APF Readiness Review

S. Vogt, M. Radovan, D. Cowley, E. Rivera, J. Burt, P. Butler, G. Laughlin,

K. Chloros, M. Peck, K. Lanclos, W. Deich, S. Allen, R. Kibrick

Table of Contents

1 Science Objectives

Perhaps no field in science – let alone in astronomy – is progressing as fast as the detection and characterization of extrasolar planets. Milestones such as the discovery of the first extrasolar planet itself (HD 114617b; Latham et al. 1989), the first extrasolar multiple-planet system (Upsilon Andromedae b,c, & d; Butler et al. 1999), the first transiting extrasolar planet (HD 209458b; Charbonneau et al. 2000, Henry et al. 2000), the first "super-Earth" (Gliese 876d; Rivera et al. 2005), and now, remarkably, the first genuinely Earth-mass planet (Alpha Centauri Bb; Dumusque et al. 2012) have all been accomplished via the precision Doppler velocity technique. The radial velocity planet surveys, which were in large part pioneered at Lick Observatory (e.g. Butler et al. 1994), have been so successful that planetary discoveries sufficient to grace the front page of the New York Times are, in a matter of a few years, routinely relegated to the status of barely meriting publication. It should be stressed at the outset, that within the span of a few short years, a scientific justification for APF that was cutting-edge and forward looking – the discovery of the first "rocky" planets orbiting low mass stars, now looks somewhat dated and provincial. The mantle has shifted from the mere discovery of planetary systems, "stamp collecting" to attaining a deep, profound and lasting understanding of the galactic planetary census.

A serious, vital question, then, is: Does APF have cutting-edge relevance given the highly advanced state of the exoplanet field? Is it worth investing the scarce resources that are still required to bring it to its full potential and to maintain it as a working instrument?

In 2008, the Geneva Planet Search Team began to make a very startling claim in their conference presentations. They asserted that the aggregate of Doppler velocity data from their survey with the La Silla-based HARPS pipeline indicates that of order 50% (or more) of the chromospherically quiet stars in the Solar neighborhood are accompanied by planets with masses $M\sin i < 30\ M_{\rm Earth}$ and orbital periods P < 100 days – worlds which have no analog in our own solar system, and which are generally called "close-in super-Earths".

The extremely large occurrence rate of super-Earths was not only unpredicted, it ran exactly counter to the expectations of theorists. If one looks at papers from the 2004-2006 pre-HARPS time frame, one finds numerous population synthesis models that could recover the then-known distributions of hot Jupiters (planets with $M \sim M_{Jup}$ and $P \sim 3$ days) and eccentric giants (planets with $M \sim$ *few* M_{Jup} and $P \sim 100$ to 1000 days). These same models, however, specifically predicted that the region occupied by close-in super-Earths would constitute a nearly depopulated "planet desert". It is therefore not surprising that the initial claims from the Geneva Team were greeted with both substantial and well-justified skepticism.

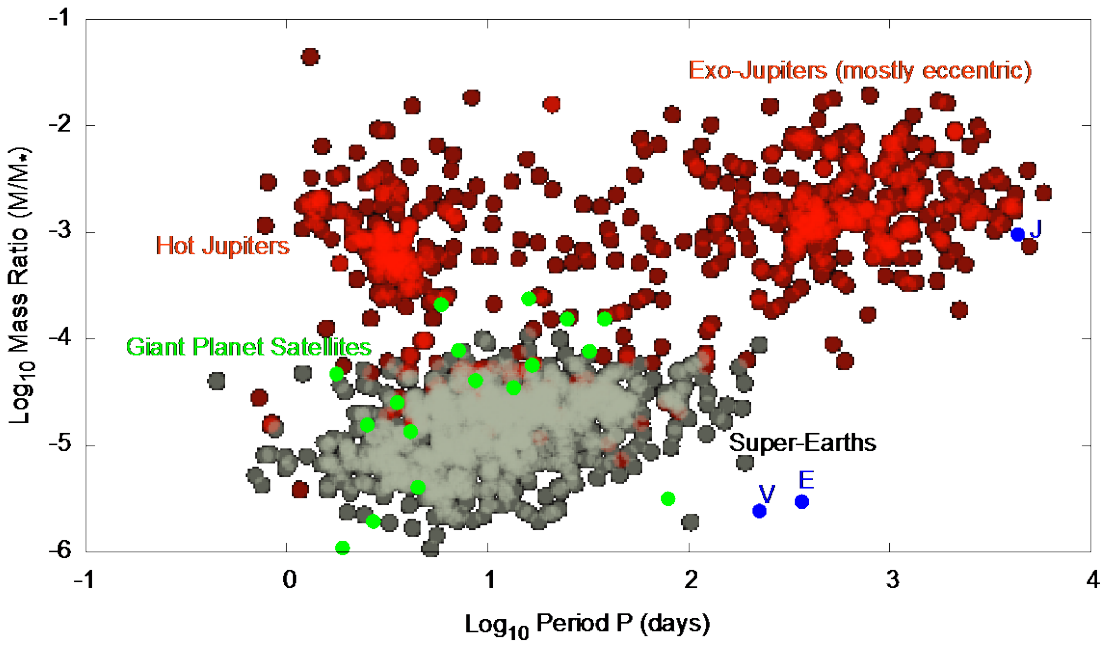

***Figure 1.1**: The current "bird's eye view" of published extrasolar planets (taken from Chiang & Laughlin 2012) in which the planet-to-star mass ratio is plotted against period. The take-away message from this diagram is that the close-in super-Earths (for which the Kepler-detected population is shown in gray) are at least 50 times more common than the well-characterized hot Jupiters, and at least 10 times more common than the eccentric exo-Jupiter population discovered and characterized by legacy Doppler surveys.*

The results of the Kepler Mission, however, have made it clear that an enormous population of close-in super-Earths does indeed exist. In the most recent catalog of Kepler planetary candidates ("KOIs") published by Batalha et al (2012), it is found that ~1,800 out of ~156,000 individual stars harbor ~2,300 transiting planet candidates – more than 80% of which have radii $R < 5R_{Earth}$. A straightforward back-of-the-envelope comparison indicates that the occurrence rate implied by the Kepler data is in broad agreement with the results from the Geneva survey. The profound, inevitable conclusion is that our Solar System – which contains no planet interior to Mercury's P=88-day orbit – did not participate in a major, if not the dominant mode of planet formation in the Galaxy. The Solar System's orbital architecture is not the norm, and a new template for discussions of planet formation appears to be needed.

The construction of a new template for the understanding of planets and the planet formation process in the galaxy hinges on two complementary efforts. The first is the statistical characterization of the planetary census. For more than a decade following the discovery of 51 Peg, the Doppler velocity surveys – with the surveys at Keck leading the way – provided statistical descriptions of the population of hot Jupiters and eccentric extrasolar giant planets.

For lower mass planets, i.e. the local galactic populations of super-Earths and Earth-mass worlds, this role will be a joint effort between space-based surveys such as Kepler, TESS, and WFIRST, and high-precision ground-based radial velocity work, such as that made possible by APF, HARPS, and HARPS-North, which will confirm orbits and measure planetary masses.

The second broad task is the detailed study of representative planets, particularly transiting planets, and multiple-planet systems orbiting the brightest nearby stars. APF, which is effectively a purpose-built instrument for obtaining radial velocities with high cadence, high throughput, and high availability, is ideally suited.

In brief, the *urgent, state-of-the-art scientific justification for APF is to discover and characterize the masses and the orbits of the nearest, most readily studied super-Earths, and to provide critical ground-based support to all-sky photometric missions such as NASA's proposed TESS mission.*

In the service of these broad roles, APF telescope will carry out a range of specific tasks. As this document demonstrates, the APF, with its fast optical system, has an end-to-end throughput that is approximately $1/5^{th}$ that of HIRES on Keck. As of early 2012, the database of Keck Iodine-cell observations of nearby stars stretches back over a decade and a half, and includes 31,190 high-precision Doppler measurements with a median instrumental uncertainty σ=1.45 m s^{-1}. These data represent 8.9×10^6 s of on-sky measurement, which with associated overheads, sums to N ≈ 370 nights – slightly more than a full year. This allocation has formed the basis for hundreds of planetary discoveries that have had a massive impact on the field. Papers based directly (or indirectly) on iodine-cell observations from Keck have generated tens of thousands of citations, and describe the discovery and characterization of well over 100 planetary systems. *Over the coming five years, APF is poised to obtain a data set of similar breadth, depth, and equal or better quality.*

2 Telescope and Dome

The APF facility is located at the Lick Observatory on Mt. Hamilton at latitude +37:20:33.1, longitude -121:38:17.7, altitude 1274 meters. The APF Telescope was designed and built by Electro Optical Systems Technologies (EOST) at their former facility in Tucson, Az.; photos and reports associated with the fabrication of the telescope at that facility, as well as its delivery to Mt. Hamilton, can be found at http://apf.ucolick.org/telescope/. Note that some of the preliminary design documents make reference to telescope drive motors that were to be friction coupled, but post-PDR the design was changed to employ direct drive motors.

The APF Dome was designed and built by Electro Optical Space Systems (EOS) of Australia. Initial assembly of the dome at the APF site was completed in May 2007. A number of fabrication and construction problems were identified as described in the document Lick APF Site Inspection Report (6 June 2007).

On-site assembly of the APF telescope was completed in May 2009. Site acceptance testing for the telescope was initiated in June 2009. While the telescope passed many of the acceptance tests at the time, acceptance tests pertaining to tracking performance were not completed until November 2010.

During the intervening two years, we have subsequently discovered a number of problems with both the telescope and dome, and these will be covered in detail in the following sections. Several of these problems are the subject of warranty deficiency claims that were filed respectively with EOST and EOS while the telescope and dome were still covered under the manufacturer's warranty. Under the terms of the contract, EOST and EOS are legally required to correct those deficiencies. Unfortunately, they have so far been either unwilling or unable to do so, and UCO appears to lack any effective means to compel them to meet their warranty obligations.

2.1 Telescope

The APF telescope has an alt/az mount and a 2.4-meter primary mirror. It has a rotatable tertiary that can direct light to either of the two Nasmyth foci. It does not have either a prime or Cassegrain focus.

2.1.1 Components

2.1.1.1 Mirrors

The APF primary mirror is a 2.4 M x 150 mm thick parabolic meniscus with a 7200 mm radius of curvature and is made from LZOS Astrosital. The blank was generated by LZOS in Russia and was polished by Rayleigh Optical Corp. (ROC) in Baltimore Maryland on September 7, 2007. The contract specifications for the mirror are governed by the "Primary Mirror Polishing Specification TS-4892-2". ROC performed a null test using a computer generated hologram (CHG) to demonstrate compliance with the contract specification the results of which are reported in the "Lick 2.4-m Primary Mirror Final Test Report" and in the "Lick 2.4 Technical Report – Lick APF Primary Mirror Test Results (TR-13108-1)". The mirror is finished in bare aluminum that was applied at Mt. Hamilton in the Shane 3 meter vacuum tank.

The APF secondary mirror is a 354 mm diameter convex hyperbola with a -1198.56 radius of curvature made from Schott Zerodur. The original secondary mirror was fabricated by Sesso to the Secondary Mirror Technical Specification and was completed in November of 2005 as indicated in the "Lick M2 Sesso final inspection" report. The SESO results were independently confirmed at UCO optical shop on the profilometer to ensure errors were not made during fabrication and testing at Sesso. Although the UCO and the Sesso results agreed, we later discovered that the assembled telescope could not meet the contract specified image quality at the proper distance from the instrument mounting plane. EOST attempted to pass the telescope off to UCO by re-spacing the optical system and did eventually meet the image quality specification, but at a position that was almost 200 mm different from the contract value. This resulted in 4-6 months of wasted effort.

When UCO became aware of the shifted focal plane Steve Vogt worked with Andrew Lowman of EOST to develop an optical prescription for a new M2 mirror to correct the problem. Because we had two independent measurements of the original M2 we were able to determine that the measured conic constant of the primary mirror was incorrect. The new secondary mirror was fabricated to "Technical Specification TS -4893-3" at the UCO optical shops and was completed in under four months, a full eight months faster than SESO was quoting for a replacement mirror. For factory testing at EOST, the UCO secondary mirror was originally coated with protected silver by an EOST subcontractor and was later re-coated in the UCO coating lab at UCSC before the telescope was installed at Mt. Hamilton.

The tertiary mirror was manufactured by Zygo and has a clear aperture of 390 mm x 275 mm on a ULE substrate as defined by the "The Tertiary Mirror Technical Specification TS-5697-2". The final acceptance tests are reported in the "Lick Tertiary Zygo Cert." Like the M2 mirror

the protective silver coating was redone at UCO prior to installation in the telescope at the Mt. Hamilton site.

2.1.1.2 Mount

The primary mirror support is a 27-point passive whiffle with a 304 stainless steel base weldment. The entire assembly was designed to be vacuum compatible so the mirror and its mount could go into the coating tank to simplify the re-coating procedure. Unfortunately, the support structure for the whiffletree became too deep to fit in the Shane coating tank, so a special lifting fixture and handling cart was provided by EOST in exchange for UCO removing the requirement that the mirror and support fit in the coating tank. The 27 flexures of the whiffletree and the lateral support assembly are attached to Invar 36 pucks bonded to the mirror back and edges. The predicted performance of M1 mirror support is summarized in the document "M1 Mirror Support Finite Element Analysis (FEA-10526-)". Additional details of the joint design and the mount details are presented in the "Passive Primary Mirror assembly-Mirror Removal/installation Procedure (MPI- 12923-1)". This document also serves as a step-by-step procedure to remove and install the mirror from the support when aluminizing.

The secondary mirror is held in place by three Invar 36 pucks which are bonded into pockets in the mirror back. The base of each puck lies on a plane that passes through the mirror center of gravity and includes a flexure which attaches to the puck via a threaded joint. The flexures provide thermal and mechanical isolation between the glass and the 1018 steel base plate. The plate has three v-grooves arranged at 120 degrees to serve as a kinematic interface between the mirror mount and the actuator assemblies (ASY-5187).

The M2 mirror is moved by three planetary ball screw actuators, each of which is driven by a harmonic drive and DC servo motor assembly (ASY-5171). The three actuators provide piston motion for telescope focus and tip /tilt correction to compensate for gravity induced alignment errors between M1 and M2. Absolute position feedback to the servo system is provided by three Heidenhain MT2501 linear optical encoders.

The M3 mirror is attached to a 45-degree frame using the same joint detail implemented for the M2 mirror assembly (ASY-5818). Like M2, the M3 assembly has a steel baseplate that provides a kinematic connection between the mirror assembly and the telescope. Unfortunately, this design does not guarantee the M2 and M3 mirror will not need to be realigned when removed from the base-plate for re-coating. Re-coating of M2 an M3 will require tooling that does not exist and therefore needs to be built to establish the optical axis of the telescope with a Davidson D-275 alignment telescope before removing either mirror.

2.1.1.3 Bearings

Both axes of the telescope use rolling contact bearings for high reliability. These require only semiannual service and should exceed the 20 year required life expectancy of the telescope.

2.1.1.4 Servo Systems (azimuth and elevation axes only)

The servo systems for the APF Telescope are digital and are implemented via a COTS programmable multi-axis motion controller (Delta Tau PMAC) board embedded in a Windows/XP-based PC. Based solely on position feedback provided by on-axis incremental tape encoders (i.e., there is no tachometer feedback), the firmware in the controller board closes the low-level position and velocity loops; high-level telescope pointing calculations are performed by the EOST-supplied software that runs on the PC.

For the azimuth and elevation axes, the controller board outputs a control voltage to an external COTS servo amplifier for the given axis, which in turn closes the current loop to the on-axis, direct drive motor for that axis.

The servo systems for the telescope performed adequately and the servo systems appeared stable (at least within the capabilities we had at that time to verify their performance) when the telescope was installed at the APF site and became operational in mid-2009. As indicated above, the telescope performance met most contractual specifications during the site acceptance tests conducted in conjunction with a representative from USNO in June 2009, although some tracking performance specifications were not met until December 2010.

The APF Spectrometer was first mated to the telescope in July 2011, and later that year we started to experience problems with the telescope servo drives for the azimuth and elevation axes. Efforts by EOST and EOS staff to solve these problems yielded some improvement but have not fully resolved them; this is also currently the case for more recent efforts by UCO/Lick Observatory staff.

The historical context of these problems, as well as the various efforts (up to mid-September 2012) to address them are described in the memo "Historical summary of APF Telescope servo performance issues". Shortly after that memo was written, UCO staff member Kyle Lanclos attended the Turbo PMAC servo training class conducted by the manufacturer (Delta Tau Data Systems Inc.) of the PMAC servo controller board that controls the APF Telescope axes; after he completed that course, UCO began its own efforts to improve the servo tuning of the elevations and azimuth axes of the telescope (see subsections 2.1.2 through 2.1.4 below).

The other hardware components that comprise the APF Telescope servo systems are described in the following subsections.

2.1.1.4.1. Encoders

Both the azimuth and elevation axes of the telescope utilize Heidenhain incremental tape encoders. According to Heidenhain technical support, the encoder tape (labeled as ERA 7x0C, Id. Number 325 691-13; see Figure 1) appears to be based on their LIDA 181 tape as described on page 26 of their 2003 Catalog for exposed linear encoders. This tape has an accuracy grade of 2.5 microns and a grating period of 40 microns, yielding 25 lines/mm; there are 133000 such lines (periods) over this 5320 mm long tape.

