# Period Variation in an Electronically Maintained Clock

Robert L. Belleville – comments/questions welcome: drbob@casaalta.com

Logging clock period via infrared optical switches at microsecond resolution shows considerable variation from cycle to cycle in my clocks. Possible sources of this variation include:

Effects of drive pulses
Ambient I/R radiation
Change in pendulum displacement
Torsional/Vibratory modes
Mechanical vibration
Air currents
Instrument errors

In this paper I will take up each of these sources. HSN 1999-3 described the system I am using to measure my 7th clock. Diagrams in that paper will be very helpful in understanding this paper. Period, temperature and barometric pressure are logged second by second using a PIC microcontroller. This is a 120 beat pendulum (248mm.) The nominal period is 1 second but it is not exactly rated to 1 second and runs fast. Clock #7 is a free pendulum and drive is only applied a few times per minute. Pendulum position is sensed optically and a drive impulse is applied only to maintain constant angular displacement. The clock is mounded on a plywood backing board and is open to the air sitting on my workbench.

Recently I made another interval timer using a Motorola 68HC11 series microprocessor. HC11 is well suited to long interval timing at a resolution of ±0.5 µsec because it contains built-in timing-specific hardware. (See diagram below.) A hardware latch captures the output of a free-running 16-bit counter on a specific edge of an input signal. This free-running counter is independent of any software uncertainties and is driven directly from the system's 8-MHz crystal. A 32.768-millisecond interval is counted before the counter overflows and starts again at zero. An interrupt is generated on each overflow so that software can increment an additional 16 bit counter. Software responds to this interrupt in far less than 1 millisecond (ms hereafter) so there is no chance of missing any timer overflow. The combined counter can measure up to one-half hour. The only problem comes from uncertainty in the relationship between the arrival of the edge that captures the free running timer and the overflow interrupt. Software removes this uncertainty. The HC11 contains 3 such input capture latches.

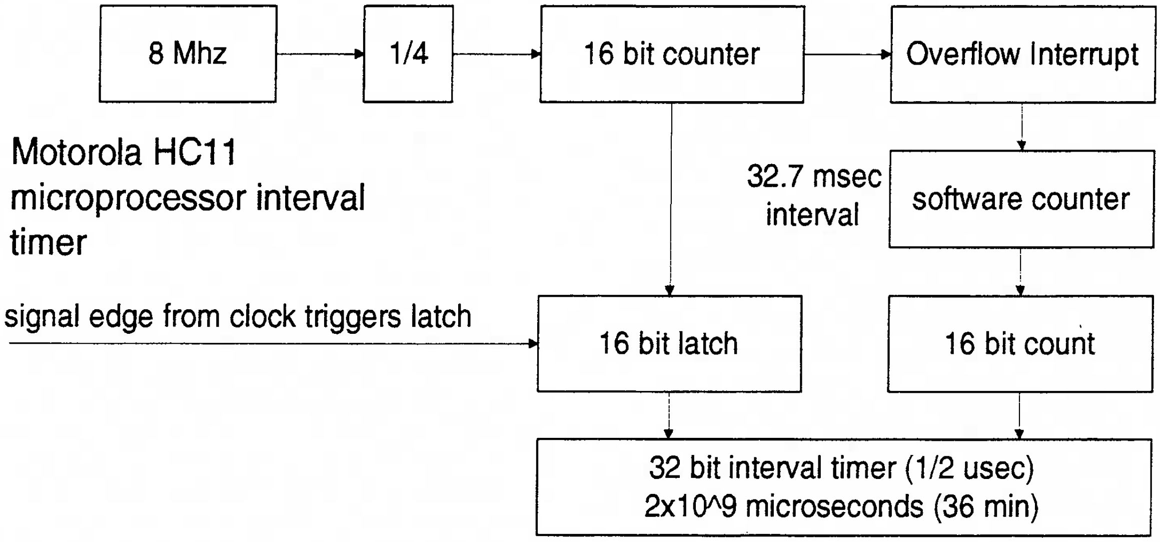

(My HC11 board is from Technological Arts – www.interlog.com/~techart. The board has everything except a 9-volt DC power supply and a case. Model Adapt-11C75DX.)

As a preliminary check I hooked the HC11 timer to a regular digital output on the PIC and sampled data to determine if any jitter is introduced by the wiring or ground system. No jitter was detected.

By running this HC11 in parallel with my PIC based monitor I can improve the reported resolution by a factor of 100. Figure 1 (figures 1 through 4 are appended to this text) shows one day of clock 7. In the third strip chart labeled 'Period', both the PIC data, reported in 52-µsec bins, and HC11 data, reported in 0.5-µsec bins, are shown. All figure 1 data streams are averaged over one-quarter hour to give the curves shown. Period correlation is nearly perfect. The offset between the two curves is artificially reduced to about 5 µsec. The actual difference between PIC's crystal and HC11's is about 30 ppm (parts per million).

Clock 7 uses closed-loop 'bang bang' control. Bang-bang can be contrasted with proportional control. Proportional control provides sinusoidal drive exactly in phase with the pendulum (Bigelow HSC1999-3). Our common household thermostats are an example of 'bang bang' control. Heat is either on or off. When temperature is a bit above thermostat set point heat is off. When temperature is a bit below set point heat comes on -- there is nothing in-between. The average temperature is then near the set point. (Or so goes the theory.) Clock #7 is given a 'bang' if it isn't swinging widely enough. An optical limit sensor 'closes the loop' to provide nearly constant angular displacement. This sensor is like a thermostat – it keeps the average angular arc nearly constant. This is a departure from free pendulum clocks built in the early part of the 20th century. These clocks were not closed-loop but got a fixed number of drive pulses each minute. In theory the pendulum would find a constant arc based on a balance between energy added and energy lost in friction. I use bang-bang because it is very simple and reliable. My #2 clock has run nearly continuously since the 1970's. Also as my previous article showed, our real problem is long term stability. Until I can get a clock to run for years without shifting period by 100 parts per million without warning or explanation my focus won't really be on fine points of drive.

