Servo Project Progress Report
Dates covered: 3/2/07 to 3/28/07
D. Clark

At last report, we had planned the following activities:

1. Find and fix the startup oscillation problem.
2. Work on gain optimization to increase the controller bandwidth.
3. Work on improving model fidelity to the actual system to more properly predict gain margins and tracking performance.
4. Simulate and verify gains and wind rejection tests.
5. Simulate and verify slewing results from the past two test runs.
6. Release off-line data reduction scripts for tracking logs to allow staff access to tracking data via the MMT web-browser interface software. Ensure wind loading data are available as part of the reduction annotations.
7. Test the latest controller updates on the telescope with the VxWorks implementation.
8. Test the latest controller iteration with the *f/9* configuration.
9. Simulate and test on hardware changing the motor amplifiers over to the Copley PWM units.
10. Release for nightly use, changing over hardware as needed to bring up the new system.

It was left to Trebisky to work on (1) and (7), while work over this period focused on gain optimizations and modeling to improve the fit of the original design models (which have always been simple approximations) over the controller bandwidth of DC to ~3Hz. Item 6 has been completed with help from Dallan Porter, and we now have nightly tracking data with wind information available at http://hacksaw.mmto.arizona.edu/plots/tracking/ and interested readers are encouraged to consider the information available there.

Simulation work at first was done in the direction of attempts to predict what the actual servo performance should be in the presence of wind disturbance, and using our new web-based reduction interface, we have the data below for comparison to simulations of wind buffeting on the new elevation controller. We still do not know for sure what the scaling of the wind gust loading to motor torque should be, nor do we have data that tells us how the wind loading varies with the relative wind to the azimuth and elevation during the tracking period, but this at least lets us get a handle on the problem.

The simulation assumes 100% loading, and is calculated over 60s, while the tracking data are not necessarily directly into the wind, over 120s. There also appears to be no difference due to changes in tracking velocity in the simulation results.

The simulation was constructed using a Davenport-spectrum filtered random disturbance scaled to a “best-guess” torque load on the telescope[1] , with the controller loop closed around our standard state-space model from the System Identification Toolbox, which is a somewhat improved version of Powell’s original model, though both are still not accurate at low frequencies. In general, the simulation shows improvement in the RMS tracking, while the peak to peak variation tends to be comparable, or even a bit larger, mainly due to lower integral gains set in the position loop.

Tracking Data File	**Wind Az**	**Wind, m/s**	**Rate**	**Az**	**El**	**RMS error**	**Pk2pk error**	**Sim RMS**
rd_20070307_044320	241	2.3	-12.7	266	68	0.088	0.926	0.061
_043536	238	1.95	-12.7	265	70	0.087	0.772	0.044
_031403	214	1.3	-5.9	207	85	0.101	1.16	0.019
_031034	209	1.65	-4.3	192	86	0.069	1.0	0.031
_030937*	210	1.6	-2.4	244	89	0.093	1.27	
_014038	102	1.75	12.6	100	78	0.084	0.772	0.035
_13711*	118	1.65	12.6	99	77	0.068	1.16	
_13524*?	115	1.75	12.6	256	89	0.112	1.19	
_225731	338	0.9	-6	207	68	0.082	1.16	0.009
_225430	195	0.75	-5.5	205	68	0.086	1.23	0.006
_201718	327	5.95	-2.3	189	61	0.074	0.88	0.408
_201422	167	3.9	-1.9	189	62	0.04	0.579	0.175

* -- Skipped due to the similarity of earlier tracking velocities
? – Funny looking data that appears to be cross coupling from the high el angle

While the new controller allows a larger peak to peak variation, the cyclical power seen in the LM628 controller is all but eliminated:

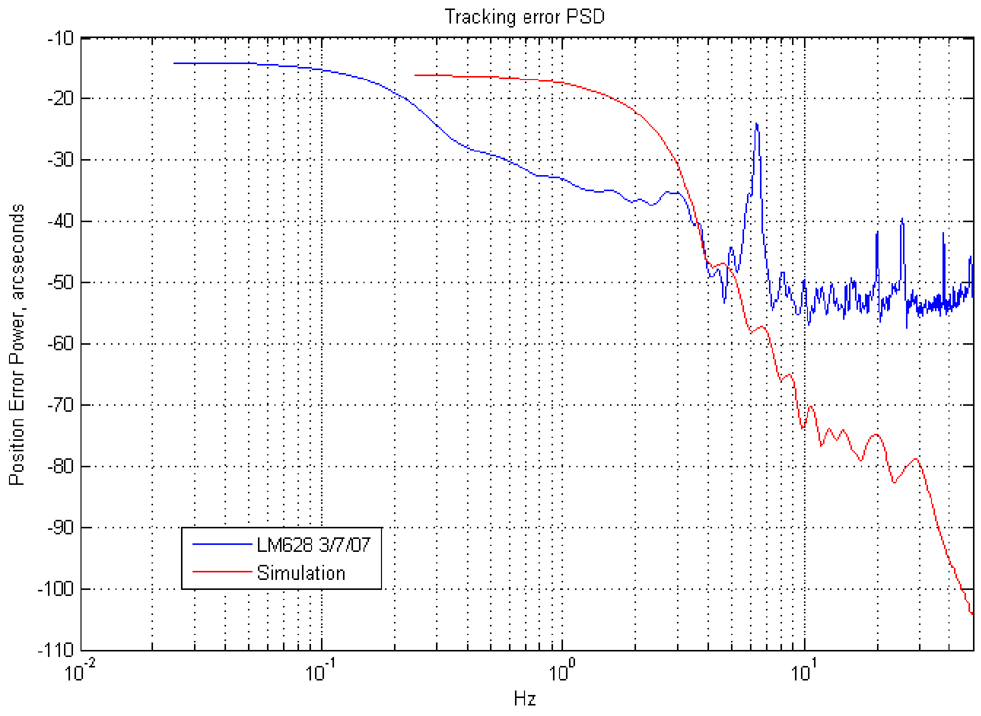

The PSD above shows that the controller (at least in simulation) allows more variation than the LM628 controller, but does not have the heretofore troublesome oscillatory motions that are particularly of concern for AO operations. The question then comes: how to improve this?

One straightforward answer is to simply increase the servo gains, particularly the position loop integral gain, as wind disturbance power is generally within the servo loop bandwidth's integral portion. Obviously, the overall closed-loop stability limits impose a constraint on this, but adjustment of the gains leads to the following predicted RMS tracking error, in arcseconds, with a given wind velocity in m/s:

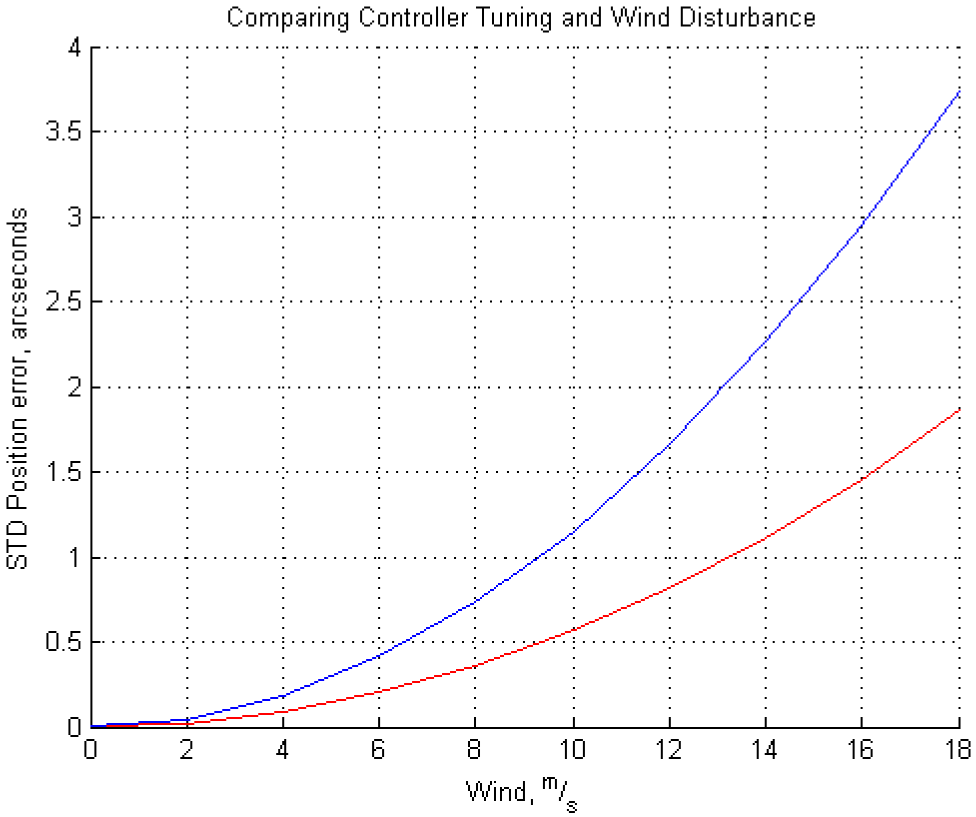

As was foreseen in the ancient MathCad wind disturbance model, wind buffeting varies with the square of the wind velocity, and leads to a primary difficulty in controller design – how to reduce the wind disturbance while maintaining closed-loop stability and its necessary constraints on servo gains.

