Tuning a 2.4 Meter Telescope... Blindfolded

Kyle Lanclos,[a] Michael Peck,[a] Michael Saylor,[a], Robert I. Kibrick,[a] Steven L. Allen,[a]

[a]University of California Observatories, Santa Cruz, CA 95064 USA
http://www.ucolick.org

ABSTRACT

Just as the 2.4 meter Automated Planet Finder (APF) commenced its final shakedown, three significant events occurred: uncontrolled telescope oscillations while tracking, liquidation of the telescope vendor's primary facility, and the expiration of the vendor warranty. Left with scant documentation, few external resources to draw upon, and limited direct local expertise, University of California Observatories (UCO) embarked on an initiative to stabilize the telescope control system at a minimal internal cost. This paper covers the problems encountered, our solutions, and the compromises made when the budget could not support a complete remedy.

Specific topics include: measurement and alignment of linear encoder signals, and custom electronics developed to enable precise alignment of the read heads and adjustment of the interpolation electronics; the use of sensitive accelerometers to isolate and diagnose sources of vibration, and to provide immediate feedback on the stability of the servo tuning; procedures used to adjust the servo control loop, and the observable effects of parameter adjustments; assessment and validation of the performance on-sky.

Keywords: Lick Observatory, Automated Planet Finder, servo tuning, linear encoder alignment, accelerometer feedback, telescope stability

1. INTRODUCTION

1.1 What is the APF telescope

The APF telescope[1] is an altitude-azimuth design with a 2.4 meter primary mirror, sited at Lick Observatory, just east of San Jose in California. The telescope drive motors are direct-drive Kollmorgen brushed servo motors, one built into the azimuth bearing, the other built into the telescope yoke around the right Nasmyth focus. A single Heidenhain LIDA 181 linear encoder tape is installed into each active bearing; there are two Heidenhain LIDA 10C (309-237-01) read heads for each encoder tape.

The telescope is controlled by a single Delta Tau Turbo PMAC controller, an ISA expansion board installed into a rack-mount computer running Microsoft's Windows XP. A traditional proportional/integral/derivative (PID) loop is used to control the telescope axes. Crucially, the raw PMAC control code is available for direct manipulation; without that access, low-cost adjustments to the servo control loop would not have been possible.

The analog drive output signals from the PMAC controller feed into a pair of Advanced Motion Controls 50A20 servo amplifiers, one for each motor. These amplifiers have their own internal settings to control the responsiveness of the amplifier to requested changes in output. The 96V input power for the servo amplifiers is provided by a pair of ganged off-the-shelf 48V power supplies.

1.2 Abbreviated history of the APF telescope

A full accounting of the construction and final acceptance of the APF telescope and dome is beyond the scope of this paper. Only the facts relevant to this paper are presented in this summary.

The design and construction of the dome was contracted out to Electro Optical Systems (EOS) Creative Technology Solutions in Australia, with the telescope design and construction handled by EOS Technologies (EOST), then in Tucson, Arizona. The telescope would support two Nasmyth ports, with one port occupied by the UCO-designed and fabricated Levy spectrometer.[1] The primary science objective was to locate extra-solar planets via precise radial velocity measurements from high-resolution spectra of individual target stars.

Send correspondence to K.L.: lanclos@ucolick.org

The contracts were issued in 2003, with the intent of producing science-grade data by 2005; on-site construction of the dome began in late 2005, and site acceptance occurred in November 2010. UCO delivered the spectrometer in January 2011, and after a delay due to a significant mechanical failure,[2] installed the instrument on the telescope in June 2011.

A variety of problems with the telescope and dome prevented significant on-sky use after the installation. Many of these problems were resolved via warranty claims. In late 2011, it became apparent that there were significant unresolved oscillations in the telescope control system; the vendor made two on-site visits to effect a remedy, and was unsuccessful. In the meantime, the facility sat largely idle. In mid-2012, the vendor elected not to remedy the outstanding warranty claims; in September 2012, UCO began its own effort to identify and mitigate the outstanding problems.

1.3 What is the blindfold?

The blindfold is a metaphor for the procedural, as opposed to technical, obstacles that impeded efforts to remedy the servo control system. Put another way, the blindfold resulted from deliberate decisions, with or without the participation of the receiving institution. There were many facets to this fundamental problem.

No servo analysis. There was no theoretical basis for any of the servo control settings, everything was strictly empirical. This left a lot of unexpected mysteries to be discovered while attempting to isolate observed problems.

Missing documentation. While UCO did receive a substantial volume of documentation from the vendor, many key aspects of the facility were not covered, including control board designs, or the procedure used to adjust and validate the servo control parameters. Many of the delivered documents contained serious errors, such as sense mismatches in wiring diagrams. The absence of key documents, and errors in existing documents, severely hindered the overall troubleshooting effort.

No source code. The control software was delivered by the vendor as installed binaries; with the sole exception of the motion control scripts in the non-volatile memory on the PMAC controller, no source code for any aspect of the control software was available for review or modification. This created a large barrier between the testing that could be performed on-sky versus testing done in an engineering mode: it was not possible to collect telemetry with high enough fidelity to diagnose problems directly while tracking stellar objects; likewise, it was not possible to effect minor changes to the servo tuning parameters with the vendor's high-level control software running.

No direct expertise. The telescope and dome were designed and constructed by external contractors. The actions taken by the vendor to remedy outstanding problems were often not shared with local staff. No formal or extended effort was made to transfer a functional technical understanding of the facility to the receiving institution. All expertise had to be developed the hard way: from first principles, leaning hard on UCO's institutional experience with large telescopes, and the sheer tenacity of UCO personnel. The APF represents at least three significant departures from existing facilities at Lick Observatory: the APF was the first (and remains the only) alt-az and the first direct drive telescope at Lick Observatory, and it was UCO's first direct exposure to Delta Tau motion controllers; going back to the mid-1980's, UCO developed motion control systems primarily using Galil controllers.

Minimal budget. Significant amounts of time and money, much of which came from external grants, went into the development of the APF as a flagship project for the future of Lick Observatory. With no small measure of institutional pride on the line, ensuring the success of the APF was seen as a way of ensuring the continued and future success of Lick Observatory. Unfortunately, the hour of need for the APF came at a time when resources were thin at UCO, and internal resources were allocated with overwhelming priority to externally funded projects. In the absence of both external funding and internal support, it was very challenging to allocate sufficient resources to apply even minimal fixes to outstanding problems at the APF.

1.4 Summary of corrective actions

Over a period of six months, attempts at improving the servo performance would interleave with work to address mechanical or electrical problems isolated in the previous iteration. Like an onion, the problems were layered on top of each other, with each layer obscuring further problems buried beneath it. Attempting to bypass the outermost layers was an exercise in chaos theory, yielding wildly unpredictable results to modest perturbations of control parameters.