In clock #7 one optical sensor detects the center of the swing and another reads end of arc at the right end. These are called 'position' and 'limit' sensors for short. Drive impulses 8-ms long are added if the limit sensor isn't occluded by a shutter which moves with the bob. The same shutter occludes the position sensor during the right half of the swing. Drive is very close to the maximum velocity point as measured by the position sensor. Drive energy is then fixed to a certain 'quantum.' If the energy needed to maintain constant angle isn't exactly an integral number of these quanta (and it never is), the system adds or drops a pulse every so often. Over an extended period however the average number of pulses per unit time is nearly constant. This could be called the 'drive ratio.' I express this here as the fraction seconds/drive-pulse.

In figure 1 the bottom strip shows drive ratio starting at about 12.75 seconds/impulse at 8:30 PM April 13 1999. Less drive is needed at 8 am when only one impulse in 13.00 seconds is required. Drive increases slowly as the day progresses and the change doesn't look related to barometer, temperature, or daylight. Currently I do not know why drive varies by this much. Note also that this variation is not strongly reflected in period either. Jaggedness is a result of 15 minutes averages. Only about 70 drive impulses are needed in each ¼ hour; so a deviation of ±1 impulse averages to 0.4 seconds/pulse. The large jagged peaks are larger than normal shock.

After tests shown in figure 1, I realized that the HC11 could measure both the left and right half-periods of the pendulum by using both the rising and falling edges of sensor output. Data showed that the position sensor was about 50-ms off center for figure 1. This has little effect on period but I adjusted the pendulum to make the left and right half period nearly identical. (Actually I achieved about ±5 ms.) A side effect of this change was a slight increase in drive ratio that you will notice in the following text and figures.

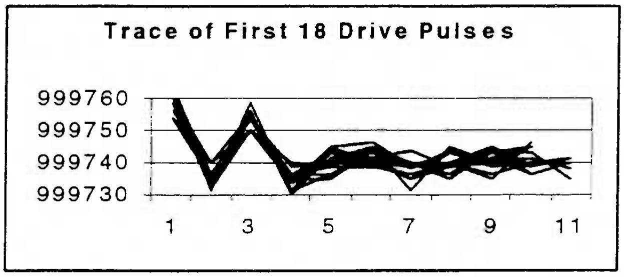

Now we can take up each jitter source effect by starting with the effect of drive impulses. As one would expect each drive impulse has a dramatic effect on period as shown in the inset figure above. Period changes by almost 30-µsec in an under-damped way. This plot starts with the cycle just after impulse and each point shows the period of each subsequent cycle. Eighteen such impulse and decay events are overlaid on this small plot. To get an exact picture of both mean period and jitter distribution I created figure 2. Basically figure 2 shows 11 normalized histograms, one for each sequential pendulum cycle following a drive impulse. About 680 impulse and decay events are shown. Each histogram bar represents the relative number of times the measured period fell into a 0.5-µsec bin (Period axis.) Each bar is normalized so that the maximum count is 7 dots wide. The width of all other bars is in proportion to the maximum. Histograms are of course usually shown rotated 90 degrees but here I wanted to be able to show 11 histograms on a common period scale and capture more than 7000 samples on a single plot. During this two-hour sample, run in the middle of the night, cooling caused a decrease in average period of almost 1-µsec. Temperature shift doesn't really change the shape of these distributions significantly since the distribution is much wider than 1-µsec.

Each cluster has a roughly normal distribution. The standard deviation is about 3-µsec. A number of outliers are seen above and below the main cluster. I take these to be greater than normal shock or vibration. (More below.)

Let me sum up drive induced variation. In sampling clock #7, or any such clock, we will from time to time see large changes in period from one cycle to the next. We could see a sample from the top of cycle 1 in figure 2 of 999.770 ms and the next from the bottom of cycle 2 of 999.722 ms. This is almost 50-µsec. In fact this isn't jitter at all but real variations in period due to drive. Actual jitter, which looks a lot like random variation, is about 15 to 20-µsec.

To see if ambient infrared radiation has much effect, I measured in daylight too. No effect was measurable. Note that direct sunlight cannot be allowed to fall on these sensors. Direct sunlight completely overloads (but doesn't damage) them.

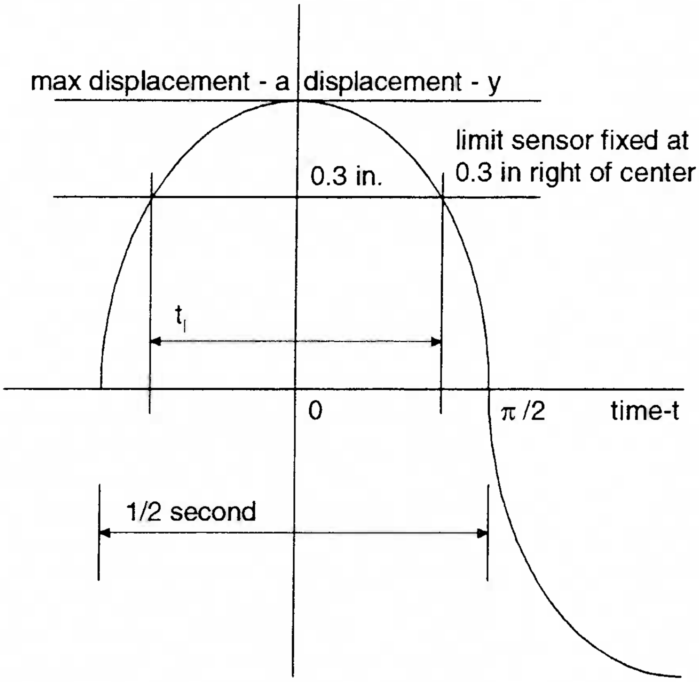

$t_l$ - time in limit sensor (seconds)

scale time so that 2 π radians is 1 second

displacement curve:
$y = a \cos(2 \pi t)$

a - maximum displacement

so

$0.3 = a \cos(2 \pi t_l / 2)$

$a = 0.3 / \cos(\pi t_l)$

Q is
$Q = \pi / \ln(a1 / a2)$

Next I measured the time interval during which the limit sensor is activated in order to determine variation in pendulum displacement. The output of the sensor phototransistor is pulled up to 5 volts by a 22KΩ resistor and feed to a 74HC14 Schmitt trigger to produce a clean logic level for the drive system. Turn-on threshold for the 74HC14 is 2.7 volts and turn-off threshold is about one volt lower at 1.8 volts. These levels mean that the phototransistor is only partially blocked by the shutter when the trip point is reached. In this range the sensor is a very short-range linear (or nearly linear) displacement detector.

