A Shortt Lesson in Allan Deviation (Tides, part 2)

Tom Van Baak, tvb@LeapSecond.com

In part 1 of this article the origin, derivation, and magnitude of lunar/solar tides was discussed. We know tides cause subtle changes in pendulum clock rate; changes in rate are instability; and an appropriate statistic for calculating clock stability is the Allan deviation. In this section, the power of Allan deviation plots to "detect tides" is explored using real data from a well-known Shortt free pendulum clock.

Shortt Number 41

The history of Shortt #41 has been covered before. We are fortunate to have access to the raw data of Pierre Boucheron's mid-1980's timing experiment with a Shortt free pendulum clock. This data was subsequently digitized and repaired by Philip Woodward and was made available to me through contacts with Bob Holmström, Bryan Mumford, and Jerry Walker. Thanks to you all.

This SH41 data set consists of approximately 8500 hours (352 days) of measurement of a Shortt clock and provides a unique, real-world example with which to explore the power of the Allan deviation statistic.

Allan Deviation

Clock accuracy (relative to some standard) can be determined with a single, instantaneous time measurement. Calculation of clock rate, on the other hand, requires at least two timing measurements separated by a fixed interval. The result is a measure of *average* clock rate; all rate measurements imply an averaging interval. Finally, it takes three or more time measurements, which is two or more rate measurements, in order to determine if a clock is running at a *consistent* rate.

Since it is assumed that clocks with time offsets can always be set and that clocks with rate offsets can be always be regulated, a truer measure of intrinsic clock quality should be neither accuracy, nor rate, but stability – which is a measure of the *consistency* of rate. Allan deviation (ADEV) is a statistical formula which computes clock stability. As more and more time or rate measurements are made at regular intervals a clearer measure of rate stability emerges. Note the phrase "Allan variance", or AVAR, was popular in the past (and that AVAR = ADEV^2).

An ADEV value is essentially a prediction, based on many averaging intervals in the past, of how far the clock rate is likely to drift during one interval into the future.

For example, for one day averages, the ADEV of SH41 is 8.72×10^{-9}. This means based on statistics from all days in the past, the clock is predicted to keep its rate constant to within 8.72×10^{-9} tomorrow (which corresponds to a time drift of 0.75 milliseconds per day).

Allan Deviation Plots

By choosing different averaging intervals, many Allan deviation values can be calculated from the same data set. For example, for hourly averages (instead of daily) the ADEV of SH41 is

5.32×10^{-8}. Even one or two point calculations like these are sufficient to compare the stabilities of one clock with another.

But individual numbers alone do make full use of all the information available. ADEV is a more powerful tool when a single plot is made with points of many different averaging times (often called ***tau***). It is conventional to make a log-log plot where the x-axis is the averaging time (tau) and the y-axis is the calculated relative stability (sigma); a so-called *sigma-tau* plot.

When this is done to its fullest an amazing and sometimes bewildering arrangement of lines and slopes and bumps are visible. A slight increase or decrease in stability may occur as averaging time gets shorter or longer. It is this set of upward, downward, flat, or curvy patterns where a log-log Allan deviation plot can act like a diagnostic "X-ray", revealing the internal workings of a clock.

We will now look into **five factors** in the creation of an ADEV plot and see what lessons it provides to those of us who wish to "see tides" in our pendulum clocks. For ease of comparison, all of the plots below use the same scale: tau (x-axis) spanning decades of averaging time from 10^3 to 10^8 seconds (less than an hour to more than a year) and sigma (y-axis) spanning decades of relative stability from 10^{-6} to 10^{-10}.

1. Sampling interval

There is no rule which states how often you have to take data from a pendulum clock. For a standard seconds pendulum one could make a measurement as often as every second or two. For a long-term experiment it might be more convenient to record a measurement only every minute, hour, or even once a day. What effect would this decision have on detecting tides?

A plot like *figure 1d* is our goal. Without further explanation here, the bumps in the first half of the ADEV plot corresponds to the effect of lunar/solar tides. Not all ADEV plots reveal tides as well as this. Why?

The ADEV plot in *figure 1a* was made from selecting **daily** samples from SH41. Notice that no points are present for averaging times less than 1 day: with one sample per day, one can only calculate ADEV for multiples of one day. By contrast the plot in *figure 1b* was made from **hourly** SH41 samples, and so there are points on the graph for multiples of one hour.

One cannot obtain stability statistics for averaging times less than the sampling interval. Tides mostly perturb stability for averaging intervals of a couple of hours to a couple of days. Thus one cannot adequately detect tides if one does not have hourly (or shorter) samples.

We observe that daily samples completely hide the effect of tides; one must take samples at least once an hour if the goal to observe the effect of tides on clock performance.