One computationally and conceptually simple way to reduce the tension between these two incompatible controller constraints is to introduce disturbance-decoupling[2,3]. This idea leverages *a priori* knowledge of the plant to be controlled. Measuring the controller outputs (i.e. plant inputs), and the actual plant outputs *while the controller is running* and comparing the plant inputs and outputs to a model of the plant with the same inputs results in a measurement of the deviation of the actual plant from the model. This deviation contains, in varying degrees, the deviation of the model of the plant, noise and sensor uncertainty (generally left unmodeled), and plant outputs from unobserved inputs, which are generally lumped together as disturbances. By compensating the plant model with a standard PID loop, the closed-loop model of the plant and the actual plant can be driven to have near-zero differences. The output of the model (“observer”) and the PID (“observer compensator”) together act together to compute an estimate of the actual disturbance signal present in the plant sensor outputs. With disturbance-decoupling, the observer compensator’s output is scaled and filtered and added to the controller’s output to the plant. Due to causality, the observer controller is incapable of completely eliminating the disturbance, but it can greatly reduce its effect, thereby achieving higher disturbance rejection without the alternative of increasing gains and the concomitant reduction in controller robustness.

Extending the previous Simulink model with an observer-compensator is a simple matter: the complete model is shown here:

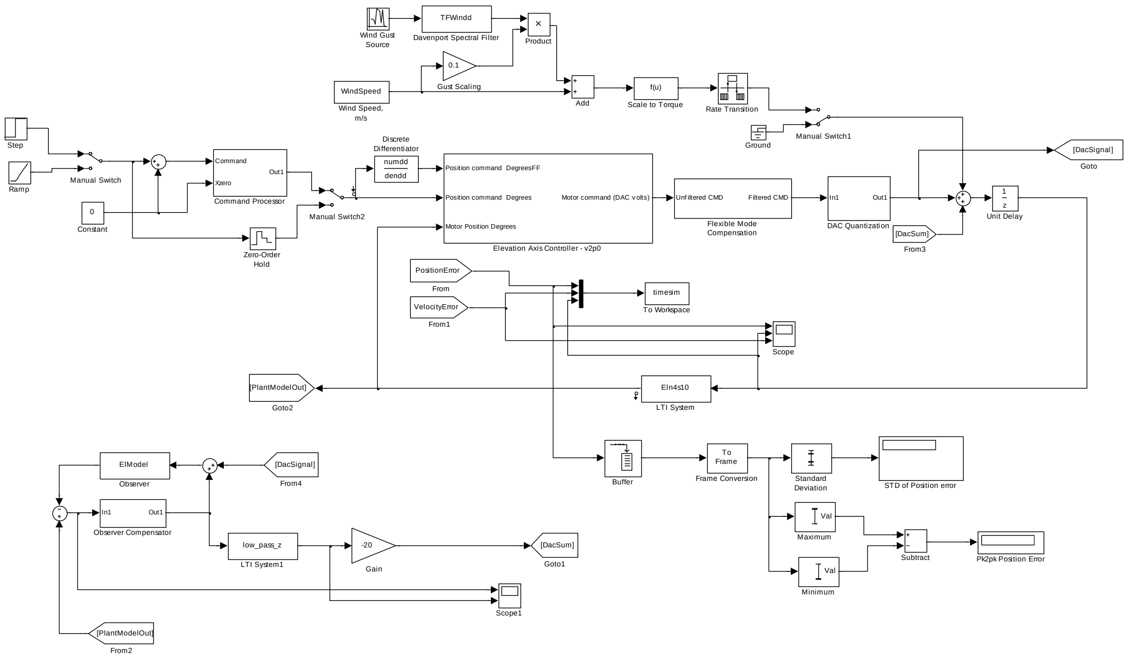

The model allows either tracking or step response to be investigated, along with the response to wind disturbance. This is essentially the same controller used on the telescope. Suppose that a wind with 6m/s velocity and 10% gust power is applied with the appropriate torque scaling to the elevation axis. The data statistics for the standard deviation and peak-to-peak variation become:

Controller Configuration	**RMS Position Error, arcseconds**	**Peak to peak Error, arcseconds**
Original gains	0.4149	2.322
Increase Ki	0.3615	2.093
Add disturbance decoupling	0.2283	1.437
Set new gains from time-response	0.2081	1.323

optimization, keep disturbance decoupling		

How does this compare with the LM628 servo? Thanks to the nightly logging, we have some comparative data with winds more or less directly into the telescope azimuth, and near 6m/s:

Logging File	**Elevation**	**Azimuth**	**Wind Az**	**Wind Speed**	**RMS**	**Peak to peak**
rd_20070315_033354	87.8	-20.18	314	6.4	0.113	0.965
rd_20070314_053123	57.66	320.3	298.5	5.05	0.115	1.23
rd_20070316_052837	79.64	116.7	106.5	6.85	0.112	1.544
rd_20070319_020228	61.24	-114.15	227.5	5.45	0.226	1.738

The difference between the simulated wind disturbance performance and the measured data from the LM628 servo implies a few things: a) the mechanical admittance of the wind velocity to torque disturbance is probably off, b) there is more work to do to improve the servo gains/disturbance rejection.

Simply cutting the input wind torque in half drives the disturbance-decoupled controller simulation output to 0.114 arcseconds RMS error, and 0.718 arcseconds peak-to-peak, as you might expect. While this is certainly in line with what is seen with the LM628 controller, the validity of the wind torque scaling is very much in question. Further work in optimizing gains and wind disturbance rejection analysis was stopped in favor of pursuing a more faithful model of the elevation axis over the frequencies of interest (e.g. DC to 3Hz) to more properly drive iterations of the controller gains. Until optimized gains with the highest possible (within the constraint of stability) integral gains are applied, it may be we shall have to accept more peak to peak tracking error in favor of eliminating the oscillatory motion if the controller is deployed with the “standard” gains.

Consider this Bode plot of the measured closed-loop bandwidth versus that predicted in simulation, where the blue and black lines are the measured data, and the green and red are the modeled results using the Powell model (red), and the System Identification Toolbox model (green):

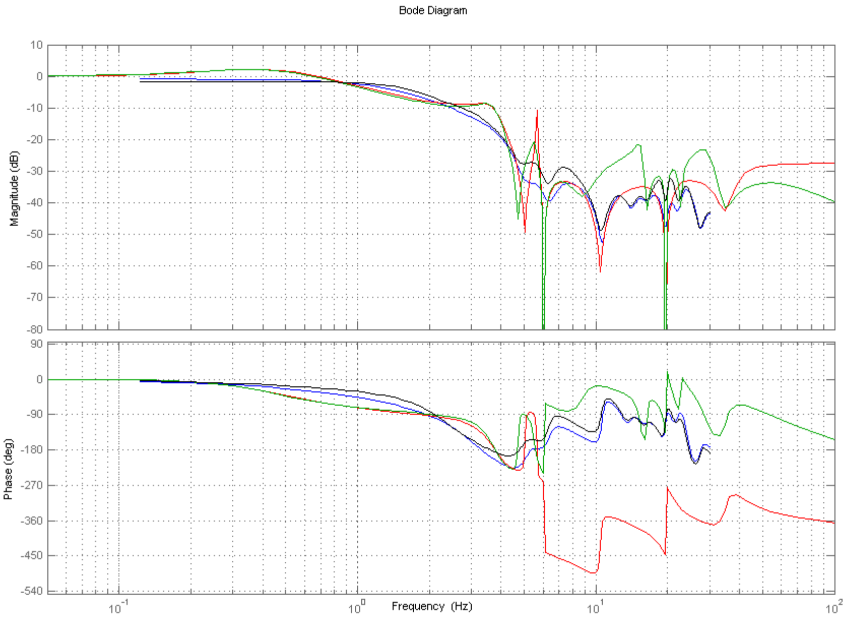

The closed-loop simulation does not equal the measured response over the range 0.1Hz to 3Hz, and this mismatch makes closed-loop analysis difficult. Consider the open-loop data compared to the design model outputs:

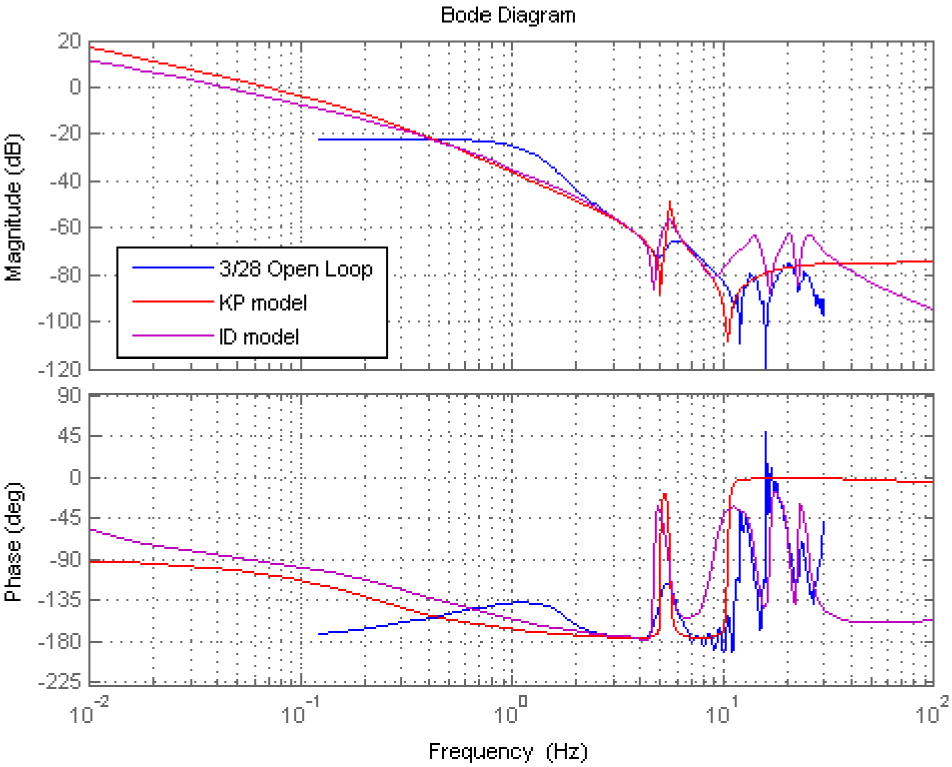

The lack of low-frequency fidelity has heretofore been left to the controller integral gain to force the proper DC low-frequency response. Since we are now in the regime where the gains in this frequency range have to be just right, it matters a great deal in terms of proper controller design.

A completely new physical-insight model of the elevation plant was constructed with much effort in order to improve the model fidelity to the low frequencies, the 6Hz modal pair, and the 20Hz companion pair. While it properly captured the phase and (generally) the magnitude response, it had positive real poles, and was consequently unstable in closed-loop form:

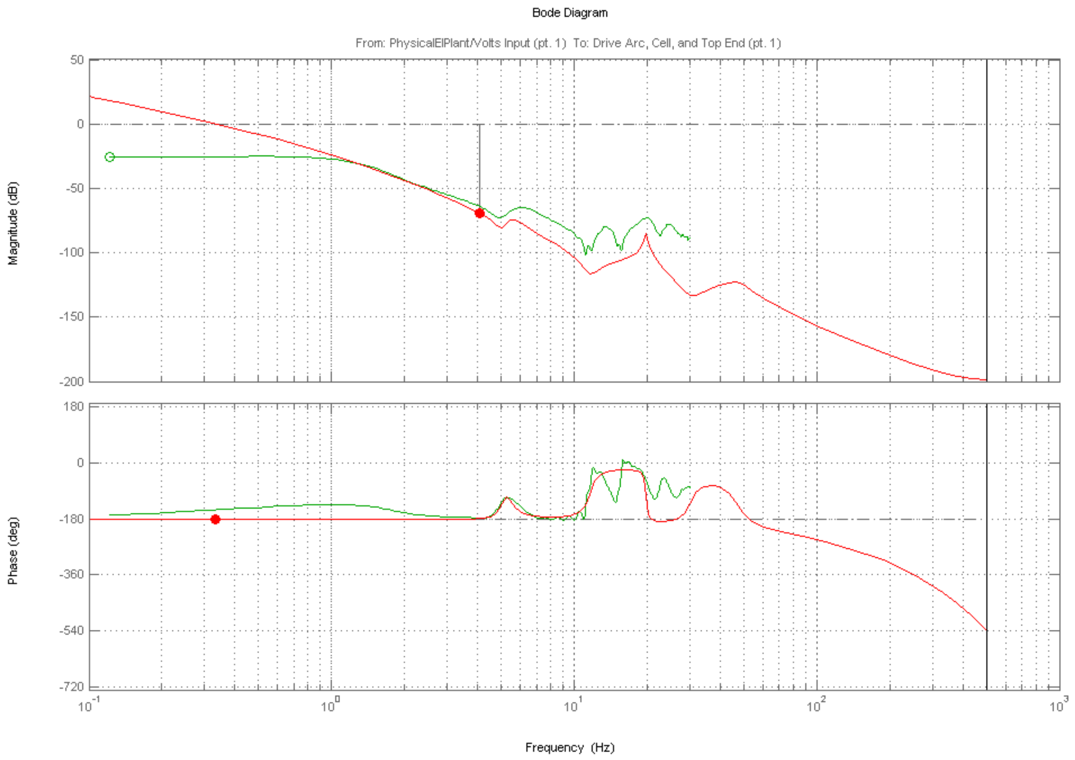

Work then turned to updating the original Powell physical-insight model, which *is* closed-loop stable. The work done on the new model was not a complete waste, however; the key insight from that work was that the long-misunderstood (and hence largely ignored) flatness of the open-loop response in the low frequencies is due to the inverted-pendulum action of the inertia of the telescope front end. The cartoon below shows the effect of gravity loading on the torque applied to the telescope: the motor torque accelerates the load with $1/Js^2$ slope, but gets “help” from the seesaw effect of the mass of the top end and cell, which combine to increase the overall acceleration to the load, and flattening the frequency response until the response once again becomes $1/Js^2$ in slope once the change in position becomes small.

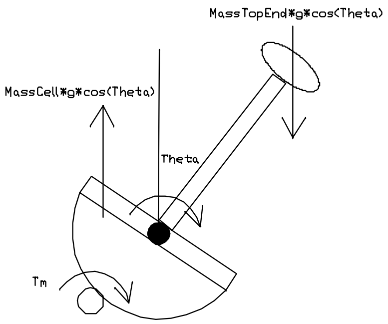

Extending Powell’s model with this idea in mind, we have this Bode plot of the new plant model, the original Powell model, and the measured data:

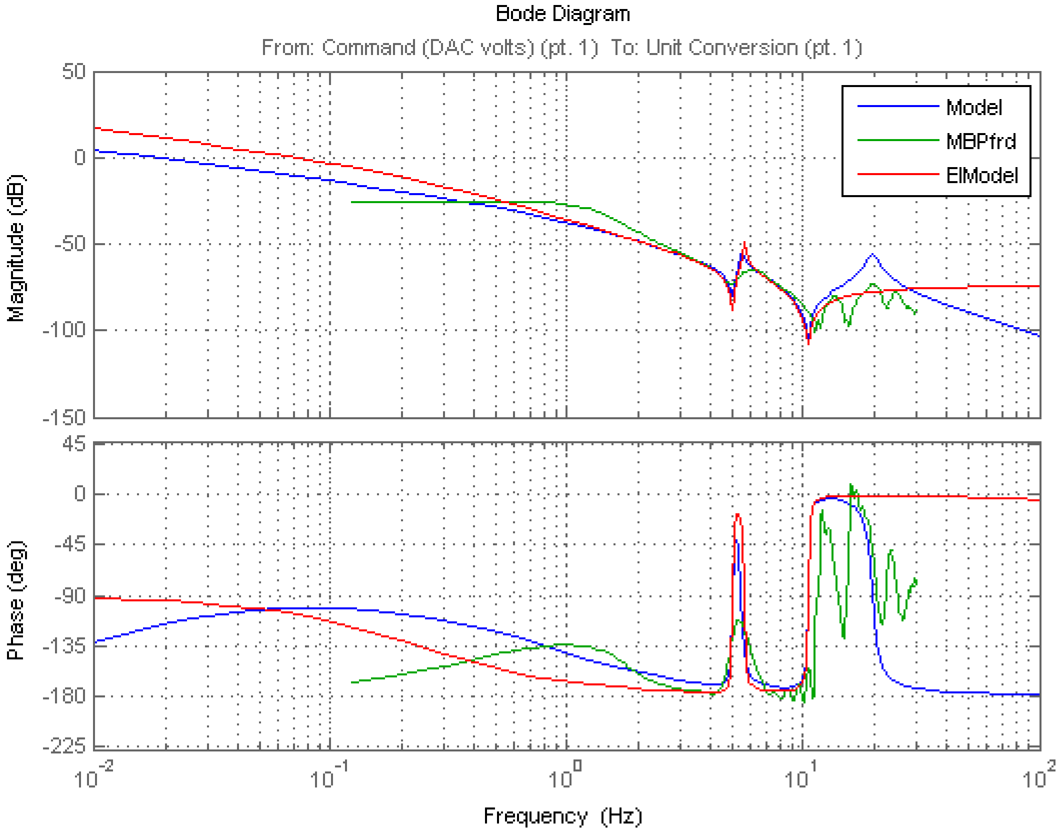

What evidence exists to support this gravity loading theory? First, all recent, and indeed archived open-loop tests with chirp signal inputs exhibit this low-frequency non-linearity. To further verify this, open loop tests were performed on 3/28 with both a chirp and band-limited noise inputs at two different elevations, 45 and 75°. The noise source was filtered with a second-order low-pass filter with a cutoff frequency of 10Hz, and scaled to keep the input torque to safe levels. First, the chirp source data results:

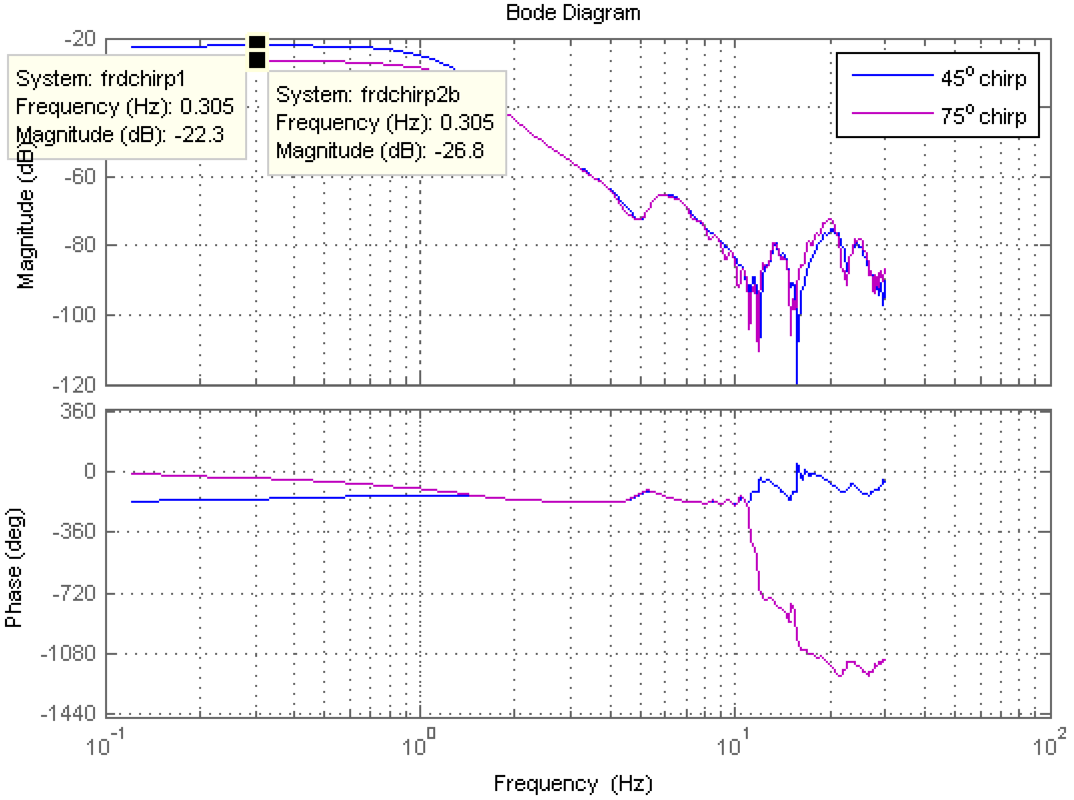

The overall modal response retains the same familiar shape. Notice, however that the low-frequency magnitude and phase have changed. The magnitude difference is 4.5db, which correlates well with the difference between $10\cdot\log_{10}[\cos(el)]$, which comes out to be 4.36db. Application of a noise input will result in lower DC content – how does it compare to the results with a sinusoidal chirp signal input? Below we have the same two responses, with the noise response at 45° (red), and 75° (yellow). The 75° data are somewhat corrupted by elevation drift due to a slight imbalance; care was taken to balance the telescope at 45°, and the data are remarkably smooth. However, the non-linear response is present for both excitation input types, and so we can safely conclude that it's for real.

The non-linear response needs to be properly modeled in order to get the controller design and its gain margins set properly; a more ideal model can only help if the disturbance-decoupled controller topology is selected. Again, it's probably worth accepting more tracking variation in return for the much-reduced overall elevation oscillation as a temporary measure until the controller design is fully optimized.

