Servo Project Progress Report
Dates covered: 3/29/07 to 4/18/07
D. Clark

At last report, we had planned the following activities:

1. Continue model evolution to increase the low-frequency model fidelity to drive controller optimization.
2. Verify soft-start behavior on the elevation axis.
3. Continue deployment of the VxWorks-based controller, with the knowledge that overall disturbance rejection *may* be lower than the existing LM628 system.
4. Optimize the controller gains and topology to improve performance and wind rejection.
5. Collect *f/15* open-loop data to ensure a complete set of responses exist for controller design and modeling.

With the help of the CAAO crew and the mountain folks, open-loop data for the elevation axis was safely and successfully collected on 4/6/07. This open-loop data was band-limited noise input, with the usual encoder outputs collected as is normally done. Using the other standard open-loop test with a sinusoidal input chirp was not done due to the potential for damaging the DM with large excursions of the telescope. The results with this test are certainly consistent with that obtained with earlier tests with *f/9* and *f/5* taken before and after the *f/15* run, respectively:

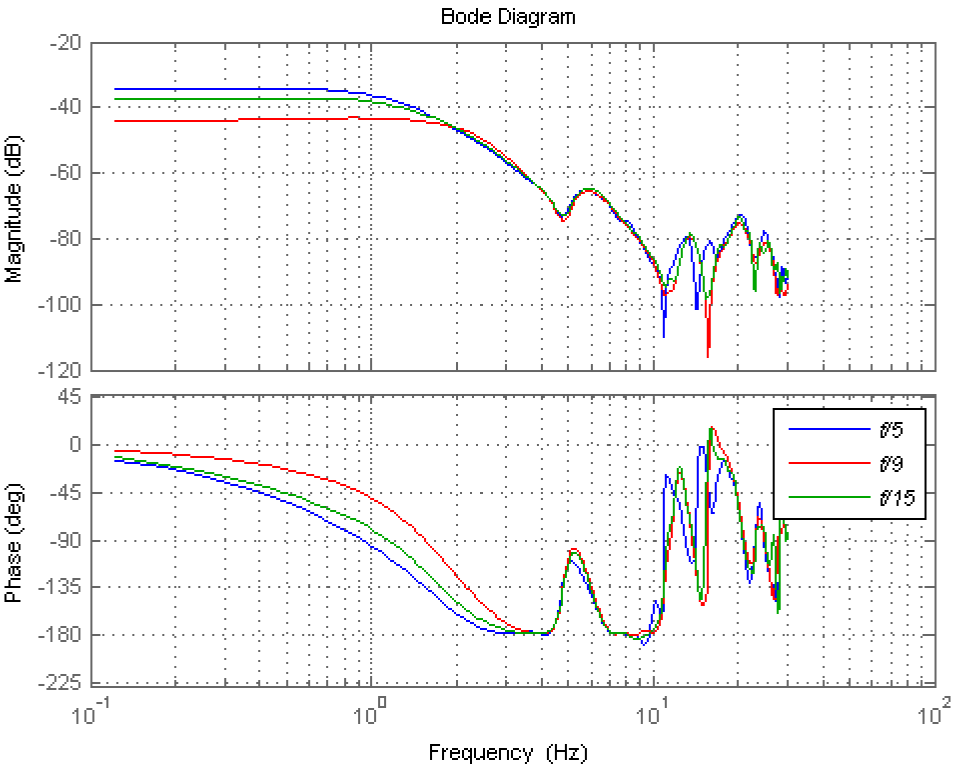

As might be expected, the changes in the different open-loop responses retain their overall modal shape, while the low-frequency response changes significantly due to the change in total inertia with configuration changes. This low-frequency difference is why robustness on the part of the new controller is of such concern; indeed, these changes may eventually require different controller gains depending on the telescope configuration for achieving the best possible performance.