2. Calculation interval

Although the standard Allan deviation formula is quite simple, there are useful variations. In *figure 1b* the averaging times calculated are 1 hour, 2 hours, 4 hours, and 8 hours, etc. This is convenient in many cases since a few calculations cover an exponential range of tau. Going up by factors of two is an octave scale. One can also go up by decades: 1, 10, 100, 1000, etc. Or a combination as in: 1, 2, 4; 10, 20, 40; 100, 200, 400, etc. All these are commonly used.

But if one uses modern computing power to compute ADEV for *every possible* integer hourly interval starting from 1 hour the plot looks quite different as seen in *figure 1c*. This calculation computes ***all tau*** instead of a small logarithmic subset of tau and the result is that the plot reveals more structure of the tides.

We observe that ADEV plots convey much more information if calculations are performed for all tau rather than just some tau. The extra computation is worth the effort if the goal is to detect the influence of tides on a pendulum clock.

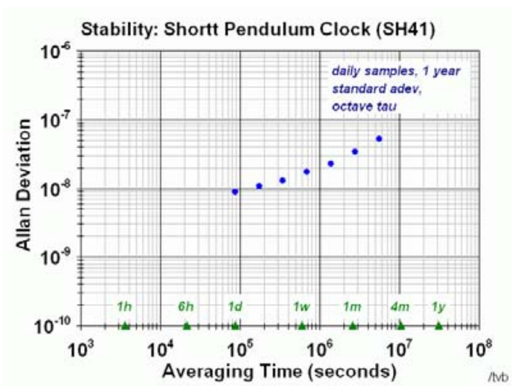

Figure 1a – daily

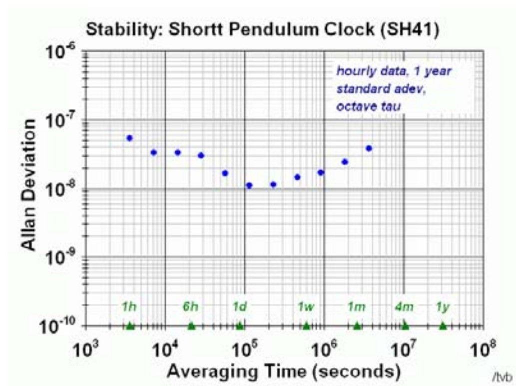

Figure 1b – hourly

Figure 1c – all tau

Figure 1d – overlapping

3. Overlapping calculations

A third trick to further clarify the influence of tides in an ADEV plot is to use overlapping samples rather than contiguous samples. For example, to calculate the ADEV for tau 1 day one simply needs to use hourly points that are one day apart; noon to noon to noon, etc. But the points 1PM to 1PM to 1PM are also one day apart, as well as 3AM to 3AM to 3AM. If one includes *all* possible pairs that are 24 hours apart, then the total number of elements in the statistic increases by about a factor of 24. The result is less uncertainly in the statistic, a better signal/noise ratio.

In *figure 1d*, an overlapping ADEV calculation gives a cleaner line, with more detail.

In summary, several calculation techniques can be combined to produce the best Allan deviation plot and maximize the evidence of tidal influence on clock stability. Use of hourly rather than daily data is mandatory. The combined use of all-tau calculations and overlapping intervals produces remarkably clear plots.

Figure 2a – full 1 year

Figure 2b – 90 days

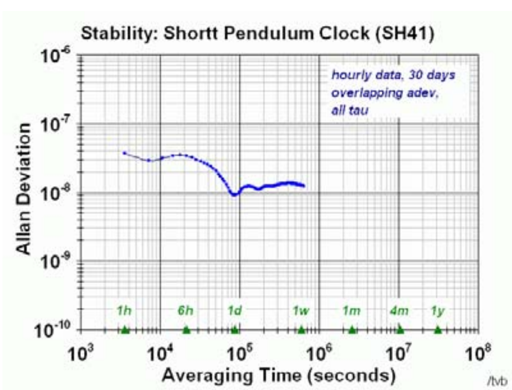

Figure 2c – 30 days

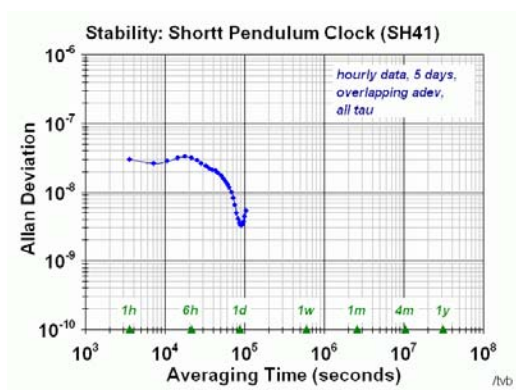

Figure 2d – only 5 days

4. Length of data set

When we use a voltmeter we expect an instant reading. If we use a thermometer we know we might have to wait for the temperature to settle. But when we measure a pendulum clock we think we need weeks, months, even years of data. Is this assumption correct?

The ADEV plot in *figure 2a* is made from almost a full year of SH41 data. By contrast, *figure 2b* is from the first 90 days only. Notice that stability points for larger tau are missing. But notice also that all points for shorter tau are identical, in spite of much less raw data.

