartz oscillator locked to a hydrogen maser. It was shaped to approximately sine form with 10 volt peak-to-peak amplitude and some white noise added. Three cases were investigated, the traditional counter, the counter preceded by a high gain limiter, and the circuit of this paper. Only single period measurements were taken. The results are expressed as a fractional frequency referred to 10 MHz and are the root of an Allen variance for 100 trials, and may be compared to results of Allan [1] and Reinhardt [2].

Case 1: $\frac{\Delta f}{f} = 1.4 \times 10^{-11}$

A slew rate of 30 volts/sec applied to an HP 5327B close to the predicted error

Case 2: $\frac{\Delta f}{f} = 3.5 \times 10^{-12}$

The HP 5327B preceded by a gain of 1000 limiter, effectively increasing slew rates to 30,000 volts/sec

Case 3: $\frac{\Delta f}{f} = 6.0 \times 10^{-13}$

Although the proposed circuit clearly does have a lower noise floor, it is hard to say by how much. These results represent a noise floor approximately 5 times higher than Allan and 50 times higher than Reinhardt. Three factors come to mind. We may be nearing the noise level of the test signal itself. Here some further measurements would be useful. Secondly, the clocking rate used in this model yields a minimum relative quantization error of 2×10^{-13}. Thirdly, good low noise circuit techniques, as exemplified in both Allan and Reinhardt, may provide lower noise floors. Nonetheless, the 10^{-12} level is quite respectable for a low cost alternative to a dedicated frequency counter.

Response to Systematic Error

This is a fairly coarse test also, introducing about a 1 volt offset for two out of every ten periods of the 1 Hz 10 volt peak-to-peak input. Care was taken to make the transitions far from zero crossings so as to avoid changing trigger conditions. For a signal with a 20 volt per second slew rate we would expect a 50 millisecond timing error based purely on the zero crossing detector technique. Only the high gain limiter alternative was measured, as it would yield the same data as for the counter alone, but with less noise. Actual values of 42 milliseconds were observed. They occurred symmetrically since the long-term average frequency did not change.

Figure 7 shows some results for a step change representing drift during the integration period. Note that the counter's "dead time" yields only half as much data. It gives three normal periods, followed by a period 42 milliseconds too long, followed by one 42 milliseconds too short, then back to three normal periods. Looking at the processor output, we see seven normal periods, one 20 milliseconds too long, one normal, and finally one 20 milliseconds too short.

Notice how the full error is present in the coarse register and half is compensated by changes in the interpolation register. These and other data show a factor of 25 suppression of DC offsets occurring outside the integration period. Actual performance may be considerably better with the much smaller perturbations one would expect in a practical system.

Note that in Fig. 7, due to the lowered clock rates, the period is determined by:

$$T = {}^{t}c_{n} + 2 \quad ({}^{t}f_{n} - {}^{t}f_{n-1})$$

Conclusion

Even in coarse testing, the circuit has achieved the goals of noise floor reduction, reduced sensitivity to systematic drifts, and no dead time. This level of performance was achieved by a prototype wherein little attention was paid to low-noise techniques. Standard CMOS and BiFET parts were used. No shielding or separation of analog and digital grounds were attempted. The output of the prototype was 7 decades of BCD information from three registers which could have easily been transmitted through the IEEE bus through an appropriate interface. All subsequent calculations and choices of averaging time could be left to the minicomputer or calculator controlling the bus. It is interesting to note that the net cost of this technique is considerably below the counter it outperformed.

Acknowledgement

This work was performed while the author was with the Hydrogen Maser Lab of the Smithsonian Astrophysical Observatory and also represents a portion of a Master's Thesis submitted to the Department of Electrical Engineering and Computer Science at M.I.T. (June 1977). The initial concept of using the A/D converter topology is due to Professor James Roberge of M.I.T.

References

[1] Allan, David W., "Report on NBS Dual Mixer Time Difference System". NBSIR 75-827.

