First look at using Streckeisen STS-2 seismometer signals to predict LIGO arm length control signals.

J. Giaime, E. Daw, R. Adhikari, W. Johnson, M. Weitz

June 26, 2001 LIGO-T010070-00-D

Contents

Abstract

Signals from Streckeisen broadband seismometers, passed through suitable linear filters, can predict the LIGO arm length signals with sufficient fidelity to allow a feedforward reduction of the microseismic peak by nearly a factor of ten.

1 Introduction

LIGO's optical lengths are controlled primarily using signals derived from the interferometric outputs, feeding them back in servo loops to the test mass magnet/coil actuators. In their low-noise mode, the test mass suspension controllers have a dynamic range of approximately $10\ \mu\mathrm{m}$. Raab and Fine [1] have estimated that length changes of LIGO's baseline due to earth tides require a range of approximately 10 times this, so a slow (bandwidth of order hertz) fine actuation system (FAS) has been incorporated in LIGO. The FAS can displace the seismic isolation system that supports each end test mass by $\pm 90\ \mu\mathrm{m}$, providing a total range of $360\ \mu\mathrm{m}$. The signal source for the FAS tide correction can either be an accurate earth tide model or an extremely small bandwidth servo loop based on interferometer signals.

Raab and Coyne [2] calculated the control dynamic range needed to track the microseismic motion, which generally appears as a peak at about 0.15 Hz in ground displacement spectra. The required peak-to-peak range estimate for LIGO Livingston is $14\ \mu\mathrm{m}$ for a 40 hour lock stretch with stationary noise of $\sigma \approx 1\ \mu\mathrm{m}$, or a worst-case range of $56\ \mu\mathrm{m}$ if significant ocean storm activity occurs during the lock stretch. If these long locking periods are indeed required for LIGO science, it may be necessary to take up some of this range using the FAS (for differential length) and the laser reference cavity (for common-mode length). It might be possible to allocate the low-frequency portion of the length feedback signal to the FAS, though this has some operational disadvantage. For example, during times when when interferometer lock is gained and lost frequently, there would likely be repeated large filter transient signals applied to the FAS, significantly shaking the suspensions.

A better method may be to use floor-mounted seismometers to measure the microseismic band motion at each building and pass these signals through a multi-input filter that estimates the changes in displacement among the buildings. Feedforward correction has an advantage over feedback in that it cannot become unstable and that it can operate independently of the interferometer. The common-mode arm length component of the estimate would be added to the laser reference cavity offset (or to the FAS), and the differential-mode component would be sent to the FAS. The transfer function of one of the FAS installations was measured at LIGO Hanford[4], and it is flat and featureless over the microseismic band, which makes it very likely to be an effective feedforward actuator.

This technical note describes some measurements made before and during the fourth engineering test run (E4) at LLO, and analysis supporting the viability of this feedforward technique.

2 Experimental test, pre-E4

2.1 Data collection and reduction

Several days before the E4 test run, three Streckeisen STS-2 broadband seismometers (purchased for use in advanced LIGO SEI development) were deployed in the PEM seismometer areas, near the test mass vacuum tanks, and covered by simple insulated (but not hermetic) tubs. During the evening of May 9, 2001, and over that night, the LLO detector was operated in recombined Fabry-Perot Michelson mode and the length control signals and seismometer signals were recorded. The detector performed well, unattended and free of human perturbation most of the night.

For the low-frequency signals considered here, the feedback gain to the test mass suspensions is sufficiently large that any external longitudinal disturbances to the test masses are nulled by the feedback so that the DARM-CTRL and CARM-CTRL DAQ channels are proportional to the sum of the external differential arm length and common-mode arm length disturbance forces. These signals are calibrated to an equivalent FP arm length change of 1.0 nm/ADC count (see ilog entry [3]).

Since the outputs of the STS-2 is proportional to the instrument's velocity, it is necessary to integrate them in time, bandpassed by the microseismic frequency range, in order to obtain a displacement signal. These displacement signals can then be summed and differenced and compared with the interferometer control signals. Unfortunately, the fidelity of this naive calculation to DARM-CTRL and CARM-CTRL is poor, and only half of the observed microseismic disturbance amplitude, on average, is accounted for.