One issue that is of interest in continuing the development of more accurate models of the telescope axis is the difference in the phase reported by the *tfestimate* routine depending on whether the i/o signals are noise- or chirp-input/output. The reader should refer to earlier reports for graphs of the standard chirp-response outputs. The chirp response phase at the low frequencies is estimated to be ~180°, while the noise response puts it at close to 0°. Reconciling these differences was a primary motivation for collection of step-response data during the M&E night of 4/11/07; the DC response can be calculated from the open-loop step response and used in driving the asymptotes in the open-loop model towards their proper gain and phase values. It's instructive to look at the time response of the chirp data, reproduced here, with the input voltage (blue), and the encoder output in counts (green):

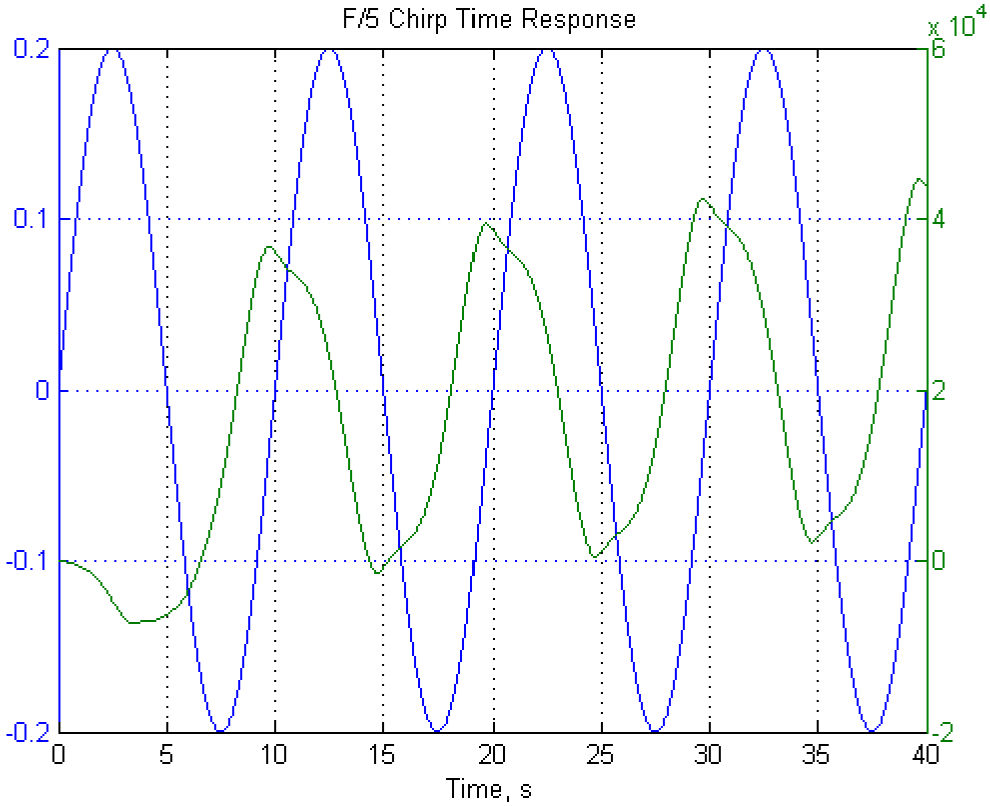

There is obvious non-linearity due to friction near the zero-crosses of the input signal, which covers the region of about ±60mV. From this data, it's clear the phase shift is really ~90°. For the best accuracy in estimating the DC response, we can use the time response from a step input to calculate the DC gain and phase and apply this knowledge to improvement of the open-loop telescope model.

To this end, step responses in both polarities of two different types were collected. One type used a 0.1V step as the input signal, switched off after 10s, the other used the same period with a 0.25V step, shaped by a 20Hz low-pass filter to limit the acceleration. This 0.25V step was switched to 0V without filtering at the end of the step input period. Unbalance torque makes the two polarities behave differently:

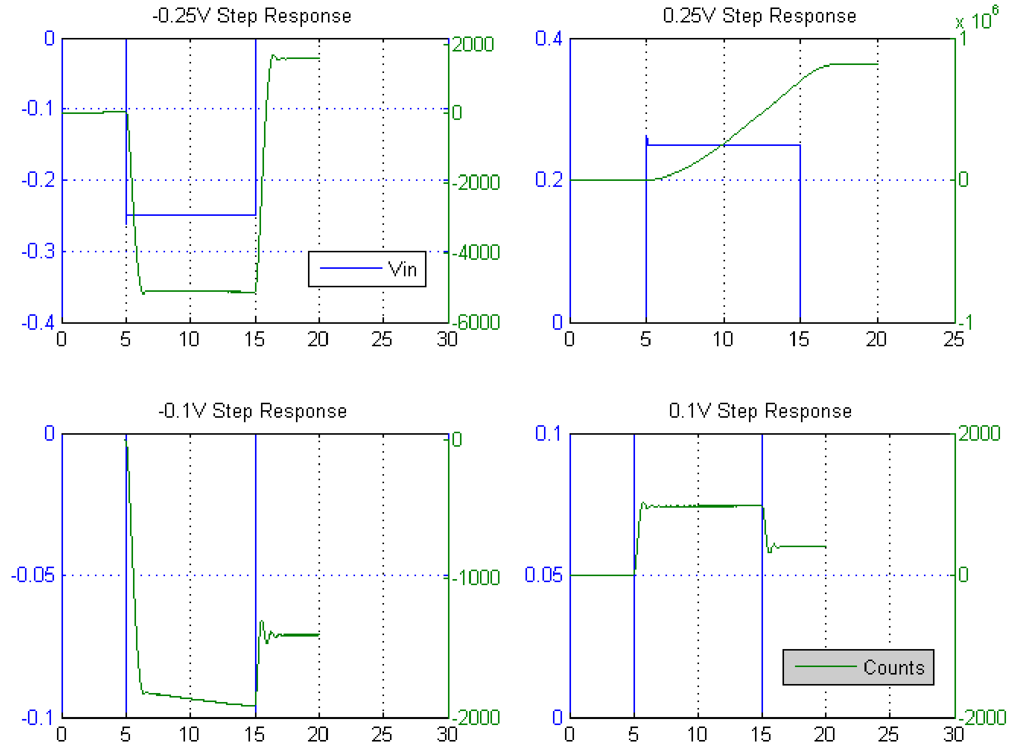

The encoder output, in counts, is the same for the other tape head and is simply inverted. We can safely conclude from this that the DC response is either 0° or 180°, depending on which encoder feedback is selected. The output of the absolute encoder tracks the above plots:

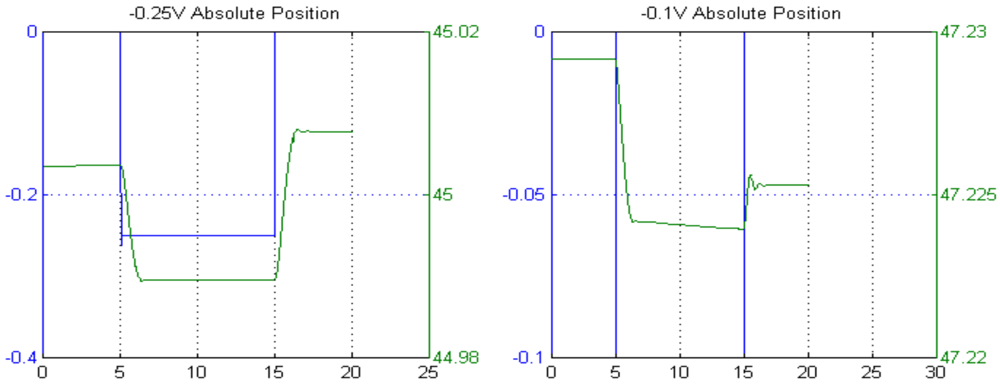

In the majority of the tests, one unexpected result is that the position “springs” back when the input torque is removed. We have old data from 2004 that shows lost motion of the elevation motor shafts when an input torque is applied that may help to explain this: the lack of anti-rotation constraints about the motor housing means that the motor body reacts to the torque by trying to “walk” along the drive arc curvature. When the motor input is relaxed, the combination of the drag link flexure and preload spring force the motor housing back to its original position, which in turn puts a counteracting force on the elevation axis, and creating a counter-force that tends to return the axis to its original position – friction makes it impossible to fully return, as can be seen.

Another important point about the step responses – when the input is removed, the velocity of the telescope quickly goes to zero. This implies some sort of frictional loading that slows the telescope down. It’s possible that a brake is dragging or a bearing problem exists in either the motor bearings or cell bearings. The time responses of the sine-input signals certainly show a frictional component, while another possibility that comes to mind is electronic braking from the motor amplifiers; since we set the amplifier output to 0A, the back-EMF of the motor as the telescope motion continues generates a counter current that must be opposed by the amplifier to maintain the zero-current condition, resulting in an increase in damping.

If the amplifier braking is the source of this damping, it has little impact on the control system design or performance; bearing problems or brake dragging are obviously of much greater concern should something be discovered to be wrong with them. Inspection of the motor brakes and drive arc pads therefore needs to be done.