Figure 2.1: *Heidenhain tape encoder on azimuth axis of telescope.*

Each tape is read by two Heidenhain LIDA 10 C (a.k.a. Id. Number 309 237-01) read heads spaced 180 degrees apart. We do not have any spares for these read heads; while they no longer appear in the Heidenhain catalog, it is still possible to purchase them from a Heidenhain distributor in the U.S. at an approximate cost of approximately $1140 with a lead

time of 4 to 6 weeks. Each read head is mounted to an EOST-designed fixture which enables the read head to be aligned with respect to the encoder tape.

See http://www.ucolick.org/~kibrick/apf/encoder_mount/ for photos of the read head and the alignment fixture.

At the end of July 2012, during his service trip to the APF site, Andrew "Jak" Gray of EOS (after first determining that the read heads were out of alignment) realigned the read heads for both of the encoders for each axis (azimuth and elevation) using Heidenhain's encoder alignment device and the procedure documented by Smith and Ritchie in December 2011. Nominal alignment to Heidenhain specifications was achieved for all 4 read heads over a limited range (about 5 degrees) of azimuth angle and over a somewhat larger range of elevation angle. Ideally, this alignment would have been optimized over a full rotation of the azimuth axis, but during much of Gray's visit, a failure of one of the Baldor Vector drives prevented rotation of the dome and thus severely limited the range of azimuth motion of the telescope.

Each read head outputs a pair of sine waves in quadrature with a nominal amplitude of 1 volt PTP and a nominal signal level of 11 uA into a 1 Kohm termination. Each read head also generates a differential reference signal that indicates when one of the distance-coded reference marks on the encoder tape passes by the read head; those reference marks are used only during the homing sequence to initialize the absolute position of each axis.

The output of each read head is connected to the input of a corresponding Heidenhain EXE 660B (a.k.a. Id. Number 331-926-01) interpolator module, which provides a 400x interpolation of the quadrature sine waves from the read head. The interpolator module outputs differential TTL square waves in quadrature at 400x the frequency of the input sine waves, and generates a reference pulse with a width of 90 degrees of the square wave period.

The differential square waves and the reference pulse generated by each interpolator module are carried via Heidenhain encoder cables to the telescope control rack on the ground floor of the APF Dome. The cable for each encoder consists of two segments: a segment that goes from the output of the interpolator to a bulkhead connector on the side of the telescope and a segment from that bulkhead to the connectors on the back panel of the telescope control rack.

On the inside of that rack, an additional cable segment routes the signals from each encoder's interpolator module to the input of an EOST-built encoder planar board (ENCP1), which then routes those signals to an output ribbon cable that carries them into a corresponding encoder channel on the PMAC board embedded in the telescope control computer.

As noted above, the azimuth and elevation axes each have two encoders (labeled "A" and "B") mounted 180 degrees apart. In the case of the azimuth "A" encoder, once the output

from its interpolator module enters the back panel of the telescope control rack, it connects to an EOST-built encoder splitter module, which splits the signals into two buffered copies.

One copy continues on to the ENCP1 encoder planar board and from there into the PMAC board in the telescope control computer. The other copy is sent via a separate Heidenhain cable into the Dome Control Cabinet, where it is conditioned by a BEI anti-dither module and then fed into the Baldor Vector drives that control the dome azimuth drive, so that the co-rotating dome follows the telescope when it rotates in azimuth (see subsection 2.2.1.4.1.3.).

2.1.1.4.2. Servo motor controller board

The programmable multi-axis controller (PMAC) board is a Turbo PMAC PC ISA Legacy board (model 79754 400-602204-trx with options 156767 301-00 PMAC-Opt 1 and 156968 302-00 PMAC-DPRAM) manufactured by Delta Tau Data Systems Inc. This board is no longer in production. At present, Delta Tau is still able to repair such boards, but how much longer they will be able to do so remains unclear. Currently, we have no spare for this board.

The relevant reference manuals for this board are:

PMAC PC Hardware Reference Manual / Programmable Multi-Axis Control

Turbo PMAC User Manual / Programmable Multi Axis Controller

Turbo PMAC/PMAC2 Software Reference Manual / Software Reference for Turbo Family

For each motor axis to be controlled, the PMAC board produces a +/- 10 volt differential output that serves as input to the servo amplifier (see subsection 2.1.1.4.5).

2.1.1.4.3. Telescope control computer

The telescope control computer (TCC) is a rack-mounted industrial PC that employs a single board computer (Model IAC-F846A) and IAC backplane (Model AP-PSL14) manufactured by Vox Technologies. This hardware is extremely old and will become increasingly difficult to obtain if replacements are needed; only one manufacturer still makes computers with ISA slots. We have no spares for this hardware. This computer runs the Windows/XP operating system and the various EOST-developed applications, including the TelescopeServer, which implements the pointing kernel and related functions.

2.1.1.4.4. GPS Time Standard

The time standard is a Model XL-AK GPS Time and Frequency Receiver manufactured by TrueTime (which is now Symmetricon). An RS-232 serial interface on this receiver connects to a serial port on the TCC, enabling the TCC to read the current GPS time from the receiver when the TelescopeServer is started. The receiver also generates a 10 MHz square wave that is phase-locked to the GPS time. This signal is buffered via the ENCP0 encoder planar board, where it is

divided by 10 to produce a 1 MHz clock that is connected to encoder channel 8 on the PMAC board, thereby providing a GPS-based clock reference within the PMAC board. The PMAC still uses its own internal crystal clock oscillator as its primary time base for all of the PMAC motion control operations, but it uses the GPS-based clock reference to "true up" its internal clock at periodic intervals. The manual for this receiver can be found here. We have no spare for this receiver.

2.1.1.4.5. Servo amplifiers

The servo amplifiers for the azimuth and elevation axes of the telescope are Model 50A20 analog servo drives manufactured by Advanced Motion Controls. These boards are still available but are reserved for existing applications only and technical support is limited. Although they are no longer in production, they can still be purchased for approximately $520, but only 30 EOS amplifiers remain at the factory. At present, we have no spares.

These amplifiers are configured to operate in current loop mode. The +/- 10 volt differential output from the corresponding channel of the PMAC board (see subsection 2.1.1.4.2) serves as the input to each servo amplifier. The output from each amplifier powers the drive motor for the given axis of motion. Note that these servo amps exhibit a noticeable zero-crossing error when the output current changes sign. Also, the azimuth servo amplifier exhibits a significant deadband and current offset when the output current changes sign.

2.1.1.4.6. Drive motors

Both the azimuth and elevation axes of the telescope employ brushed DC direct drive motors manufactured by Kollmorgen. The azimuth axis uses a model QT-23515 drive motor and the elevation axis a model QT-23514. The datasheets for these motors can be found here. These motors appear to be out of production. However, since they are built directly into each bearing for each axis, they would be extremely difficult to replace in any case. For more details, see EOST report DN-6581-1 Direct Drive Preliminary Design Note.

2.1.2 Servo System Tuning

Tuning the servo systems for APF's azimuth and elevation axes involves two steps. The first step is analog tuning of the servo amplifiers (by adjusting the values of hardware components) and the second is digital tuning of the PMAC board (by adjusting PID and related parameters) as defined in the PMAC firmware file ALLFILES.PMC. In addition, before commencing with such tuning, it is important that the telescope encoder read heads are properly aligned with respect to the encoder tape for each axis and that the telescope tube is adequately balanced.

2.1.2.1 Historical context

As confirmed by former EOST employee Kevin Harris, EOST never carried out a full servo analysis of the APF Telescope drive systems nor did they develop a transfer function for the overall system. Rather, when tuning the telescope servo systems, they employed an ad hoc, iterative procedure.

Servo tuning was first performed when the telescope was assembled in EOST's facility in Tucson. The initial versions of the PMAC servo tuning firmware file, ALLFILES.PMC, were implemented by Adrian R. Loeff (EOST's principal architect for their telescope pointing kernel) and Kevin R. Harris (one of EOST's lead engineers) between October 2005 and November 2008. Unfortunately, the only record we have off these previous iterations of the PMAC servo tuning (and some of the analog tuning of the servo amplifiers) are the revision history comments found at the start of the ALLFILES.PMC file.

As can be seen from those comments, version 2.1 of the ALLFILES.PMC file, implemented by Kevin Harris (KRH) on Nov. 7, 2008, was the last version of the PMAC servo tuning that was performed before the APF Telescope was delivered to Mt. Hamilton in April 2009. Around that time, Harris, Loeff, and Bruce Cook, who were a key part of the senior technical staff at EOST, left the company.

All of the PMAC servo tuning that was performed by EOST after the telescope was in operation on Mt. Hamilton (versions 2.2 through 2.4) was carried out by Elwood C. Downey (ECD), who appears to have joined EOST sometime during 2009; Downey was not part of the original EOST team that designed and built the APF Telescope, but he had worked on the commissioning of the EOST-built telescope at Magdalena Ridge Observatory (MRO).

Except for the final version (2.4) of the PMAC servo tuning that Downey generated in January 2012, his other versions (2.2 and 2.3) were produced without UCO staff able to observe what he was doing. In fact, EOST management (Stubbe) demanded that our staff not be present, arguing that this would be a distraction and would interfere with Downey's ability to

work efficiently. We were also instructed by EOST not to attempt any servo tuning ourselves, since we did not have the tools, documentation, or training needed to do so. We were also concerned as to whether our making adjustments to the servo tuning ourselves would jeopardize UCO's warranty rights under the terms of the contract with EOST.

The earliest version of the ALLFILES.PMC for which we have a copy is Version 2.2, which was generated by Downey in August 2009 at the APF site. Unfortunately, while the revision history comments for that version indicate that Downey changed the tuning for the azimuth and elevation axes, they provide no details of which parameters were changed, nor how, nor why.

The final version of the ALLFILES.PMC file generated by EOST staff was version 2.4, which was generated by Downey and Mark Croft during their service visit to the APF site on January 11-12, 2012; UCO staff Vogt, Radovan, and Kibrick was permitted to observe this servo tuning effort. This was the first time we were given access to EOST's servo tuning procedures and tools or given explicit authorization by EOST staff to make our own adjustments to the servo tuning.

The January 2012 service visit by Downey and Croft was in response to a formal warranty deficiency claim (TELE-002) that UCO filed with EOST in December 2011 as a result of serious telescope oscillations that were first observed at the beginning of that month. An MPEG-format video generated at that time using images from the guide camera clearly illustrates the extent of these oscillations. Each frame of this video is from a 1-second integration on the guide camera. To provide a sense of scale, a 2” x 2” box has been overlaid onto these images.

The final servo tuning performed by Downey and Croft during their January 2012 service visit reduced the severity of these telescope oscillations but did not eliminate them. During that visit, they informed us that EOST was in the process of shutting down their facility in Tucson and that both of them were about to be laid off.

In the months that followed, we experienced additional telescope oscillations of varying amplitude. Some of these produced no audible symptoms but resulted in star images (on the guide camera) that were elongated along the position angle corresponding to either the azimuth or elevation axes. Other oscillations resulted in low to moderate buzzing of the respective drive motors during tracking and/or slewing, and a few oscillations resulted in violent oscillations that forced us to rush for the E-stop button.

In the meantime, the remaining management at EOST indicated that our warranty deficiency claim was still pending and that they would try to do something about it, but this discussion dragged on for the next several months. The end result was that in July, responsibility for addressing these warranty issues was transferred from EOST to EOS, and that led to Jak Gray's service visit to the APF site at the end of July.

As noted earlier, during that site visit, Gray re-aligned the read head assemblies for the Heidenhain tape encoders for the azimuth and elevation axes. Unfortunately, as noted above, the re-alignment of the read head assemblies for the azimuth axis was performed only over a very limited range of azimuth rotation. Gray also adjusted the analog tuning of the servo amplifiers as noted in subsection 2.1.2.3 below, but he did not make any adjustments to the PMAC servo tuning due to lack of time.

Following Gray's service visit, after several exchanges of emails with EOS, it became clear that they had no clear plan for how to resolve our servo tuning issues, despite our still pending warranty deficiency claims. They also stated that Gray would be on leave for 6 weeks starting October 1.

Accordingly, UCO decided that it could not wait any longer for EOST or EOS to address our outstanding warranty claims, and that we would have to assume responsibility for solving these servo tuning problems ourselves. As a result, we are now in the process of trying to obtain the necessary training for our staff and the diagnostic hardware and software needed for us to carry out this task. Since the APF budget never included funding for this activity, we are financially constrained as to how quickly we can move forward given our current lack of resources.

2.1.2.2 Balancing the telescope tube

During their service visit in January 2012, Downey and Croft checked the balance of the APF Telescope tube by confirming that with the servo drives powered off and the tube manually positioned to various elevation angles, the tube would remain still when they let go. Based on that test, they concluded that the balance of the tube was adequate.

With the transfer of warranty responsibility from EOST to EOS in mid-2012 came new advice regarding the importance of telescope tube balance to the successful operation of the servo drives for the APF Telescope. According to Ian Ritchie, Operations Manager for EOS, the servo performance of EOST-built telescopes is significantly compromised if the tube is inadequately balanced. Both Ritchie and Gray recommended that we check the telescope balance. Gray described a procedure for doing this by plotting the drive current versus the elevation angle of the telescope while moving back and forth between the zenith and the horizon. For a properly balanced telescope tube, the result should be a horizontal line of data points.

Using that test, in August, we determined that there was still a residual imbalance which could be reduced to negligible levels by first adding new weights to the existing weights at the top end of the telescope tube and then re-distributing the weights between the front and back of the top end. First, 5 pounds of weight were added symmetrically to the top end of the tube.

Then, a total of 11.25 pounds was redistributed from the front of the top end (i.e., the side that faces the dome shutter) to the back of the top end.

Note that this balancing operation must be done with the mirror covers open, since closing those covers makes the telescope quite bottom heavy. For the same reason, all servo tuning activities must also be carried out with the mirror cover open.

2.1.2.3 Tuning the servo amplifiers

As noted earlier, the Advanced Motion Controls model 50A20 servo amplifiers are operated in current loop mode. The datasheet for this amplifier includes a block/diagram schematic on page 2. (Note that the 100 Kohm resistor that is shown on the schematic in parallel with R28 and the 0.01 capacitor that is shown in parallel with C73 both appear to have been removed from both of our servo amps). The current loop tuning procedure involves adjustments to the values of the following 4 components:

2.1.2.3.1. Proportional gain resistor (R28)

As described in AMC's procedure, the value of this resistor is adjusted to provide a slightly under-damped response to a step input. There is a DIP switch on the servo amp (Switch 7) which can be set to temporarily short out the integral gain capacitor (C73) while the value of R28 is being adjusted.

During his service visit in late July, Gray determined that the value of R28 should be 440 Kohm for the elevation servo amp and 220 Kohm for the azimuth servo amp. Unfortunately, this change resulted in significant buzzing of the elevation drive motor during tracking. As a result, in mid-August, R28 for the elevation servo amplifier was changed to 221 Kohm. In early October, UCO personnel analyzed the response of the system with different values for R28, and arrived at a preferred value of 200 Kohm. They also rechecked the R28 value for the azimuth axis, but left it set to 220 Kohm.

2.1.2.3.2. Integral gain capacitor (C73)

The integral gain capacitor C73 is adjusted after the value of R28 has been optimized with C73 shorted out. The goal is trim out any residual ringing in the step response. During his visit in late July, Gray determined the value of C73 should be 40 nF for the azimuth servo amp and 10 nF for the elevation servo amp. In mid-August, when the value of R28 for the elevation servo amp was changed from 440 Kohm to 221 Kohm, the value of C73 for that amplifier was changed to 47 nF. At that point, the analog servo tuning for both servo amplifiers was quite similar.

Subsequent discussions between our electronics engineer, Michael Peck, and the support engineers at Advanced Motion Controls, resulted in the decision to set DIP Switch 7 to the “ON” position for both servo amplifiers so that C73 remains shorted out at all times. Note

that we also tested the use of a 1 uF capacitor for C73 (and with Switch 7 set to “OFF”) but did not find that it provided any benefit. At present, Switch 7 is set to the “ON” position on both servo amps such that C73 is shorted out.

2.1.2.3.3. Reference gain potentiometer (pot 3)

This potentiometer determines the overall gain of the servo amplifier and thus establishes the mapping between the control voltage input and the output current of the amplifier driving the motor. In adjusting this potentiometer, the goal is to ensure that they full dynamic range of the control voltage is utilized.

For each motor axis channel, the PMAC board uses a 16-bit signed DAC to generate a +/- 10 volt control voltage. The reference gain potentiometer on the servo amplifier should be adjusted so that +10 volt control voltage input from the PMAC results in close to the maximum positive output current from the servo amplifier, while a -10 volt input results in close to the maximum negative output current.

In mid-October we checked the settings of the reference gain potentiometers for both servo amplifiers. While this potentiometer was correctly adjusted on the azimuth servo amplifier, the gain for the elevation servo amplifier was found to be set a factor of 9 too high. As a result, only about 11% of the dynamic range of the PMAC DAC for that channel was being utilized, which resulted in significant quantization of the control voltage generated by the PMAC board for that channel. The reference gain for potentiometer for the elevation servo amp was then re-adjusted to reduce the gain by a factor of 9.

2.1.2.3.4. Offset potentiometer (pot 4)

This potentiometer adjusts the zero point offset of the servo amplifier. It is adjusted so that a control voltage input of zero volts yields an output current of zero amps. In mid-October, we determined that this potentiometer had been misadjusted for both amplifiers, and those misadjustments were corrected.

2.1.2.4 Tuning the PMAC board

The PMAC board has a large number of parameters that impact the servo tuning. In addition to the classical proportional, integral, and derivative terms, the PMAC provides velocity, acceleration, and friction feed forward terms, along with terms for implementing low-pass filters and notch filters.

As we will likely continue to adjust these PMAC servo tuning parameters during the time between the date on which this report will be issued (Nov. 9) and the date of the Readiness Review (Nov. 16), the parameter values described in the following subsections are subject to change, and updated values will likely be presented at the review.