Problems that could be thoroughly decoupled from other aspects of the control system were the first to be approached. The telescope tube was not balanced (section 2) along its elevation axis; restoring balance required only that the telescope slew at constant velocity, precise tracking was not required. Determining the zero point offset of the motor controller and the servo amplifiers required only those two respective pieces of equipment, no motors or telescope were necessary. Properly tuning the current loop (section 4) within the servo amplifier only required a change in the digital-to-analog converter (DAC) output from the motor controller, and like the offset, no motors or telescope were necessary.

After peeling back these layers, but most especially after the adjustments to the current loop tuning, the full magnitude of the poor performance of the servo control parameters was apparent. Experimenting with the basic PID parameters along with low-pass filters yielded some improvement, but each incremental improvement would run into a point of instability well before achieving the desired performance. Each time, the boundaries of the solution space had to be expanded by electronic or mechanical means: some fixes were small, such as decreasing the gain in the servo amplifiers, thereby decreasing the quantization on the output signal; some fixes were large, such as properly aligning the encoder heads (section 5), thereby decreasing the error on the encoder feedback signals.

After addressing enough of these issues, the servo control system started to behave more like a simple, small servo motor, in that the PID parameters could be adjusted (section 6) to predictable outcomes, and secondary terms in the servo control loop such as velocity feed-forward started to have an observable and beneficial effect on the overall performance. In all early attempts, the feed-forward terms had either no impact, or a negative impact. In this phase, the mechanical resonances of the telescope structure began to dominate as sources of instability. As it was clear that no funding was available to address these major mechanical issues, our effort shifted to finding a set of 'lucky' parameters that avoided exciting these resonances.

The key parameter in this case was the frequency of the low-pass filter. A minor adjustment from 12 to 14 Hz would yield a very different pattern of resonances, amplified to varying degrees by the mechanical resonances. After limited trial and error, a frequency was discovered that avoided the worst of the mechanical resonances. The telescope began taking world-class data immediately afterward.

2. BALANCING THE TELESCOPE TUBE

An on-site visit by the vendor in July 2012 was prompted by an inability to 'home' the telescope. The root source of that problem was a failure in one of the Baldor controllers used with the dome drive motors, but that determination was not made for several months. One suggestion offered by the vendor was that the telescope tube was not sufficiently well balanced, and the balance should be checked and adjusted as necessary before proceeding with further diagnosis.

2.1 Procedure provided by vendor

The procedure used by the vendor involves slewing the telescope in elevation from the zenith to the horizon, and back again, at a constant speed. With a multimeter on the elevation axis servo amplifier output, one technician would call out the absolute position of the telescope, and the multimeter reading would be recorded at each call out.

Upon plotting the amplifier output against the elevation angle, the two lines should be straight. If they are not, and if comparing the absolute output values, the two lines will intersect at a specific elevation angle. That angle is the balanced point. Having made that determination, weight would be added or removed from the top end of the telescope tube to bring the telescope better into balance.

When attempted in July 2012, this procedure did not succeed, due largely to the failing Baldor controller; while the failing controller was for an entirely different subsystem of the telescope, it was pumping enough noise into the electronics cabinet that it was not possible to use a multimeter or oscilloscope for any diagnostic purpose.

2.2 Procedure devised by UCO staff

Before the vendor concluded their on-site visit, we suggested an alternate procedure, with the intent of minimizing by-hand data collection. Nearly all of the telemetry offered by the vendor-supplied telescope control software is captured by UCO-authored software, and distributed as Keck Task Library (KTL)[3,4] keywords, the common application programming interface (API) for all software at Lick Observatory. These KTL keywords, in turn, are collected in a history database,[5] which provides on-demand access to any collected data set. By using software to collect the quantitative data, it was possible for a single technician to carry out the basic procedure and analysis.

The three parameters of interest were the actual telescope position, the telescope velocity, and DAC output, all for the elevation axis. The velocity was inspected to ensure it was constant for the duration of the collected data set.

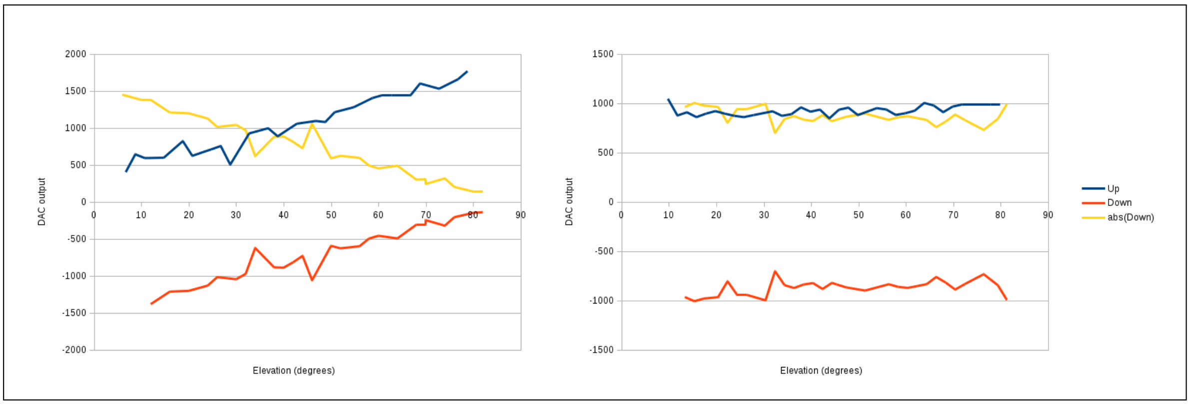

Figure 1. DAC output at constant velocity, before and after balancing the telescope

The initial test of this procedure demonstrated an observable imbalance, as seen in figure 1. The slope of the line indicates the nature of the imbalance, if any; when the telescope is heavy on the bottom, the lines will have a positive slope; when heavy on the top, a negative slope; when perfectly balanced, the slope will be zero.

Each adjustment to the weights at the top end of the telescope tube was followed by a data collection pass. After several iterations, the telescope was considered successfully balanced; the resulting data set is shown in figure 1. A small slope is present in the final data, indicating that the telescope remains ever so slightly heavier on the bottom side. This residual imbalance is not sufficient to accelerate the telescope when it is unpowered, and is easily handled by the PID control loop.

3. THE UTILITY OF ACCELEROMETERS

While the telescope is in operation, the vendor-provided control software publishes telemetry roughly twice per second. While sufficient for extended tasks such as telescope balancing procedure in section 2, this data stream was not of high enough fidelity to assess, much less diagnose, oscillations in the servo control system.

To address this lack of visibility, a pair of high-resolution Dytran 3191A1 accelerometers were installed on the telescope, one near the base of the telescope tube, the other on the telescope yoke. The software provided with the accelerometers gives live feedback on the acceleration of the telescope, and a power spectrum highlighting the frequencies that are most active. The data can also be captured on-demand, and analyzed after the fact. Qualitative statements of "the star is making Lissajous figures on the guider" or "it's buzzing" or "it's shaking itself senseless" could now be matched with hard data. The accelerometers also provided an essential common view between different control modes of the telescope, in that tests could be developed with low-level software to match the accelerometer signature of an oscillation experienced while tracking a star.

