Time Dependence Study of Length Control Calibration During the April 2000 Engineering Run at LHO

Peter Shawhan (LIGO/Caltech)
May 30, 2000

This note describes a study I did to check the stability of the length control calibration procedure used during the "one-arm engineering run" conducted at LHO on April 3-4, 2000. Various information about the run may be found at the LHO Engineering Run Home Page, `http://blue.ligo-wa.caltech.edu/engrun`.

The basic calibration procedure is described in an electronic log entry at 16:52:19 on April 3. The technique was to use the GDS excitation engine to drive the end test mass sinusoidally at 30, 300, and 1000 Hz, and to measure the resulting magnitudes of these lines in the power spectrum for the length control error signal. This test was done roughly every four hours while data was being collected; each time (except one), a single value was recorded for each frequency. The values are listed in an electronic log entry at 20:13:16 on April 4. There is some variation over time; it is tempting to assume this is due to natural long-term drifts in the transfer function.

I repeated the calibration procedure offline using the data stored on "fortress", with a few changes. First, I made an effort to subtract the baseline noise under the peaks by measuring the power at frequencies just below and just above the peak; this was a correction of 2-3% at 30 Hz and around 6% at 300 and 1000 Hz. Second, I performed multiple calibrations during the time interval when the excitation was turned on, which typically was several minutes. The procedure consumes 44 seconds of data, so I spaced the calibrations a minute apart to avoid overlap. My goal was to determine how repeatable the calibration procedure is over short time scales, i.e. to try to determine whether the variability is due to: 1) long-term drifts; 2) short-term drifts; or 3) point-to-point measurement error.

Figure 1 shows my calibration results for three different time intervals during which the excitation was turned on. Details about the time intervals, etc., are given in the figure caption. I draw the following conclusions:

- The first time interval, during which the arm lock was stable, shows that there is some real point-to-point variability which is larger than I would have expected. It is also significantly different at the three frequencies: the rms is about 1% at 30 Hz, 3% at 300 Hz, and 1.5% at 1000 Hz.
- The second and third time intervals give two different pictures of how the transfer function changes just before the arm loses lock. In one case, there is no sign that lock is about to be lost; in the other case, the transfer function changes dramatically in the last few minutes prior to loss of lock. Perhaps two different lock-loss mechanisms were at work in the two cases? It is also interesting that the transfer function drift in the latter case is upward at 30 and 300 Hz but downward at 1000 Hz.
- The second time interval shows that the transfer function is a strong function of the ASC centering/alignment servo gains (or perhaps is influenced by the source of the ITM position input, which is different during locking?). The shift directions at the three frequencies follow the same up/up/down pattern as was seen before loss of lock during the third time interval.