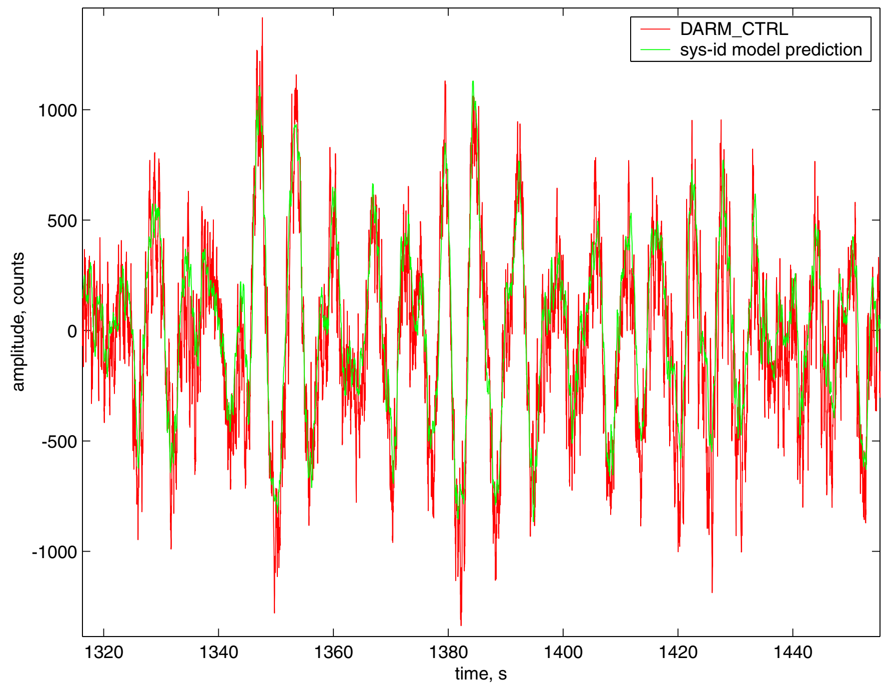

Figure 1: System identification fidelity: DARM-CTRL (red) and the filter's output (green) during several minutes of the "seg7" lock segment. 1 ADC count is equivalent to the control signal that corrects for 1.0 nm of differential arm length disturbance.

2.2 System identification

To make a filter that accounts for more subtle effects, the MATLAB system identification toolbox is used. Good results are obtained by fitting the data to a "black box" state space realization (using the `pem` function) of a multi-input filter that takes the seismometer signals as input and predicts the control signals as its outputs. The algorithms in this toolbox adjust a set of model parameters in order to minimize the variance in the difference between the data and the statespace outputs. To produce an effective model, it is helpful to hand-tune a bandpass filter around the microseismic band to prefilter the input data and in some cases post-process the outputs to reduce extraneous noise from transient low-frequency events. This procedure was carried out using about 10% of the data from lock segment "seg7," and a filter was produced. Figure 1 shows a comparison between the filter's output and the actual DARM-CTRL signal it is intended to predict.

Figure 2 contains a bode plot of the transfer functions from each input of the filter to its output. Its form is basically as outlined above, a bandpass bracketing the microseismic peak, with an integrator around the center of the peak. An integrating filter has a $90°$ phase advance; the "most effective" transfer functions in Fig. 2 seem to be advanced approximately another $40°$. There are some possible sources of phase shift between the ground displacement and the interferometer length signal, though none are completely satisfactory. The test mass suspension local damping loops were operating during the time this test was done, and they can introduce several degrees of phase change. The inputs used by the filter are the x and y channels from the corner station seismometer, the x signal from the LIGO X arm, and the y signal from the Y arm. While this