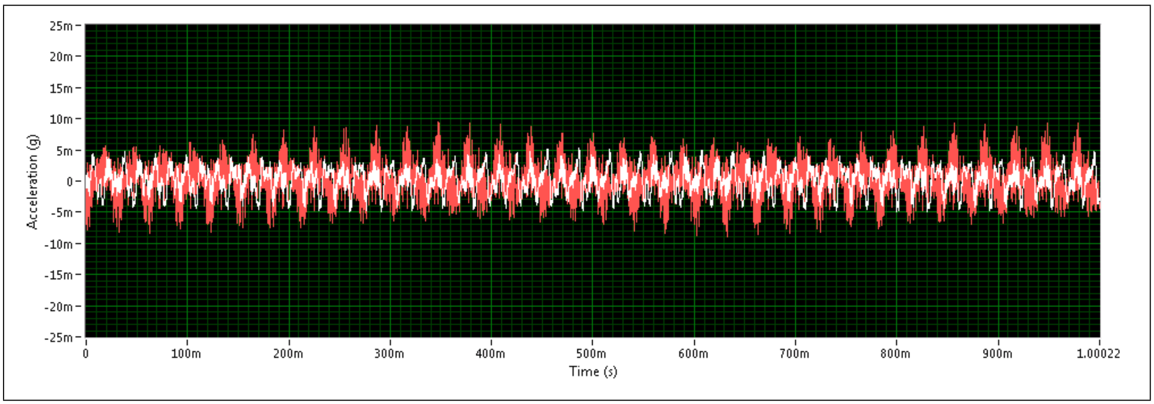

Figure 2. Acceleration signature of an oscillation while holding position

A simple early oscillation is shown in figure 2. Each axis is represented by its own color: the elevation axis is in white, and the azimuth axis is in red. When the acceleration is graphed over time, a strong oscillation shows what appear to be standing waves, harmonic structure in what would ideally be a flat line. This specific oscillation occurred while the telescope was servoing to hold position while pointed at the zenith.

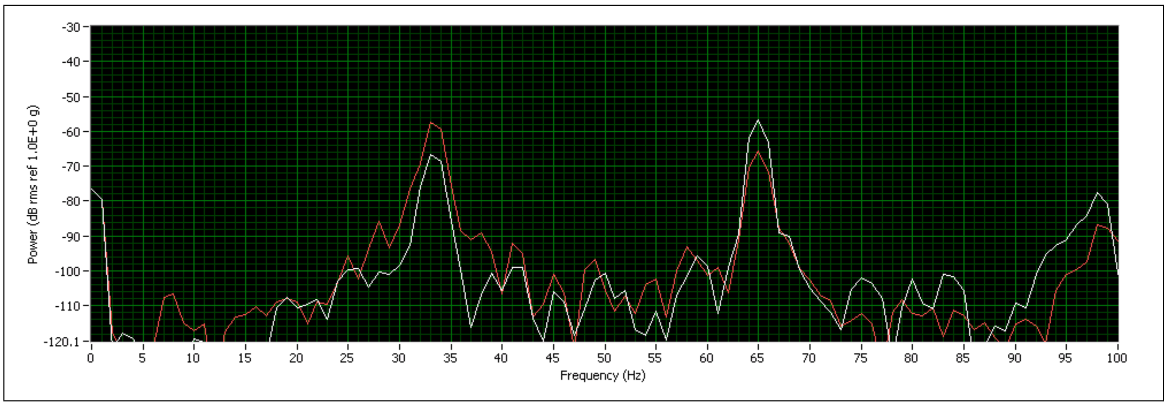

Figure 3. Power spectrum of an oscillation while holding position

The power spectrum for the same oscillation is shown in figure 3. The oscillation appears to be driven by the azimuth axis in the 30-35 Hz range, and is coupling into the elevation axis; the entire structure is then resonating at or near even multiples of the base frequency. With a peak power above -60 decibels, this oscillation was audible to personnel in the dome. Though this oscillation was being driven by the servo control loop, it was amplified by a mechanical resonance in the telescope yoke.

Empirical identification of mechanical resonances was also possible with the accelerometer measurements. The procedure, though crude, was effective: walk around the telescope, and bang on different components of the structure. When something rings, check the power spectrum to identify the resonant frequency. Two major resonances were identified in this fashion: one was a large steel plate at the base of the telescope, just above the azimuth bearing; this plate exhibited a strong resonance at 45 Hz. The second resonance was a matched pair of large steel plates in the telescope yoke, on either side of the telescope tube (see figure 4), with a strong resonance between 25-35 Hz. These resonances would have been readily identifiable via swept sine analysis, had a full servo analysis been performed.

4. TUNING THE CURRENT LOOP

The Advanced Motion Controls 50A20 servo amplifiers used with both the azimuth and elevation axes are smart motor driver/controllers, with the ability to operate without an external controller, including options for setting proportional gain, bandwidth filtering, and integration.

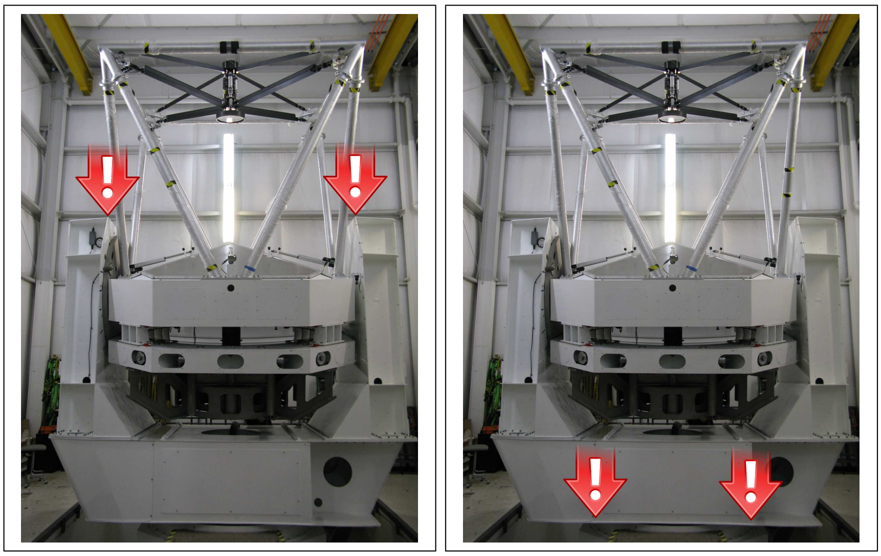

Figure 4. Location of metal plates subject to resonances at 25-35 Hz (left) and 45 Hz (right)

In the as-delivered configuration, both servo amplifiers were set up in such a way as to duplicate some of the servo control functionality performed by the Delta Tau PMAC controller. In particular, the servo amplifiers had filtering, integration, and high gain settings that negatively impacted the performance of the overall servo control system.