Moving on to verification of soft-starting the controller, at last report a Simulink session produced a soft-start diagram that a) limited the velocity estimator’s output during startup, and b) soft-started the input signal to the flexible-mode filters. One activity during the 4/11 M&E was to verify that the new soft-start works properly. Using the xPC Target test machine, this verification happily matched up well with the linear simulation (which does not include non-linear friction):

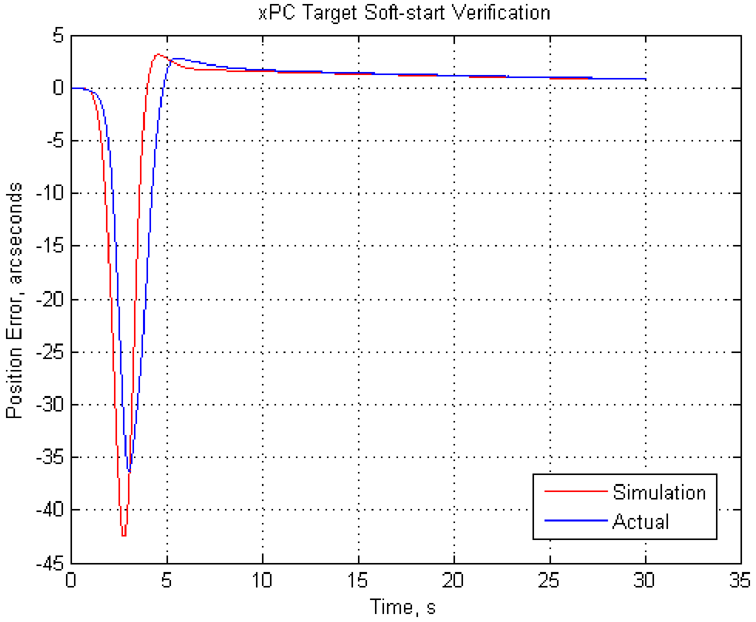

This is a vast improvement on the rail-to-rail DAC output that occurs without the soft-start installed.

Testing of the servo controller in its VxWorks incarnation was done after this initial work was completed; Trebisky can comment on his work in a separate report.

However, additional testing was done after other engineering activities early on the morning of 4/12 to collect additional data on the wind shake response of the telescope and controller using the xPC Target test machine, and it’s worth noting the difference between the earlier VxWorks-based controller wind shake and the xPC Target results.

Using the standard tracking logger data, we have a representative tracking-error plot from (1 of 3) objects that were acquired with the VxWorks controller. The wind at the time was in the high 20s, with gusts to over 30mph from the SW. The telescope azimuth is to the NE.