Figure 3 shows a great deal of information about the decay of the pendulum as measured by the limit sensor. This sensor is only tripped for a short time near the right end of the swing. Like figure 2, I show a series of histograms so that the distribution of sample data can be clearly seen. Note how smoothly the interval decays and how closely clustered these samples are around the maximum. These time intervals can be accurately converted to displacement by looking at the inset figure below. Since we know that the displacement curve is a sinusoid it is relatively easy to convert the time interval clipped off at the top of this cosine curve into maximum displacement. The table in the lower left corner of figure 3 makes use of this relationship to give the actual amplitude at each cycle of the pendulum. The total change is only 0.002 inches. These displacements also give us an idea of period variation due to displacement change. As the table shows this is very small – only about 0.5-µsec maximum. In fact closed-loop control keeps mean displacement at about cycle #5. The actual error then is only about 0.1-µsec (±1 cycle) or less. Over a period of an hour or so this error is reduced to near zero. Variation in period due to variations in displacement is a tiny fraction of jitter seen.

As a bonus, cycle by cycle Q can be given by noting the exact mean decay from one cycle to the next. Figure 3 also shows Q. The mean is about 4500.

(The shutter for I/R switches needs to be heavy enough to block the beam from the LED. Paper, even heavy card stock, isn't enough. I use 1/16-inch aluminum.)

To the extent that drive impulses aren't exactly aligned with pendulum motion an unwanted torsion force is applied to the system. This torsion will (if it is anywhere near the torsional resonance point of the system) cause the shutter to be driven out of the plane of motion and wave about in the sensor aperture. This is another possible source of jitter. In addition impulses could setup in-plane vibratory modes too. This could be about the center of mass of the bob for example. Clock #7 is relatively small and ridged. A force of only ½ gram or less is applied by the electromagnet. While both of these jitter sources are surely present at some level, I judge by the smooth decay shown in figure 3 (cycle to cycle change of only about 0.00015 inches) that these sources are small. But below we'll see how important small vibrations are.

I did two experiments to get a rough idea of the effect of air currents and mechanical vibrations on period. Figure 4 is typical of all histograms for both experiments. In the first experiment I placed a fan 10 feet from the pendulum and directed the flow so that it didn't actually fall on the pendulum but so that an occasional breath of air would touch the shaft and bob. The fan was on low. In the second I placed the fan, also on low, under the pendulum's table so that the breeze would bump the legs and the motor would vibrate the floor. There really isn't any real difference between these two groups. In both cases jitter increases greatly from 15-µsec, shown in figure 2 with quiet conditions, to 150-µsec or more. This would increase the standard deviation to about 25-µsec. The distribution isn't normal but it does show a strong central tendency which is skewed to the right. Clearly period jitter is very sensitive to any form of mechanical disturbance.

Instrument errors refer to jitter introduced into a digital interval-timing system through time-base error, ±1 count error, and electrical noise. Time base error refers to error introduced by short-term variations in the oscillator used to drive the counter. In my case this an 8-MHz crystal driving the HC11. Over intervals of an hour or two I assume this error is very much smaller than 0.5-µsec. I ignore this error for now. (There is also an error related to the absolute frequency of the crystal. I know from experiment that this crystal is at least 30 ppm off. This effects the absolute period measurement of course. But since for this work I am only interested in variations of relatively short periods, I am also ignoring this error in this paper.)

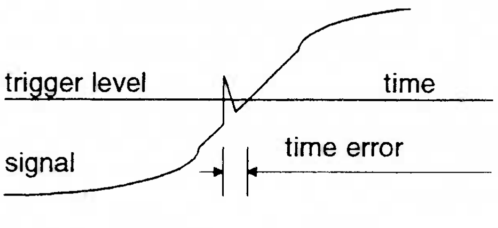

All sampling digital interval timers have a ±1-count error. Hewlett Packard's application note #200 "Fundamentals of the Electronic Counters" (www.hp.com) has a complete description of this error. Simply put it is impossible to know, on the digital counter side of things, whether or not we just missed a change in the measured signal. In my case this error is fixed at ±0.5-µsec.

Finally if electrical noise, or in my case infrared radiation noise, is present it may trigger the counter a bit too soon or a bit too late. The inset figure shows how this happens (below left.) This error is related both to the amount of noise and the rate of voltage change near trigger threshold. Higher noise or slower signals give higher error. With an oscilloscope I measured voltage change of the position infrared sensor as it enters the Schmitt trigger to be 625 volts/second. As an example, peak to peak sensor noise of 0.001 volt would result in a ±1.6-µsec error. In fact I don't see that large a noise component; however, this is rather hard to measure. So even if we accept noise at this level, total error is ±2.1-µsec. (±0.5 and ±1.6) (Again the HP paper for many more details.)

(Power supply noise is introduced identically into both the Schmitt trigger and the sensor. Power supply noise is then found directly in the threshold levels of the 74HC14. So power supply noise rejected. This is call 'ratiometric' sampling.)

Instrument errors account for only about one quarter of the jitter seen in figure 2 (4.2 of 15-µsec.)