We adjusted the settings on the servo amplifiers to get linear response and lower gain so that we could do all the loop tuning digitally in the PMAC controller. The specific changes were:

1. Turn switch 7 "ON" to deactivate integration
2. Turn switch 3 "OFF" to increase the current loop proportional gain
3. Turn potentiometer 1 "Loop Gain" fully counter-clockwise for the minimum gain
4. Adjust potentiometer 3 "Gain" for full scale output current with the PMAC input at 90% of full scale.
5. Adjust potentiometer 4 "Offset" such that the amplifier had zero output when the PMAC output was commanded to zero

At this point, the output from the servo amplifiers responded linearly to commanded changes. The servo control loop tuning could now be handled exclusively and deterministically within the PMAC controller itself.

5. ALIGNING THE ENCODER READ HEADS

Of the various telescope oscillations, one of the most reproducible had an interesting character, in that the peak frequency of the power of the oscillation, as observed on the accelerometers, was linearly related to the velocity of the telescope in either azimuth or elevation. With the velocity reported in degrees per second, the scale factor was approximately 370. This relationship was confirmed with velocities between 0.01 and 2.0 degrees per second, with corresponding peak frequencies between 3.7 and 740 Hz. The velocity dependent nature of the oscillation suggested a problem with the encoder subsystem;[6] a quick check of the encoder tape properties confirmed this hypothesis.

The encoder tape itself has a grating period of 40 microns, or 25 lines per millimeter; over its full length of 5,230 millimeters, there are 133,000 grating periods, or lines. The tape is affixed end-to-end within a circular groove inside the azimuth bearing. Hence, for each one degree rotation of the azimuth axis, the read head will traverse 369.444 (133,000 / 360) lines on the encoder tape.

After discussing the observed behavior with technical staff at other facilities using similar encoders, the consensus opinion was that the likely cause was a misaligned encoder read head. The non-optimal sine waves generated by a misaligned read head cause the encoder's interpolation electronics to produce interpolated square waves that are not evenly spaced over a given period of the encoder tape, resulting in a nonlinearity in the encoder signal. Even though the square waves were not evenly spaced, the interpolator module still produced the correct number of squares waves over each grating period. As a result, there is no net change in the total number of interpolated encoder counts and thus no cumulative effect on the overall telescope pointing, although the varying width of these square waves will result in localized, very small-scale pointing and tracking errors. Thus, if considering only telescope pointing accuracy, the encoder appears to be working correctly.

Since the PMAC controller is trying to close the position and velocity loops to the encoder counts that result from the quadrature decoding of the interpolated square waves, and since it is attempting to achieve a velocity profile calculated to single encoder counts, any periodic error in the width of the square waves will translate directly into a periodic error in the commanded velocity profile; with the PID loop constantly trying to correct for these errors, the observable oscillation is triggered.

5.1 Encoder hardware considerations

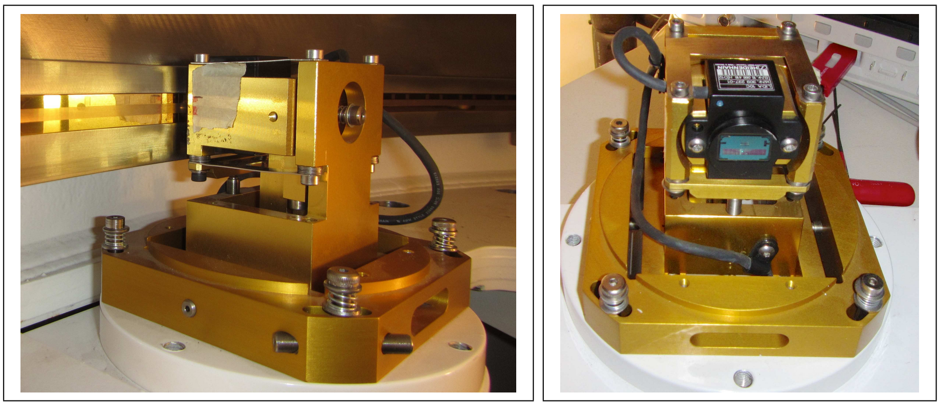

Figure 5. Heidenhain LIDA 10C read head on its custom EOST mounting assembly

5.1.1 Heidenhain LIDA 10C optical encoder read heads

The Heidenhain LIDA 10C read heads are precision optical instruments that view about 600 encoder tape lines at a time, generating two output channels in quadrature as it traverses the tape. The optical averaging of 600 lines makes for very uniform sine wave outputs that are insensitive to localized dust, dirt and scratches, and even the seam between the butted ends of the linear encoder tape. Both the azimuth and elevation axes use a pair of encoder read heads mounted 180 degrees apart to average out any signal variances that could occur due to imperfections in the machining of the mounting surface for the encoder tape.

The read heads are sensitive to accumulated buildup of dust, reducing the contrast of the tape. This in turn reduces the amplitude of the sine wave outputs, a degradation that is reasonably well tolerated by the system, but the decrease in amplitude in the index channel causes the output of the index pulse to not cross zero volts so that index marks may not be properly detected. The dust buildup is not a problem for the azimuth tape which is mounted vertically within the azimuth bearing. The dust buildup on the elevation tape is a problem for the upward-facing section of the tape, which changes depending on the current elevation of the telescope. This implies that dust buildup will be heaviest on the portion of the tape facing upwards while the telescope is stowed pointing at the zenith. While it is possible to clean the encoder tape, doing so did not affect the degradation.

5.1.2 EOST encoder read head mounting assembly

The EOST encoder mounting assembly (see figure 5) is designed to place the read head behind a bulkhead and facilitate precise adjustments to the alignment of the read head from the bottom of the mechanism outside the plate. Three flaws in the encoder mounts had to be addressed in order to successfully align the read heads. The first is that the vertical positioning springs return to center, and the adjustment screw can only push the head up, away from the center; to work around this, the spring plates were bent so that the resting position was below center, allowing a full range of adjustment through the center. The second flaw is that the yaw position at the head is held by friction, and will move when adjusting pitch and roll. The yaw adjustment is made by rotating the whole base of the mount in its roughly six inch mounting hole; this adjustment is made with opposing set screws that push on one of the four mounting bolts. Making yaw adjustments slides the base around, in the process changing the distance between the encoder tape and the read head, which then requires readjusting the distance between the read head and the encoder tape. There was no direct fix for the second flaw, thus requiring another check of the yaw after adjusting the pitch or the roll. The third flaw had several of the read heads hard against the lower spring plate, preventing both pitch and roll adjustments; in one case, this required the installation of washers to create a gap for the read head, in others its was sufficient to loosen the two screws holding the read head, push the read head away from the spring, and re-tighten.

5.1.3 EXE 660 B quadrature interpolation hardware

The output of each read head is connected to the input of a corresponding Heidenhain EXE 660B (331-926-01) interpolator module, which provides a 400x interpolation of the quadrature sine waves from the read head. The interpolator module outputs differentially driven square waves in quadrature at 400x the frequency of the input sine waves, and generates a reference pulse with a width of one quadrature count.