2.1.2.4.1. Adjusting the DAC offset (Ixx29)

Similar to the offset potentiometer on the servo amplifier, this adjustment is used to trim out any residual zero point errors in the control voltages generated by each channel of the PMAC board, so that a commanded velocity of zero results in a control voltage output of zero volts for the given axis of motion. The DAC offset is set by PMAC parameter Ixx29, where 'xx' is replaced by the motor number; note that for the APF Telescope, PMAC motor #9 corresponds to the azimuth axis, #13 to the elevation axis.

In early October, we determined that these DAC offsets had never been set in the ALLFILES.PMC file that is used to set all of the PMAC parameters. We also determined that there were small zero point errors that did need to get trimmed out. Accordingly, we determined the following values for the DAC offsets: elevation axis (I1329) = 440, azimuth axis (I0929) = 110.

2.1.2.4.2. PID Parameters

The following PMAC servo tuning parameters for the elevation and azimuth axes of the telescope correspond to those specified in Version 2.5 of the ALLFILES.PMC file installed on October 16, 2012:

Table 1. Key PMAC servo tuning parameters for the elevation axis (Motor #13)

PMAC Register	**Description**	**Value**
I1329	DAC Offset	440
I1330	PID Proportional Gain	1117245
I1331	PID Derivative Gain	1000
I1332	PID Velocity Feed Forward Gain	1500
I1333	PID Integral Gain	120000
I1334	PID Integration Mode	0
I1335	PID Acceleration Feed Forward Gain	500
I1336	PID Notch Filter Coefficient N1	0
I1337	PID Notch Filter Coefficient N2	0
I1338	PID Notch Filter Coefficient D1 (35 Hz low pass)	-0.4413770643
I1339	PID Notch Filter Coefficient D2	0
I1363	Integration Limit	50000
I1368	Friction Feed Forward	50

Table 2. Key PMAC servo tuning Parameters for the azimuth axis (Motor #9)

PMAC Register	Description	Value
I0929	DAC Offset	110
I0930	PID Proportional Gain	500000
I0931	PID Derivative Gain	350
I0932	PID Velocity Feed Forward Gain	3800
I0933	PID Integral Gain	20000
I0934	PID Integration Mode	0
I0935	PID Acceleration Feed Forward Gain	500
I0936	PID Notch Filter Coefficient N1	0
I0937	PID Notch Filter Coefficient N2	0
I0938	PID Notch Filter Coefficient D1	0
I0939	PID Notch Filter Coefficient D2	0
I0963	Integration Limit	50000
I0968	Friction Feed Forward	0

In previous operation (as tuned by EOST), the servo amplifiers were operated in a non-optimal state. The de-tuning of those amplifiers masked the overall performance of the PMAC control loop; once the servo amplifiers were correctly tuned, the problems inherent in the servo system were quickly exposed.

Our servo tuning started with the interactive tuning tool provided by Delta Tau as part of their PEWIN software suite (see section 2.1.3.1.3). This allowed for rapid iterative testing of

different basic parameters; after determining a reasonable set of values, we would move on to tests at constant velocities using the AxCON tool (see section 2.1.3.1.4). This allowed for validation of the tuning parameters as well as modest adjustments to remedy problems observed at tracking or slewing speeds. The final tests occurred using the normal telescope server software, and involved tracking stars with or without the dome shutter open. While tracking, perturbations would be applied to the telescope structure in order to see how well the system responded to an external disturbance.

Our initial attempts at adjusting the tuning parameters were foiled by poorly set maximum velocities and accelerations for both the elevation and azimuth axes. After setting those parameters to rational values, each successive pass at tuning the overall system resulted in improved performance. The major improvements to the tuning are as follows:

- Significant increases in the servo update rate for both axes (see section 2.1.2.4.3)
- Significant increases in the proportional terms, increasing the "stiffness" of the control behavior
- Large reductions in the integrator limit terms, eliminating at least one class of runaway oscillation for both axes

2.1.2.4.3. Servo update rate

The servo update rate for each axis is determined by the value of the corresponding Ixx60 parameter (where 'xx' is the motor number), which sets the servo cycle period extension register. Currently these are set as follows:

Axis	**PMAC Register**	**Register Value**	**Update Rate**
Azimuth	I0960	3	566 Hz
Elevation	I1360	12	174 Hz

When EOST last worked with the APF telescope in January 2012, the servo update rates for both axes were reduced from 377 Hz to 98 Hz, presumably to introduce additional damping into the control system. This had the side-effect of amplifying the tendency of the telescope to go into oscillation, as it introduced a significant phase lag between an impetus and the response of the PID loop. This specific frequency was an unfortunate choice, as it pumped oscillations around 32 Hz, which were also being pumped by a periodic errors in the signal from the azimuth Heidenhain encoders at specific tracking speeds (see section 2.1.7.7).

Another side effect of having these low servo update rates is that the proportional term is not as effective, especially when trying to move at a constant velocity. The integral term had been increased to negate this effect, but that increase also increased the tendency of the system to drive itself into oscillation, especially with a large integrator limit.

Quick testing with the elevation axis demonstrated improved overall performance of the system with a modest increase in the servo update rate, from 98 Hz to 174 Hz. Beyond the benefits mentioned in the previous paragraph, this also prevented the control loop from interacting with the periodic error in the azimuth encoder.

More extensive testing with the azimuth axis similarly demonstrated improved overall performance up to a significantly higher update rate, 566 Hz. At this high frequency, the control system is effectively prevented from pumping mechanical resonances or other periodic errors, and enabled the reduction in the integrator limit for the azimuth axis, from 5,000,000 down to 50,000.

Planned near-term work will include increasing the servo update rate of the elevation axis to match the azimuth axis.

2.1.3 Existing Tools for Measuring the Servo System Response

Here we briefly describe the various tools we currently have available for monitoring the performance of the telescope servo systems. Some of these tools can only be used with personnel present in the dome, while others can be used remotely.

2.1.3.1 On-site off-line tools

At present, using these monitoring tools requires having at least one staff member present in the APF Dome, either to monitor the output of the tool directly or to be able to promptly (and reliably) engage an E-stop if a dangerous oscillation develops. Two of these tools cannot be used when the telescope is online (i.e., able to operate on-sky), because they require that the TelescopeServer process be stopped before these tools can be run.

2.1.3.1.1. Oscilloscope monitoring of servo amp current

To enable monitoring of the servo amplifier output currents (for the azimuth and elevation drives) using an oscilloscope, we fabricated a set of breakout cables that provide a location to which an oscilloscope probe can be solidly attached. Oscilloscopes screen shots can be captured to memory and transferred to a flash drive for subsequent analysis and documentation. At present, we have no ability to monitor or operate the oscilloscope remotely.

2.1.3.1.2. Microphone to monitor sounds from drive motor

We have installed a microphone in close proximity to the drive motor for the elevation axis. This microphone is connected to a set of amplified speakers located near our control desk

on level 1 of the APF Dome. This enables an observer working from that location to monitor the sounds coming from that motor. This audible feedback provides a very sensitive and rapid means of alerting the observer to servo system instabilities in that axis so that they can take immediate corrective action. While we would like to be able to provide this audio feedback to APF observers who are working remotely, to date we have been unable to identify the resources needed to make that happen.

2.1.3.1.3. Delta Tau's PEWIN servo tuning tool

The PEWIN tool is Delta Tau's primary tool for configuring the parameters of the PMAC board and for interactively testing those parameters on the servo systems that are attached to and controlled by the board. This tool runs under Windows/XP on the telescope control computer (TCC). To run this tool, the TelescopeServer process must first be shut down.

Typically, this tool is used to adjust PMAC parameters, generate various test patterns (e.g., a step input or a parabolic velocity profile) that are output to the servo system, and to measure/display the response of the servo system to those test patterns. The user interactively adjusts the PMAC parameters until the observed responses to various test patterns are determined to be satisfactory. For additional details, see Section 4.3.1 of the Lick 2.4m Telescope Software User Manual SUM-13498-1.

The PEWIN can either be run locally on the TCC console, or remotely over the network, provided that there is a staff member in the APF dome who is able to monitor the telescope and engage an E-stop if needed. In the later case, the PEWIN session on the TCC is typically run via rdesktop from another (typically Linux) host inside the APF Dome. That rdesktop session is run within the context of a VNC desktop that is shared between the staff member in the dome and the staff member working from the remote site (e.g., UCSC). The remote user's connection to the VNC desktop is via an ssh tunnel.

2.1.3.1.4. EOST's AxCON servo tuning utility

The AxCON utility is EOST's tool for configuring and testing a subset of the PMAC servo tuning parameters. In addition to setting and displaying the current values of those parameters, AxCON can be used to command the telescope to jog at specific velocities for a specific angular distance.

For the selected axis of motion (e.g., azimuth or elevation), AxCON provides a real-time graph of both the PMAC's DAC output value and the servo following error for that axis. The user interactively adjusts the PMAC parameters for that axis and tests the servo system over a wide range of tracking and slewing speeds until the real-time graph of the following error shows that the settling time and following errors for that axis are acceptable.

While the AxCON program can issue commands for both axes in parallel, it can only display the results for one axis at a time. Like PEWIN, the AxCON program can be run locally or remotely, following the same approach and restrictions as described above. The TelescopeServer process must be stopped before AxCON can be run, and AxCON can't be run at the same time as PEWIN. For additional details, see Section 4.3.2 of the Lick 2.4m Telescope Software User Manual SUM-13948-1.

2.1.3.2 Remotely-accessible online tools

The following tools can be used at any time, including when the telescope is tracking an object on sky. These tools are accessible both locally and remotely. No one need be present in the dome in order to use these tools.

2.1.3.2.1. Accelerometers

To provide a more quantitative means of assessing the frequency and amplitude of the oscillations that intermittently afflict the servo systems for the azimuth and elevation axes, we have installed single-axis accelerometers (Dytran Model 3191A) on both telescope axes; additional details can be found here. These accelerometers (which are on loan from another project) are interfaced to a PC running the National Instruments LabView Signal Express software under Windows 7. That software is able to display (in real-time), record, and play back the measurements from each accelerometer, and to display those measurements in both the time and frequency domains.

Using the same approach as described above for the PEWIN software, the LabView Signal Express displays are routed (via rdesktop) to a shared VNC desktop that enables both local and remote users to simultaneously monitor the accelerometer outputs in real-time.

2.1.3.2.2. 1 Hz monitoring of telescope keywords

EOST's TelescopeServer process publishes periodic updates (typically at around 1 Hz) for a variety of its internal variables, such as telescope position, velocity, and servo following errors. EOST provides a DeviceBrowser tool which enables one to display in a tree-structured format all of the variables that are published by the telescope. However, this tool provides no capability to log specified variables to a time-stamped log file for later analysis or to generate live, time-history plots of such variables.

UCO has developed a layer of software that maps individual keywords to corresponding variables that are published by the TelescopeServer and related EOST/EOS servers. This software layer utilizes the infrastructure provided by the Keck Task Library (KTL) and associated tools that are widely deployed at both Keck and Lick Observatories.

Using this keyword layer, we can easily capture most of the interesting variables (e.g., servo following error for each axis) published by the TelescopeServer and either record these

into time-stamped log files or display them via live, time-history plots. In addition, all such keyword broadcasts are now archived at all times to a searchable database. This has proven to be an extremely useful capability for monitoring telescope performance during observing and for diagnostic purposes.

2.1.4 Servo System Response

The following description of the servo system response for the azimuth and elevations axes corresponds to the measured response obtained during October. As noted earlier, additional efforts to tune the servo systems are ongoing. Thus, the responses illustrated below show the types of responses we can measure, but they do not reflect the most current tuning of the system. Updated results will be presented at the review.

2.1.4.1 Response to step input

Shown in Figure 2.2 is an example profile of a step move for the elevation axis between two discrete encoder positions. The large initial oscillation is due to noise from the servo amplifier that occurs whenever you transition from a positive output voltage to a negative voltage. This zero-crossing rarely occurs when tracking, although it does occur in the elevation axis when tracking through the meridian and can occur in the azimuth axis when tracking circumpolar stars. Also visible in this graph is a modest amount of high frequency noise at the servo update rate (174 Hz), likely an amplification of encoder noise by the derivative term in the control loop.

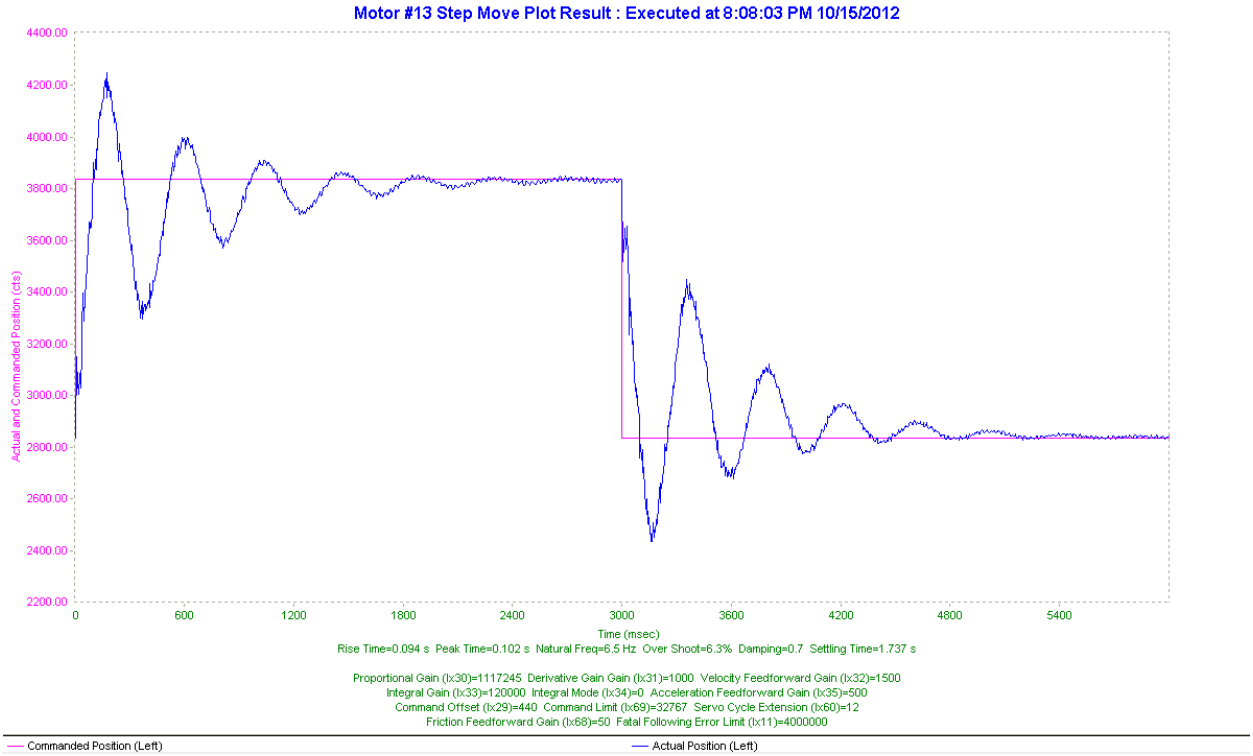

***Figure 2.2**: Response to a step input.*

2.1.4.2 Response to parabolic velocity input

Shown in Figure 2.3 is an example profile of a parabolic move for the elevation axis, likewise between two discrete encoder positions, but with gradual acceleration and constantly changing velocities. By comparison, the step move attempts to accelerate at its configured maximum, and will attempt to arrive at the destination in the minimum amount of time allowed. The variation in the commanded velocity (pink) is a bug in the graph (and not in the actual generated velocity profile), resulting from the sampling rate exceeding the servo update rate; in normal conditions, the commanded velocity vs. time is shaped like a pure sine wave for this test. The 174 Hz noise noted above is also visible here, most readily in the graph of the following error (green). Also visible are the zero-crossing errors at roughly 500 ms and 4600 ms into the test.

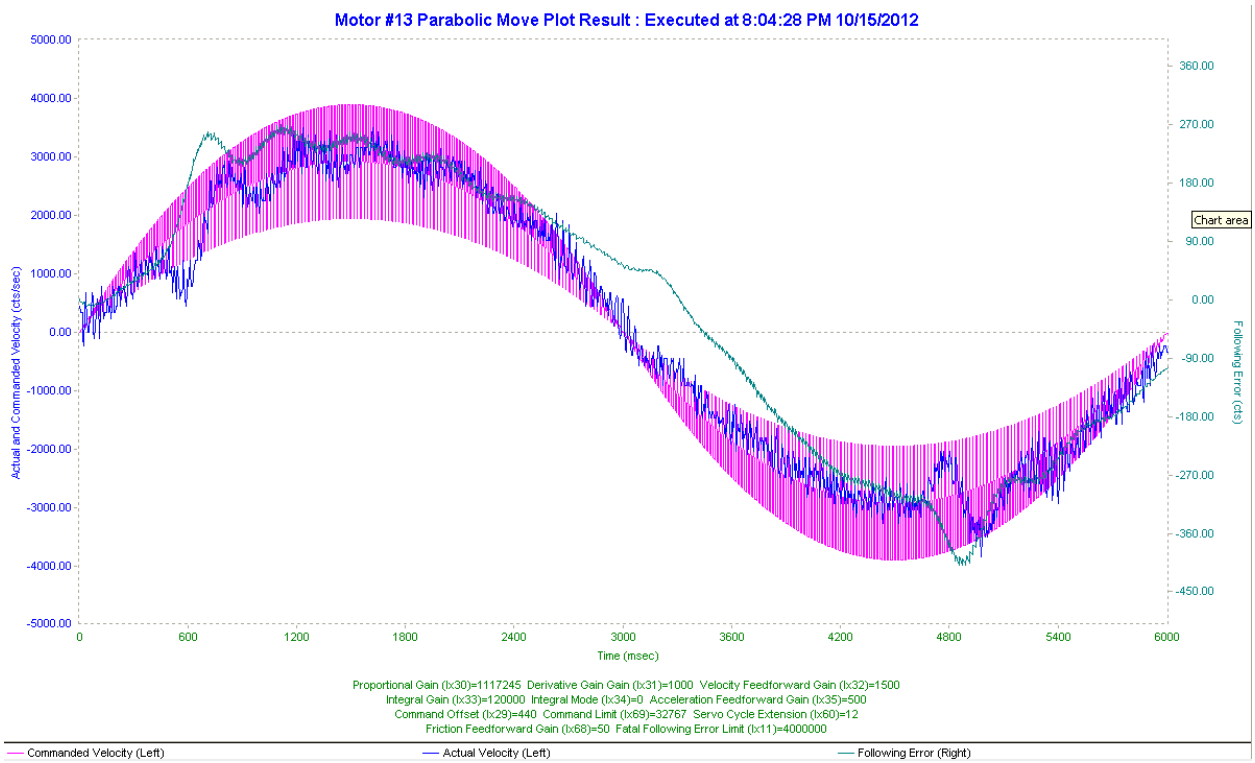

***Figure 2.3**: Response to parabolic velocity input.*

2.1.4.3 Following-error at select velocities

Shown in Figure 2.4 is an example of the behavior of the elevation axis when receiving an external impetus while tracking at a constant low velocity. The following error (FE) is in arcseconds, and the DAC output is in 16-bit DAC counts (+/- 32767). This figure is an excellent example of the lack of critical damping in the system, and highlights one of the key areas of conflict in the PID tuning process: an inability to introduce sufficient damping via the derivative term without introducing high-frequency oscillations due to noise amplification. This specific example was a fairly large impetus, and the axis fully recovered in three seconds.

