Time Dependence Study of Length Control Calibration During the April 2000 Engineering Run at LHO (Revised Version)

Peter Shawhan (LIGO/Caltech)
June 14, 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. (This is a revised version, thanks to feedback from Michael Landry and Stan Whitcomb on my earlier version, dated May 30.) 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 spectra for the length-control error signal (H2:LSC-AS_Q_TEMP) and the control signal fed back to the position of the input test mass (H2:LSC-AS_I_TEMP). This test was done roughly every four hours while data was being collected, and a value was recorded at each frequency for each channel. (In one case, sets of values were recorded just before and just after re-locking.) 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, but that is not the only possible explanation.

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 the control signal for three different time intervals during which the excitation was turned on. Details about the time intervals are given in the figure caption. At each frequency, the individual values have been divided by the average of all 22 values for that frequency, to make it easier to read off fractional changes from the plot. The error bars are a somewhat naive estimate of the measurement error—basically, a generous estimate of the uncertainty in subtracting the baseline noise, but without attempting to evaluate the statistical error in the measurement of the peak height, which is thought to be small. There are a few conclusions which can be drawn from these measurements:

- The first time interval, during which the arm lock was stable, shows that there is point-to-point variability which is larger than the estimated measurement error. 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. Based on this figure alone, it is unclear whether this variability is due to real short-term variation or whether the true measurement error is larger than the estimate.

- 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. This pattern is not entirely unreasonable, since the unity gain frequency is something like 600 Hz, and frequencies above and below it could be affected differently.
- 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 to the servo, 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 2 overlays the calibration values for both the error signal and the control signal. The time variability is remarkably well correlated between the two, indicating that the measurement errors are indeed small and the gain of the length-control servo is quite stable. Thus we may conclude that the observed variability is due to real minute-to-minute variation of the plant transfer function (end test mass to error signal).

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 3, 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 magnitudes of the ETM-motion shifts do not really match the magnitudes of the shifts for the error signal and length control 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.)

As a final check, I looked at the shadow-sensor signals for the *input* test mass to see whether the excitation applied to the ETM could be inducing motion of the ITM. There was no sign of a peak at 30 Hz (the only frequency I could check, since the ITM shadow sensors were sampled at 256 Hz). Therefore, the excitation applied to the ETM represented the true length variation of the arm cavity at those frequencies.