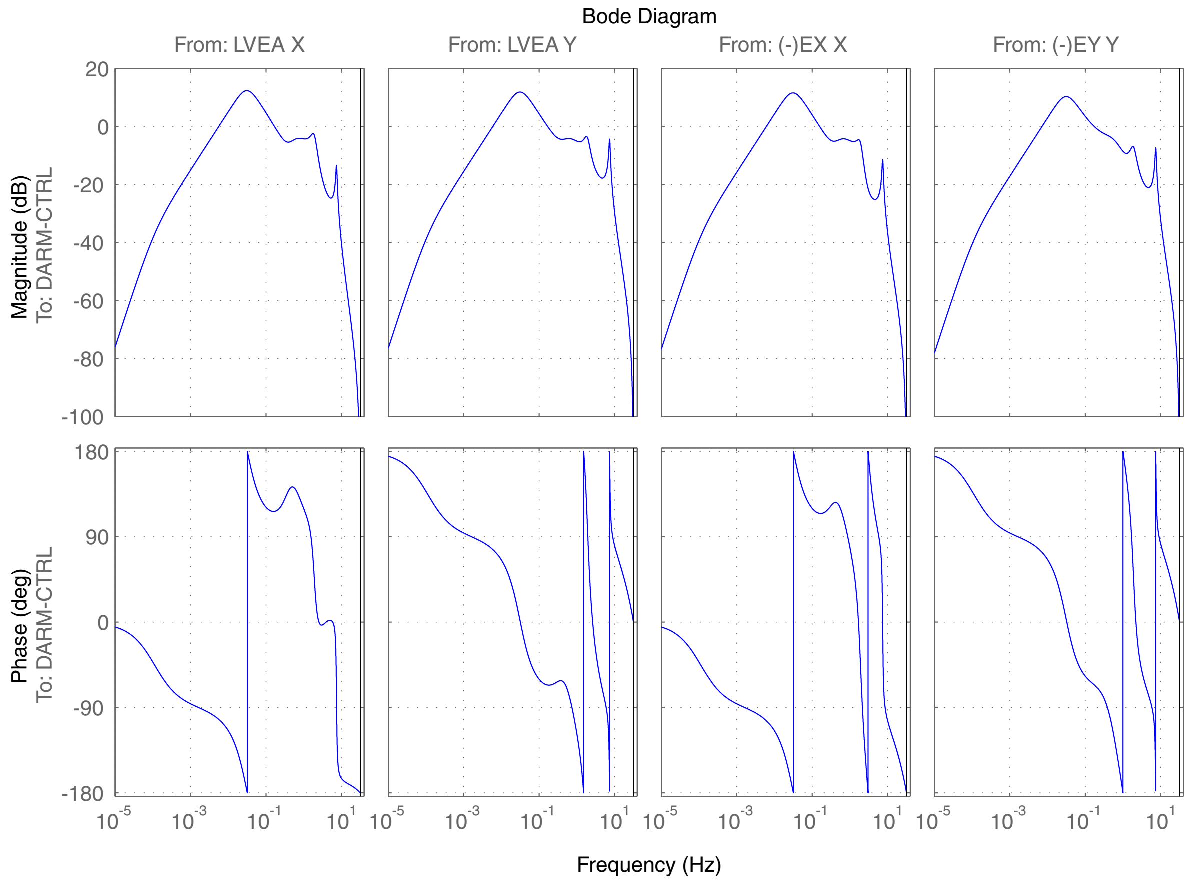

Figure 2: Bode plot (magnitude and phase transfer functions) of the four-input, one-output, statespace filter that was used to produce a candidate feedforward signal for the differential arm length in the 4 km recombined FP-arm Michelson interferometer at LLO. This was obtained by minimizing variance between the DARM-CTRL and the filter output using MATLAB on a short subsegment of "seg7," and then hand-tuned slightly.

ought to cover the most important ground motion, in principle the three technical slabs have 18 degrees of freedom; it is possible that the microseismic waves are coupling to test mass motion significantly through other DOFs, and as a result the phase relationship between the DOFs we measured and the length signals may vary. Similarly, Fig. 3 shows a filter separately constructed to correct the CARM-CTRL signal.

To estimate how well the filtered seismometer signals would do as a source of feedforward control, data from nine segments of interferometer lock, 20 - 40 minutes long each, are used as input to the filter. The filter's output is then subtracted from the control signal to simulate what would remain as disturbance for the interferometer length servos with feedforward active. Welch estimates of the displacement spectral densities of difference signal and the original control signal are plotted together in Figure 5 (for DARM-CTRL) and in Fig. 4 (for CARM-CTRL). It appears that the feedforward technique can reduce the microseismic disturbance by almost a factor of ten in displacement, without adding noise appreciably outside of its band.

