Improving Elevation Axis Disturbance Rejection, Part I of ??

D. Clark

October 5, 2007

During M&E time in September, we found that the peak-to-peak tracking error of the new Simulink-generated elevation controller was of order 2X worse than with the LM628 controller under (more or less) similar wind disturbance, though the new controller was superior in terms of modal frequency suppression. To investigate improvements to this performance deficiency, more work was done to further optimize the controller's disturbance rejection properties.

The first question about the new controller's performance was, "Was the servo output saturating and causing the servo gain to drop during tracking?". Recall that there is a finite amount of a) DAC voltage output, and b) Amplifier current available for driving the telescope. The DAC saturates at ±5V, while the servo amplifier gain of 3.5A/V and the limited bus voltage internally leaves the maximum current output at 12A (which means the DAC limit is really 3.4V). For the figure below, and all other related graphs in this report, the servo telemetry file sv_20070919_2320403 was used; this was approximately a 10m/s wind with the telescope azimuth to the SW. The position error PSD and time-series output was graphed in a previous report, which noted a 8.95 arcsecond peak-to-peak error, resulting in RMS tracking of 1.39 arcseconds over a 60s period. The graphs below make use of the final 30s of the data in this file, below is the DAC output during this period:

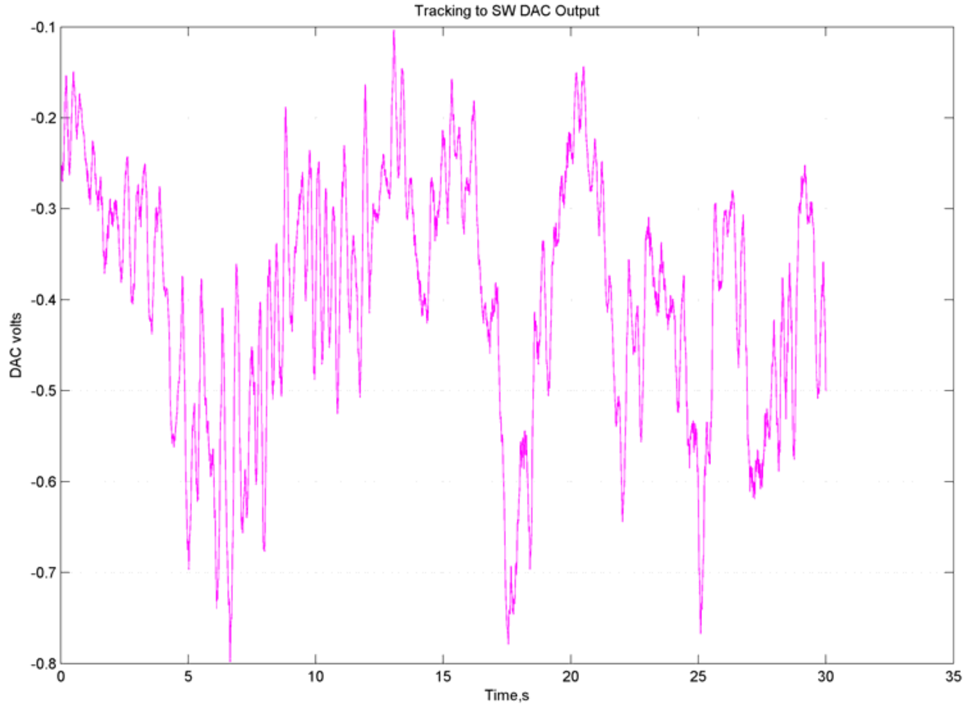

Note that the DAC voltage is well under the limit of 3.4V; motor amplifier saturation was not at issue during this tracking period.

The next step was to generate a time-series simulation of the wind loading with the as-used servo parameters for the purposes of comparing in an experimental way methods of improving the tracking performance. Below we have the simulation output versus the telemetry data:

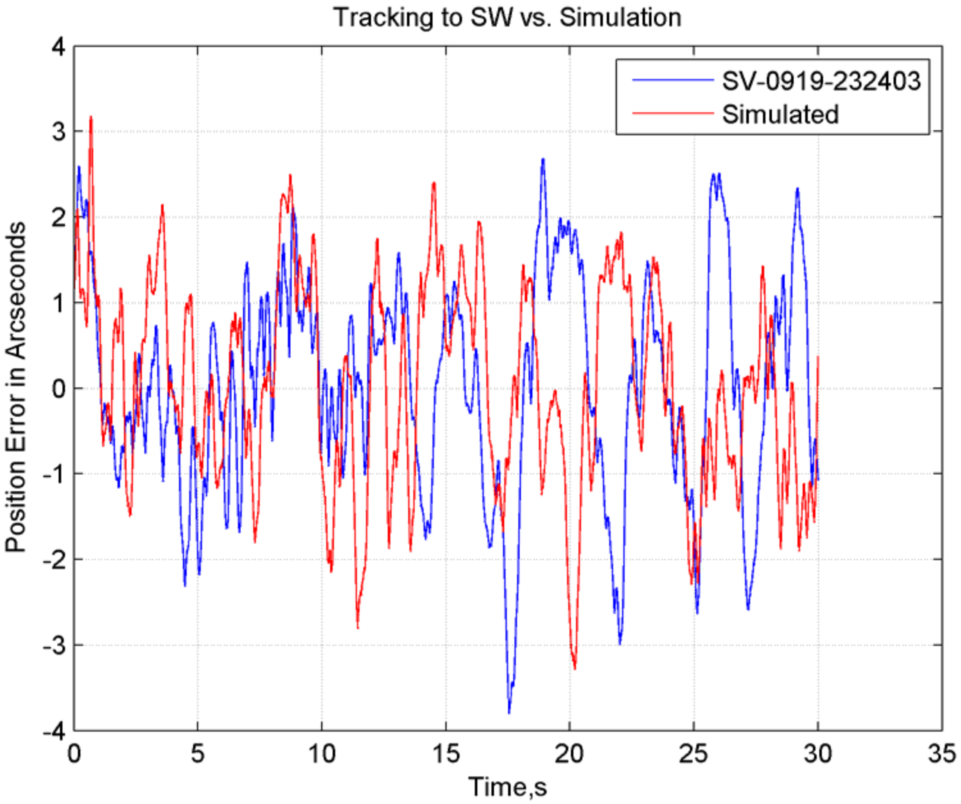

The file data have a peak-to-peak value of 6.49 arcseconds, with and RMS of 1.26 arcseconds. The simulation output peak-to-peak is 6.38 arcseconds and an RMS of 1.25 arcseconds. The data is in excellent agreement; no further work in scaling the simulated wind data was (or is) contemplated.

Next, the effect on the disturbance rejection of adjusting the position-loop gains was investigated. Without regard to their effect on the controller's stability or command-signal bandwidth, the position loop integral and proportional gains were arbitrarily increased to show in a quantitative way their effect on the disturbance rejection.

For the purposes of control design, disturbance rejection is defined as the ratio of the controlled-variable output (i.e. the telescope's position) to a torque signal input that comes from outside the control loop (traditionally summed into the drive amplifier input). Ideally, therefore the response is -∞ at all frequencies. In a practical system, infinite rejection is not possible; loop integrators make the rejection high at DC to some low frequency, while the proportional gain interacting with whatever modal frequencies result in the rejection lowering in low-to-intermediate frequencies, and the controller's derivative gain and the filtering effect of the load inertia combine to again increase the disturbance rejection at higher frequencies.

The controller as tested implemented a variable-gain integral term in the position loop. What was its' effect on the loop disturbance rejection? Below is a Bode plot of the disturbance rejection for the minimum and maximum integral gain. Since it varies with the absolute value of the position error, the instantaneous disturbance rejection will lie between the two lines shown:

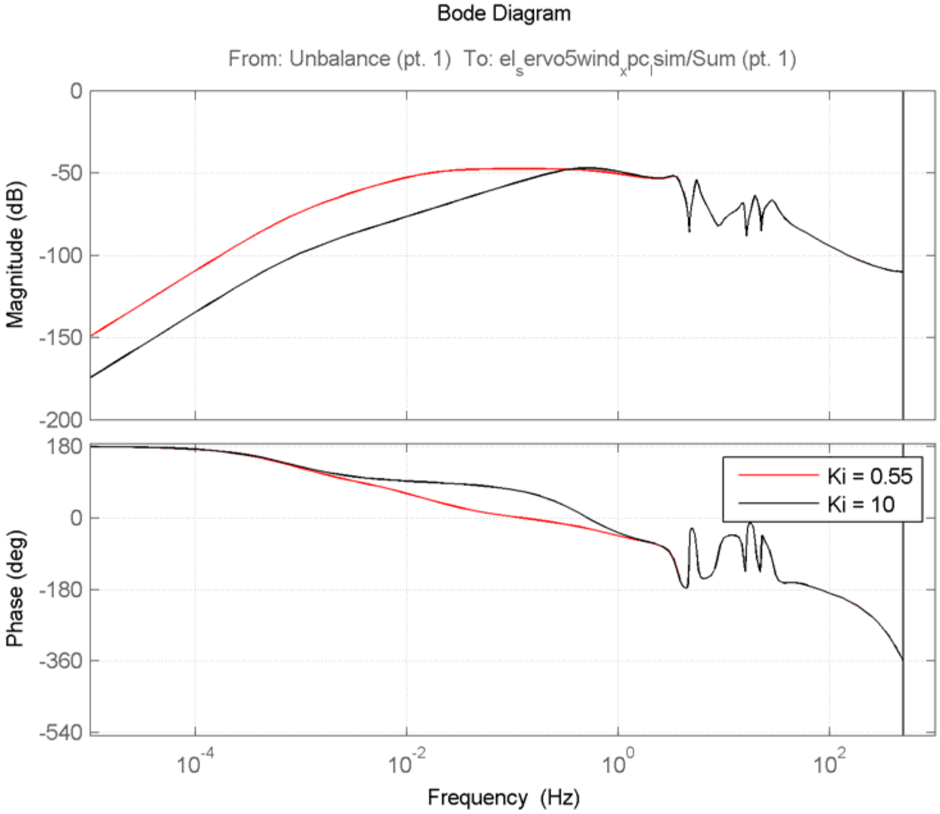

The effect of the increased integral gain is surprisingly ineffective at frequencies approaching 1Hz, right where the wind-buffeting rejection is needed. However, it significantly increases the servo stiffness at the low frequencies, so there is some benefit to having it.

The fastest path for the error signal in the controller is the proportional gain, and as been mentioned, it tends to operate in the intermediate frequencies. Again, without regard for stability concerns, this gain was raised arbitrarily:

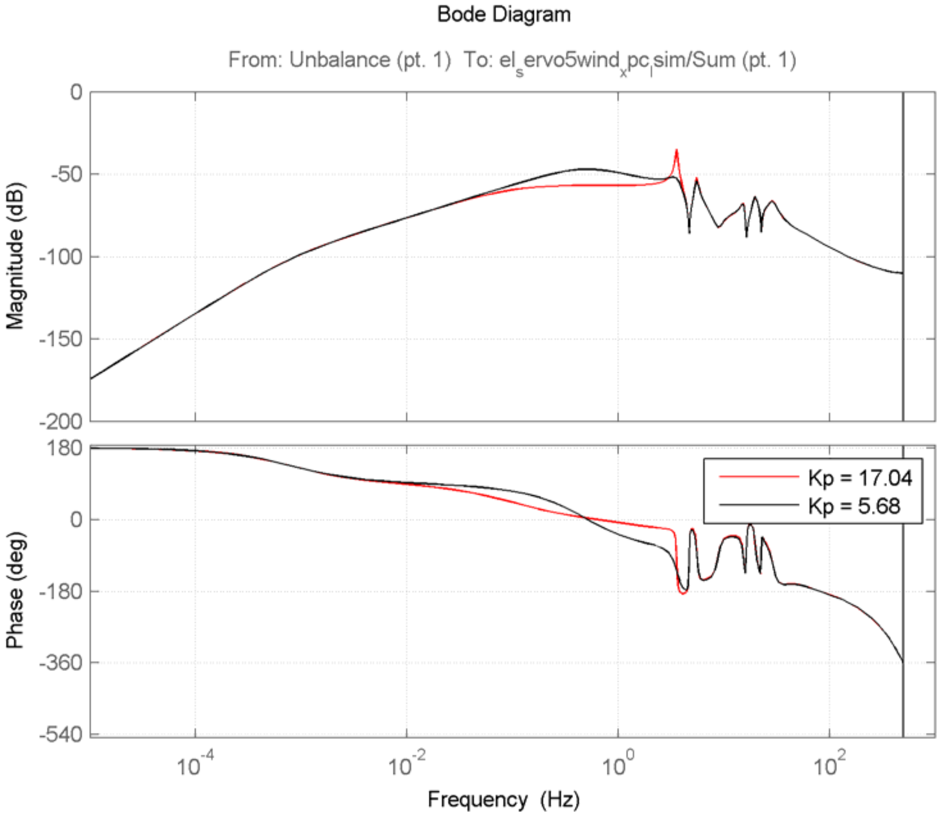

The proportional gain clearly improves things in the frequency region of interest. This is achieved at the unavoidable cost of creating gain peaking at ~ 3Hz. In other words, the control loop stability is compromised, and is likely to oscillate at 3Hz, especially if an unwanted excitation is provided (e.g. a tuning fork). This is confirmation of the gain peaking evidenced in the LM628 data at ~ 2.5Hz being indicative of a marginally-stable control loop. The fact that the open-loop plant includes a modal pair at 5 and 6Hz (notch and peak, respectively) means that the closed-loop bandwidth can only approach, and never equal, the notch frequency; indeed, the servo loop will be non-linearly closer to instability as the bandwidth is raised towards the notch frequency.

Given that reasonable adjustments to the position loop gains can improve both the control-signal bandwidth and the disturbance rejection, some work in adjusting those (and the velocity loop gains, as the loops interact) was performed:

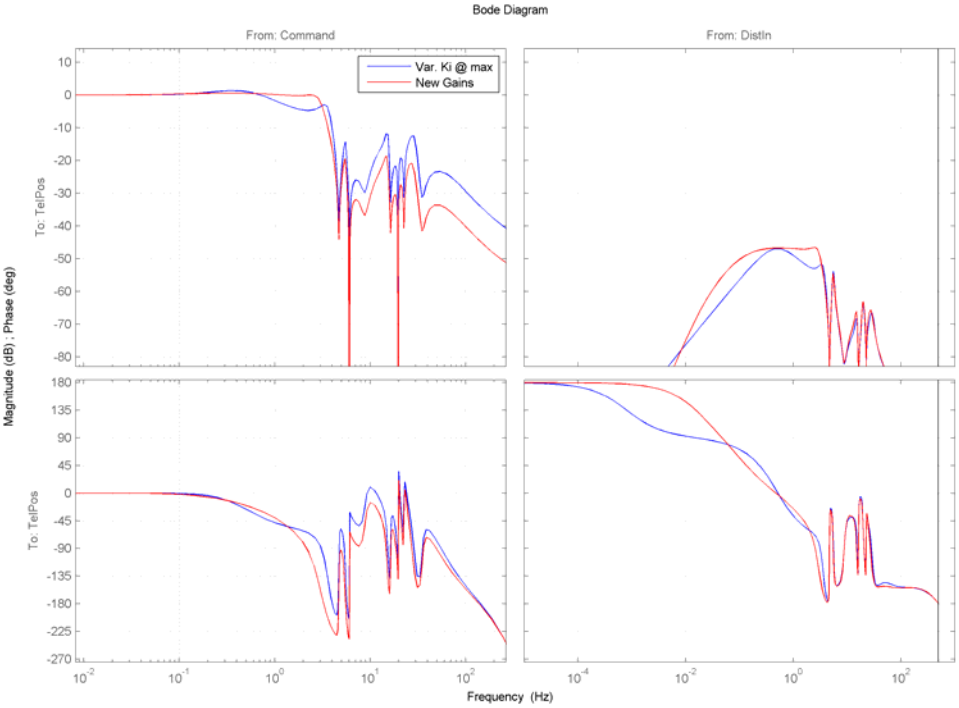

As can be seen, the control loop command signal path is much improved, at moderate cost to the disturbance rejection (i.e. little to no change in the overall disturbance rejection). The time-series simulation confirms this result:

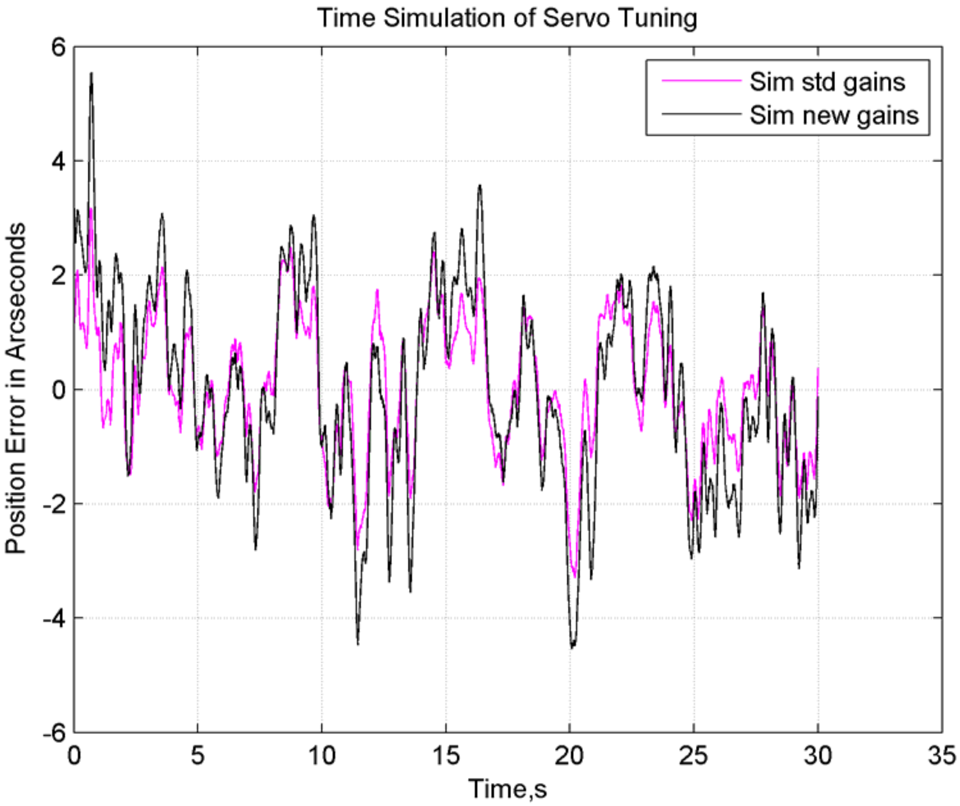

The new gains result in somewhat degraded position loop error in the face of the wind disturbance simulation. The improved control-signal bandwidth makes it worth testing, however, to confirm (or disprove) the simulation results.

Why is the wind disturbance so difficult to remove? A couple reasons: 1) conservative control design, avoiding gain peaking and seeking to preserve robustness when millions of dollars worth of equipment and safety of personnel are at stake means we aren't as aggressive as could be when applying potentially unstable servo gains, and more compellingly, 2) the error signal is very small at the 0.1 arcsecond level. To explain, consider—

> The controller is designed to operate for a number of scaling reasons in units of degrees. This means that for a 1° signal level, we have for a 0.1 arcsecond error, an error signal at 1/36000, or 2.778×10^{-5}, which in decibel terms is -45db. This implies that the disturbance rejection must be a minimum of 45db, and probably higher, to safely reject signals down below a few arcseconds. The controller gains in the position and velocity loops that most quickly operate on this small signal are the proportional gains, and if we temporarily ignore the (time-lagged) effect of the integral and derivative terms, multiplication of this signal by the gains, the servo amplifier output gain, the motor torque constant, and the mechanical advantage results in a final telescope torque of 19.924 ft-lbf. This is too small to do much against a wind load that can be in the hundreds of ft-lbf.

The -45db disturbance rejection factor is in fact about what we see in the simulation Bode plots. What is desired is to *increase* the disturbance rejection, without resorting to controller gains that make the servo stability questionable – especially when we routinely change the telescope inertia by changing the optical and instrument loads. How to do this?

One advanced method is to introduce disturbance decoupling. This leverages that fact that given a sufficiently accurate model of the plant to be controlled, comparing the actual plant to the modeled plant in real time results in a measurement of the plant's deviations from the model; those deviations can be attributed to errors in the model, or plant inputs that cannot be accounted for by the known inputs to the plant, i.e. disturbances. Scaling the measured deviations and injecting them back into the plant causes the disturbances to be minimized, up to limits imposed by the model's accuracy and the phase lag associated with the disturbance-estimation calculation. Here is a block diagram illustrating this controller topology:

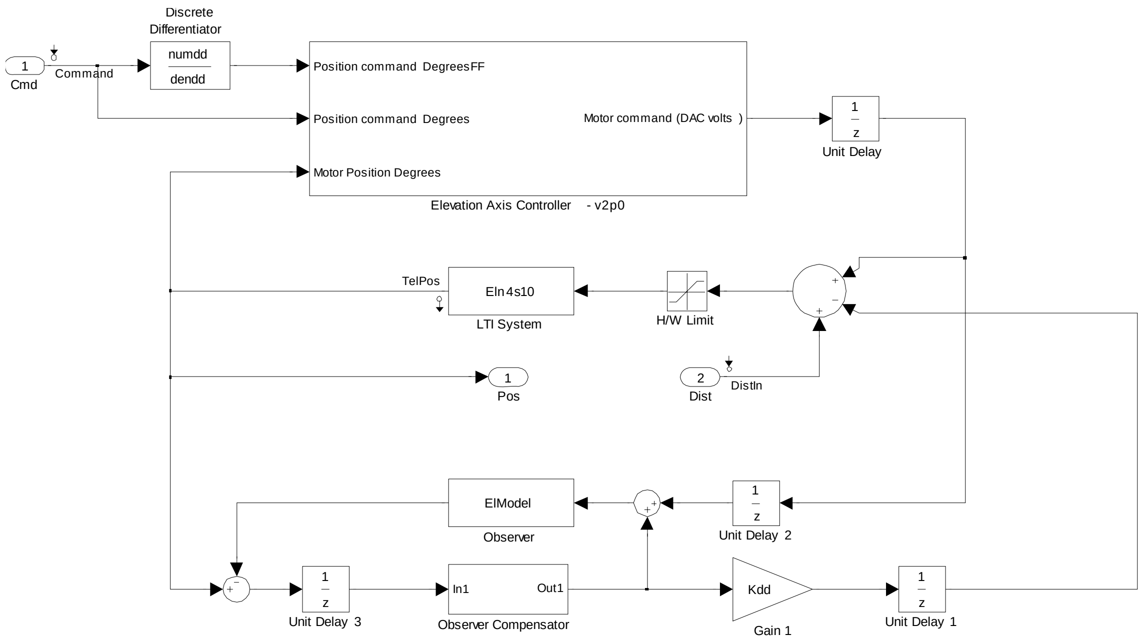

This is the simple Simulink diagram used for evaluation of the closed-loop controller command and disturbance response, augmented by a Luenberger Observer with disturbance decoupling. The large block marked "Elevation Axis Controller" is in fact the controller we've been using to drive the telescope. The telescope itself is simulated by the LTI System block, 'Eln4s10', which is the identified model of the telescope that includes the 20Hz modal frequency. A sum block on its input allows the disturbance input to be modeled. To add disturbance decoupling, the plant input before the disturbance is used to drive the telescope model, Powell's original telescope model based on the early servo design work. The difference between the two models is compensated by a simple PID and is fed back to the telescope model and scaled and applied (in a negative sense) to the telescope – it is this last signal that represents the estimated disturbance. The disturbance-estimation authority can be scaled (within reason) to provide an adjustable counteraction to the disturbance.

This approach is wholly dependent on the accuracy of the system model for best performance and robustness. In simulation, two different models for the plant and observer help to investigate the deleterious effects of model error on the controller robustness and performance. In the real world, the controller is also affected by noise, scaling errors, and nonlinearities in addition to modeling error. If these error sources are sufficiently small, the controller gains can be safely raised, without compromising the control loop bandwidth or stability.

Consider the Bode plot of the closed-loop control-signal bandwidth and disturbance rejection:

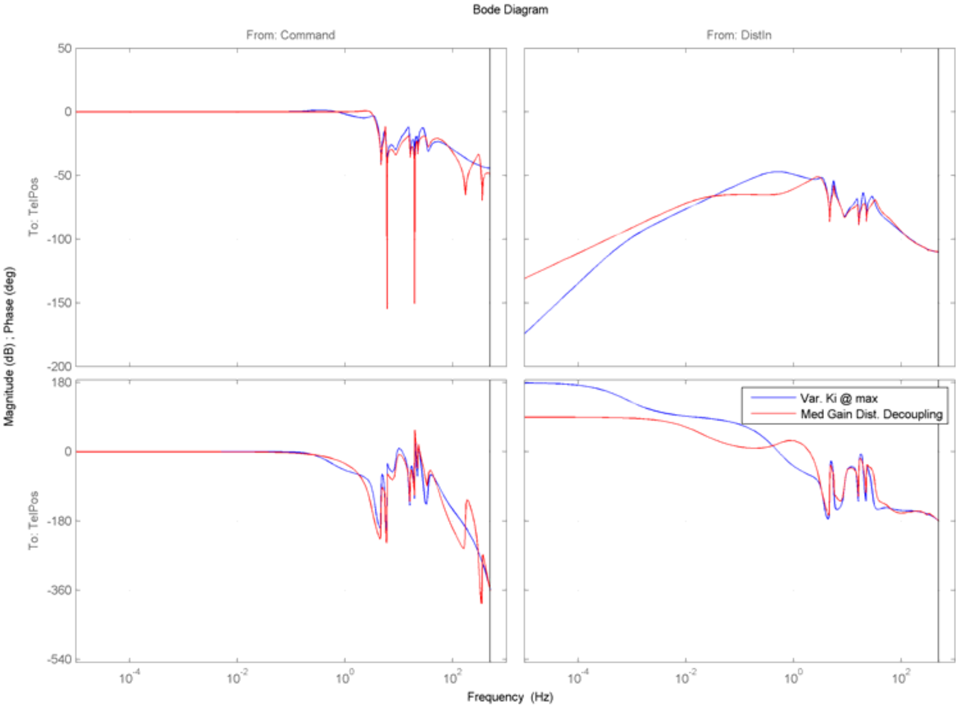

As you see, the control-signal path compared to the un-augmented controller with new gains is nearly identical, while the disturbance rejection shows much improvement. This is also seen in the time-series simulation, which compares well to the servo telemetry results. The simulated RMS is 0.23 arcseconds, with a peak-to-peak value of 1.45 arcseconds. The effect of the disturbance decoupling is dependent on the allowed gain of the decoupling signal (in the below case it is 10; for the above Bode plot, it was 5). We will need to test this controller topology on the telescope, starting with small decoupling gains to explore the allowable parameter space. If the servo is stable and well-behaved with the Luenberger Observer in place, this may well point the way to a much-improved controller.

More test and verification await.

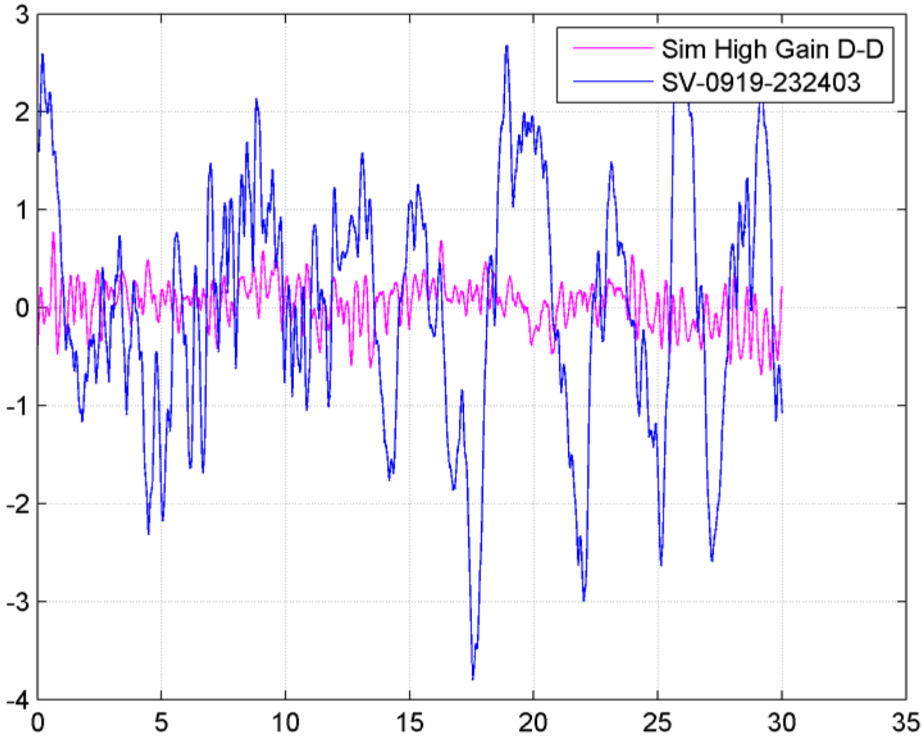
Sim High Gain D-D
SV-0919-232403
3
2
1
0
-1
-2
-3
-4
0
5
10
15
20
25
30
35