Ideally, if the encoder read head moves at a constant velocity past the successive lines on the encoder tape, the 400 quadrature square waves generated by the interpolator for a single line pair on the encoder tape should be of constant width. When the PMAC motor controller subsequently performs quadrature decoding of those square waves, the resulting encoder “counts” should occur at a constant rate that corresponds to the actual velocity of the read head relative to the tape.

In order for the interpolator to deliver such linear performance, the mechanical alignment of the encoder read head relative to the encoder tape must first be optimized over the full length of the tape. Because the alignment will never be perfect, there will always be residual nonlinearities; the remaining nonlinearity can be addressed by the adjustment of two potentiometers inside the interpolator itself. A custom electronics breakout board (see section 5.3) was developed in order to assess the magnitude of the nonlinearity.

5.2 Adjustments using Heidenhain PWM-9 Diagnostic Kit

After explaining the observed behavior and suggested remedy, EOS agreed to provide a Heidenhain PWM-9 Diagnostic Kit to use on a temporary basis for our alignment effort. The PWM-9 connects directly to the outputs from the encoder read head, and provides a pass-through connector to use for the existing wiring to the interpolator. This allows for direct inspection and analysis of the quadrature signals from the encoder read head.

The diagnostic kit came with a user’s manual with detailed procedures for proper alignment of the encoder read heads. This was a labor intensive process, consuming multiple hours for each read head. Before making any alignment changes, all four encoder read heads had quadrature offset errors between 2 and 8 degrees; the PWM-9 manual states that the deviation must be less than 20 degrees, and that a deviation less than 10 degrees is desirable. After several iterations adjusting the read head mounts, the deviation was reduced to between 0.5 and 2 degrees. While the improved alignment did reduce the magnitude of the nonlinearity, and therefore the amplitude of the oscillations it induced, the error remained strong enough to significantly deteriorate the performance of the telescope at tracking speeds.

5.3 Measuring and removing the nonlinearity

In January of 2013 we developed a digital breakout board to be connected in series with the outputs from the interpolation box. In addition to offering test points to analyze the interpolator output, the board was intended to remain permanently installed to provide deglitching or dedithering of the interpolator output, but those glitches were successfully addressed by changing the interpolator’s internal dip switch setting 5 to ‘open’, thereby enabling what is labeled as ‘hysteresis’.

In figure 6, the scope is set up to alternately trigger on two points 180 degrees out of phase with each other, so that the different pulse widths are superimposed. Note that the sine wave crossings of the overlaid

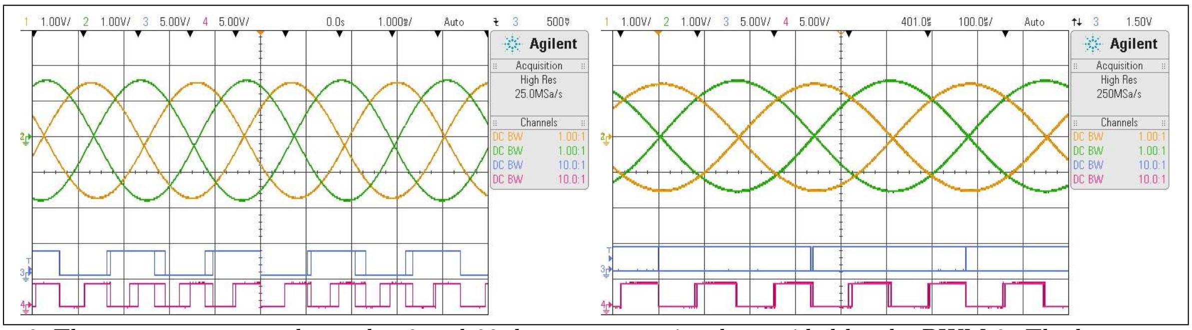

Figure 6. The upper traces are the analog 0 and 90 degree output signals provided by the PWM-9. The lower traces are the 400x quadrature output of the interpolator, divided down to give an observable number of transitions on the scope. These graphs show readings from before and after adjustments to the mechanical alignment.

quadrature signals are not necessarily occurring near zero, some occur well below zero. This is direct evidence of the nonlinearity from the encoder read head. The square waves shown are from our digital breakout board, and represent the interpolator outputs divided down by one of several fixed factors. In the case of figure 6, the square wave at the bottom of each image is the interpolated signal divided down by a factor of 100. Thus, one cycle of the encoder signals corresponds to four cycles of the divided-down interpolated output. Any misalignment of the square wave transitions is likewise direct evidence of the nonlinearity from the encoder read head.

With the encoder read heads at their best adjustment there still was a significant nonlinearity out of the interpolator. Documentation was not readily available on the function of the two potentiometers inside the interpolator; when contacted, Heidenhain declined to clarify the matter, stating that if adjustments were required, the devices should be shipped back to the factory for service. Experiments conducted by EOS at their facility in Australia indicated that it should be safe to proceed with adjustments to the potentiometers, though it was not obvious there would be any significant benefit.

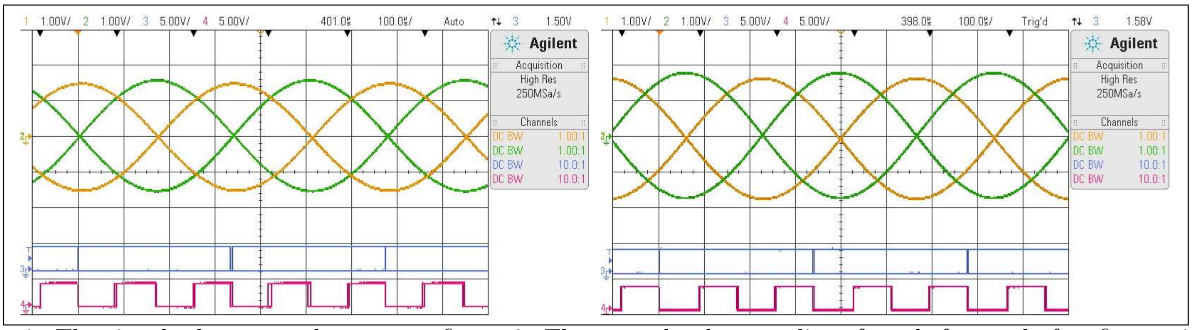

Figure 7. The signals shown are the same as figure 6. These graphs show readings from before and after fine tuning of the interpolator potentiometers.

The two potentiometers are labeled '0 degrees' and '90 degrees'; in some of the interpolators, these had been set far from the center position. Empirical testing showed that adjustments did have an impact on the measured nonlinearity. Our final procedure was to center these potentiometers, mechanically align the head for best readings on the PWM-9, and then tweak the potentiometers to remove the residual nonlinearity observed via the divided interpolator output. Adjusting these potentiometers while watching the non-linearity on the scope allowed a nearly perfect linearity out of the interpolator (see figure 7). This success was backed up by the results under servo control, with residual oscillations occurring at such low power that they did not impact the overall performance of the telescope in any significant fashion.