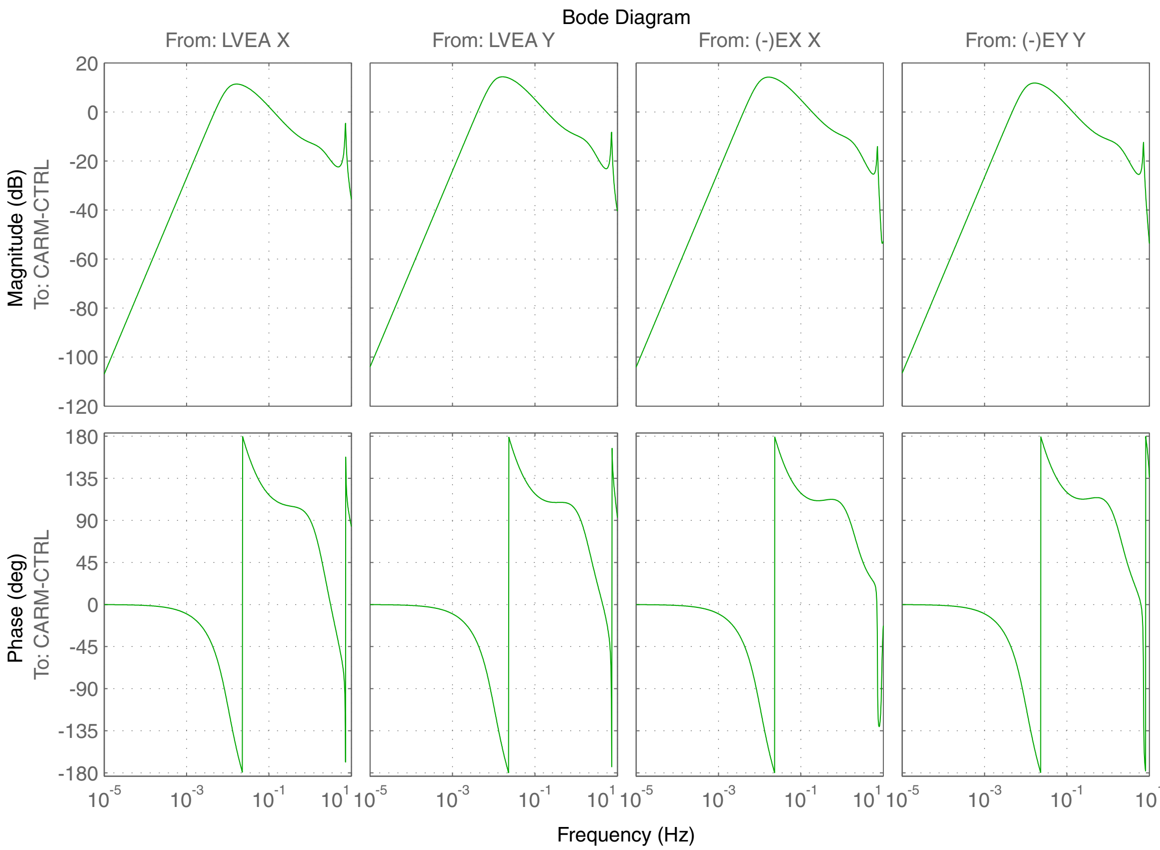

Figure 3: Bode plot of the four-input, one-output, statespace filter that was used to produce a candidate feedforward signal for the common-mode arm length.