Similar observations in the absence of an external impetus demonstrate the overall performance of the system at tracking speeds. For example, at low speeds, the graph shown in Figure 2.5 (note: this measurement was taken before the amplifier gain was correctly set) is dominated by periodic noise, with measured following errors on the order of 0.1 arcseconds. The magnitude of the following error scales with the commanded velocity.

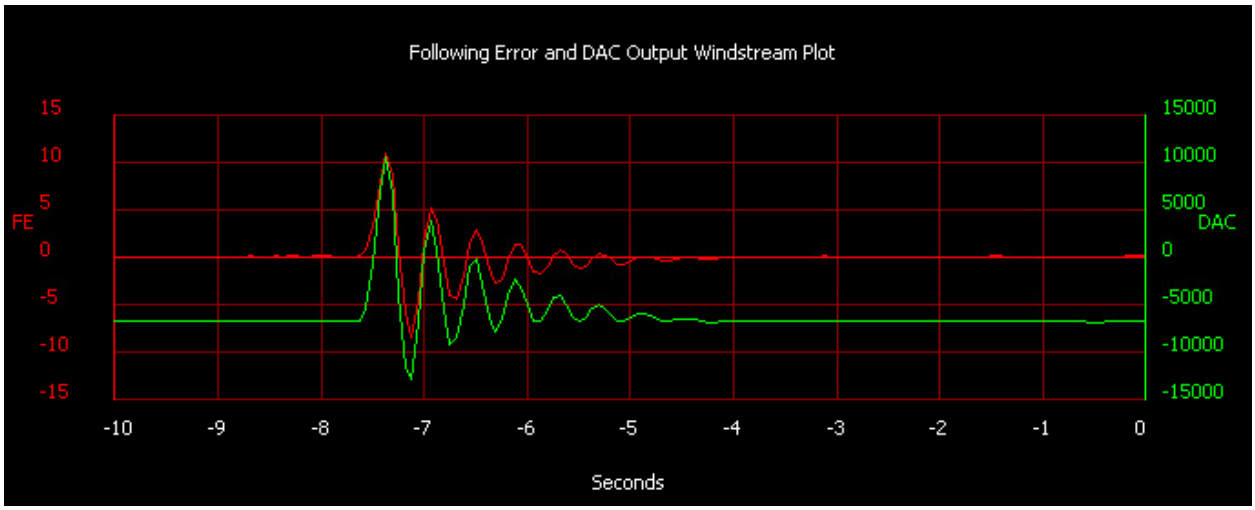

***Figure 2.4**: Response to an external impetus while tracking.*

A large following error at slew speeds is not considered a significant flaw, so long as the error does not grow unbounded. Delta Tau recommends that the integral term be disabled while the system is accelerating or moving at a constant velocity, and that it only be used to holds position; the APF telescope requires the use of the integral term in order to control the following error when moving at a constant velocity. If the integral term is undersized, the following error can take full minutes to converge to zero.

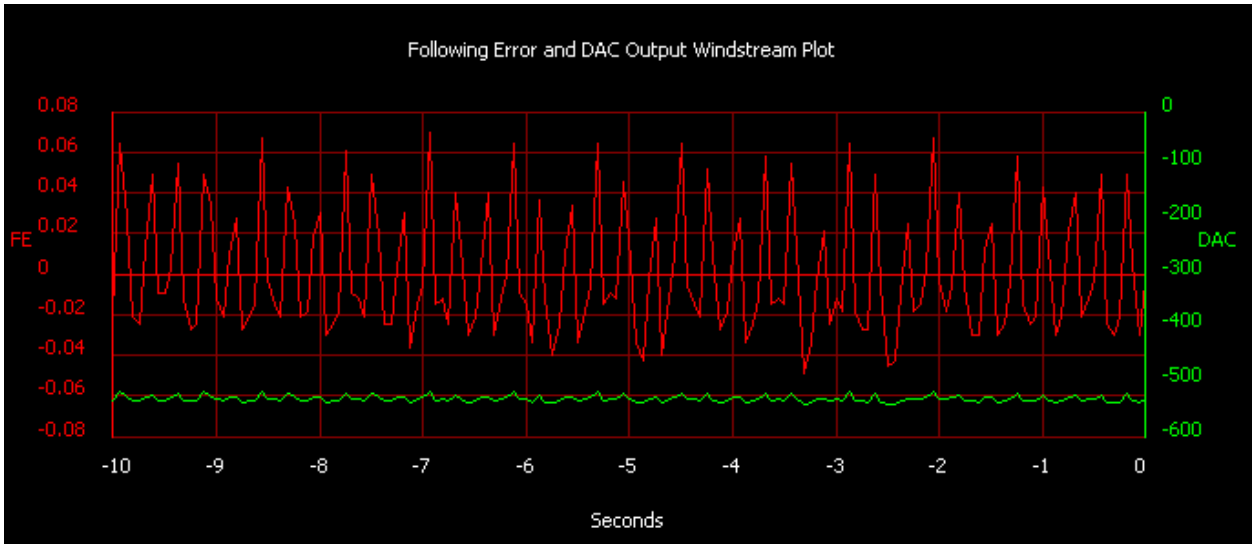

***Figure 2.5:** Following error during tracking.*

2.1.4.4 Accelerometer measurements

Many of the oscillations that we observe with the accelerometers are difficult to reproduce, indicative of a servo system that is only conditionally stable. As we continue our efforts to improve the servo performance, the character of the observed oscillations varies as

well. Accordingly, we will present in this report only some example accelerometer measurements. We will present the most recent accelerometer measurements at the review.

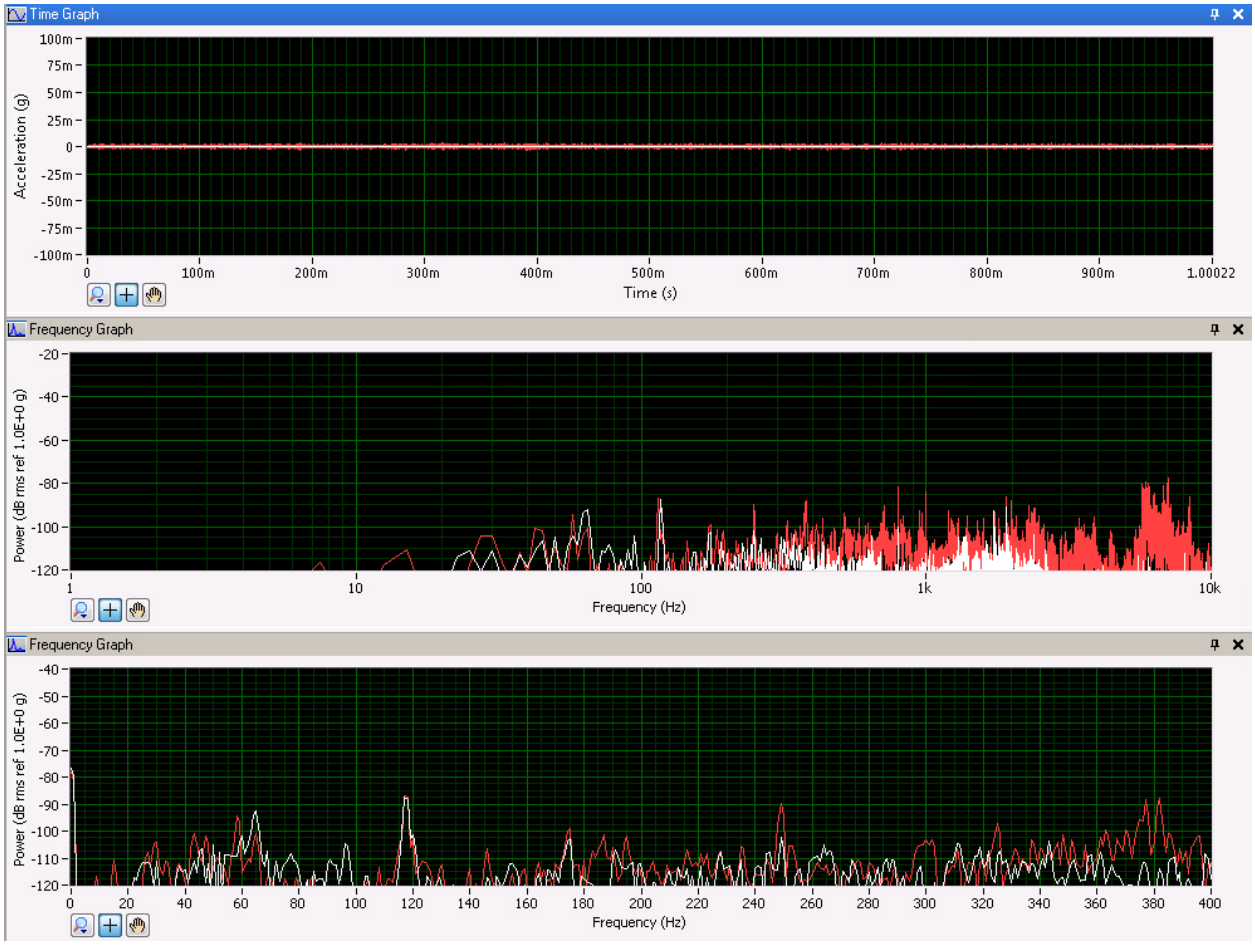

Figure 2.6: *Accelerometer measurements for telescope powered off.*

Figure 2.6 shows the observed accelerometer measurements from the APF Telescope when its servo drives are unpowered (the azimuth axis is shown in red and the elevation axis in white). The top plot is amplitude versus time, the other two show amplitude versus frequency. Most of the low-frequency (less than 200 Hz) noise in the graph is due to the air circulation fans on the third level of the dome. While those fans are normally left running during the daytime, they are usually turned off during nighttime observing. The feature that appears at 118 Hz is due to vibration in the instrument coolant supply line (that attaches to the manifold on the side of the spectrometer) and which is fed by the NESLAB chiller on level 1 (see section 3.2.16).

Figure 2.7 shows a similar set of graphs, except that this was captured when the telescope was experiencing a mild oscillation. The measured power is not peaking significantly above -60 dB, and there is significant structure visible in the graph: in this case, low-frequency

peaks at 34 and 141 Hz. The former is due to the periodic encoder error (see discussion following Figure 2.7), and the latter is likely an interaction with the servo update rate of 174 Hz on the elevation axis. This helps to illustrate the tight coupling between the two axes, in that sources of error on either axis are quickly propagated to the other axis, and potentially amplified. Also visible in this graph is a high-frequency "forest" from about 400 Hz to 2000 Hz. This series of peaks is usually the result of an interaction with a high derivative term.

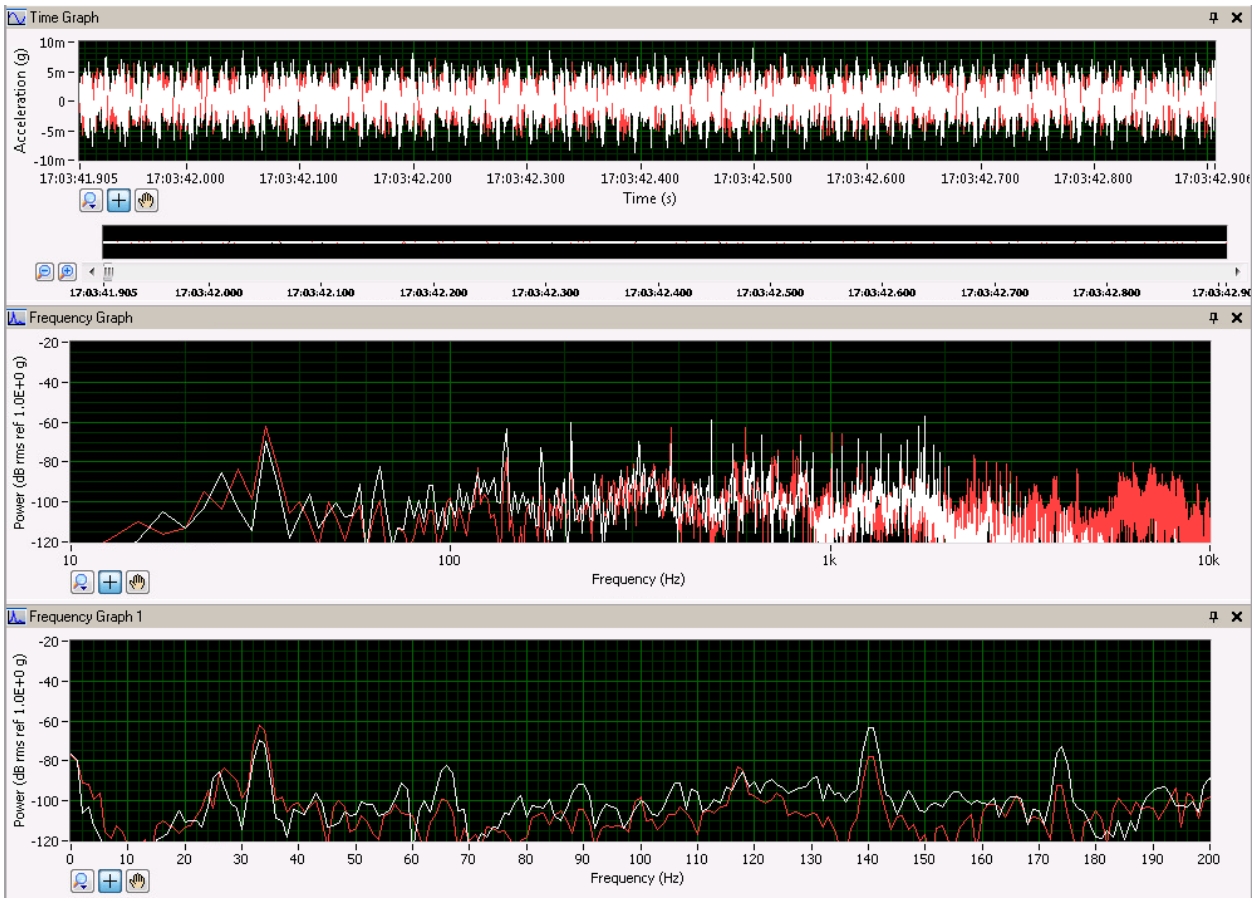

Figure 2.7: *Telescope oscillating while tracking*

There is one type of oscillation that is extremely reproducible: an oscillation in the azimuth axis whose frequency is linearly related to the commanded azimuth velocity according to the relationship:

oscillation frequency ~= 370 * commanded azimuth velocity

where the oscillation frequency is measured in Hz and the commanded azimuth velocity is measured in degrees/sec. We believe this oscillation is probably caused by a problem with the mechanical alignment of the read heads for Heidenhain tape encoder for the azimuth axis.

As noted in subsection 2.1.1.4.1, the encoder tape has a grating period of 40 microns (or 25 lines/mm), and has 133000 such periods (or lines) over its 5320 mm length. The tape is

affixed, end-to-end, in a circular groove within part of the azimuth bearing. Thus, the 133000 lines on the tape correspond to a full 360-degree rotation of the azimuth axis. Hence, for each one degree rotation of the azimuth axis, approximately 370 of these lines on the encoder tape (i.e., 133000/360 = 369.44444) are rotated past the encoder's read head.

After discussing this behavior with technical staff at other observatories where similar Heidenhain tape encoders are used, the consensus opinion is that the oscillations described by the relationship listed above are indicative of a non-optimally-aligned encoder read head. The resulting non-optimal sine waves generated by the 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.

Even though the square waves are not evenly spaced, the interpolator module is still producing the correct number of squares waves (i.e., 400) for each period of the encoder tape. As a result, there is no net change in the total number of interpolated encoder counts and thus no cumulative effort 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, from the standpoint of 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, if there is a periodic error in the width of those square waves, that will result in a corresponding periodic error in the commanded azimuth velocity, which in turn will cause an oscillation at the frequency described in the relationship above.