[2] Reinhardt, Victor and Donahoe, Theresa, "A Simple Technique for High Resolution Time Domain Phase Noise Measurement". Proceedings of the 8th Annual PTTI Meeting. NASA, GSF.

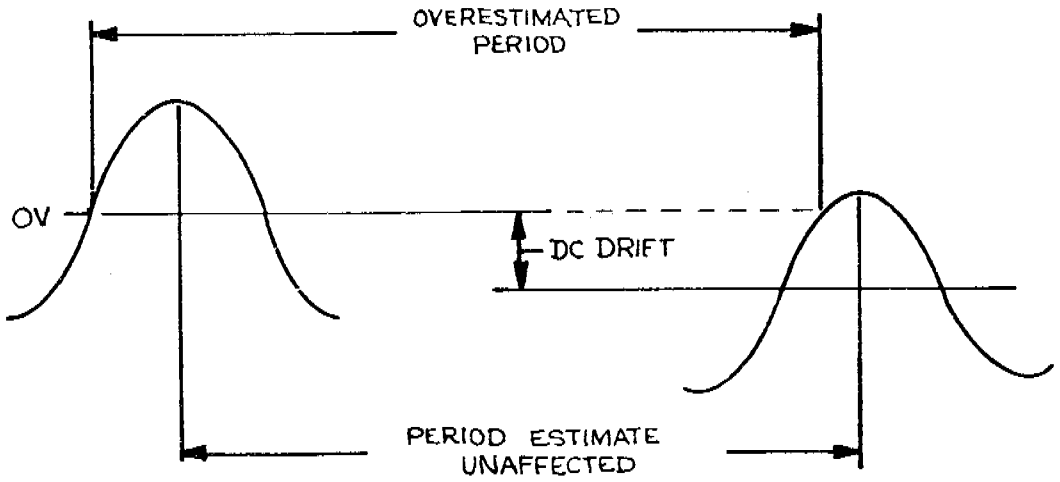

Figure 1. EFFECT OF DC OFFSET

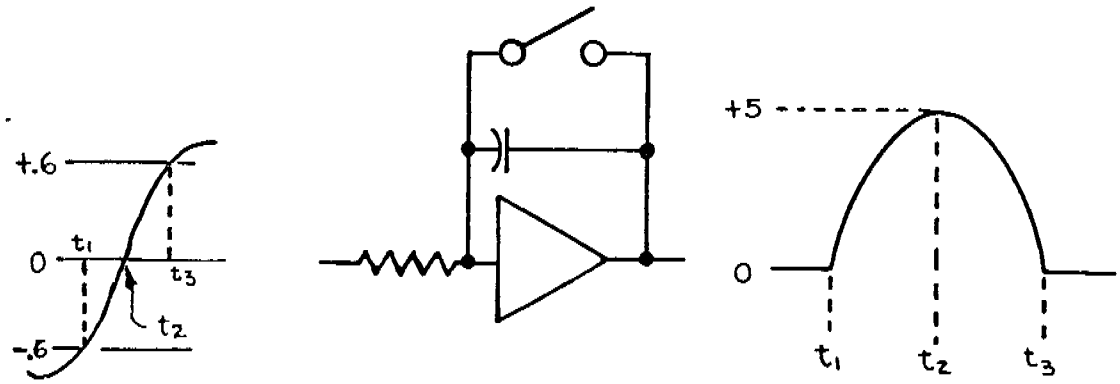

Figure 2. SWITCHED INTEGRATOR

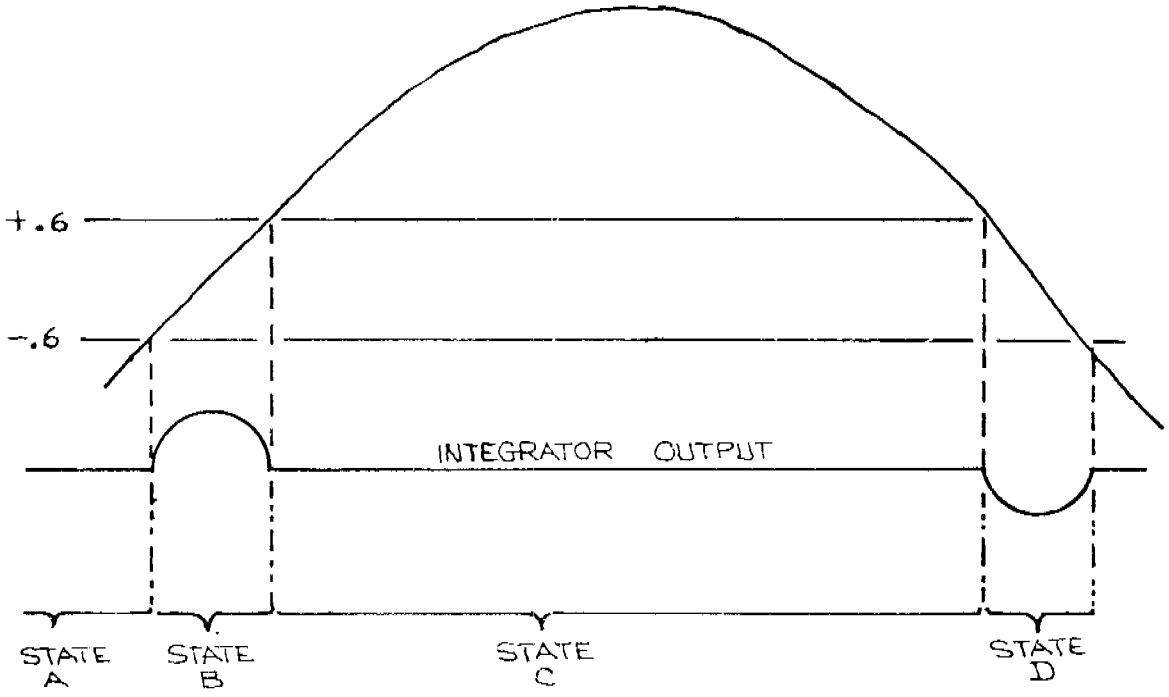

Figure 3. SYSTEM WAVEFORMS

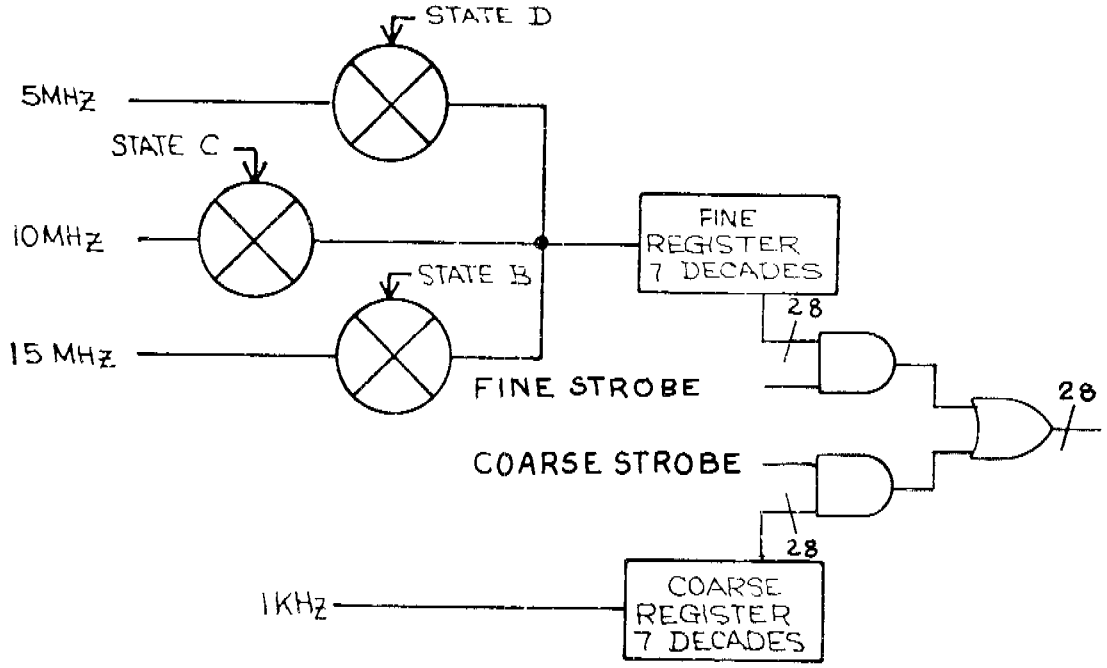

Figure 4. GATING SCHEME

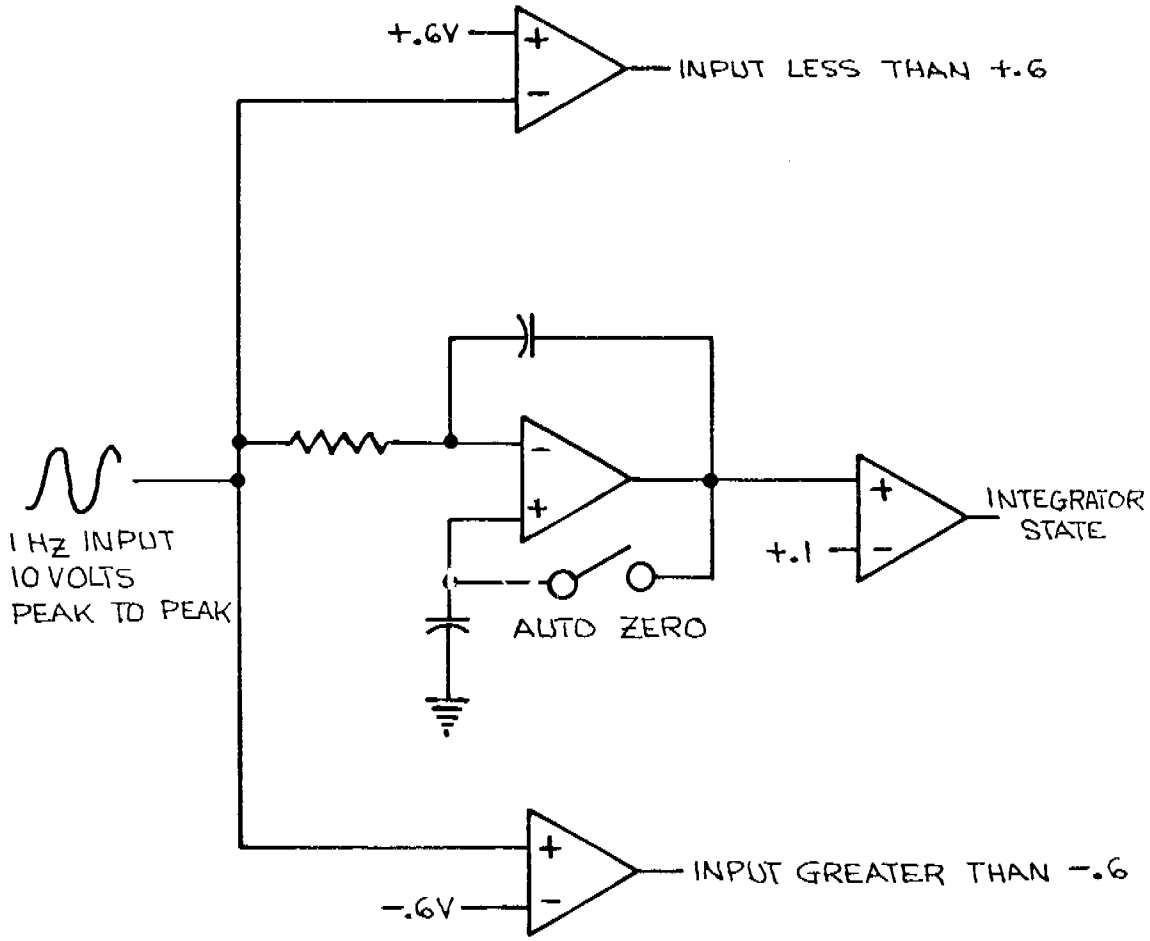

Figure 5. ANALOG SECTION

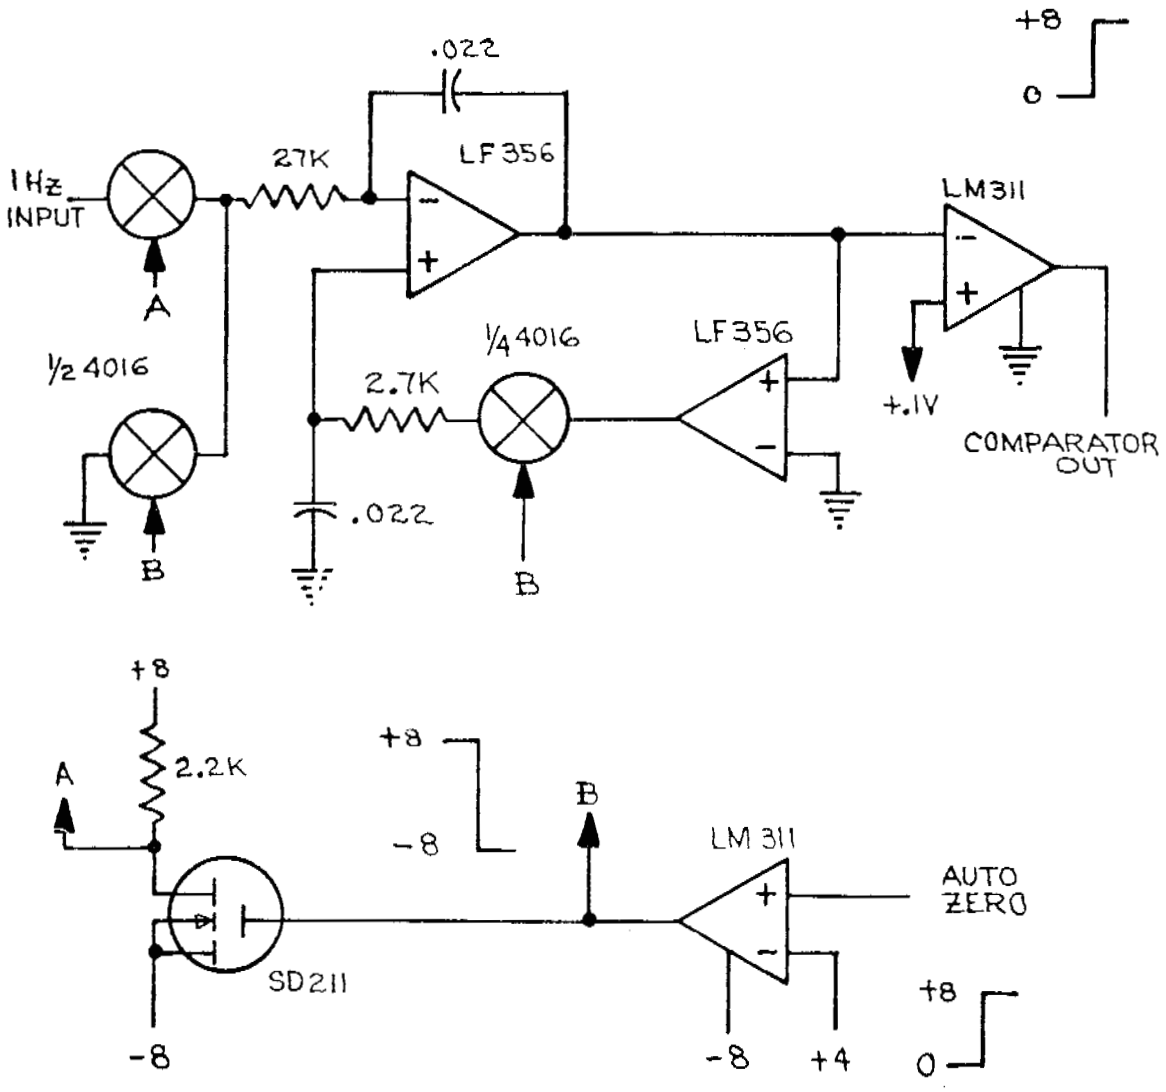

Figure 6. INTEGRATOR SCHEMATIC

COUNTER WITH LIMITER

DATA		Δt from 1.0 second in μs
9999903	T x 10^{-7} sec	-9.7
9578161		-42183.9
10422557		+42255.7
10000136		+13.6
9999904		-9.6
10000134		+13.4