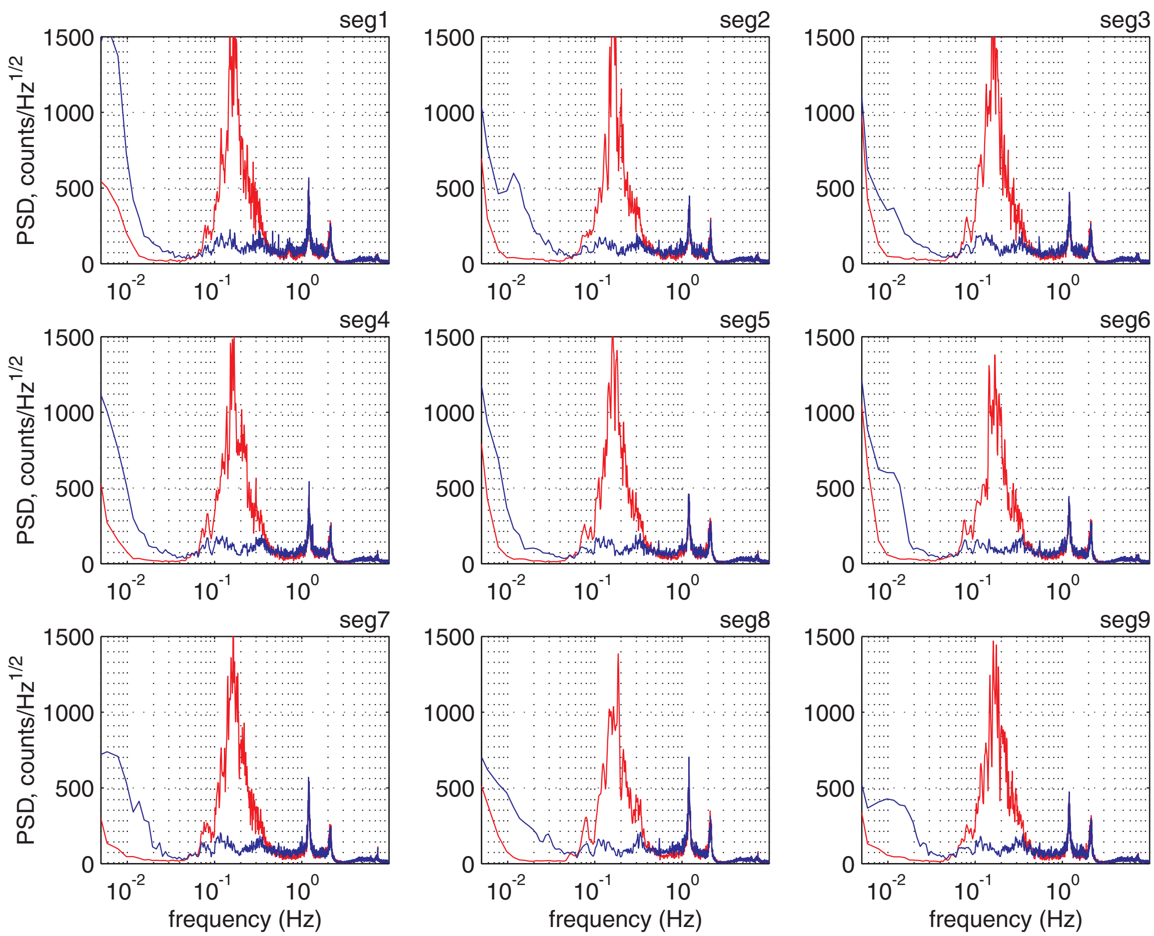

Figure 4: Equivalent displacement spectral density plots of nine segments of data from a pre-E4 run (5/10/2001) with the interferometer locked. The segments are each typically 20 - 40 minutes long. The red trace is the uncorrected DARM-CTRL channel, calibrated to be 1 nm/ADC count, and the blue trace is what remains after the prediction from the seismometer signals is subtracted from DARM-CTRL. Typically, the microseismic peak in the differential arm length is reduced in height to below $2 \times 10^{-7}\ \mathrm{m}/\sqrt{\mathrm{Hz}}$.

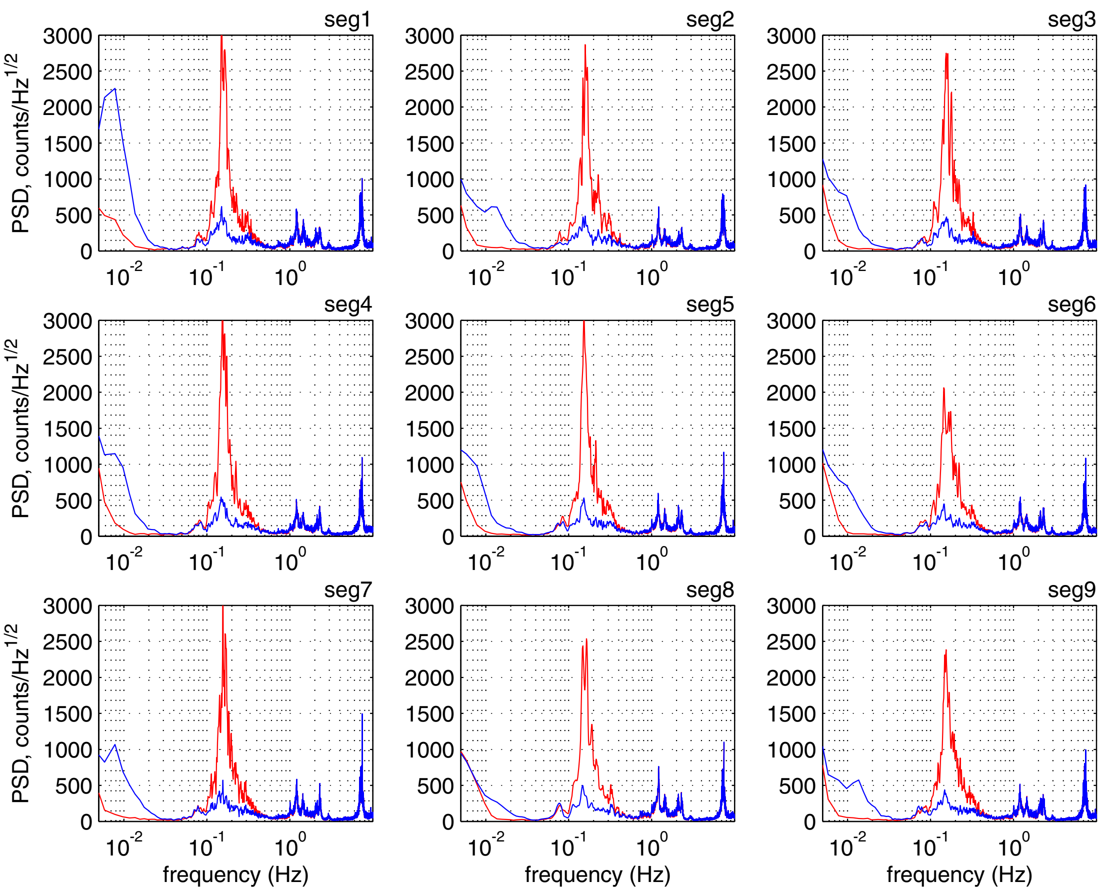

Figure 5: Equivalent displacement spectral density plots of nine segments of data from the pre-E4 run. The red trace is the uncorrected CARM-CTRL channel, calibrated to be 1 nm/ADC count, and the blue trace what remains after the prediction from the seismometer signals subtracted from CARM-CTRL.