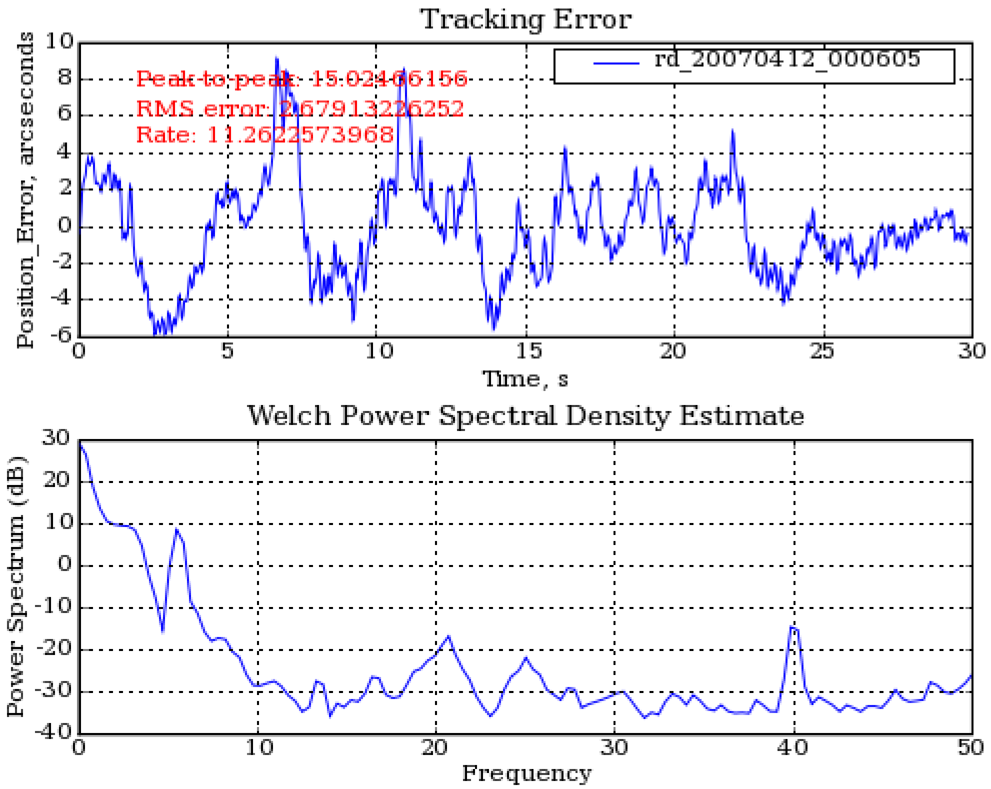

The PSD plot clearly shows that the modal filters are not working; in addition, the overall tracking error is much larger than expected from simulation, probably because the controller bandwidth is much lower than it should be. In addition, it does not come close to the performance of the LM628 servo from the same night:

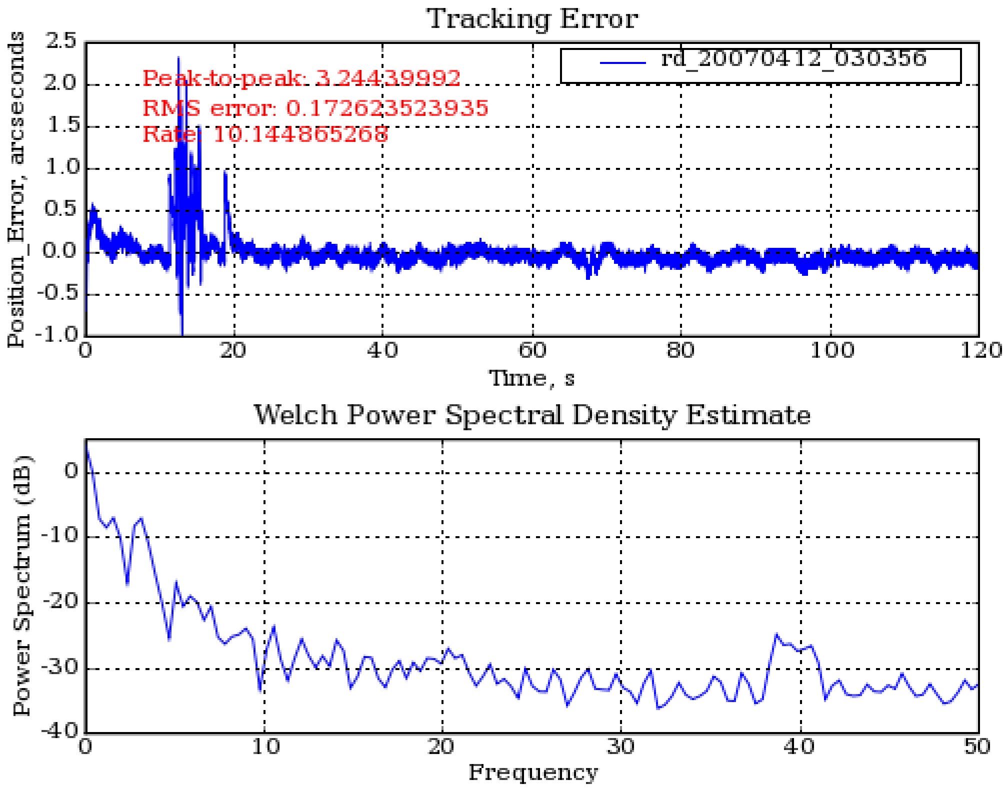

The tracking logger script apparently went away sometime late on the night of 4/11, as it has not updated the logging data since. However, other data collected since the tracking error webpage went live is available.

A table summarizing the tracking data is available in the last report, and will not be reproduced here. It's interesting to average the RMS error values for the LM628 servo from that table, however, and compare it to the average amount of oscillatory power from some additional tracking logs to get an estimate of how much of the total tracking error is due to oscillatory modes in the servo (all values in milli-arcseconds):

Average RMS Error	Average Peak to Peak Error	6Hz Mode	20Hz Mode	35Hz Mode
81.17	1008	52.4	2.82	1.17

Admittedly, the sample size for the tracking error data is small, and perhaps we could set up some standard MySQL scripts to reduce the tracking error logs to build a better and more representative database to refine these numbers. It's clear from this early data, though, that the 6Hz mode contributes the majority of the overall tracking error, and *cannot* be reduced with the LM628 servo system. Indeed, suppression of this and the 20Hz mode are two of the motivating factors for even working on improvement of the servos in the first place.