Those who have aligned Heidenhain encoder read heads on similar EOST-built telescopes state that the read head needs to be aligned looking at the full length of the encoder tape (by rotating the axis through its full 360 degrees of motion) rather than a small section of the encoder tape. One has to pick the best alignment/performance for the entire range of motion.

Unfortunately, when the azimuth encoder read heads were re-aligned by Gray in late July, we were unable to rotate that APF Dome during many of the days that he was at the APF site due to a failed Baldor Vector drive controller in the dome azimuth drive system. As a result, the re-alignment of the azimuth read heads was performed over the very limited range of azimuth rotation (about 5 degrees) set by the end-of-travel limits on the “trailer hitch” (see subsection 2.2.1.4.1.2) assembly that is part of the dome azimuth drive system.

These results suggest that the mechanical alignment of the azimuth encoder read heads needs to be checked, and if needed, re-aligned. Since we do not have the Heidenhain test equipment needed to perform that task, we are working with EOS to have them loan us their equipment so that we can perform this alignment check (and possible re-alignment) ourselves.

However, if we re-check the mechanical alignment of the read heads and find that it remains within specifications for the full 360 degrees of azimuth rotation, then the other possibility is that the electronics inside the read heads have drifted out of alignment. If that proves to be the case, then we might either need to remove the read heads and ship them back to Heidenhain for re-calibration at the factory or purchase a set of replacement read heads. Alternatively, we might be able to post-process the encoder signals to reduce this periodic error; see subsection 2.1.8.1.2).

2.1.5 Telescope Pointing

Pointing calculations are performed in the TelescopeServer process that runs on the Telescope Control Computer (TCC). The server uses a pointing model that is derived from a large set of pointing observations of bright stars with well know positions. The various components that affect the pointing model are described below.

2.1.5.1 Homing the telescope axes

Since the Heidenhain tape encoders on the azimuth and elevation axis are incremental, an absolute position reference needs to be established whenever the telescope control system is first powered up. The homing sequence for each axis employs a mechanical home switch that establishes a coarse home position for the axis.

After that coarse home position has been reached, the telescope is jogged at moderate speed while the control system attempts to detect, in the correct sequence, three successive absolute position reference marks on the Heidenhain encoder tape (see Figure 2.1). The incremental distance between each set of marks is distinct, and enables the control system to verify that it has detected the expected sequence of marks, enabling it to establish a precise and stable absolute position for the axis.

Once the telescope has been homed in all axes, it need not be re-homed unless there is a loss of power, the TelescopeServer is restarted, or the telescope is moved manually at a speed that exceeds the maximum count rate of the encoder system.

2.1.5.2 Time reference

The telescope's time reference is provided in combination by the internal clock on the PMAC board and the GPS Time reference. See subsection 2.1.1.4.4 for details. The current estimate for UT1-UTC is obtained from the same IERS Bulletin A that is consulted for the current estimate of polar motion.

2.1.5.3 Polar motion

The TelescopeServer accounts for polar motion via the IERS Bulletin A notices that are issued regularly by USNO and which can be downloaded from http://maia.usno.navy.mil/ser7/mark3.out. The continued availability of these bulletins from

this site and in this format is not guaranteed. Should the format for these bulletins change in the future, we will need to develop tools to convert from the new format to the current one.

2.1.5.4 Automated collection of pointing observations

We have developed a set of software scripts that can be run unattended (except for the need to monitor the telescope for servo instabilities) and which will collect a set of pointing observations for about 215 well-known bright stars that are relatively evenly spaced over the portion of the celestial sphere visible from the APF site.

For each such pointing star, the telescope is slewed to the catalog coordinates for that star, which results in the image of that star landing somewhere within the roughly 1' x 1' field of view of the guide camera. The autoguide software is then enabled to offset the star onto the center of the spectrometer slit. Once that star image is centered on the slit and its position is stable, the telescope's encoder coordinates are recorded into a file along with the catalog coordinates. That file can then be input into a pointing analysis program such as TPOINT (see subsection 2.1.5.7.1).

Depending on whether the maximum azimuth slew speed (see section 2.2.3.2) is set to its nominal value of 4 deg/s or reduced to its cold-weather setting of 3 deg/s, collection of a full set of pointing observations takes either about 2.5 hours or 4 hours.

2.1.5.4.1. Guide camera

The guide camera is a PhotonMAX-512B made by Princeton Instruments (subsequently acquired by Roper Scientific). This camera has 512 x 512 imaging pixels which are each 16 x 16 microns. It employs a back-side illuminated e2v frame-transfer CCD featuring both a conventional and a charge multiplying readout amplifier (only the conventional readout amplifier is used in normal operation). The CCD is thermoelectrically cooled to -70 C and the heat from that cooler is removed from the camera by a liquid coolant circulation pump that is located outside the spectrometer. The vacuum is maintained using metal seals that are pinched off, eliminating the need for periodic re-pumping, which would require opening up the spectrometer.

The camera employs a Camera Link interface that connects to a PCI-bus frame grabber board manufactured by Princeton Instruments (PI) and sold as part of the camera; PI also provides a Windows/XP device driver for this board, along with its PVCAM applications programming interface and its WinView application that can be used to operate the camera. While we used PI's WinView software for initial configuration and testing of the camera, that software is not part of the operational APF software. However, it can be used to capture windowed images from the camera at very high (e.g., 45 Hz) rates.

The PCI-bus frame grabber board, device driver, PVCAM libraries, and WinView software are all installed in the Autoguide Computer (AGC), which was a deliverable from EOS. The AGC computer runs Windows/XP and is located inside the telescope control rack on level one of the APF Dome.

Because the distance between the guide camera (mounted inside the APF Spectrometer) and the telescope control rack exceeds the maximum distance over which the CameraLink interface can operate, a set of CameraLink-to-Fiber converters (sold by PI) are used to overcome that limitation. There is one such converter located just outside the APF Spectrometer and another located just outside the telescope control rack; a dual fiber optic cable connects between the two converters, while a standard CameraLink cable connects the guide camera head (inside the spectrometer) and the frame grabber board (inside the AGC) to their respective converters.

The APF guide camera has been in operation since 2006 and has been extremely reliable. Since its installation onto the telescope in 2009, it has been operating continuously, except when turned off during maintenance operations on the spectrometer or the AGC. The original coolant circulation pump developed some problems with mineral deposits in 2010(?) and it was replaced by a new circulation unit. We were able to repair the original unit and it now serves as a spare. We also have a spare frame grabber board.

Unfortunately, we have no spare for the guide camera itself, and this particular model is no longer produced. However, PI has a newer model that has the identical mechanical form factor and CCD, but which employs a Gigabit Ethernet interface in place of the CameraLink interface; that would eliminate the need for the PCI-bus frame grabber board and the CameraLink-to-Fiber converters. The cost of the new model camera is approximately $35K.

Were we to upgrade to the newer model, some software changes would be needed. However, since PI now has a version of the PVCAM software that runs under Linux, such an upgrade might enable the AGC to be eliminated, such that UCO's guide camera software (see 2.1.5.5), which runs on a Linux host, could interface directly to the guide camera via the Linux-based PVCAM software layer that could be installed on that host. Such an architecture would be easier to maintain and would likely have a longer operational lifetime.

2.1.5.4.2. Guide camera mounting and optical path

The guide camera is mounted to a linear stage used to adjust the camera's focus. The camera is mounted in an up-looking orientation, so the motion of the focus stage moves the camera up and down relative to the spectrometer's optical breadboard, which is mounted horizontally.

As shown schematically in Figure 2.8, the beam from the telescope's tertiary (entering from the bottom left) first passes through the two prisms of the instrument's atmospheric

dispersion compensator (ADC), which is located within the elevation bearing. After passing through the ADC the beam hits a pellicle (actually, a polished and coated flat piece of glass) which passes about 96% of the light straight through to the iodine cell and then on to the spectrometer's adjustable slit stage, while reflecting about 4% to a set of focal reducing optics for the guide camera. The output side of those optics sends that beam to a folding flat mirror (not shown) that then directs it downward into the (up-looking) guide camera. There is also a corner cube (at the right edge of the figure) that sits behind a normally closed-shutter. When that shutter is open, the guide camera can image the slit.

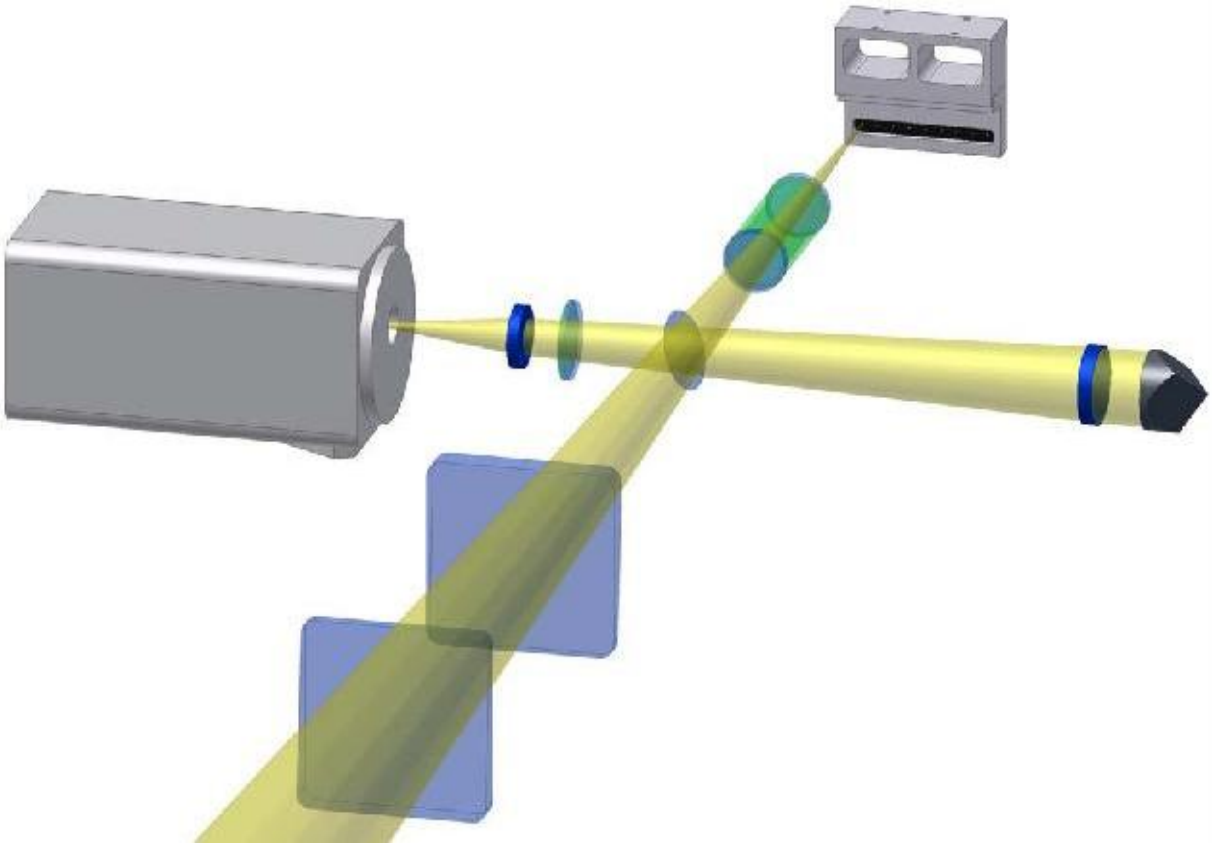

***Figure 2.8**: Guide camera optical path.*

The original design called for using an actual pellicle to pick off the light sent to the guide camera, and several different pellicles were used within the spectrometer during the first few months after its installation onto the telescope in mid-2011. Unfortunately, the pellicles proved extremely fragile and were easily damaged. In addition, when we first started observing oscillations of the star images from the guide camera, we suspected those might be due to acoustic oscillations of the pellicle. As a result, the pellicle was replaced by a 3.175 mm thick flat polished glass plate, the back side of which was anti-reflection coated (with Solgel) to reduce ghost images resulting from reflections off of the back side. While such reflections do result in

such ghosts, they are sufficiently faint and separated from the primary image so as to not interfere with the target acquisition and autoguiding software.

In its current configuration mounted within the spectrometer, the plate scale of the guide camera is 0.108 arcsec/pixel. Using observations of sets of bright stars densely grouped within an open cluster, we derived a plate solution and FITS world coordinate system (WCS) for the guide camera. As a result, when images from the guide camera are displayed using the ds9 image display, ds9's compass rose correctly displays the correct orientation for the currently-selected WCS and either the RA/Dec or alt/az coordinates will be displayed for the pixel in the guide image at which ds9's cursor is pointing.

It should be noted that the guide camera performs two distinct functions within the APF Spectrometer. First, it is used for target acquisition and guiding. Second, it is used as the detector for the instrument's exposure meter. So that the exposure estimates generated by this exposure meter would more accurately reflect the exposure levels in the Iodine region of the APF spectra, in late 2011 a 1mm BG-18 and a 3mm GG-475 filter were inserted into the optical path to the camera. While these filters improve the performance of the exposure meter, they significantly reduce the total flux that reaches the guide camera (especially for redder stars), thereby reducing the camera's faint limit.

Also note that prior to the installation of the APF Spectrometer onto the telescope in mid-2011, the guide camera was bolted (in a side-looking orientation) to a temporary optical breadboard that was attached (in place of the spectrometer) at the left Nasmyth focus of the telescope. In that configuration, the beam from the tertiary mirror bounced off of a folding flat mirror which directed all of the light from the telescope through an ND1 filter and into a different set of focal reducing optics that fed the light into the guide camera. The guide camera plate scale in that configuration was 0.322 arcsec/pixel. That temporary configuration was used for all of the telescope pointing and tracking tests that comprised the site acceptance testing performed from mid-2009 through mid-September 2010, when the ADC was installed into the elevation bearing. The temporary breadboard that held the guide camera remained in place until early 2011, when it was removed to make way for installation of the spectrometer.

2.1.5.5 Guide camera software

EOS wrote the AutoguiderServer (a Windows executable that runs on the AGC) which provides a communications layer that sits above the PVCAM libraries. The AutoguiderServer formats the image stream from the guide camera for transmission over the local area network using the EOS software framework, thus enabling those images to be relayed from the Windows-based AGC to a Linux-based host.

EOS also wrote an image display client that can be used to display the guide camera images transmitted by the AutoguiderServer. While this client has very limited display options, it

is sometimes useful for testing the low-level connectivity of the various elements of the guide camera software. However, for normal observing, that image display client has been superseded by the ds9 image display software (see section 2.1.6.2.1).

UCO wrote its own guide camera service (known as 'eosgcam') which runs on a Linux host; it receives the guide camera image stream transmitted by EOST'S AutoguiderServer that runs on the Windows-based AGC. The eosgcam service merges that image pixel stream with the relevant telemetry from the camera (e.g., exposure time, binning, windowing, camera gain, etc.) and telescope (e.g., RA, Dec, UT, etc.) in order to construct a valid and well-documented FITS-format file for each image from the camera. As noted earlier, these FITS files contain a WCS that describes the current orientation and scale of the guide camera image.

The eosgcam service writes each successive guide camera image to disk as a 2D format FITS file. Each such file is automatically loaded into the guide camera's ds9 image display (using ds9's xpa protocol) and is also read by the autoguider software. Keyword-settable options in the eosgcam service determine whether these 2D format FITS file remain on disk or are deleted once they have been loaded into ds9 and processed by the autoguider.

In addition, the eosgcam service can write a sequence of successive guide camera images to disk as a 3D format (i.e., datacube) FITS file; this enables the recording of “movies” of guide camera images that can then be replayed for subsequent analysis using ds9. This recording option is turned on and off via a KTL keyword. This is an extremely useful tool for monitoring telescope tracking performance or diagnosing problems with the telescope drive system. It was this capability that was used to document to EOST the telescope oscillations that were first observed in December 2011.

2.1.5.6 Atmospheric Dispersion Corrector (ADC)

The ADC is mounted inside the left elevation bearing and consists of two prisms: one of which is fixed and the other is mounted to a linear stage that moves parallel to the optical axis of the telescope. The fixed prism faces the tertiary and also serves as the entrance window to the spectrometer. The movable prism is moved (so as to adjust the spacing between the two prisms) as a function of the telescope elevation angle in order to match the dispersion compensation to the air mass corresponding to that elevation. The ADC is able to provide full compensation down to an elevation angle of 14 degrees.

While the installation of the ADC into the optical path has a negligible impact on the plate scale of the guide camera, it does have a significant impact on the telescope pointing. Changing the position of the movable prism results in an offset (along the elevation axis vector) of the star on the guide camera. This effect needs to be included in the telescope pointing model.

2.1.5.7 Constructing the telescope pointing model

The telescope pointing model is derived from a set of pointing observations of bright stars whose positions are well known and which are distributed as evenly as possible across that portion of the celestial sphere that is visible from the APF site on a given night. Each pointing observation consists of the catalog position of the star along with the telescope's azimuth and elevation encoder readings for the instant in time at which that star image was “observed” to be centered on the spectrometer's slit.

Using a pointing model composed of several geometrical terms (e.g., telescope axis misalignments, non-perpendicularity of the axes, encoder zero point errors, etc.), pointing model software performs a non-linear least squares fit of the observed encoder readings to their corresponding catalog positions. That fit determines the optimal values for the various terms that are included in the pointing model. Additional terms (e.g., telescope tube flexure) can be added to the model in an effort to improve the fit. For APF, two different pointing model programs can be used.

2.1.5.7.1. TPOINT

TPOINT is the telescope pointing software used by most major observatories (including Lick) to develop their telescope pointing models. It was developed by Patrick Wallace in the mid-1970s and has been maintained and updated for over 35 years. UCO has a site license that enables it to use TPOINT for any of the Lick Observatory telescopes on Mt. Hamilton.