5.4 Alternate approaches to remove nonlinearity

The measurement of the nonlinearity was intended as the first step towards developing electronics to actively remove the nonlinearity from the quadrature outputs of the encoder interpolator.[6] This promised to be an expensive and time consuming process, and measuring the magnitude of the problem was seen as a key step towards making an informed judgment on whether to pursue such a remedy. Because the alignment of the read heads and fine-tuning of the interpolators was so successful, actively scrubbing the nonlinearity from the encoder signals was no longer pursued.

6. TUNING THE SERVO CONTROL LOOP

Our procedure for tuning the servo control loop involved three discrete phases: fundamental tuning in Delta Tau's PMAC Executive software; slew and tracking tests using EOST's AxCon software; and finally, testing on-sky in real observing conditions. As mentioned in section 3, the one consistent source of feedback in all three phases were the accelerometers; each software component used in each discrete phase required exclusive access to the Delta Tau PMAC controller in order to function.

Our objectives were to minimize the incidence of oscillations, and to achieve a tracking performance with tracking errors no worse than 0.1 arc seconds on sky. For the latter spec, that translated to an error budget of 30 encoder counts on either axis, though the on-sky magnitude of a following error in azimuth is reduced by the cosine of the elevation angle. In addition to the following error, the DAC output was the other key parameter monitored throughout these exercises. Occasional sanity checks of the DAC output with an oscilloscope were helpful in ensuring that our modifications were having the intended effect.

6.1 Tuning with Delta Tau's PMAC Executive

Delta Tau's PMAC Executive software ships with its own interactive tuning tools, allowing fairly rapid iteration between adjusting servo control parameters and tests of instantaneous response. The two tests used in our procedure were a step-response, which commands the telescope to a single discrete position at maximum acceleration, and a parabolic-response, which commands the telescope along a constantly varying velocity profile at low acceleration.

Figure 8. Step response test, showing ringing and high-frequency instability

The tuning tools also include calculators for low-pass frequency filters, which proved early on to be extremely effective in knocking down high-frequency oscillations observed on the accelerometers. The frequency selected for the low-pass filter was essentially a free parameter; it had to be selected before any further adjustments were made, and if it was changed, it proved necessary to start the tuning procedure from scratch. After probing a variety of filter frequencies between 4 Hz and 40 Hz, we settled on 9 Hz for the azimuth axis, and 12 Hz for the elevation axis.

Another free parameter was the servo update rate. The default servo update rate for a Delta Tau Turbo PMAC controller is 2.252 kHz. The telescope vendor lowered the servo update rate multiple times in an effort to reduce the incidence of catastrophic oscillations; after the last site visit, the frequency had been reduced to 98 Hz. After extended experimentation with other values, we determined that the default servo update rate, which was also the maximum for our controller, enabled the best performance for the APF telescope.

Lowering the servo update rate had three major observable impacts: it added ‘virtual’ damping to the system, it decreased the effectiveness of the proportional term, and it increased the likelihood of triggering mechanical resonances via beat frequencies. Our decision to implement a low-pass filter mitigated the need for ‘virtual’ damping, in that the low-pass filter takes responsibility for not responding to high-frequency changes, as opposed to ignoring those same changes by sampling at a lower frequency.

At lower update rates, the control loop could absorb less of an increase in the proportional term before oscillating; in order to approach nominal performance, it was necessary to increase the integral term by relatively large amounts. Doing so would allow the telescope to slew successfully, but the following error would not converge to zero during the slew. With the lower proportional term, the overall system was not stiff enough, and responded poorly to sudden impulses, such as wind gusts. These compromises were not necessary when operating with the maximum servo update rate.

The remaining parameters were the proportional, integral, and derivative terms, and the velocity feed-forward term. The friction feed-forward and acceleration feed-forward terms did not yield measurable improvements in performance, and were left at zero.

Each tuning iteration began with these four parameters set to zero. The proportional term would be increased until the telescope began ‘ringing’ around a target position when performing a step-response test (see figure 8). The derivative term would then be increased until the ringing was eliminated. At this point, the step response would generally include an overshoot, with varying amounts of success at converging on the commanded target position. The proportional term would then be pushed further, until a single ‘bounce’ from the initial overshoot began to develop. The integral term would be increased until the telescope converged quickly on the target position.

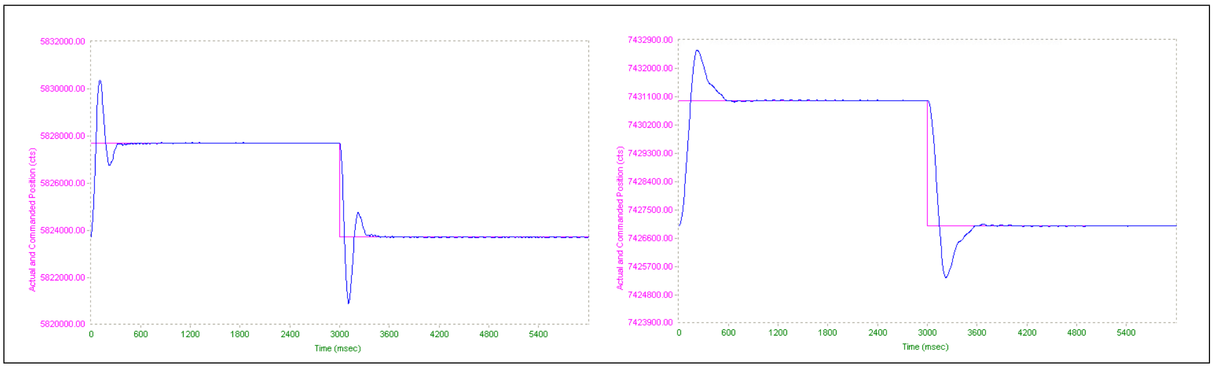

Figure 9. Final step response profiles, azimuth and elevation axes

Further iteration would then occur on the derivative and proportional terms, trying to increase the proportional term as much as possible without inducing ringing. At high enough values, the derivative term introduced its own instabilities, usually in the form of high-frequency noise (see figure 8). The derivative term is also sensitive to quantization of encoder signals and encoder noise in general; properly setting the gain on the servo amplifiers and mitigating noise on the encoder feedback signals was crucial to opening up the parameter space.

The suitability of the parameters determined via the step response tests would then be probed on a parabolic response test (see figure 10). The two parameters that saw the most adjustment in the parabolic response test were the integral and velocity feed-forward terms, with the objective of minimizing the following error as the telescope moved through its commanded velocity track. The results of the parabolic response tests were particularly interesting for our needs, as the nature of this test more closely mimics the use case that we care about: precise velocity tracking at low and constantly varying speeds.