As was reported earlier, a wind shake model was constructed for deriving controller optimization in terms of disturbance rejection. Using data collected with the xPC Target test machine during this last M&E run was an excellent opportunity to collect data that compares directly to the results with the LM628 and “new” VxWorks controller under virtually identical conditions.

For collection of the wind shake data with the xPC system, 3 tests were conducted. First, the controller was started and left running while the amplifiers were left off and all the telescope brakes were set. This gives a direct measure of the mechanical admittance of the structure to a fairly stiff connection to the yoke. Second, the controller was started and left idling with all brakes released, simply holding its position in the wind. Last, the controller was started and tracking a ±10 arcsecond/second was commanded via the Java GUI written for communication with this version of the controller. The wind during this testing was mid-20s, gusting to over 30mph, with the telescope oriented away from the wind to the NE, and the elevation set to 45°.

Here are the time data plots of the position error signal for these tests, less the startup and halt portions:

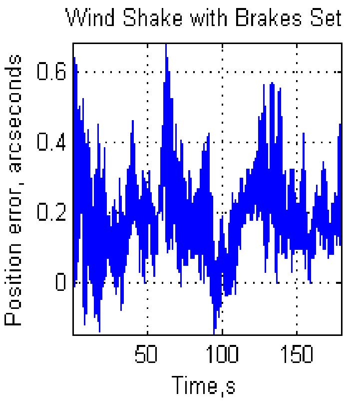
Wind Shake with Brakes Set

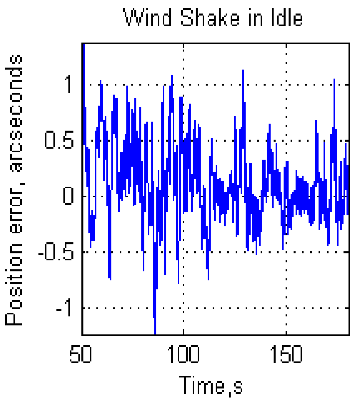
Wind Shake in Idle

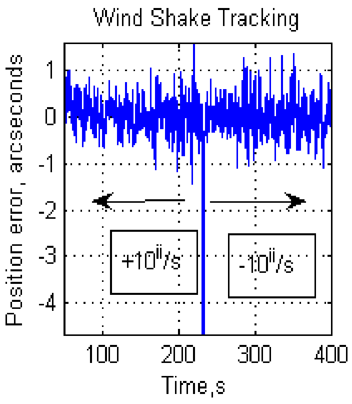
Wind Shake Tracking

The relevant statistics for these plots are, in arcseconds:

Test	RMS Error	Peak to Peak Error	Mean Error
Brakes set	0.11	0.83	0.19
Holding position	0.334	2.354	0.087
Tracking +10 arcseconds/second	0.303	2.168	0.053
Tracking -10 arcseconds/second	0.361	2.429	0.011

Note that with the brakes set, the position error is nearly the same as the ultimate tracking goal. Note also that it contains a DC and variable component, providing evidence that the current wind-disturbance model is correct. Small imbalance torques also influence the results somewhat.

The time data above were also reduced to PSDs, for comparison to the data from the LM628 and the VxWorks controller:

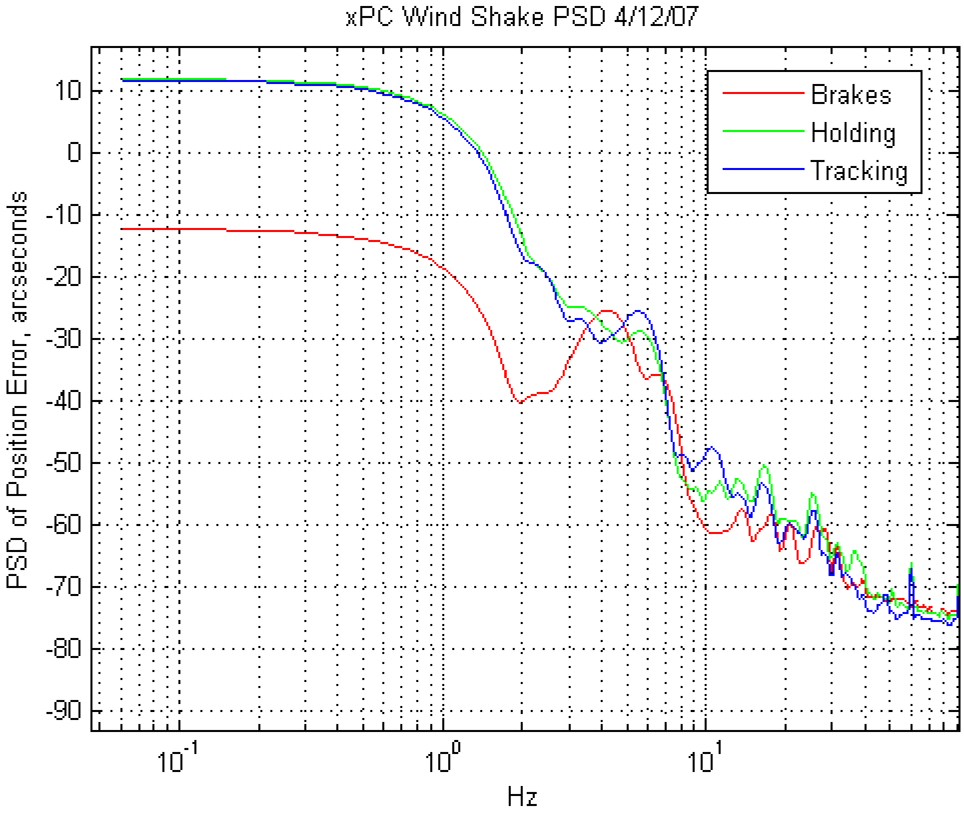

In linear terms, the 4.2Hz peak, which is equivalent to the 5.5Hz peak when the servo is running, is down at 2.8milliarcseconds. Indeed, the notch filter suppression is practically as good as having the telescope brakes set! The higher-frequency modes are all below the primary modes by a factor of 25db, or 300 in linear terms. This is *significantly* better than the LM628 servo.

Taking this data and comparing it to the wind shake model output, it turned out the wind shake input torque had to be *reduced* by nearly 50% to rescale them to equivalence, but

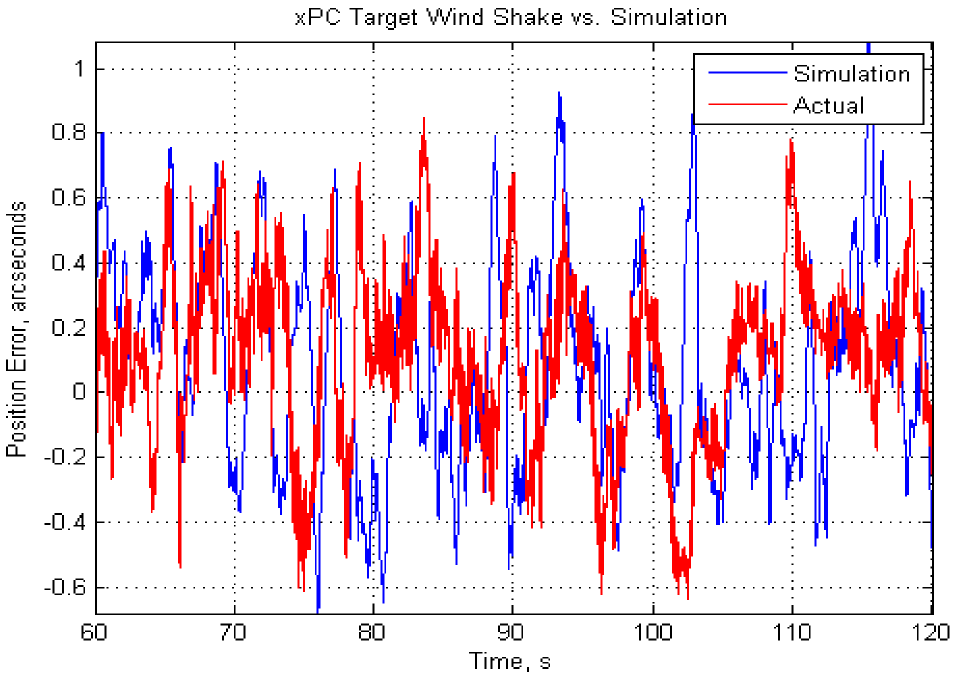

once this is done the two are in decent agreement.