Unfortunately, while TPOINT can be used to generate pointing models for APF, the model terms used by TPOINT are not identical to those used by EOST's telescope pointing kernel in the TelescopeServer. Thus, while TPOINT can be used to characterize the overall pointing behavior of the APF Telescope, its output can't be fed directly into EOST's software. Instead, one needs to use EOST's StarCal application (which is their own re-implementation of TPOINT) to generate pointing model terms that can be used directly by EOST's TelescopeServer software.

For a given set of pointing observations and corresponding model terms, both TPOINT and StarCal will yield similar results in terms of all-sky RMS pointing errors. The most recent set of pointing observations for the APF Telescope were obtained on August 29, and the pointing model derived from those observations is what is currently installed. Using TPOINT, the following pointing model was determined for that set of pointing observations:

	coeff	change	value	sigma
1	IA	-0.000	-32061.17	1.891
2	IE	+0.000	-1077.02	0.488
3	AW	+0.000	+15.82	0.194
4	AN	+0.000	-6.30	0.200
5	CA	+0.000	+33.19	2.457
6	NPAE	-0.000	+9.25	1.888
7	TX	+0.000	-24.69	0.333
8	HECE2	+0.000	-19.91	0.646
9	HECE9	-0.000	+1.16	0.277
10	HACA2	+0.000	-4.97	0.296
11	HESA2	-0.000	-0.69	0.237

Sky RMS = 2.41

Popn SD = 2.47

2.1.5.7.2. EOST-supplied software

2.1.5.7.2.1. StarCal

StarCal is EOST's re-implementation of TPOINT, but with some significant differences. Unlike TPOINT, it is not strictly an off-line analysis tool. Rather, StarCal has built-in features for interactively acquiring a set of pointing observations by commanding slews and offsets of the telescope; it assumes that there is a human in the loop, and for each and every pointing target to be acquired, it requires the human observer to navigate through multiple pop-up windows and menu options. When acquiring a large set (i.e., > 200) of pointing observations, this proves to be a labor-intensive, inefficient, and error prone process, and one that is totally unnecessary. That is why we developed our own automated scripts to acquire such large sets of pointing observations.

Although StarCal uses a different nomenclature and sign convention for its pointing terms, for a given set of pointing observations, it will yield similar results to TPOINT, i.e., all-sky (i.e., from 89 down to 15 degrees of elevation angle and from -105 to +305 degrees azimuth) RMS pointing errors of about 2.5 arcsec, with worst-case pointing errors well under 10 arcsec. Below is the comparable StarCal pointing model based on the August 29 set of pointing observations; this is the model currently being applied by the APF Telescope Server.

```
       Std DevValue
Term   arc secarc secName                              Formula
-----  --------      ---------     ---------------------------- -----------------
 01      0.78   -32057.39     Azimuth Encoder Offset              A=1
 02      0.37     1076.98     Elevation Encoder Offset     E=1
 03      0.12       16.08     X Axis Tilt                         A=tEcA; E=-sA
 04      0.12       -6.32     Y Axis Tilt                         A=tEsA; E=cA
 05      1.13       28.29     Transverse Misalignment             S=1
 06      1.00       12.78     Az-El Non-Orthogonality             A=tE
 07      0.25      -24.75     Tube flexure                 E=t(PI/2-E)
 08      0.48       19.82     Cos(2*Elevation)in Elevation E=c(2E)
 09      0.21       -1.27     Cos(9*Elevation)in Elevation E=c(9E)
 10      0.20       -4.76     Cos(2*Azimuth) in Azimuth    A=c(2A)
 11       0.18      -0.60     Sin(2*Azimuth) in Elevation  E=s(2A)
Sky RMS = 2.45
```

2.1.5.7.2.2. Mount Model Tool (MMTool)

As noted above, the StarCal software assumes that a given set of pointing observations was generated using StarCal's interactive acquisition process, and it further assumes these observations are recorded in its own proprietary, undocumented binary format. Further, StarCal has no ability to import pointing observation files written in the public ASCII text file format used by TPOINT. This made it impossible for us to use our automated scripts to generate a set of pointing observations in a format that StarCal could import directly for analysis.

To overcome this obstacle, EOST agreed to provide us with a copy of a software tool, MMTool, that they use in house to address this problem. Among its other capabilities, MMTool can input a set of pointing observations in the public TPOINT format and convert them into the proprietary binary format required for input into StarCal.

2.1.5.8 Pointing requirements

The contract with EOST requires that the pointing accuracy be better than 3 arcsec RMS to 70 degrees zenith angle (i.e., down to 20 degrees elevation angle) after mount modeling correction.

2.1.5.9 Measured pointing performance

The currently achieved telescope pointing performance (i.e., better than 2.5 arcsec RMS all the way down to 75 degrees zenith angle) not only meets but exceeds the contract requirement.

In more practical terms, for successful robotic observing, the APF needs to be able to automatically acquire the relevant targets of interest. Most of the candidate target stars for the APF exo-planet search effort are relatively bright (e.g., magnitude < 10) and have no nearby neighbors of comparable brightness. In those cases, the automatic acquisition process is relatively trivial. First, the telescope is slewed to the target coordinates, and barring thick clouds or other such obstructions in the light path, the image of that star will always land within a radius less than 10 arcseconds from the center of the spectrometer's slit. Second, the autoguider software locks onto the brightest target in that area and offsets the telescope to center the star image onto the slit center.

Unfortunately, there are a few candidate stars that have nearby (i.e., less than 10 arcseconds away) neighbors of either comparable or greater brightness than the candidate star. In such cases, the simple acquisition process will not guarantee that the candidate star will be centered on the slit, either because the neighbor star is brighter or because the separation between the candidate and its neighbor(s) is less than the worst-case pointing error.

For such stars, we use a two-step process in which we first use the default (simple) acquisition process to acquire an automatically-selected nearby pointing reference star (i.e., a bright star with a well-known position and no nearby neighbors of comparable brightness). Acquiring such a reference star greatly reduces the worst-case error for any small slews that are made within that local region of the sky. We then make a small slew from that reference star to the candidate star, and restrict the autoguider to only lock onto the target object that is within a few arcseconds of the slit center; that will be our candidate star. Candidate stars that require this two-step process must be flagged in the target list used by the automatic acquisition scripts.

2.1.6 Telescope Tracking

2.1.6.1 Open-loop

2.1.6.1.1. Tracking requirements

The contract specifications for telescope tracking rates are:

- > 0.5 degs/sec in azimuth
- > 15 arcsec/sec in elevation

The specifications for open-loop tracking accuracy (not including atmospheric seeing effects) are:

- <= 0.1 arcsec RMS over 10 seconds of time
- <= 0.50 arcsec RMS of 10 minutes of time
- <= 2.50 arcsec RMS over 1 hour of time

Note that these last contract specifications reflect a relaxation by a factor of 2 in the tracking accuracy requirements specified in the original APF request for proposals.

2.1.6.1.2. Measured open-loop tracking performance

The tracking performance was measured in 2009 and 2010 as part of the site acceptance testing and documented in late 2010. In that document, two tests were noted as not meeting the specification due to windshake; both tests were repeated under better conditions and the specification was met.

2.1.6.2 Closed-loop operation

Except for the periodic tracking errors that are caused by the periodic errors in the Heidenhain tape encoders for the azimuth axis, APF's open-loop tracking performance is so good that the autoguider software quite often has little to do. In fact, in those cases where displacements of the star image are primarily due to atmospheric seeing effects rather than telescope tracking errors, open-loop operation may perform as well (if not better) than closed-loop. See section 3.2.2 for further discussion of closed-loop guiding.

2.1.6.2.1. Autoguider software

The autoguider software for APF processes the FITS format files that the guide camera software (eosgcam) creates from the images read from the APF guide camera. Following the completion of a slew (by both the telescope and the ADC) to a new target, the autoguider software will lock onto the image of the target star and will offset the telescope by the requisite amount so as to align the centroid of that star image with the center of spectrometer's slit. As each subsequent image is read out of the guide camera, the autoguider will repeat that process so as to keep the star image centered on that slit. Because the guide camera is fed from a pellicle that picks off the light ahead of the slit (see Figure 2.8), the autoguider software sees the full star image, even when the star is centered on the slit.

In addition to its basic autoguiding functions described below, the autoguider service also implements the functionality of APF's exposure meter. For each guide camera image, the autoguider software computes the sum of the counts in the pixels comprising a rectangular region that corresponds to the spectrometer's slit; that sum estimates the star+sky flux that went down the slit during the period of time corresponding to that guide camera image. The software also computes the sum of the counts in the pixels comprising an identically-sized region well-separated from the position of the slit in the image; that sum estimates the sky background. The difference of those two sums estimates the star flux that went down the slit during that guide camera exposure.

These frame-by-frame measurements of the star flux going down the slit are then integrated to determine the current exposure level for the science exposure currently in

progress. When the desired exposure level is reached, the science exposure is terminated and read out. These frame-by-frame measurements are combined with the time at which each guide camera image was read out so that the photon-weighted midpoint of the science exposure can be determined and logged; this is an important component of the radial velocity calculations for the science target.

The pixels that comprise the rectangular region that corresponds to the spectrometer slit are determined via a one-time calibration process. The spectrometer is opened up and an LED light source is temporarily inserted so as to back-illuminate the spectrometer slit. The normally-closed shutter that sits in front of the corner cube (see Figure 2.8) is temporarily opened so that the guide camera is able to image the front of the slit. As a result, the guide camera sees an illuminated rectangular region that corresponds to the slit.

A best-fit rectangle is mapped to that slit image using ds9's region-definition feature, and that definition is saved to disk as a ds9 region file. This mapping process is repeated for each of the different-sized slits on the movable slit/decker plate. Whenever a different slit is selected by moving the position of the slit/decker stage, the corresponding ds9 region file for the selected slit is loaded into ds9 (so as to mark the location of the slit in each image) and into the autoguider software (so that the correct region of pixels in used in the exposure meter calculations).

2.1.6.2.1.1. Source extractor

The autoguider software uses the Source Extractor (a.k.a. sextractor) image analysis package to process each image frame read from the guide camera. For each object that is detectable in the image within a specified search radius of the slit center, sextractor computes a number of image statistics, including the image centroid, its semi-major and semi-minor axes, and an estimate of the flux. The autoguider software then sorts those results in order of brightness, and returns the statistics for the brightest object in its list of detected objects; it also broadcasts those statistics as KTL keywords, enabling them to easily be displayed, plotted, or recorded to log files. For example, plotting the semi-major axis, semi-minor axis, and flux versus time provides a very useful diagnostic for adjusting the telescope's M2 focus.

The search radius that sextractor uses for detecting objects in the guide camera image depends on the autoguider's operating mode. Immediately following a slew to a new target, that mode is set to 'acquire' and the search radius is typically set to about 25 arcseconds. Once the target object has been successfully centered on the spectrometer's slit, the mode is switched to 'guide', and the search radius is reduced to 2.5 arcseconds.

2.1.6.2.1.2. Image statistics keywords

As each image is read out of the guide camera, the autoguider software broadcasts keyword values for each of the following sextractor-generated image statistics for the brightest object within the specified search radius of the spectrometer's slit:

- SOURCEX X centroid
- SOURCEY Y centroid
- FLUX flux
- S_A semi-major axis
- S_B semi-minor axis
- S_FLAGS sextractor flags (e.g., “ANALYSIS OK”, “SATURATED”, etc.)
- S_THETA position angle
- S_X2 X second moment
- S_Y2 Y second moment
- S_XY XY moment

In addition to exposing the statistics computed by sextractor, the autoguider software also calculates and broadcasts statistics of its own, including average and RMS tracking errors and the computed telescope offsets that were applied.

The autoguide software also provides keywords that control its operation. For example, the MAXRADIUS keyword is used to set the search radius that sextractor will use when analyzing the next image. Other keywords enable one to specify if several guide camera frames should be co-added together before processing by sextractor.

2.1.6.2.1.3. DS9

The DS9 software (a.k.a. SAOImage DS9) is an astronomical imaging and data visualization application in use in many observatories. It is maintained by Bill Joye at SAO. For APF, we are currently using version 6.2 of this software.

2.1.7 Remaining Problems and Concerns

Aside from the problems enumerated below, the APF Telescope hardware and software is fully operational and capable of supporting the pointing, tracking, guiding, and exposure-meter functions required for local, remote, and robotic operations. Most of the remaining problems are either issues with the telescope encoders or the tuning of its servo systems. Once

these problems are resolved, the APF Telescope should be ready for a human observer to execute his or her exo-planet search program each and every night that weather permits.

2.1.7.1 Oscillations in the elevation axis

When the elevation axis is operated by itself, its servo tuning appears to provide acceptable settling times and following errors with little tendency for the axis to go into oscillation on its own. Unfortunately, when oscillations develop in the azimuth axis, these very often couple mechanically into the elevation axis, causing it to start oscillating as well.

2.1.7.2 Oscillations in the azimuth axis

As noted earlier, periodic errors with the Heidenhain tape encoders for the azimuth axis are causing periodic velocity errors which result in oscillations of the azimuth axis whose frequency is a function of the commanded azimuth velocity. Until those encoder issues are resolved, these oscillations will make it difficult to achieve an acceptable set of servo tuning parameters for the azimuth axis.

2.1.7.3 Long settling times at slow tracking rates

This is entirely a function of the servo tuning. Our recent servo tuning efforts have yielded significant improvements over the long settling times we observed with the last set of tuning parameters that EOST installed in January 2012. Unfortunately, some combinations of PMAC parameters that yield reduced settling times also appear to reduce servo stability.

2.1.7.4 Susceptibility of telescope to windshake

Like most telescopes, the APF Telescope is susceptible to windshake, especially when the dome shutters are wide open, which has been our default mode of operation. Depending on the operating mode of the dome shutters (see subsection 2.2.2), it is possible to operate one or both of the dome shutters in a wind-shielding mode, so as to reduce the exposure of the telescope to the wind (see section 2.2.2.3).

Unfortunately, use of the wind-shielding mode raises some operational concerns with respect to the safety of the primary mirror (i.e., moving the dome shutters across an uncovered, up-looking M1). Possible options include:

- Closing the primary mirror covers whenever the dome shutters are moved over an up-looking mirror
- Leaving the mirror covers open but pointing the telescope at the horizon prior to moving the dome shutters over the mirror

The first option is our current practice, but this is not time-efficient because the primary mirror covers take a long time to open and close. Whether the second option is a satisfactory alternative is not yet clear.

2.1.7.5 Stability of M2 focus with temperature and elevation

The TelescopeServer software provides automatic compensation of the M2 focus as a function of temperature and telescope elevation. The temperature compensation assumes linear expansion and contraction of the telescope tube as a function of the telescope tube average strut temperature; that average is derived from a temperature sensor on each of 4 of the 8 struts that comprise the tube. While this model appears to work to first order, residual focus shifts are still observed as a function of temperature.

We suspect that there may also be residual focus shifts as a function of the telescope elevation angle, but we have not yet attempted to make a rigorous measurement of this effect.

At present, we still have a human observer in the loop who monitors the telescope focus and adjusts it if necessary. With the transition to robotic operation, either the calibration of the automatic focus compensation mechanisms provided by the TelescopeServer will need to be improved, or we will need to implement an auto-focus script that the observing robot can run periodically during the night.

With regard to establishing the M2 focus, note that since the guide camera does not image the star on the spectrometer slit, some other means is needed to determine when M2 is properly focused with respect to that slit. This determination is made manually by means of a knife-edge test at the slit, using the movable slit/decker stage as the knife-edge. Performing this one-time test requires opening up the spectrometer and using an inspection mirror to enable one to see the back of the slit. Once best M2 focus is achieved at the slit, the guide camera focus stage is first adjusted for best focus of the slit outline, then the fold optics are adjusted to yield the best focused star image.

2.1.7.6 Heidenhain azimuth encoder periodic errors

These are discussed in subsection 2.1.4.4, and they represent the most urgent problem we need to resolve in order to achieve acceptable servo tuning for the telescope.

2.1.7.7 Heidenhain encoder interpolation dither

In addition to the periodic errors we observe from the Heidenhain tape encoders for the azimuth axis, we also observe interpolation dither on the interpolator outputs for all of the telescope axes. As seen in Figure 2.9, this interpolation dither results in fast (of order 1 usec) glitches on the quadrature square waves generated by the encoder interpolator modules.

While the PMAC boards have some ability to ignore such glitches, at certain speeds they may have difficulty discriminating these glitches from the actual square wave transitions. Also, these glitches may be causing problems for the Baldor Vector controller for the dome's azimuth drive, since that drive receives a copy of the encoder pulses from the telescope's azimuth A encoder.