3 Experimental test, E4

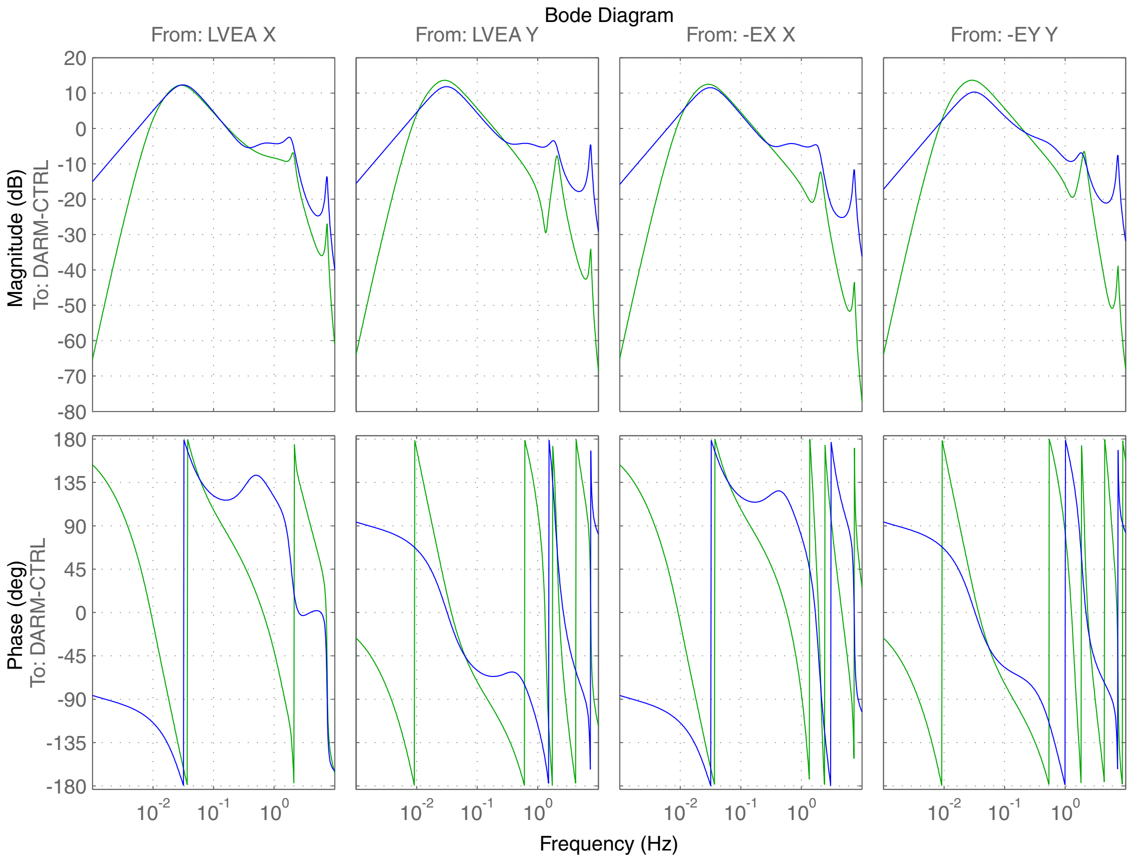

Figure 6: Bode plot comparing the DARM-CTRL prediction filter that is effective for the pre-E4 run (blue) and the one that is effective for the E4 data (green). Note that the magnitudes are nearly identical at the microseismic peak, but that the phases are considerably different.