9578080		-42192.0
10422487		42248.7
9999854		-14.6

PROPOSED CIRCUIT (with 1 volt drift during one period)

DATA		T	Δt (μs)
1724543	T_f in 2 x 10^{-7} sec	.9999978	-2.2
001000	T_c in 10^{-3} sec		
1724554		.9999998	-.2
001000			
1724555		.9792412	-20075.8
001040			
1828349		.9999816	-18.4
001000			
1828441		1.0208344	+20834.4
000960			
1724269		.9999426	-57.4
001000			
1724556			

Figure 7. EFFECT OF SYSTEMATIC OFFSETS

APPENDIX: Interpolation Interval

1) let t_o mark the entrance of State B

t_1 mark the entrance of State C

t_2 mark the entrance of State D

t_3 mark the end of State D

Note that the system is synchronized so t_o also marks an epoch of the coarse clock.

2) Since State B runs from t_o to t_1, its mid-point lies at

$$t_+ = 1/2\ (t_o + t_1)$$

This is the best estimate of the time of the positive zero crossing (referred to t_o).

3) Identically State D gives the best estimate of the negative zero crossing

$$t_- = 1/2\ (t_2 + t_3)$$

4) The "peak" is estimated by the mid-point of the two zero crossings

$$t_p = 1/2\ (t_+ + t_-)$$

$$t_p = 1/2\ [1/2\ (t_o + t_1) + 1/2\ (t_2 + t_3)]$$

$$= 1/4\ (t_o + t_1 + t_2 + t_3)$$

$$= 1/4\ (t_1 + t_2 + t_3)$$ Since we set t_o=0 as our reference point.

5) We have, however, measured the durations of B, C, and D and not their absolute timing.

Duration of State B = $t_B = t_1 - t_0 = t_1$

Duration of State C = $t_C = t_2 - t_1$

Duration of State D = $t_D = t_3 - t_2$

Substituting: $t_1 = t_B$

$t_2 = t_B + t_C$

$t_3 = t_B + t_C + t_D$