This is an encouraging result. The torque input scaling most likely will need to be adjusted upward by a factor of 2 or 3 to correctly simulate the response when pointing into the wind, but the simulation RMS as-is 0.33 arcseconds with 1.766 arcseconds peak-to-peak output, which is very similar to the field data collected.

Gain optimization and other tweaks to the controller will only improve this; adding the disturbance-decoupled Luenberger observer/compensator without changing the gains gives the following simulation output, for example:

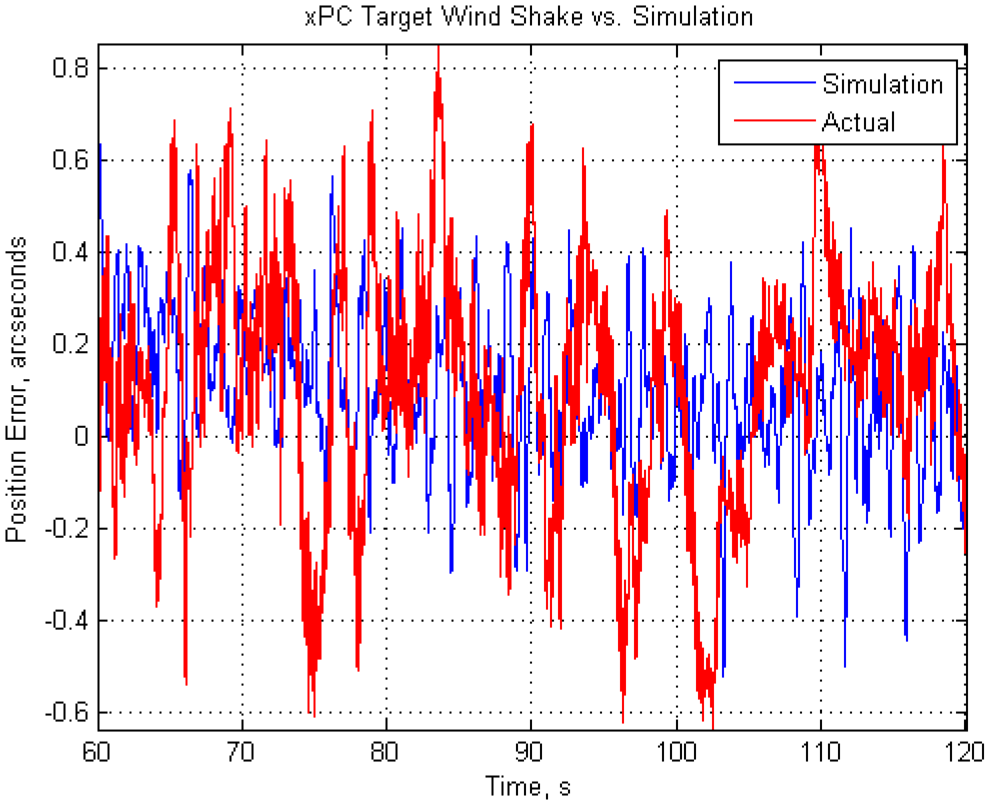

The position RMS is then 0.167 arcseconds, and 1.157 arcseconds peak-to-peak, for an almost 2:1 improvement. A notable aspect of the disturbance-decoupled simulation model is that the closed-loop telescope model and observer model are not the same, and both are indeed not of the highest fidelity to the measured open-loop data; both will degrade the controller performance in the simulation. Having better models for improvement in the controller design remains a primary design goal.

The good agreement between the collected wind data and the model gives confidence in moving forward with continued wind-rejection modeling. It’s important to continue to monitor the LM628 tracking performance, as well, to determine what improvements, if any, are being achieved.

The VxWorks implementation and the xPC Target controller apparently not producing the same performance is the most problematic issue we face. Keith Powell has been somewhat freed up from his CAAO responsibilities, and is now assisting with improvements to the servo modeling and general optimization efforts. Without an implemented controller that is 100% equivalent to the simulation or the xPC Target test system, succeeding in getting the control system to work is very much in doubt.

The upcoming activities are:

1. Continue model evolution to drive improvements in controller design.
2. Inspect the motor and drive arc brakes to ensure there is not a problem with them dragging.
3. Measure the motor currents during the step response switch-off period to determine if electronic damping is present and eliminate bearings as a suspect in the unexpected damping.
4. Collect wind shake data as the opportunity becomes available to increase the statistical validity of the wind shake modeling, from both the nightly logging sources and independent tests with the xPC test system.
5. Debug and verify proper VxWorks controller implementation.