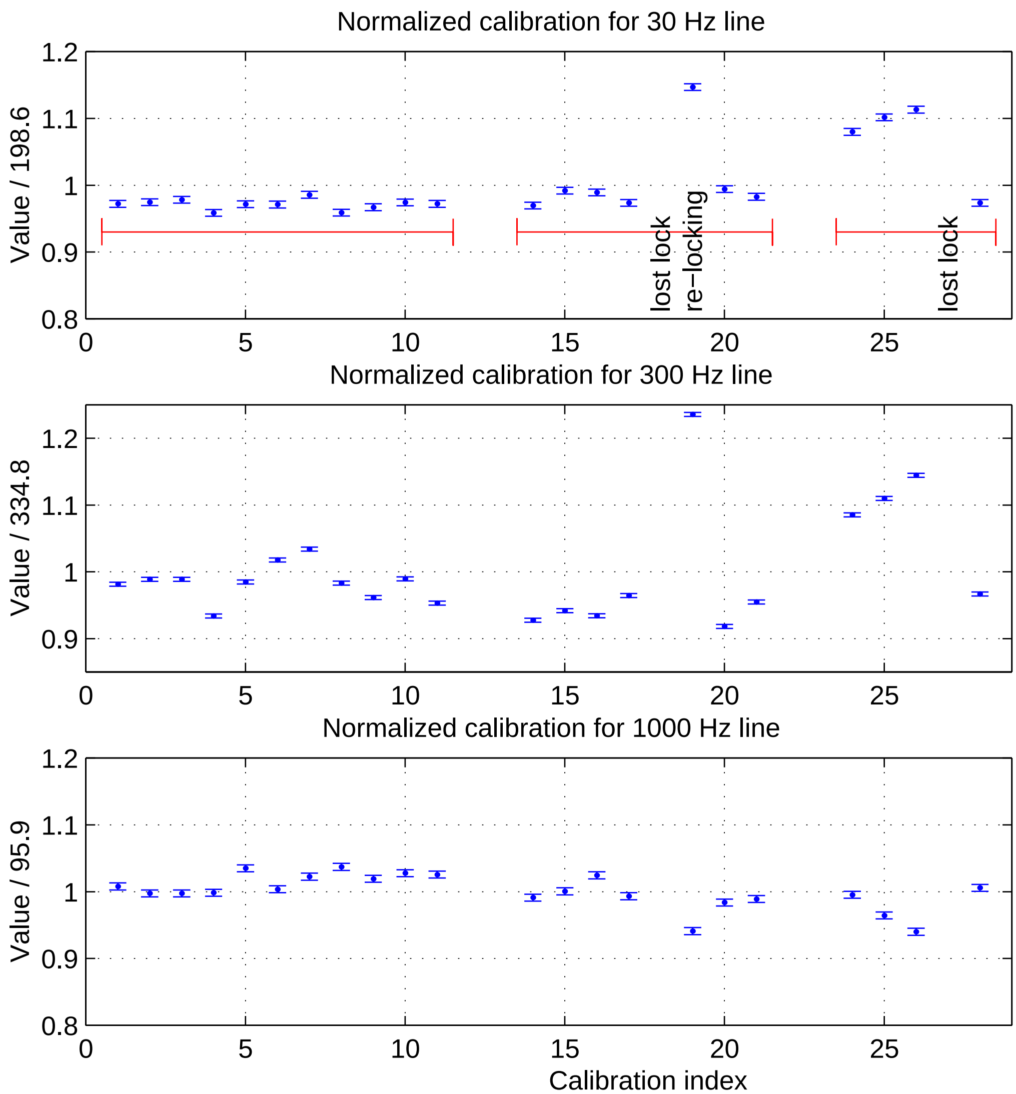

FIGURE 1: Measured calibration values for the control signal at each frequency, normalized to the average for that frequency. The horizontal bars in the top plot group calibrations taken during each of three time periods when the excitation was turned on: April 3, 20:14-20:25 PDT; April 4, 0:00-0:09 PDT; and April 4, 9:14-9:23 PDT. 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. The calibration labeled “re-locking” was done immediately after lock was re-acquired, while the auto-locking script was still adjusting the ASC alignment/centering gains; this had a big effect on the transfer function, as shown.