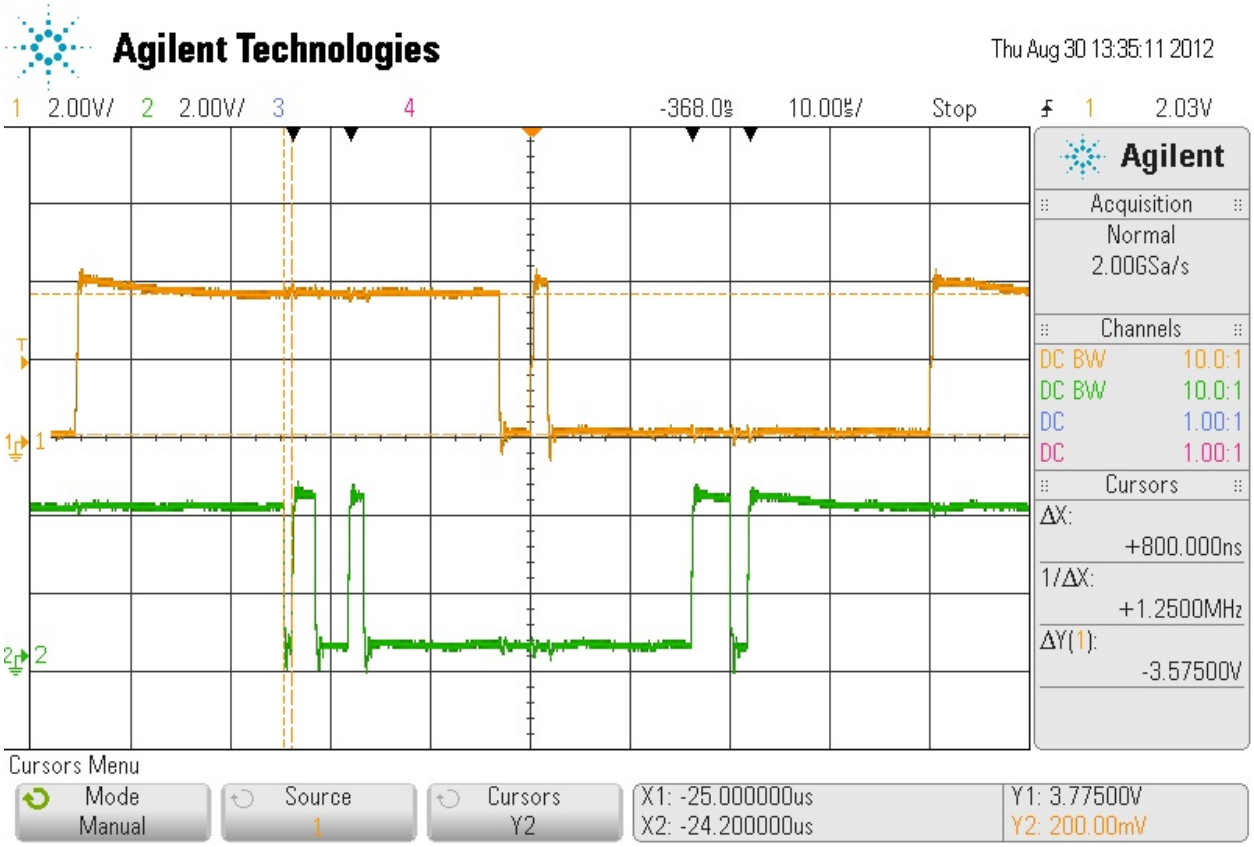

Figure 2.9: *Glitches caused by encoder interpolation jitter*

2.1.8 Next Steps

Below we list the next steps we plan to take to address the problems and concerns enumerated immediately above.

2.1.8.1 Construct break out boards for signal capture

The telescope control rack does not provide convenient test points to enable monitoring of the relevant signals (e.g., encoder signals, servo amp drive currents, etc.) for the telescope servo drives. To overcome that lack, we have already fabricated some breakout boards and will be building additional boards this month.

2.1.8.1.1. Heidenhain encoder breakout board

We have designed and are in the process of fabricating an encoder breakout board that will provide convenient access (at the back of the telescope control rack) of the quadrature square-wave signals from the Heidenhain encoder interpolator modules (both the “A” and “B” encoders) for the azimuth and elevation axes.

In addition, this board will contain some active logic that will provide some pre-processing of these signals. This logic will be implemented using an FPGA, and the implementation will occur in two phases. In phase 1, this pre-processing logic will:

- remove the interpolator dither glitches (see Figure 9).
- generate a virtual tachometer output by using the signals from the encoders to generate a signed analog output voltage that is proportional to the telescope's velocity in the given axis

In phase 2, we plan to implement a scheme that will enable this breakout board to dynamically compensate for the periodic errors in these encoders, particularly the ones for the azimuth axis. This compensation will be based on a method developed and used at the SOAR Telescope and the Blanco 4-m Telescope.

There are two characteristics of the Heidenhain encoders that may be affecting the tracking errors.

1. Observed dither/jitter/glitches at the decision points out of the Heidenhain interpolators.

2. There is a known non-linearity on the output of the Heidenhain encoder interpolator. This non-linearity shows up as a velocity determined oscillation in the following error (at a rate of 370 cycles per degree). This error has become the dominant tracking error contributing as much as +/- .3 arcsec error.

Proposed fixes:

1. We have designed and are building a circuit board that goes inline with the Heidenhain encoder signals to the PMAC and removes the decision point dither. It is hoped that this correction will reduce the low level noise in the following error. This board also provides a diagnostic tachometer output from the Heidenhain encoder signals. This board is also a learning platform for dealing with the encoders on the telescope is a building block for the second and much more complex PCB. The cost of the hardware was less than $200 and will require approximately two man weeks of time to program and test.

2. We propose to design, build a second PCB with substantially greater computational power that will allow for correction of the encoder interpolator non-linearity. The concept was presented in an SPIE paper on the SOAR telescope in 2008. We will implement a digital phase lock loop that watches the non-linearity of the previous 10 cycles (*400 interpolation = 4000 encoder counts) from the magnetic tape magnetic lines and removes that average profile from the current cycle. This technique should

remove the majority of the non-linearity and eliminate the error as the dominant tracking error source. The hardware cost is estimated at less than $2k to fabricate and will require approximately two man months to program and test.

2.1.8.1.2. Servo amplifier breakout board

We have already fabricated and used extensively passive breakout boards for the azimuth and elevation servo amplifiers. Each board sits in series with the input signal cable for the respective servo amplifier, and provides test points to which oscilloscope probes can be attached securely to monitor each of the relevant signals, e.g., the control voltage generated by the PMAC controller for that axis, the monitor test point for the servo amplifier current output, etc.

We have recently fabricated, but have not yet had an opportunity to test, a version of this breakout board which provides a buffered, single-ended copy of the PMAC's differential control voltage. This will enable us to monitor the PMAC control voltage using the same USB-based digitzer module and LabView Signal Express software that is currently used to monitor the outputs from the accelerometers, so that these signals can be plotted using a common time base.

2.1.8.2 Explore options to buy/rent/borrow signal analyzer

As noted in 2.1.2.1, a full servo analysis has never been performed for the APF Telescope. We believe it would be useful to do so, but we currently lack the equipment and other resources needed to perform this task. We also currently lack access to the requisite signals, but that lack should be remedied once we complete the fabrication test of the breakout boards described above.

We have been exploring options for obtaining a signal analyzer similar to the one cited in SPIE papers that describe the servo analysis performed for the SOAR Telescope. Refurbished models of that particular signal analyzer can be purchased for approximately $3200, and such units can be rented for a two-week period for about $1600. At present, the APF budget has no funds budgeted for such equipment.

2.1.8.3 Obtain resources needed for full servo analysis

None of the staff at UCO has performed a full servo analysis of a telescope like APF. While we are in the process of enabling access to the relevant signals (via installation of the breakout boards) and trying to obtain the needed test equipment, it is likely that we will require guidance from an outside consultant who has expertise in this area. At present, the APF budget has no funds budgeted for that.

2.1.8.4 Continue to refine PMAC tuning parameters

We are continuing our efforts to improve the servo tuning by adjusting the relevant PMAC parameters for each axis and those efforts have overlapped the writing of this report. In fact, the parameters listed in Tables 1 and 2 have recently been changed and further updates may occur prior to the Review. The latest parameters and results will be described at the Review.

2.1.8.5 Consider adding dynamic PID tuning to PMAC code

Users of other EOST-built telescopes (e.g., PanStarrs) have described an alternative approach to adjusting the PMAC parameters; that approach involves dynamically adjusting the PMAC's proportional gain (Kp) term for each axis based on the combined velocity of the elevation and azimuth axes. A small value of Kp is used during slew motions, and higher values during tracking; a functional curve or look-up table is developed empirically to map a given combined velocity to its corresponding Kp value, with separate mappings determined for the azimuth and elevation axes.

While we are not currently planning to implement such a scheme for APF, we are keeping this approach in mind should our current attempts to improve the servo tuning prove inadequate.

2.1.8.6 Research direct-drive tuning strategies used elsewhere

We have also been searching the literature for papers that describe servo tuning strategies employed on other telescopes that employ direct-drive motors to see what other approaches might prove useful for APF. We would appreciate receiving pointers to any such papers of which the Review Committee members are aware.

2.1.8.7 Experiment with dome-shutter wind shielding options

Once we have achieved adequate tuning of the telescope servo systems and addressed the safety concerns enumerated in 2.17.4, then, contingent on the availability of telescope engineering time on nights with moderate to strong (but not excessive) winds, we plan to conduct tests to evaluate the extent to which telescope wind-shake can be reduced by operating the dome shutters in a wind-shielding mode, so as to reduce the area of the dome shutter opening that is exposed to the wind.

2.1.8.8 Collect/analyze M2 focus data; refine autofocus model

If staff resources permit and if telescope engineering time is made available for this purpose, it would be helpful to collect measurements of how the M2 focus varies as a function of telescope tube temperature and telescope elevation angles. Using such data, it may be possible to refine the relevant look-up tables in EOST's telescope control software which already provide some level of compensation for these two effects. Developing an automated M2

focusing script (which is likely needed for robotic observing) would greatly simplify (and might enable the automation of) these measurements.

2.1.8.9 Working with EOS to adjust encoder read heads

As noted earlier, the Heidenhain encoders for the azimuth and elevation axes currently exhibit two distinct problems:

- interpolation dither pulses on the quadrature square waves generated by the interpolation electronics, and
- periodic errors in the electronic interpolation of the quadrature sine wave signals from the encoder read head.

One or both of these problems is likely due to mechanical misalignment of the encoder read heads relative to the Heidenhain encoder tape. The interpolation dither pulses may also be due to non-optimal adjustment of the gain and bias potentiometers within the electronics contained inside the encoder read heads.

EOS is currently assisting us with our efforts to improve the mechanical alignment of the read heads for the azimuth axis. We have recently completed the paper work that will enable them to loan us their Heidenhain encoder test equipment that is needed to perform that alignment. Once we receive that equipment, we plan to re-check the mechanical alignment of those read heads, and re-align them if necessary.

Using their own spare Heidenhain encoders, EOS has set up a small test apparatus at their facility to investigate the effect of adjusting the potentiometers within the Heidenhain read head electronics and they have indicated they will report to us the results of their tests. Note that while Heidenhain advises against adjusting these potentiometers and insists that this should only be done at their factory, it appears that other users of such Heidenhain encoders may have ignored that advice and made such adjustments themselves.

2.2 Dome

The APF Dome is a co-rotating EOS Ice Storm 2 enclosure. Various documents are available which describe the design of this dome, including:

Automated Planet Finder Enclosure Design Analysis Report DN-07753-01

Lick Automated Planet Finder Enclosure Technical Description Document

Lick Automated Planet Finder Observatory Operations Manual EMB07782-01

Icestorm 2 Mechanical Maintenance Manual HBO07753-01

The documents for the Preliminary Design Review for the APF Dome can be found here.

The APF dome consists of a stationary outer ring wall (which supports the rotating hemispherical dome) and a stationary central concrete pier (which supports the telescope). In between the two is a three-level, annular-shaped carousel that rotates with the hemispherical dome. The ground floor of this rotating carousel contains all of the control electronics for the dome and telescope, as well as the electronics and coolant pumps for the spectrometer; the ground floor is maintained at +20 degrees C and is thermally isolated from the upper two floors by an insulated ceiling and an insulated door to the staircase to the next floor. A curved staircase, which rotates with the carousel, provides access between the first and second floor; that floor provides access to the dome bogey wheels, the dome azimuth drive motors and the stationary portion of the azimuth bearing. A separate curved staircase (which also rotates with the carousel) connects between the second and third floors; that floor provides access to the telescope and spectrometer.

All three floors of the APF Dome co-rotate with (but are isolated from) the telescope. This provides convenient access to the spectrometer from the level three dome floor, even while the telescope is tracking an object. It also means that the control cables and plumbing (that connect between the telescope-mounted spectrometer and the control electronics and pumps on level one) do not need to be routed through the dome's azimuth cable wrap.

Note that because the outside door to the facility remains fixed (since it attaches to the stationary outer ring wall) while the floors of the dome carousel rotate, this can be somewhat disorienting to new users of the facility. A series of glowing EXIT signs with arrow heads pointing towards the shortest path to the outside door are attached to the inside of the stationary ring wall to help users find their way out of the dome as quickly as possible.

2.2.1 Components

The major components of the APF Dome are summarized here.

2.2.1.1 Shutters (Front and Rear)

The dome shutter consists of a front shutter panel and a rear shutter panel which can be driven independently. Each shutter panel is driven by motors contained within the panel that engage stationary chains that are attached to the arch girders of the dome. This design is illustrated in slides 31 through 41 of the Dome Preliminary Design Review Session 1 presentation.

To close the dome shutters, the two shutter panels are driven towards each other until the rear edge of the front shutter panel and front edge of the rear shutter panel contact each other at their respective "home" positions and engage their respective weather seals. To open the shutters, the two panels can either be driven apart, or can be kept mated together and moved in unison towards the rear of the dome; see section 2.2.2. The shutters can be operated under control of the Dome Control Computer (DCC) or via manual control pushbuttons located inside the dome.

2.2.1.2 Vent doors

Aligned with level three of the dome are four independently-operable vent doors which can be opened at the start of the night to increase the airflow through the dome so that the air temperature inside the dome can be brought more rapidly into equilibrium with the outside air. On windy nights, it may be advisable to leave one or more of these vent doors closed to reduce the exposure of the telescope to strong winds. The vent doors can be operated under control of the DCC or via manual control pushbuttons located inside the dome.

2.2.1.3 Dome azimuth cable wrap and limit switches

The dome azimuth cable wrap is located on level one of the dome and the stationary portion of this assembly is mounted to the fixed concrete pier in the center of the dome. This cable wrap primarily carries various utilities (i.e., electrical power, chilled water lines from the external chiller, and external network cables) from the stationary central pier to the rotating portion of the dome. Those external utilities enter the dome through the outer ring wall and then pass underneath the metal plate and the rotating level one floor before emerging at the central pier. The motive force for rotating the dome azimuth cable wrap is provided by the dome drive motors, not the telescope drives.

2.2.1.4 Servo systems

The servo system electronics for the dome azimuth and shutter drives are located in the Dome Control Cabinet on the first floor of the dome.

2.2.1.4.1 Dome azimuth drive

The dome azimuth rotation is “loosely” slaved to the azimuth motion of the telescope, such that the azimuth of the dome always remains within a few degrees of the azimuth of the telescope; this is necessary to prevent collisions between the telescope and structures within the dome. (See slides 8 and 9 of Session 2 of the APF Dome Preliminary Design Review).

The dome and telescope are intended to be mechanically isolated from one another, so that motion of the dome does not perturb telescope tracking; prior to the installation of the APF Spectrometer in mid-2011, for the most part, (see 2.2.1.4.1.2) they were so isolated. As a result of that installation, there is now a large cable bundle (consisting of electrical cables and coolant hoses) that hangs between the spectrometer (which is itself firmly attached to the right yoke of the telescope) and one of the vertical columns on the rotating dome carousel.

Because this cable bundle creates a new connection between the telescope and the dome, it is critical that this bundle remains sufficiently flexible and slack so that the rotating dome does not impart significant torques or vibrations to the azimuth axis of the telescope. This photo of this cable bundle was taken in October 2011; since that time, the amount of slack in this bundle has been significantly increased. However, this bundle remains a cause for concern, especially on cold winter nights, when the cables and hoses in this bundle are likely to become stiffer than normal.

2.2.1.4.1.1 Telescope azimuth encoder/encoder splitter

The rotation of the dome is slaved to the azimuth motion of the telescope by feeding a copy of the signals from the telescope's azimuth “A” encoder into the servo drive system that controls the dome rotation. The signals from that encoder are replicated by the encoder splitter module in the telescope control cabinet. If the rotation axis of the dome and the azimuth axis of the telescope were perfectly concentric, such an arrangement might be sufficient to keep the dome and telescope in sync. However, because they are not, following errors will develop, and an additional sensor is needed to measure the residual position error between the telescope and the dome.

2.2.1.4.1.2 Linear potentiometer and “trailer hitch” limits

The dome rotation servo system uses a linear potentiometer (with 450 mm stroke and a self-aligning bearing) to sense any residual misaligment of the dome rotation relative to the telescope azimuth. The body of this potentiometer (blue cylinder at right edge of photo below)

mounts to a horizontal I-beam (EOS calls this the “Interface Beam”) that connects between two of the vertical columns that are part of the rotating multi-level carousel.