As I mentioned above these sensors actually have a short range in which they can be used to measure displacement directly. This is because most of the phototransistor surface is the photosensitive base structure. Infrared radiation from the IR Led floods this small square surface. As the shadow of the shutter crosses the beam, photocurrent is linearly reduced. As photocurrent drops below the transistor's saturation point, collector current drops in proportion transistor gain. This current flows through a 22K resistor so we get a linear voltage change. We can calibrate the position sensor in the following way. A pendulum, of the sort of clock #7's, has a maximum velocity at bottom of swing of 47.8 millimeters/second. The linear part of the position sensor's voltage waveform is 625 volts/second. Dividing we get 0.0765 millimeters/volt. Why is this interesting?

Mechanical vibration (air currents, frame motion, torsion, etc.) is coupled to the position sensor by the shutter. To account for the remaining 10-µsec of jitter only (0.0765mm/v x 625v/s x 10-µsec) 0.0005 millimeters (20 millionths of an inch) of vibration is necessary! This vibration can come at any frequency. (This makes it next to impossible to see on an oscilloscope.)

Finally a summary:

Under quiet conditions observed jitter is about 15-µsec. This jitter is approximately normally distributed about the mean. This is good news because as more and more samples are combined the average tends toward the true mean. (The statistical law of large numbers.) Only a thousand or so samples are needed to get a good estimate of the mean. Since the standard error of the mean is proportional to the square root of the sample size, 20 minutes sampling gets to within 3% of the true mean. Within such a short sample, temperature and other effects are not very important.

In my free pendulum with drive only about every 11 seconds, drive impulses pull the period well away from the long-term average. The mean period right after a drive impulse is applied is about 18-µsec longer than average. Figure 2 shows that the first 4 cycles after drive are +18-µsec -4-µsec, +14-µsec and -6-µsec from free period shown in the last 7 cycles. (+22-µsec. altogether) On average over 11 cycles this is 2-µsec/second or 2 ppm. Any variation in this impulse and decay pattern will change average clock period. These experiments do show that the pattern remains fixed for long periods of time.

All forms of mechanical noise whether against the bob or against the frame have a huge effect upon jitter. (Drive impulse induced vibratory oscillation has to be considered part of overall mechanical noise.) Jitter is increased by a factor of 10 even for quite modest vibration. I have shown how sensitive I/R detectors are to this noise. Various mechanical and electric measures can be taken to reduce noise. If only reasonable care is taken jitter is reduced to 15-µsec (as shown in figure 2.) I do not currently know how much this number could be reduced if extreme measures were taken. This jitter seems to be random with (perhaps) normal distribution. Based on my assumptions of electrical noise above, it doesn't seem likely that jitter can be reduced below about 4-µsec even if there was no mechanical disturbance. My working hypothesis remains that mechanical disturbance is the major non-systematic (not temperature, barometer, circular error, or drive variation) period jitter in free pendulums measured with infrared sensors.

Comments or questions: drbob@casaalta.com.

Figure 4. Typical Period Histogram with Increased Vibration/Air Currents

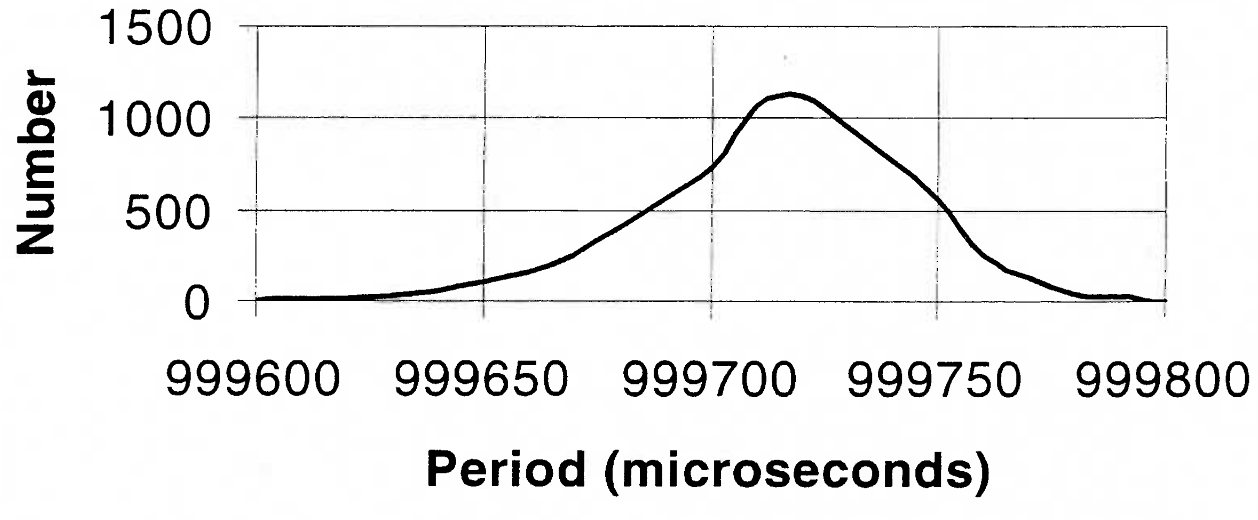

Figure 1. Clock #7 24 Hours - Period Check and Amount of Drive - 4/13/99 hr#10 is 6AM