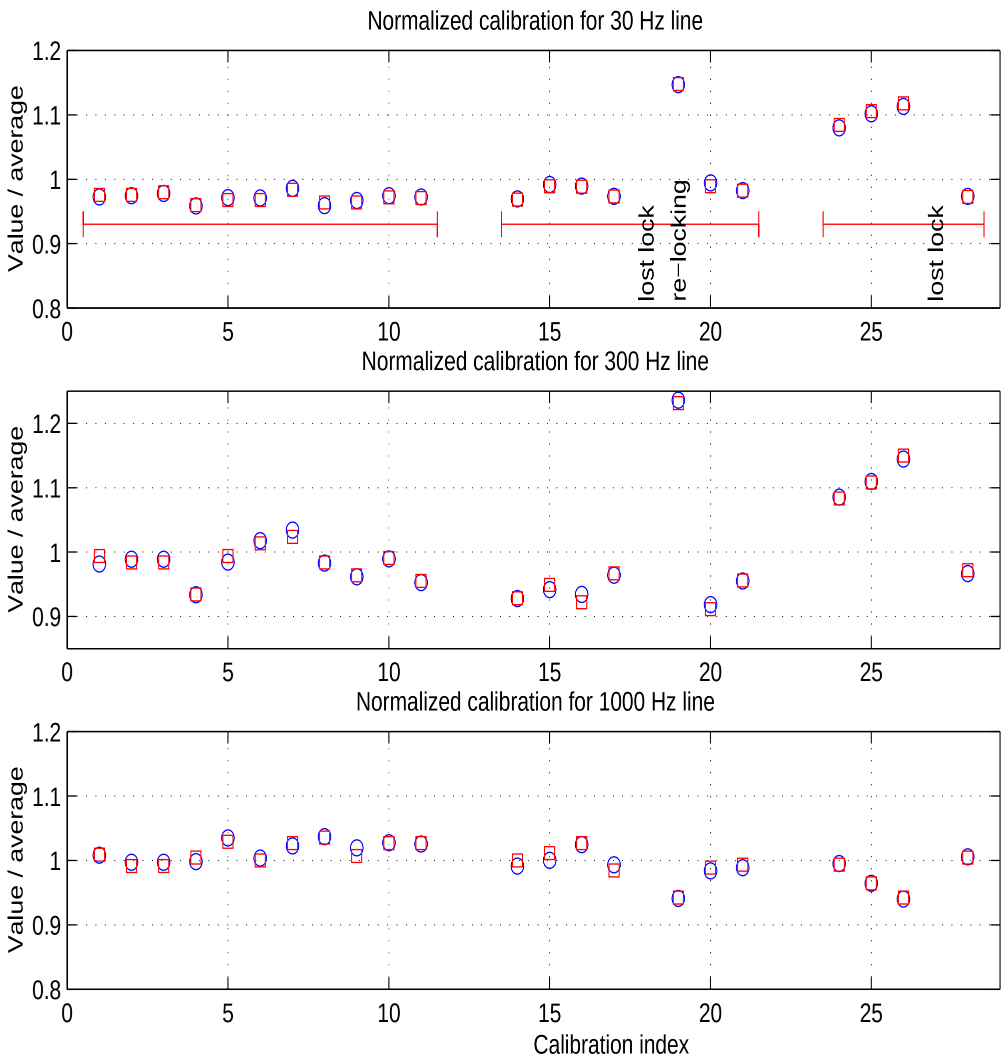

FIGURE 2: Same as Figure 1, but includes both the length-control error signal (red squares) and the control signal (blue circles).

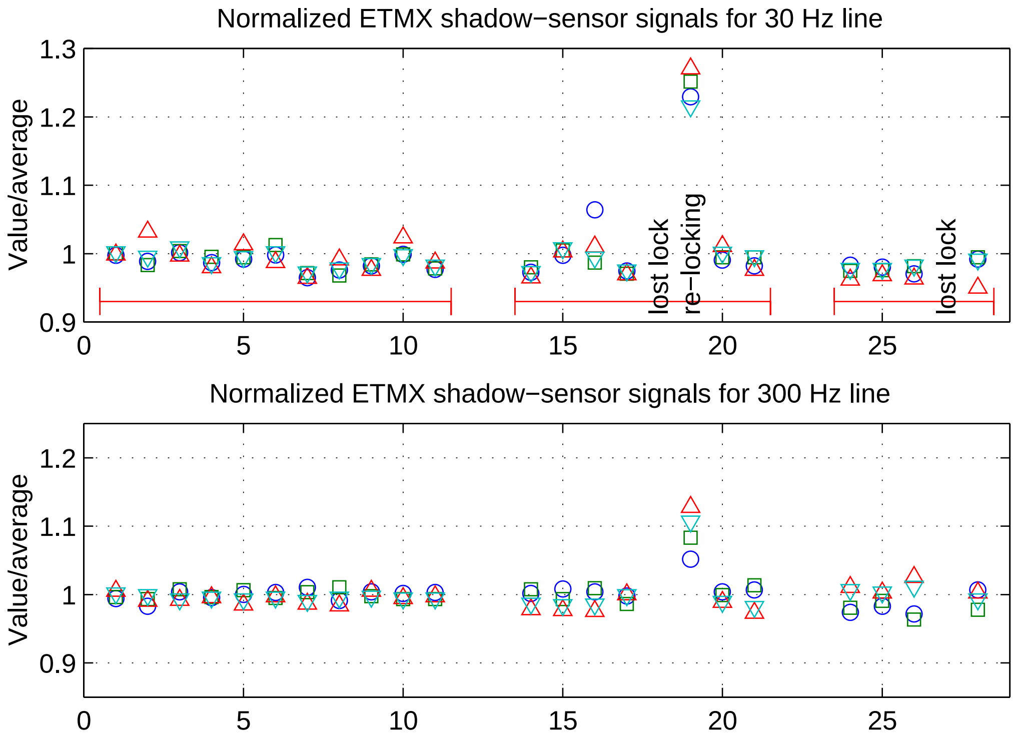

FIGURE 3: Normalized peak heights for 30 Hz and 300 Hz lines in shadow sensor signals from the end test mass. The four different symbols represent the four shadow sensors at the corners of the end test mass, each normalized to its own average. The time periods are the same as in Figure 1; the vertical scales are the same as well, though shifted in the case of the top plot.