Throughout most of the tuning effort, adjustments to the velocity feed-forward term did not have the expected effect. For simple motor systems, the velocity feed-forward term is expected to be roughly the same as the derivative term. When set too high, the velocity feed-forward term inverts the phase of the following error vs. time plot, such that the following error is out of phase with the commanded velocity profile. It was not until the very last round of tuning that the velocity feed-forward term provoked a phase change on the parabolic response

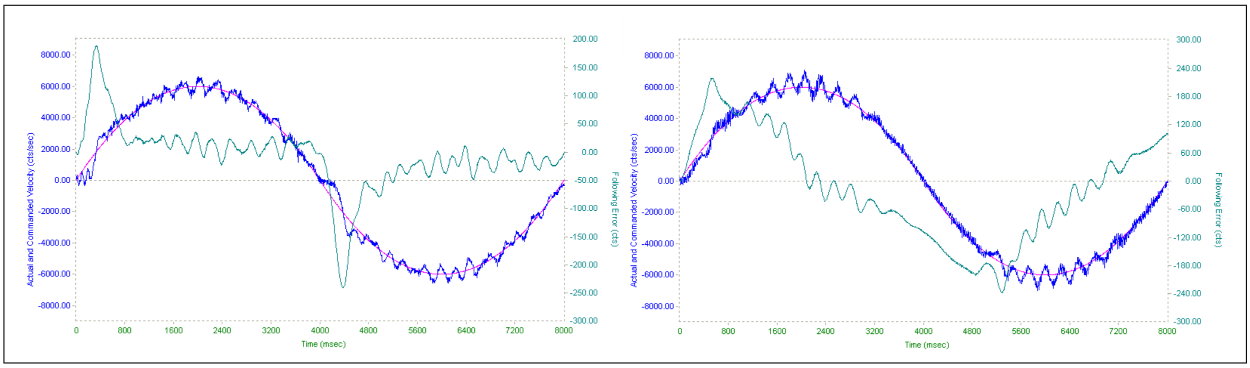

Figure 10. Final parabolic response profiles, azimuth and elevation axes

test, and indeed, the following error was minimized with a term value some 20% higher than the final derivative term, at least for the azimuth axis. The elevation axis received one less round of optimization, as its performance had already reached an acceptable level. Note in figure 10 how the following error for the azimuth axis hovers around zero for most of the move, while the following error for the elevation axis takes longer to converge to zero; this is largely due to the absence of a velocity feed-forward term for the elevation axis.

Other features in the parabolic response curves warrant explanation. The azimuth axis shows a large increase in the following error when it accelerates from a velocity of zero; this is due in part to static friction, but also due to the zero-crossing error in the servo amplifiers. There is a significant dead band around zero where the amplifier does not generate an output signal; once the DAC output is high enough to escape the dead band and overcome the static friction, it is too high for the low commanded velocity, and the telescope overshoots slightly before correcting. The elevation axis shows a similar pattern, but it is less pronounced, as the inertial load on the elevation axis is much lower. The ringing about the commanded velocity is characteristic of the elevation axis's response to the sudden 'knock' resulting from this zero-crossing error.

6.2 Tuning with EOST's AxCon

With the basic tuning complete, the next step involved tests at constant speed, both slew speeds (2-3 degrees per second) and tracking speeds, typically between 0.1 and 0.0001 degrees per second. These two modes put pressure on the control loop in different ways: the tracking performance tended to be better with a high proportional term, but with those same parameters the telescope would be driven into oscillation at slew speeds.

AxCon is a Windows application provided by EOST; it is designed to interface with EOST's on-board programs on the Delta Tau PMAC controller, allowing the telescope to be slewed through its full range of motion at its nominal speed and acceleration. AxCon also exposes the basic servo control parameters, allowing them to be changed on-the-fly while the telescope is in motion; at the same time, it provides high fidelity feedback of the DAC output and the following error as seen by the motion controller.

The first series of tests would probe the overall stability of the telescope for slew speeds, starting first at low speeds, and ramping up gradually to full slew speeds. The accelerometer displays would provide the first clue if something were amiss, well before any audible shaking would set in. If a set of parameters passed at slew speeds, testing would proceed to tracking speeds. The proportional, integral, and derivative terms would be increased in an attempt to improve the overall tracking performance; after reaching a point of diminishing returns, the adjusted parameters would be re-tested at slew speeds, and backed off as necessary to ensure stable performance.

Each axis can be individually controlled within AxCon, and while both motors can be active simultaneously, only one axis can be monitored and manipulated at a time. The ability to drive both axes simultaneously proved essential for testing in the later phases of our overall effort, as it was discovered that increases in the rotational inertia of the telescope about the azimuth axis had an amplifying effect on the mechanical resonances; this amplifying effect was stronger when the servo loop was closed on the elevation axis. After this discovery, tests of the azimuth axis were universally performed with the telescope servoing to hold position near the horizon, maximizing the impact of the mechanical resonances.

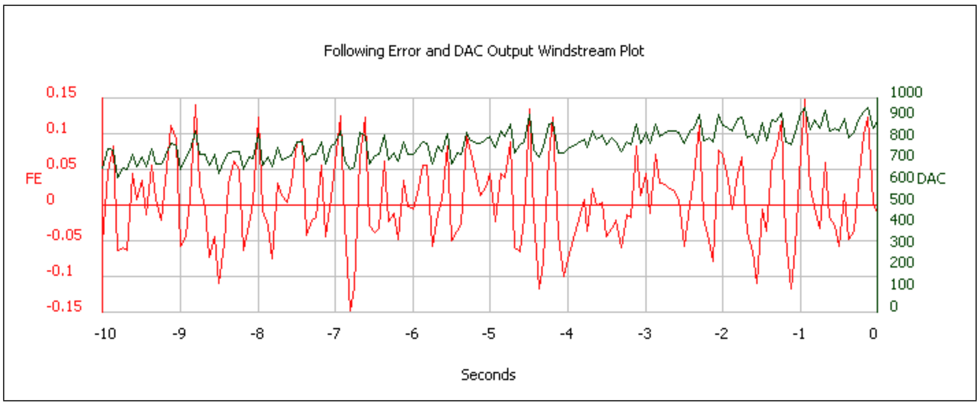

Figure 11. Azimuth tracking performance at 0.01 degrees/sec; following error (FE) is in arc seconds

While tracking at slow speeds, both axes would be tested for their stability in response to a sudden impulse. For testing purposes, this usually involved a technician giving the telescope a strong push aligned with one axis; the performance under these conditions gave an idea of how the telescope might respond to a gust of wind. When shifted from its calculated position track, the telescope should quickly and decisively return to its programmed track; in terms of the following error, any sudden increase in the following error should be immediately corrected such that the following error returns to zero. Poor tuning parameters demonstrated a variety of failures in this test: taking too long to zero out the following error; overcorrecting for the error, ringing around zero error (see figure 12); amplifying the error, making corrections in phase with the oscillation of the telescope.