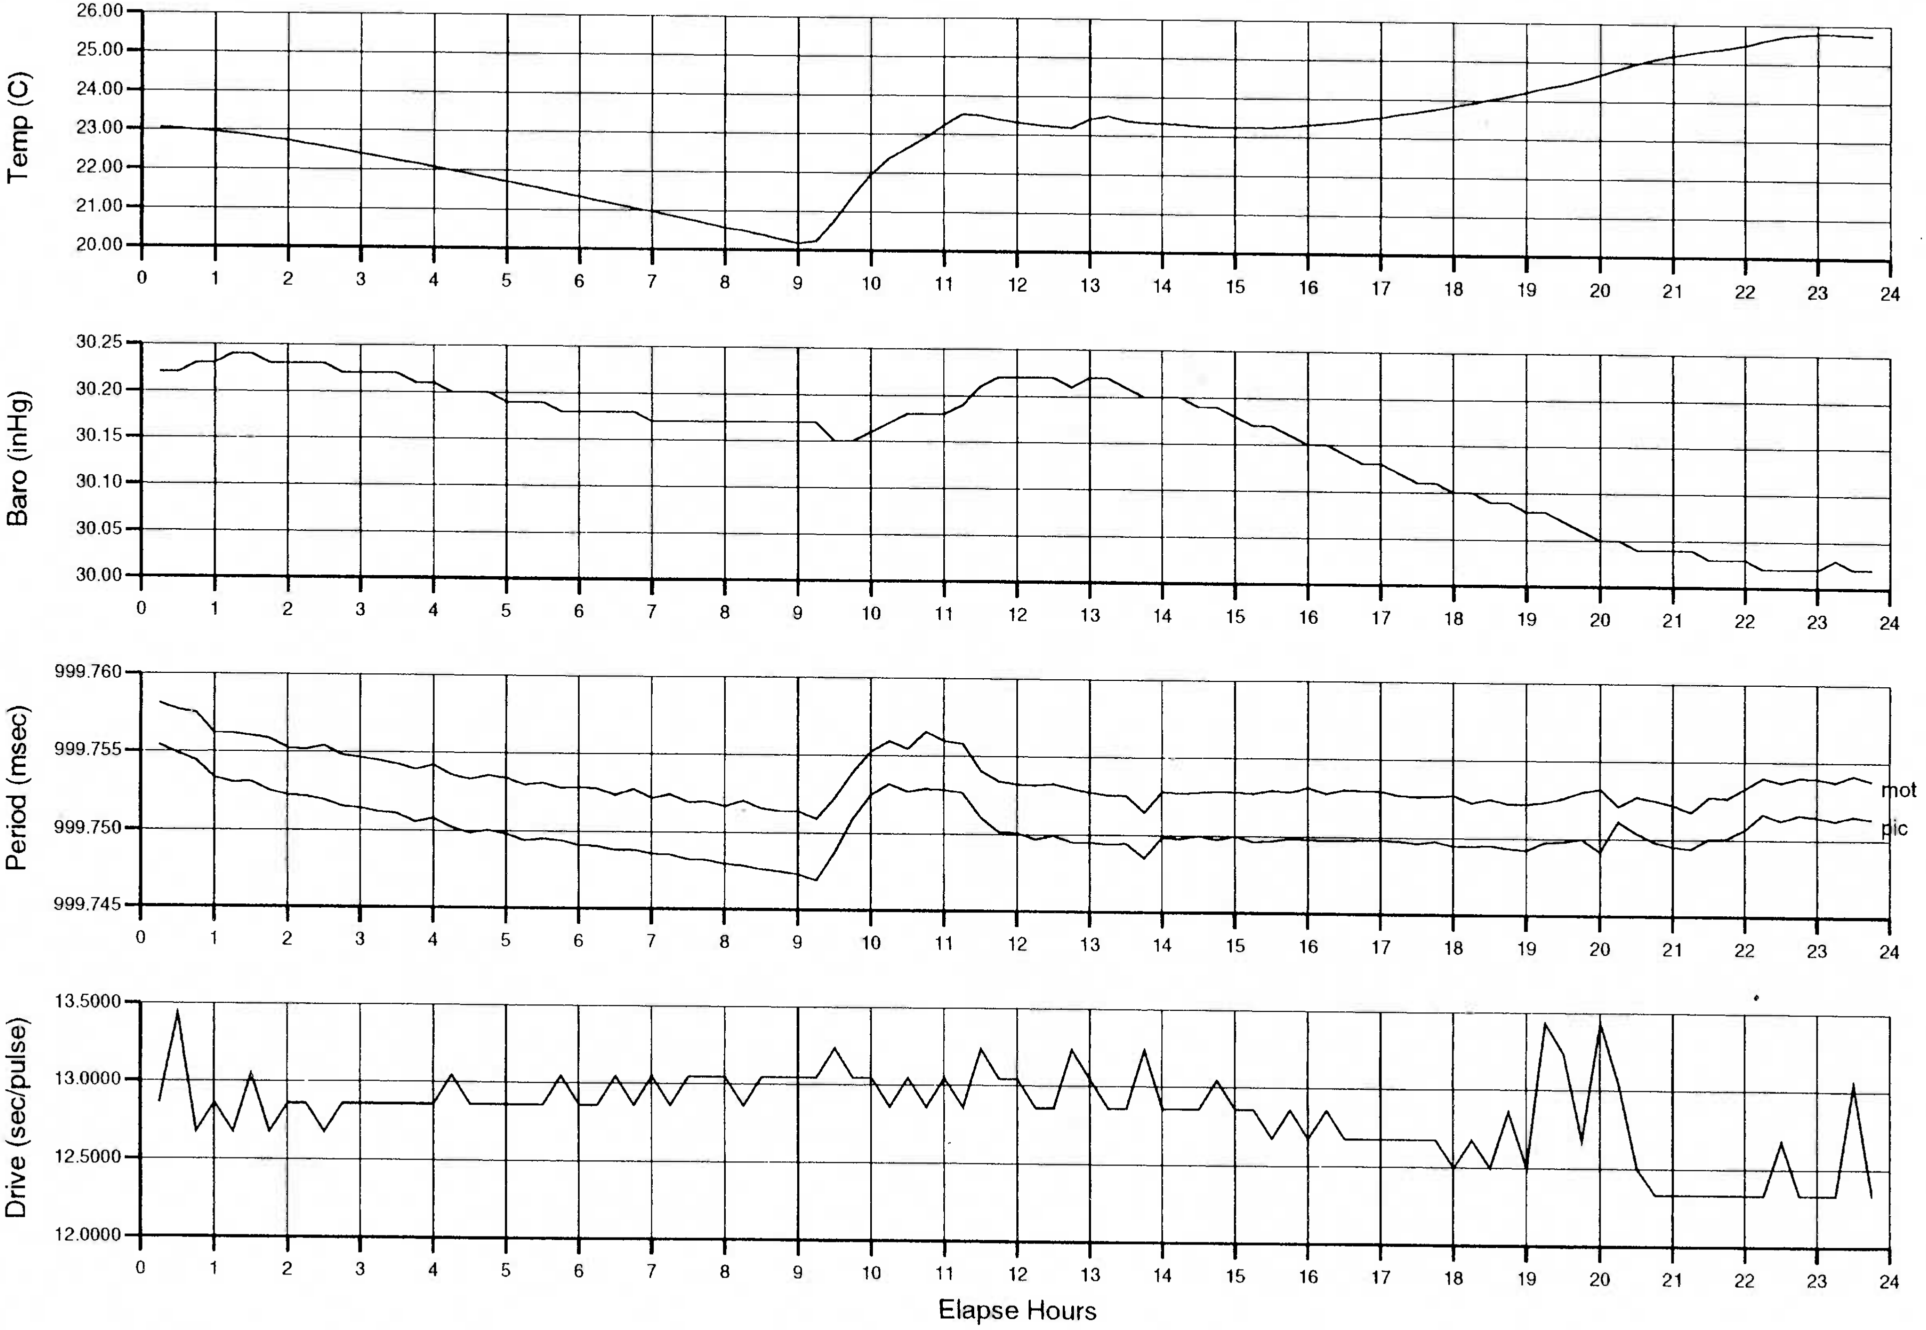

Figure 2. Period Histograms for each Cycle following Each Drive Pulse

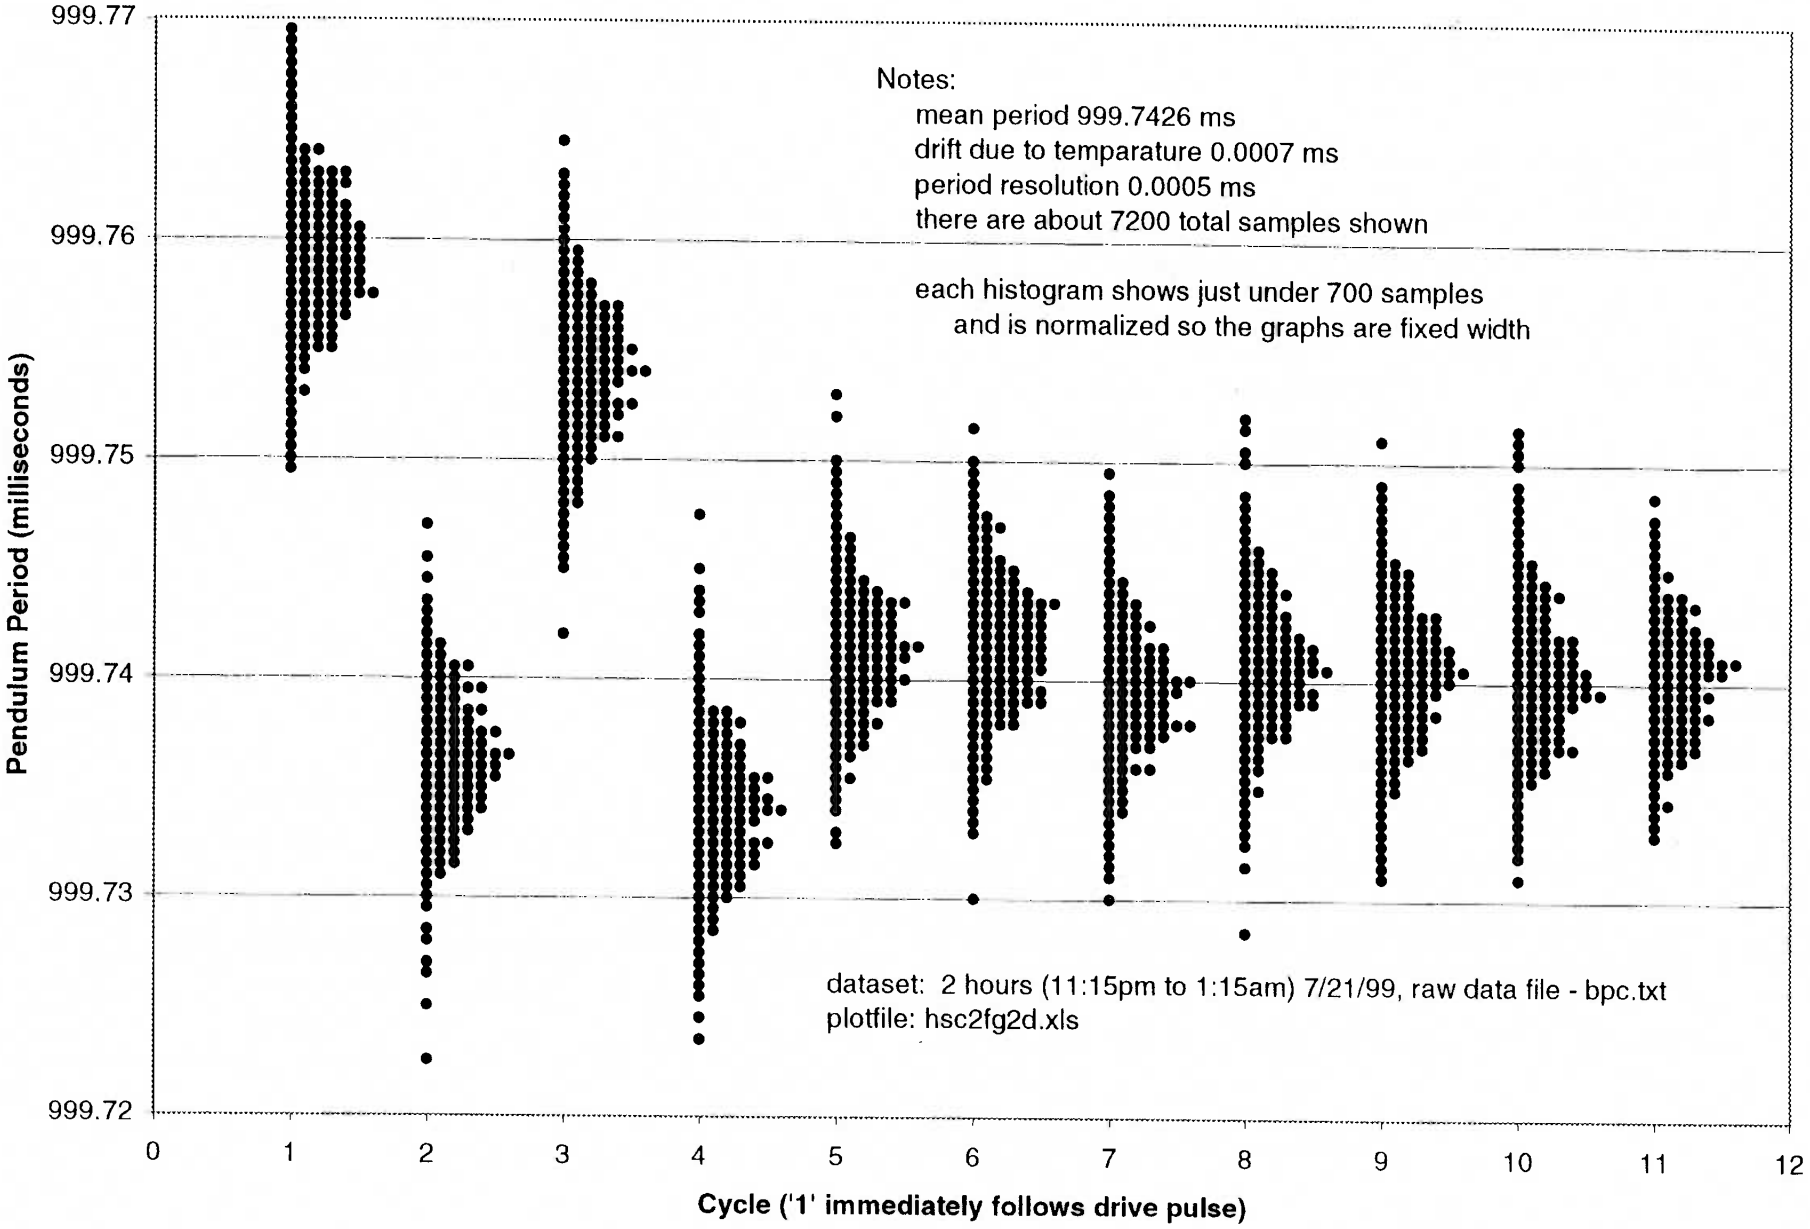

Notes:
mean period 999.7426 ms
drift due to temparature 0.0007 ms
period resolution 0.0005 ms
there are about 7200 total samples shown

each histogram shows just under 700 samples
and is normalized so the graphs are fixed width

dataset: 2 hours (11:15pm to 1:15am) 7/21/99, raw data file - bpc.txt
plotfile: hsc2fg2d.xls