Similarly, *figure 2c* is made from 30 days of data. And finally, *figure 2d* is made from just the first 5 days (120 hours) of the SH41 data set. The number of points plotted on the right side of the plot (longer averaging times) is dependent on the length of the data set used. But it is also instructive to see that the shape of the ADEV plot on the left (shorter averaging times) is quite immune to the length of the data set.

There seems to be an expectation that one must collect a massive amount of data in order to obtain precise results. But ADEV calculations can show if a clock is stable enough to detect tides in a matter of days. One does not need weeks or months or years of data.

The length of the data set has more to do with how far to the right an ADEV plot extends (long tau) than it does with how an ADEV plot looks on the left (short tau). This is expected when we think about it. And it implies seeing tides is not something that requires extended run times.

Figure 3a – all 5 digits

Figure 3b – 4 digits

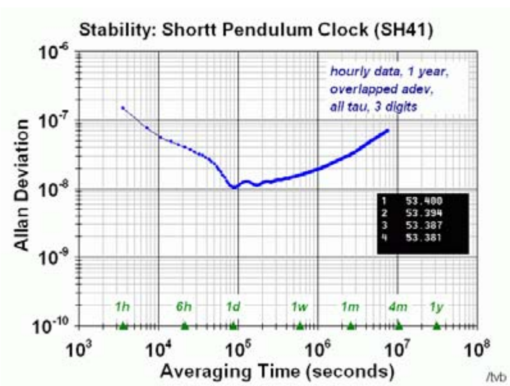

Figure 3c – 3 digits

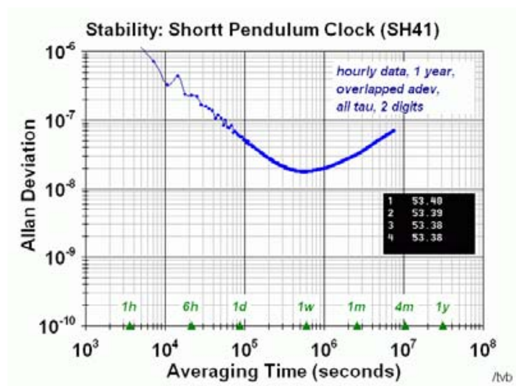

Figure 3d – just 2 digits

5. Measurement resolution

Lastly, we look into the nature of the raw data itself. Allan deviation is based on periodically spaced time error measurements between the clock being tested and some reference clock. How accurate or how precise do these measurements need to be? Is it important?

Each hourly point in the SH41 raw data file has 10 µs (microsecond) resolution; that is 5 decimal places, as shown in *figure 3a*. We can edit Boucheron's raw data to remove the 5th digit of every point and re-run the calculations to produce *figure 3b*; we now have only 4 decimal places in each measurement, 0.1 ms (millisecond) of resolution. Interestingly, it makes almost no difference in the plot.

In *figure 3c*, 2 decimal places have been deleted leaving only 1 ms resolution. We can see that most of the tidal hump is now gone; a few slight wiggles remain. Finally, *figure 3d* shows the ADEV of SH41 as if it were measured with only 10 ms resolution. Same good clock, but the poor measurement resolution has completely hidden any trace of tides.

This set of Allan deviation plots strongly conveys the importance of measurement resolution. Too much resolution adds little or nothing to the stability plots. Too little resolution can partially or completely hide subtle effects such as tides. One can downgrade high-resolution data to but if data is collected at low-resolution nothing can be done to improve it.

Conclusion

With the best of conditions the effect of tides on pendulum clocks can be seen directly in rate or progressive time error charts (e.g., Fedchenko gravimeters via George Feinstein). We have also seen the effect of tides through Fourier spectral analysis (e.g., Philip Woodward). In this section we explored how an Allan deviation plot can be used to see the effect of tides on a precision pendulum clock.

In summary, the effect of tides is more easily revealed in an Allan deviation plot if certain techniques are followed: short sample rates, plotting all tau, and calculating overlapping samples.

Moreover, the tidal signature is revealed rather quickly; only a few days of data are required for reliable statistics. Longer run times improve the quality of the analysis but not significantly.

Finally, there is a minimum requirement on data resolution in order for tides to be detected. Data with higher resolution adds little value to the analysis. On the other hand, data without sufficient resolution cannot ever reveal tides, no matter how long the run is.

A rule-of-thumb for an experimenter to detect tides in their precision pendulum clock is to collect hourly samples with 0.1 millisecond of resolution, or better. If tides are not visible in Allan deviation (or spectrum analysis) plots it is safe to conclude that either the clock itself is intrinsically too unstable to detect tides or the clock itself is fine but the measurement system is not recording precise enough samples.

In the next section we will explore in detail the Allan deviation of two other clocks: a so-called *Dream* pendulum clock, and the infamous *Littlemore* clock.