6) Substituting into the result of 4)

$$t_p = 1/4\ [t_B = (t_B + t_c) + (t_B + t_c + t_d)]$$

$$= 3/4\ t_B + 1/2\ t_C + 1/4\ t_D$$

hence the ease of computation by measuring the state durations with clock rates of 3:2:1.

DR. VICTOR REINHARDT, NASA Goddard Space Flight Center:

I think it is a very interesting approach, especially the summation of the positive and negative edge. We had some later data from our system--an equivalent dual-mixer system. In short term we did beat you, but we did notice diurnal variations equivalent to 20 picoseconds due to temperature effects. We traced it down to the mixer itself and not any of the components--the DC offsets in the mixer. So if Bob Vessot is going to give us a 10^{-17} oscillator, we certainly need approaches similar to yours to get out the drifts in the measurement system.

DR. BLOMBERG:

Thank you.

MR. HERMAN DAMS, National Research Council, Canada

You mentioned the problem of dead time in the dual balanced mixer system. In the system we built at NRC about three years ago, which was similar to Dave Allan's, we overcame that problem by simply having a counter with a dual front end so that you count continuously. You take your readout, of course, at the end of the one-second period from the front end counter, which at that time has stopped. The other one was still counting.

DR. BLOMBERG:

I mostly brought that up as a practical issue for systems that would be implemented with commercially available counters. I didn't mean to connect this with that effort.