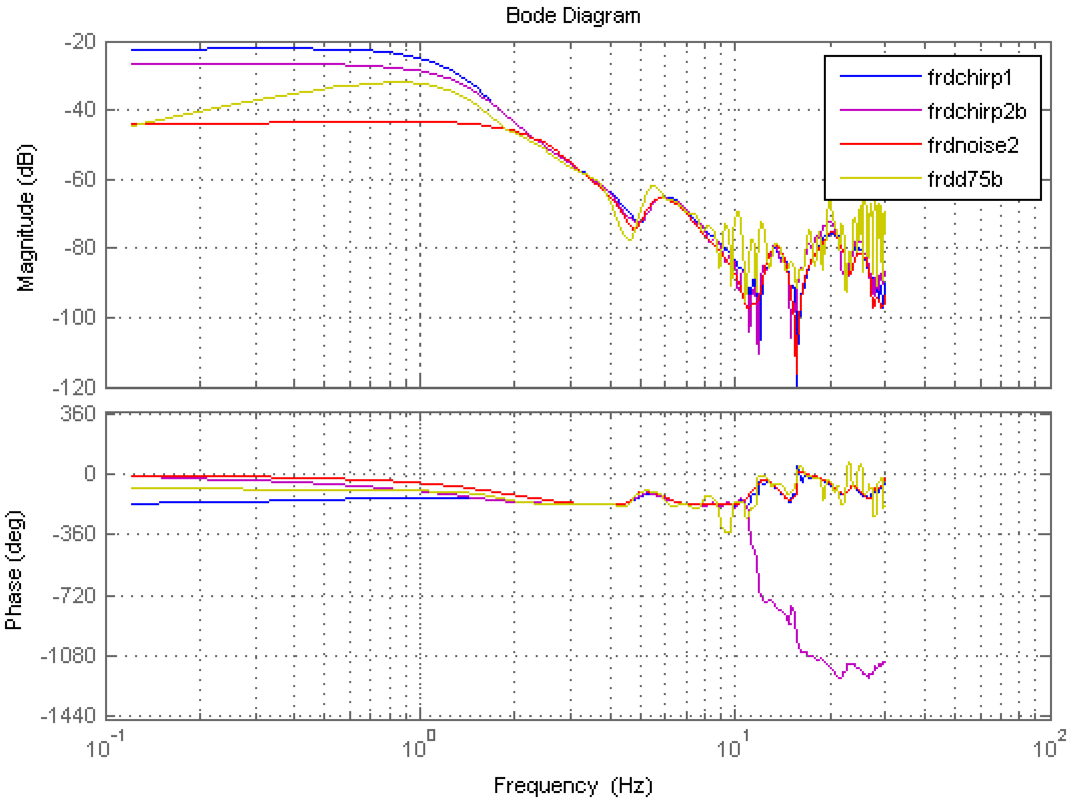

Some attention during the M&E night of 3/28 was also paid to simulation of the startup oscillation, which continues to plague us. There were several timing issues in the initialization period for the controller in Trebisky's version of the Simulink diagram. A version of this Simulink diagram was quickly built, and the loop closed on the standard elevation axis model. As expected, starting the controller in simulation causes the controller output to rail for several seconds in an oscillatory manner, as is observed on the telescope.

Further investigation revealed that large transients were being applied to the velocity loop from the velocity-estimation filter during startup, and these large transients, multiplied by the velocity loop gains, were causing large transient signal inputs to the flexible-mode filters, some of which are IIR (Infinite Impulse Response) types. An IIR filter tends to have undamped output signals over infinite time when an impulse signal is applied; the startup transient is therefore officially Not Good. A test fix was to apply a time-based soft-start to the output of the velocity estimator and the input to the flexible-mode filters. This reduced the large output transient to a 70-arcsecond, gentle transition as the controller gains came fully online over the course of ~3-5s. It remains to be tested on the telescope, but it is believed that this is comparable to startup transients seen with the existing LM628 servo.

It is critical, at this juncture, to get the controller working in the VxWorks system. All testing (including closed-loop) discussed here has been done with the xPC Target test controller, which has never been designated as a permanent replacement controller. Now that we understand how to successfully soft-start the controller, it should be a much safer proposition to run the controller and get it working on the telescope. The next critical path is twofold: a) create a high-fidelity model, and b) design the controller gains and topology to optimize overall performance. With a good model in place, we should be able to fully predict the disturbance-rejection properties of the new controller, and can confidently investigate alternate controller topologies, such as the disturbance-decoupled Luenberger observer discussed above, or even more exotic controllers, if desired.

The activities to be pursued are:

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.

Sources

1. W. Gawronski, *Modeling Wind-Gust Disturbances for the Analysis of Antenna Pointing Accuracy*, IEEE Antennas and Propagation Magazine, Vol. 46, No. 1, February 2004.
2. W. Gawronski, *Antenna Linear-Quadratic-Gaussian (LQG) Controllers: Properties, Limits of Performance, and Tuning Procedure*, IPN Progress Report 42-158, Jet Propulsion Laboratory, August 2004
3. G. Ellis, *Observers in Control Systems*, ISBN 0-12-237427-X, 2002.