**Figure 3. Decay of Pendulum Displacement after Drive Pulses**

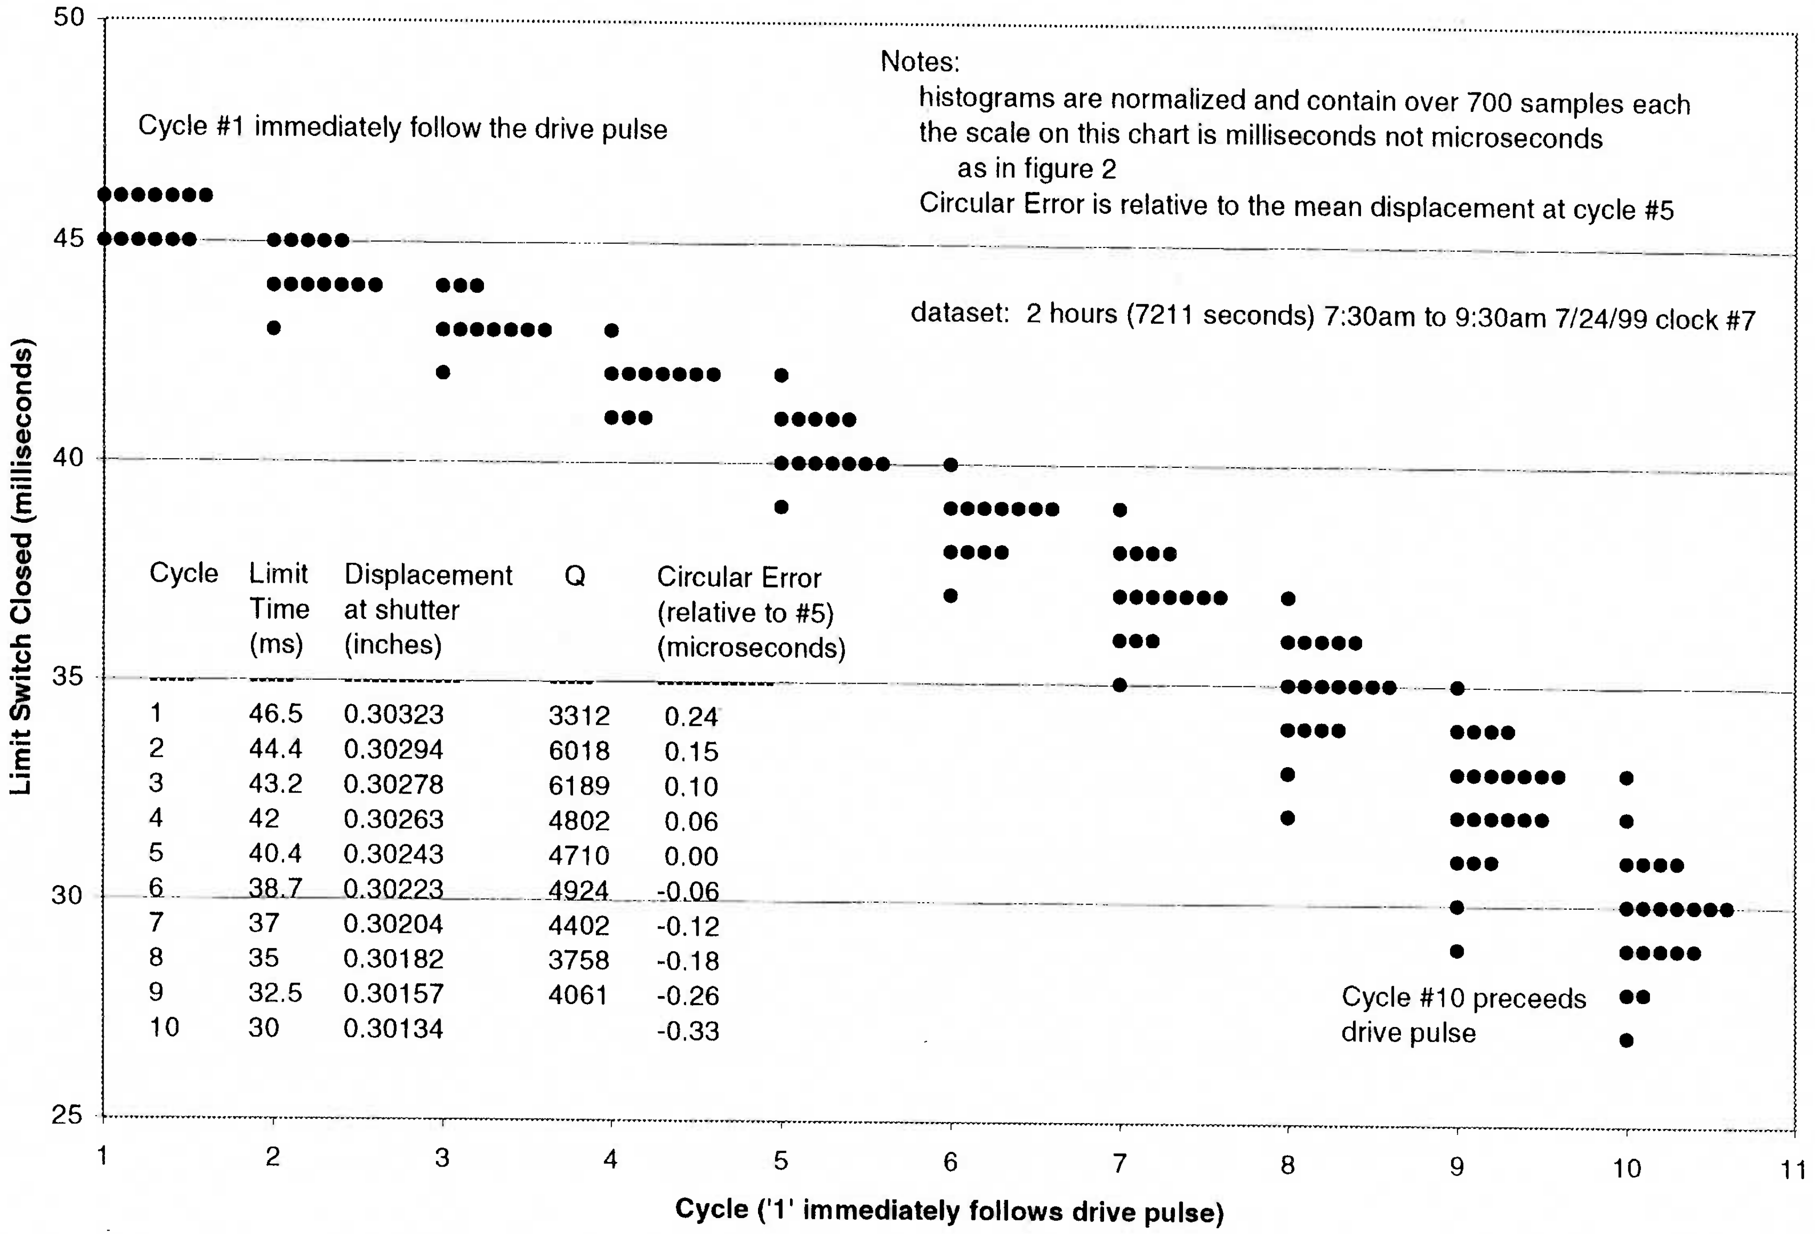

Cycle #1 immediately follow the drive pulse

Notes:

histograms are normalized and contain over 700 samples each

the scale on this chart is milliseconds not microseconds as in figure 2

Circular Error is relative to the mean displacement at cycle #5

dataset: 2 hours (7211 seconds) 7:30am to 9:30am 7/24/99 clock #7

| Cycle | Limit Time (ms) | Displacement at shutter (inches) | Q | Circular Error (relative to #5) (microseconds) |
|---|---|---|---|---|
| 1 | 46.5 | 0.30323 | 3312 | 0.24 |
| 2 | 44.4 | 0.30294 | 6018 | 0.15 |
| 3 | 43.2 | 0.30278 | 6189 | 0.10 |
| 4 | 42 | 0.30263 | 4802 | 0.06 |
| 5 | 40.4 | 0.30243 | 4710 | 0.00 |
| 6 | 38.7 | 0.30223 | 4924 | -0.06 |
| 7 | 37 | 0.30204 | 4402 | -0.12 |
| 8 | 35 | 0.30182 | 3758 | -0.18 |
| 9 | 32.5 | 0.30157 | 4061 | -0.26 |
| 10 | 30 | 0.30134 |  | -0.33 |

Cycle #10 preceeds drive pulse