A few days after the test run described above, another long set of data was collected during the LIGO E4 run at Livingston, with the detector in approximately the same configuration. Unfortunately, the filters designed from the pre-E4 data did not do a very good job predicting the E4 control signals; the disturbance could only be reduced by a factor or two or so, not the desired ten. So, the same filter-construction procedure was carried out with a short segment of E4 data, and after some hand-tuning an effective filter was found. The result is shown in Figure 6, plotted together with the pre-E4 filter for comparison. Spectral density plots showing the disturbance reduction that may be realized are shown in Fig. 7

The reason for the difference between the two effective filters is not fully known. There was somewhat more transient low-frequency noise during E4 (possibly due to the increased human and vehicular traffic), so more stringent bandpassing was added to the filter, causing a phase dispersion around the microseismic peak. Still, the “effective” transfer function phase value right at the peak center differs considerably from the pre-E4 filter, even though the magnitudes are almost identical

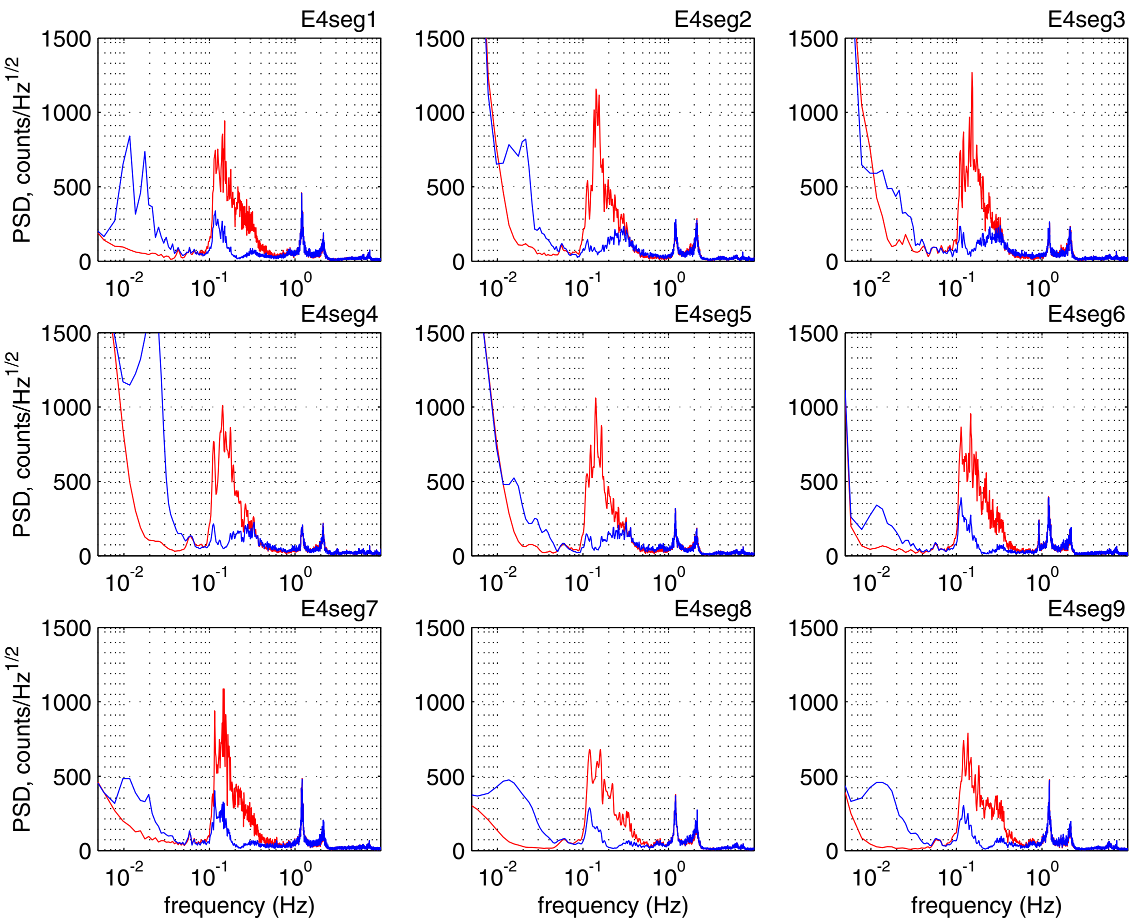

Figure 7: Equivalent displacement spectral density plots of nine segments of data from the E4 run. The red trace is the uncorrected DARM-CTRL channel, calibrated to be 1 nm/ADC count, and the blue trace what remains after the prediction from the seismometer signals subtracted from DARM-CTRL.

in the two filters. This phase discrepancy may be due to a changed microseismic disturbance vector; the level and shape of the peak are constantly different during E4, which was a period of very low microseism, and it is plausible that the waves are of a different character. It is also possible, though not likely, that the detector settings were changed in a way that has not been documented between the runs.