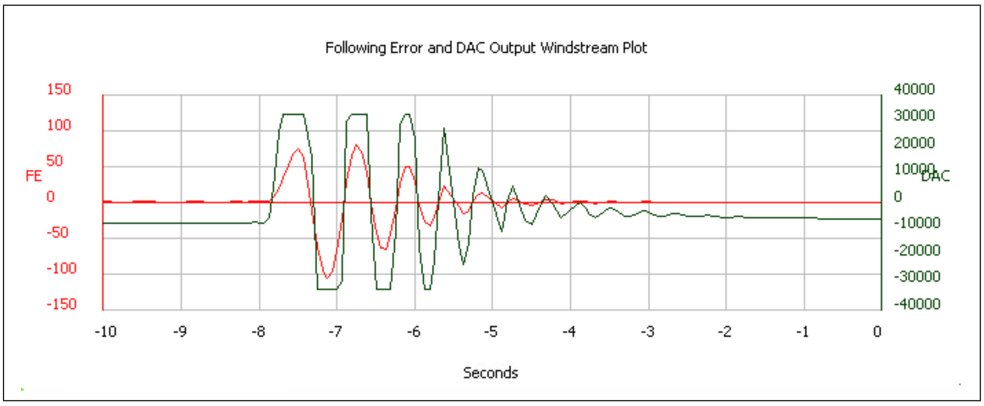

Figure 12. Poor correction of a large impulse to the elevation axis; note the DAC output clipping at +/- 32k

6.3 Validation on stellar objects

The final test for any tuning adjustments was to assess the performance of the telescope on-sky. The APF telescope operated for several months in a shared risk mode, where a science team would attempt to take data with the facility, inevitably uncovering new failure modes. In this way, on-sky testing provided valuable feedback motivating the expansion of the daytime testing procedures in addition to direct performance validation.

With the exception of the accelerometers (as mentioned in section 3), the telemetry available while observing was not sufficient to provide direct insight into the nature of any performance problems. The guide camera and the science camera could only be used to prove the absence of problems, but in that respect, they were the tools that mattered: if a low-level oscillation did not impact the quality of the science data, it did not need to be fixed.

Unfortunately, this validation process was dominated by qualitative rather than quantitative measurements. Elongation of star images, or variation in a star's apparent position between guide camera frames could have provided feedback on which telescope axis was oscillating, and to what degree, but a quantitative analysis was never established. Apparent star motion that would drive the star outside the boundaries of the selected aperture at the focal plane should have reduced the radial velocity precision for a given observation, but making that assessment requires an extended sequence of successful exposures, and is not at all practical for immediate feedback.

Open dome testing quickly made it apparent that it is very difficult to emulate the effect of wind shake while inside a closed dome. The turbulence from the recirculation fans was too regular; literal hands-on manipulation of the telescope could in only the most limited fashion mimic the chaotic persistence of the wind. On-sky testing clearly showed successive performance degradation as the telescope's exposure to gusts of wind increased, by any combination of opening vent doors, opening the dome shutter, and rotating into the wind. In normal operation, the telescope is operated with the vent doors closed and the dome shutters minimally opened in a wind shielding mode, thus minimizing the influence of the wind.

7. CONCLUSION

When our effort to remedy the telescope control system began in September 2012, the project was primed for failure. Despite the scientific potential of the facility, there was no budget to support an extended repair effort; the facility needed first to demonstrate its potential, to prove its scientific merit in order to attract future support. Such a proof was not possible while the telescope was considered a threat to its own safety, in that it was not even possible to home the telescope without triggering uncontrolled oscillations.

In spite of a diversity of political and technical obstacles, with no budget to speak of and limited staff time to dedicate to the effort, UCO personnel (both active and retired) were successful in remedying the fundamental control of the telescope, meeting and exceeding the original specifications for tracking performance. There were many lessons learned from this effort, many relating to the layers of the blindfold described in section 1.3.

Documentation is key. Institutional memory can substitute for solid documentation for a short period of time, but it cannot last; when the facility is created by outside contractors, there is no institutional memory to trivially draw from when problems arise. In both cases, but especially in the case when outside contractors are involved, having complete and accurate technical documentation is absolutely key to the ongoing maintenance of the facility.

Documentation includes source code. If the facility ships with software, and the software will be used in an operational capacity, the ability to inspect and modify the behavior of that software is absolutely essential to the long-term health of the facility. Set aside any questions about improving the software or expanding its functionality; there will inevitably be mistakes in the source code, or questions about its behavior that require inspection at the lowest level. Without the source code, bridging the knowledge gap between the documentation and the hardware becomes orders of magnitude more difficult. Without the source code, bringing the control software forward to the next generation of computing equipment becomes a practical impossibility; this last point only becomes more true as the complexity of the software increases.

High fidelity telemetry is essential. In order to diagnose behavior at the sub-arc second level, some visibility into the high-frequency behavior of the control system is required. The average performance may look very stable when sampled at 1 Hz, where harmonic structure may be visible when sampled at 10 Hz, and the full picture only becomes apparent at 100 or 1000 Hz sampling rates.

Extended retention of telemetry is essential. Low fidelity telemetry has enormous value when retained indefinitely.[5] Even if the full picture is not apparent, the signature of a problem may be apparent, and can guide further troubleshooting. Some trends may only be apparent in an extended data set, and are not practical to observe with a high fidelity data stream. Being able to mine the stored telemetry is invaluable when answering questions that were not asked at the time the problem occurred.

Divide and conquer. A problem that can be solved in isolation is one less free variable to consider when approaching broader issues. In many cases, eliminating isolated problems clarified the boundaries if not the causes themselves of the problems that remained.

REFERENCES

[1] Vogt, S. S. et al., "APF - The Lick Observatory Automated Planet Finder," *PASP* **126**(938), pp. 359–379 (2014).

[2] Radovan, M. V. and Cabak, G. F., "Buckling failure of the Automated Planet Finder spectrometer determinate spaceframe," *Proc. SPIE* **8446**, 84468A–84468A–12 (2012).

[3] Conrad, A. R. and Lupton, W. F., "The Keck Keyword Layer," in [*Astronomical Data Analysis Software and Systems II*], R.J. Hanisch, R.J.V. Brissenden, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **52**, 203–207 (Jan. 1993).

[4] Lupton, W. F. and Conrad, A. R., "The Keck Task Library (KTL)," in [*Astronomical Data Analysis Software and Systems II*], R.J. Hanisch, R.J.V. Brissenden, & J. Barnes, ed., *Astronomical Society of the Pacific Conference Series* **52**, 315–320 (Jan. 1993).

[5] Lanclos, K. and Deich, W. T. S., "A complete history of everything," *Software and Cyberinfrastructure for Astronomy* **8451**(0Y), Proc. SPIE (2012).

[6] Warner, M. et al., "Adaptive periodic error correction for Heidenhain tape encoders," *Ground-based and Airborne Telescopes II* **7012**, 70123N–70123N–8, Proc. SPIE (2008).

Kyle Lanclos is a member of the Scientific Programming Group at University of California Observatories. He received his bachelor's in computer science with a minor in astrophysics from the University of California at Santa Cruz in June 2000, and began working with UCO in December 2000.