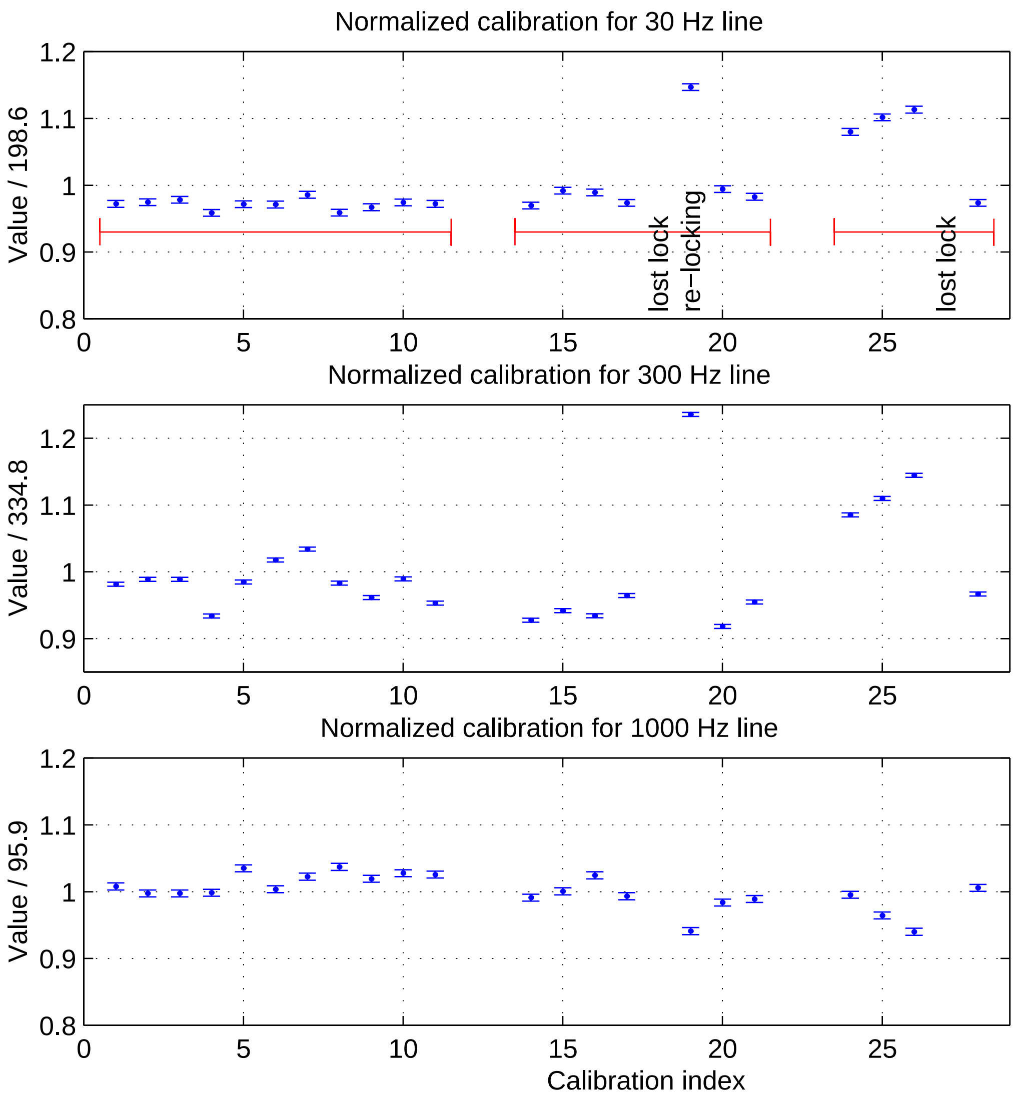

FIGURE 1: Measured calibration values at each frequency, normalized to the average for that frequency (to make it easier to read off fractional changes). The error bars are a naive estimate of the measurement error (basically, a generous estimate of the uncertainty in subtracting the baseline noise, but without including any estimate of the statistical error in the measurement of the peak height). The horizontal bars in the top plot group calibrations taken during each of three time periods when the excitation was turned on: April 3, 20:14-20:25 UTC; April 4, 0:00-0:09 UTC; and April 4, 9:14-9:23 UTC. Consecutive calibrations are separated in time by one minute, except when the arm lock was lost during the second and third intervals and it took a few minutes for the auto-locking script to restore lock. I also did a calibration during the second time interval immediately after lock was re-acquired, while the auto-locking script was still adjusting the ASC alignment/centering gains, which had a big effect on the transfer function, as shown.

After performing all of the calibrations, I went back and checked whether the actual motion of the end test mass could have a time-varying amplitude and could cause the apparent time variation of the calibration numbers. I did this by measuring the signals on the four shadow sensors at 30 Hz and 300 Hz. (Although the shadow sensors are sampled at 2048 Hz, the "Fourier Tools" software refuses to calculate a power spectrum all the way up to the Nyquist frequency.) The results, shown in Figure 2, indicate that the actual motion of the end test mass is quite stable, except during the one calibration when the auto-locking script was still adjusting the ASC alignment/centering gains—and in this case, the magnitude of the shifts do not really match the magnitudes of the shifts for the length control error signal, which is curious. The most interesting observation is that no shift in the shadow-sensor peak heights is observed during the third time interval just before losing lock, unlike the length-control-error-signal peak heights. (For this study, I did not take the time to estimate and subtract baselines point-by-point, but I did make sure that the baselines were no larger during the "re-locking" measurement than during the immediately following measurement. In fact, they were somewhat smaller.)

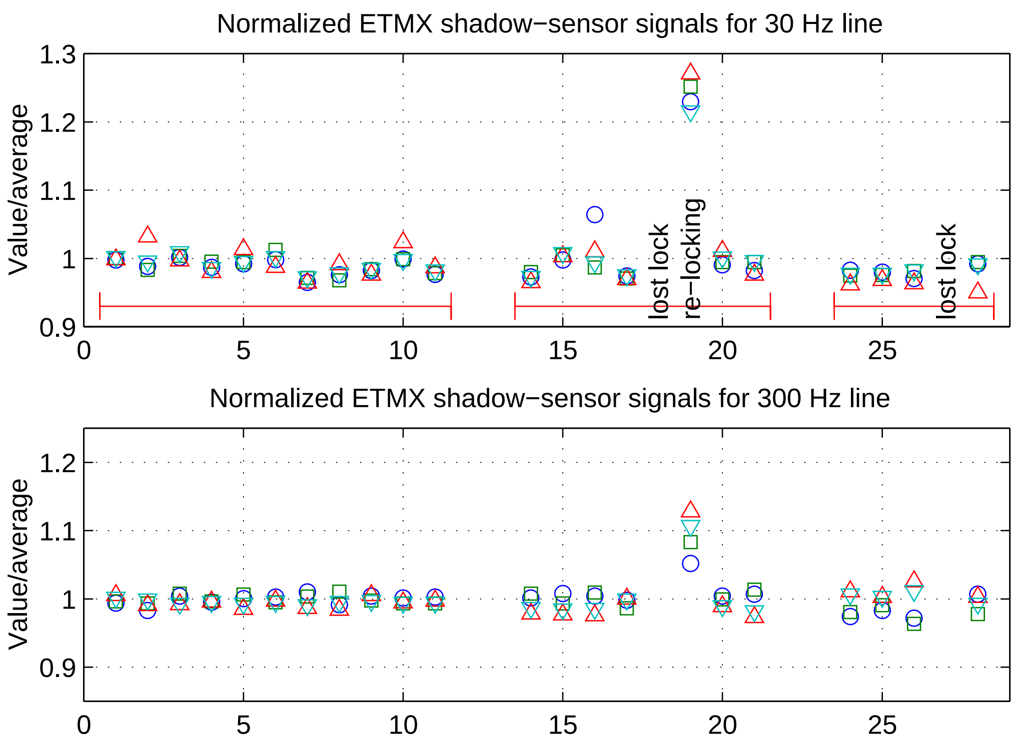

FIGURE 2: Normalized peak heights for 30 Hz and 300 Hz lines in shadow sensor signals from the end test mass. The four different symbols represent the four shadow sensors at the corners of the end test mass. The time periods are the same as in Figure 1; the vertical scales are the same as well, though shifted in the case of the top plot.