4 Discussion and recommendations

Since whatever is changing in the coupling from ground motion to length disturbance appears to do so very slowly, it is very likely that a fairly simple procedure can be devised to maintain an effective feedforward algorithm. Though signals appear on the seismometers from time to time, perhaps due to local human activity or wind gusts, that can add transient excess low-frequency noise, on balance it appears that the operational stability of LIGO can be improved using feedforward in this way, at

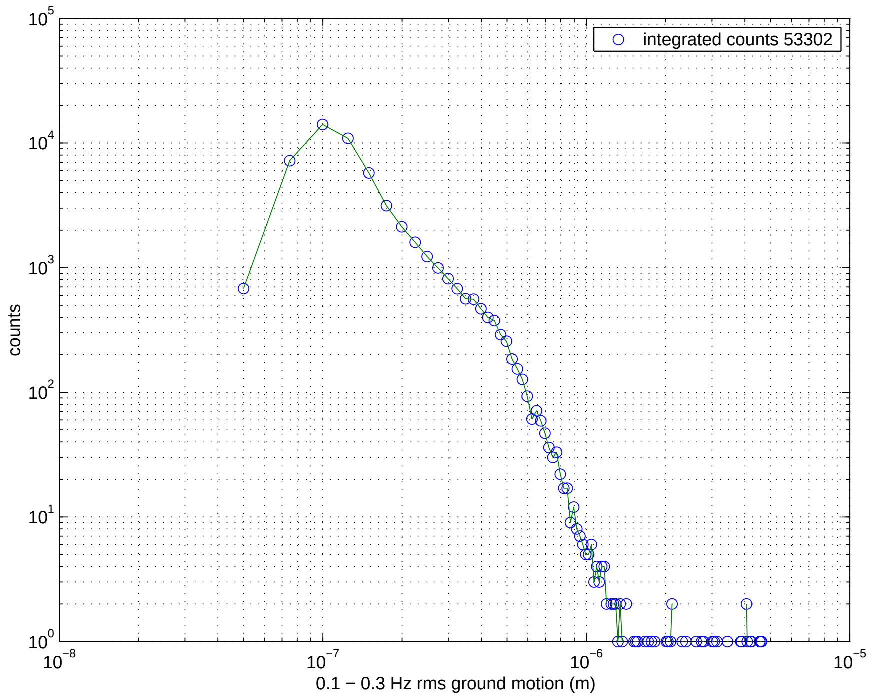

Figure 8: Histogram of the RMS of displacement in the 0.1 - 0.3 Hz band in the corner station at LLO, between May 15 and June 22, 2001. This period includes the intense weather activity associated with Tropical Storm Allison. Each point represents a 1 minute segment, and a handful of extreme outliers were excluded. (The 0.1 - 0.3 Hz band covers the center of the microseismic peak, though not all of it.)

least during conditions similar to those during the data runs described here.

Whether it is worth the additional complexity over simply feeding back to the test mass suspensions depends on the size the microseismic excitation and the desired interferometer locked segment time. This question was considered in [2], which includes (as mentioned in the Introduction above) an estimate of $14\ \mu\mathrm{m}$ as the needed control range for a 40 hour lock. Figure 8 shows a histogram of slab microseismic motion at LLO during six weeks that included both very quiet periods and the passage of a locally intense tropical storm. The peak in the distribution ($0.1\ \mu\mathrm{m}$) is only 10% of what was assumed in [2], though there are a significant number of minute-long segments with RMS motion above $1\ \mu\mathrm{m}$. During the segments of interferometer lock analyzed for this note, peak-to-peak DARM-CTRL values of about $3\ \mu\mathrm{m}$ are typical over a half-hour, if the earth tides are subtracted. Further study of the statistics of the needed dynamic range would be desirable if we would want to operate with only test mass suspension feedback, since it is possible ground noise levels that were seen in [2] may return.

Recommendations:

- LIGO should implement (as planned) the microseismic motion feedforward to the FAS.
- We should replace the Guralp seismometers (which are too noisy for this use) with Streckeisen STS-2 units. (these cost about $13,000 each.)

- A fully-functional feedforward system should be tested at LLO this Summer.
- The tests described here should be repeated at LIGO Hanford.

References

[1] Raab, F., and Fine, M., "The Effect of Earth Tides on LIGO Interferometers," LIGO-T970059-01-D, 2/20/97.

[2] Raab, F., and Coyne, D., "Effect of Microseismic Noise on a LIGO Interferometer," LIGO-T960187-01-D, 2/20/1997.

[3] http://www.ligo-la.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=05/07/2001&anchor_to_scroll_to=2001:05:08:01:11:21-rana

[4] http://blue.ligo-wa.caltech.edu/ilog/pub/ilog.cgi?group=detector&date_to_view=03/21/2000&anchor_to_scroll_to=2000:03:22:00:20:36-daniel (Measurement carried out by V. Sannibale, D. Sigg, and others.)