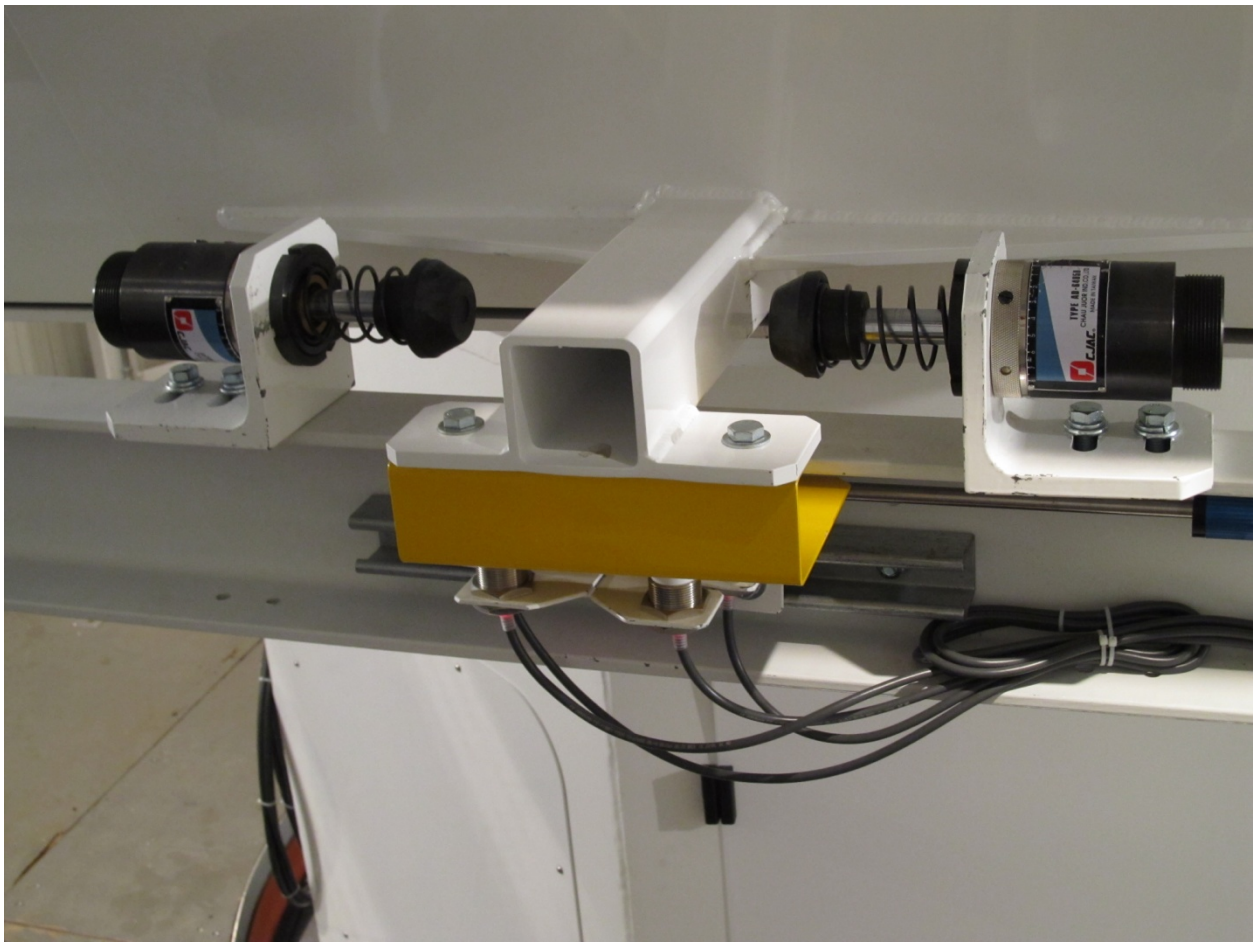

***Figure 2.10**: Linear potentiometer.*

The potentiometer's tap (silver rod extending out of the blue cylinder) attaches to the yellow bracket that is bolted to the white “trailer hitch” box beam which projects from the telescope. If the telescope is offset in azimuth relative to the dome, this trailer hitch box beam will move primarily left or right relative to the horizontal I-beam, thereby pulling or pushing on the potentiometer's tap so as to change its resistance.

Also attached to the I-beam are a set of Hall Effect sensors (Omron E2E-X18MEI proximity switches) which provide a set of soft and hard limits for the dome rotation system. If the telescope and dome get too far out of alignment, a soft limit is tripped, causing both the dome and telescope to be commanded to stop. If the soft limit is exceeded, a hard limit is tripped and power to the dome and telescope drives is disabled; in that case, before drive power can be re-enabled, the telescope needs to be pushed by hand in order to re-center the trailer hitch so as to clear the hard limits.

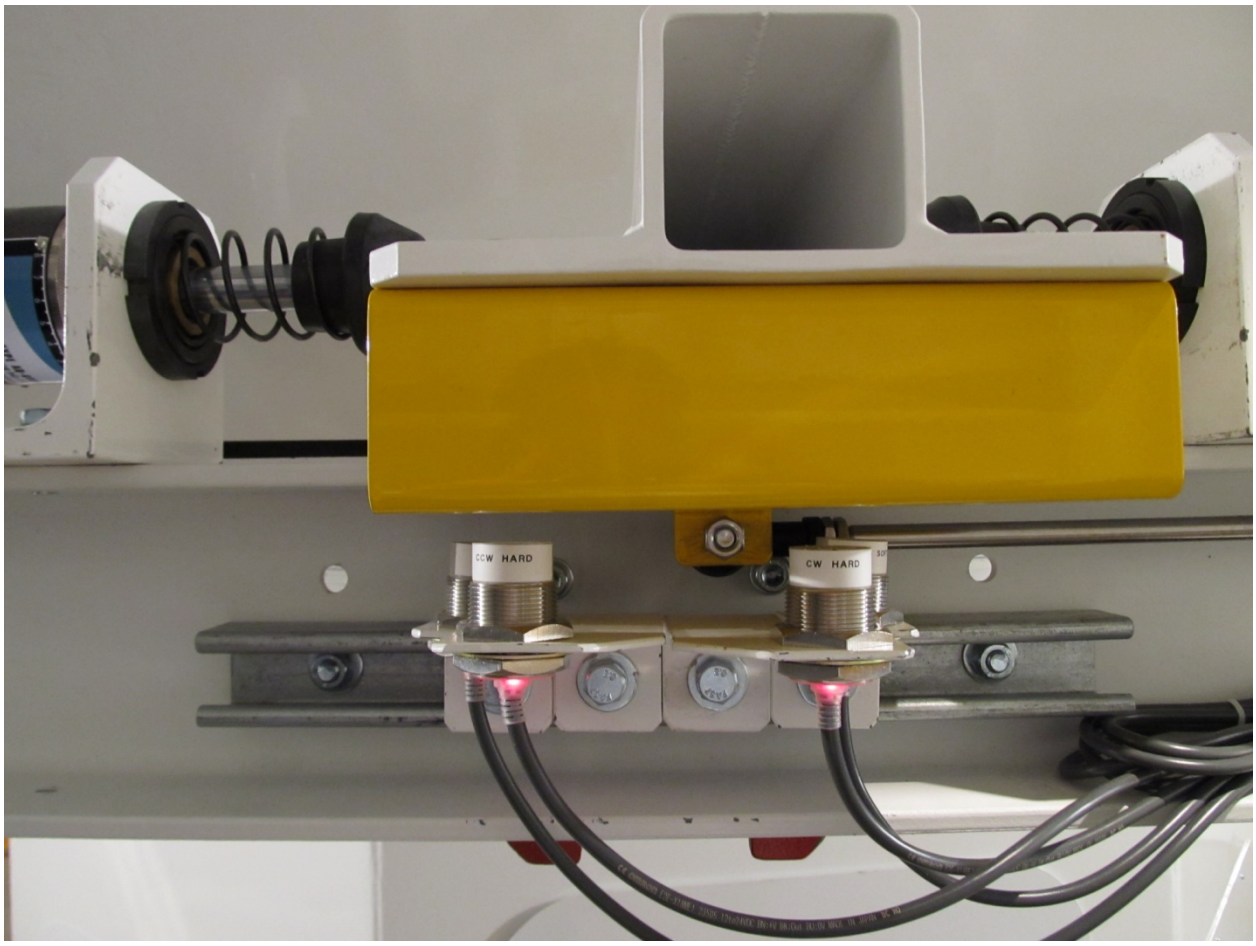

Figure 2.11*: Shock absorbers.*

In addition, a pair of shock absorbers (CJAC AD64050) is located on either side of the trailer hitch and will cushion the impact if there is a large speed differential between the telescope and dome at the point where one of the limit switches is tripped.

Note that prior to the installation of the spectrometer in mid-2011, the linear potentiometer was the only mechanical link between the telescope and dome. To ensure that the rotating dome does not couple any significant torque or vibration into the telescope, it is important that the potentiometer's tap continues to move smoothly in and out of its barrel. This should be part of the periodic preventive maintenance checks.

Also note that when power is first enabled to the dome azimuth drive controller (see 2.2.1.4.1.3), it will check the reading from the linear potentiometer to see if the dome and the telescope are within a specified centering tolerance. If not, the controller will command a small rotation of the dome in the appropriate direction until the dome is centered within that tolerance. This power-on procedure assumes that there is no significant mechanical coupling between the telescope and the dome, so that a small rotation of the dome will not induce *any* rotation of the telescope in azimuth.

However, if the potentiometer's tap were to get stuck in its barrel, or if the large cable bundle loop that feeds into the spectrometer (see 2.2.1.4.1) becomes too rigid, then this power-up procedure can lead to a runaway condition, causing the dome and telescope both to rotate in azimuth, presumably until the dome slams into one of its end-of-travel limits. Because the stuck tap or the overly-rigid cable loop creates (an unintended) mechanical linkage between the telescope and the dome, when the dome rotates (in an effort to improve its centering relative to the telescope), it effectively drags the telescope along with it, so that the centering error is never reduced, and this cycle continues until the end-of-travel limit is tripped.

2.2.1.4.1.3 Baldor Vector drive controllers

The dome azimuth servo system employs two Baldor Model ZD18H203-E Vector drive controllers, one for each of the two dome drive motors. One motor controller is configured as the master, the other the slave. Both controllers are located in the Dome Control Cabinet on level one of the dome. As configured, these controllers can only be configured using the membrane keypad and LCD display on their front panels. Operational status can also be monitored via that display.

Since their installation in 2006, each controller has failed once (the slave in May 2011, the master in July 2012) and each was either replaced or repaired; at least one of the refurbished replacement units had an extremely noisy power supply, which radiated noise throughout the dome control cabinet, causing problems for other circuitry in that cabinet. Although Baldor no longer manufactures this model, there is a third-party vendor (Practec) that still repairs this model for about $500 and sells refurbished units for about $1000. At present, we do not own any spare controllers. Baldor makes a replacement unit, the model VS1GV63-1B, which costs $1165. Were we to replace our obsolete Model 18H drives with the newer model, we would order it with an Ethernet interface so as to enable it to be configured and monitored remotely, rather than via the front panel keypad and LCD.

2.2.1.4.1.4 Dome drive motors

There are two dome azimuth 3-phase AC drive motors spaced 180 degrees apart and operated in a master/slave configuration. One end of the motor shaft connects to a reduction gearing whose output connects to the drive wheel, while the other end connects to an encoder via a flexible coupling; the encoder provides motor velocity feedback to the Baldor controller which operates the motor.

The motors are Baldor 2KW 1750 RPM flux vector motors ZDNM3661 wound for 3-phase 240-volt power and fitted with Mary M32 brakes. The motors are coupled to MR2180 16.0:1 coaxial gear reducers connected to Rossi MRC21140-112 gearboxes.

2.2.1.4.2 Dome Shutter Drive Controllers

The dome shutter servo system employs two Baldor FDH2A07TR-RC20 Flex Drive II AC servo drive controllers, one for the front shutter and the other for the rear shutter. We have no spares for these controllers.

2.2.1.5 Dome control computer

The overall control of the various dome systems is provided by software written by EOS which runs on the Dome Control Computer (DCC), which is located on the level one dome floor just behind the stairs to level two. The DCC is an industrial, rack-mounted PC that runs Windows XP. It contains a CANbus interface board that interfaces to the various pieces of CANbus hardware (e.g., temperature, humidity, and dew detectors, the three fan cooling units, etc.) that are located throughout the dome. We have no spare for either the DCC or the CANbus interface board it contains.

2.2.1.6 Dome Cooling

Dome cooling is provided via chilled water that is fed to the dome from an external Multistack chiller unit located just below the Astrograph Dome. The chilled water is distributed to 3 fan cooling units, one for each of the three levels of the dome. The fan cooling unit for level 1 is mounted on the outside of the dome's ring wall and has its air intake and output located on the inside of the ring wall. The chilled water for level 2 goes up through the central pier and connects to a fan cooling unit located inside the metal structure that sits atop the concrete pier and which supports the azimuth drive motor and the stationary portion of the telescope's azimuth bearing. The chilled water for level 3 goes through the azimuth cable wrap and feeds a fan cooling unit that hangs down from the inside top of the rotating hemispherical dome.

In addition, on level 3 of the dome are two large fans, which can be used to improve air circulation within that level of the dome. Both of those fans and all three fan cooling units are controlled by EOS software that runs on the Dome Control Computer.

2.2.1.7 Dome UPS

Most of the critical equipment within the dome is powered by an APC Symmetra LX constant-online UPS. This UPS is sized to provide sufficient backup power so that the dome shutters can be closed in the event of an extended power outage; if external AC power remains off for more than 3 consecutive minutes, unless manually overridden, the dome control computer will automatically command the dome shutters to the closed position.

The UPS also provides power conditioning for the rather noisy and distorted single-phase AC power that feeds the APF Dome, and it enables observing to continue uninterrupted despite the occasional brief power outages and brownouts that sometimes hit Mt. Hamilton.

Note that all of the Lick Observatory facilities on Mt. Hamilton are served by a 1+ Mw diesel power generator that automatically starts up within a few seconds if the power from our local utility (PG&E) fails. Provided that this generator and its electrical switch gear is functioning correctly, the key role for the APF UPS is to keep the facility powered during the several seconds that it takes for Mt. Hamilton generator to come online.

The only APF equipment that is not backed up by the UPS are the various fan cooling units, the external Multistack chiller that provides the chilled water to the dome, and the NESLAB chiller that provides chilled coolant to the spectrometer. As a result, if the facility is run from UPS power for an extended period (e.g., because the generator fails to start up), level 1 of the dome can rapidly become overheated because of the loss of cooling coupled with the considerable amount of heat generated by the multiple racks of electronics. In such situations, because the NESLAB chiller will also be off, the temperature inside the spectrometer can also start rising.

Once utility power is restored, both the NESLAB chiller and the external Multistack chiller will normally restart on their own. However, the fan cooling units inside the APF Dome will not restart until they are explicitly commanded to do so; at present, this requires manual intervention, which can be done remotely. Our APF monitor task software needs to be updated so that if it detects that the fan cooling units have tripped off and that cooling is needed, it will command those units to restart once utility power has been restored.

2.2.1.8 Dome lighting and surveillance cameras

Overhead lighting is provided on all three levels of the APF Dome. The lighting for each level can be controlled individually using toggle-style light switches located at various points (typically near the stairways) in the dome.

All of the lighting on all three levels can be globally switched on or off via the EOS dome control software that runs on the DCC, thus enabling the lighting to be controlled remotely; that software does not provide individual control of the lighting on each level, nor can it detect whether or not any of the lights in the dome have been switched on or off.

UCO has also installed eight closed-circuit video surveillance cameras at strategic locations throughout the dome, and these can be used to provide direct feedback as to whether the lights inside the dome are turned on or off. These cameras also enable us to view:

- the inside of the dome shutter (where the panels meet when closed)
- the primary mirror cover
- the “trailer hitch” assembly

- the areas in the path of the telescope tube
- the outside door

Six of the eight cameras are black-and-white IR cameras equipped with IR LEDs that illuminate their respective fields of view and which can produce useful images at night when the dome's interior lights are turned off. However, since the IR LEDs create light pollution within the dome, these cameras are only switched on if there is a need to check on the status of the areas that those cameras monitor. The other two cameras are color optical cameras and produce useful images only during the daytime or when the dome interior lights are turned on.

Two of the cameras have microphones attached to them and enable us to monitor nearby sounds. One microphone is mounted close to the dome shutters, while the other is mounted adjacent to the elevation drive motor.

The images from all eight cameras are fed to a PC-based digital video recorder, which maintains a circular buffer of the images from the last 24 hours. By replaying images recorded to the DVR, we can check for unusual events that might have occurred during the previous period. This DVR provides a web-based interface that enables us to remotely view and listen to either the live or recorded video and audio.

2.2.1.9 Dome Safety Systems (E-stop)

There are eight Emergency Stop (E-Stop) buttons located at strategic locations throughout the dome. These are large, red, illuminated buttons that are easy to hit, and they remain physically engaged until the button is pulled back out; while engaged, the button will flash on and off. In addition, there is a Software E-stop that can be activated either by clicking on the Software E-stop button within the Dome Client GUI or by writing to the ESTOP keyword in the eosdome keyword service, either of which can be done locally or from a remote location.

Once any E-stop button (including the Software E-stop) is tripped, this puts the dome and telescope into an E-stop state and removes all power from any motorized devices that are part of either the dome (dome azimuth drive, shutter drives, and vent doors) or the telescope (azimuth drive, elevation drive, M2 actuators, tertiary drive, and M1 cover) . At present, the E-stop state does not disable any of the motorized devices that are contained within the APF Spectrometer, nor does it disable the various cooling systems or fans.

Once the dome and telescope are in the E-stop state, that state cannot be cleared until the E-stop button that triggered that state has been manually reset; once that is done, the E-stop state itself must be manually reset, either via the Dome Client GUI or the by writing to the appropriate keyword. Once the E-stop state is cleared, then the dome azimuth drive power

must explicitly be re-enabled (either via the GUI or by writing a keyword) before any of the motorized devices on the dome or telescope can be operated.

Since there are currently no brakes (nor reduction gears) on either the azimuth or elevation axes of the telescope, those axes are not restrained by their respective drive motors and are thus free to rotate once an E-stop has been set. If the primary mirror cover is open, the telescope tube is well-balanced, but if that cover is closed, then the tube is quite bottom heavy. In the unusual case where the mirror cover is closed and the elevation axis of the telescope is servoing to a position near the horizon, if an E-stop occurs, the top of telescope tube will rapidly accelerate toward the zenith (because of the tube imbalance) and will run into the end-of-travel shock absorbers for that axis. To avoid this situation, if a maintenance activity requires positioning the telescope to the horizon with the mirror cover closed, the telescope should be pinned at that position during that activity. Similarly, if deteriorating weather conditions prompt closing of the mirror cover and dome, the telescope should be slewed to the zenith prior to closing the mirror cover. See Section 5 for additional details regarding APF Safety Systems.

2.2.1.10 Dome lightning suppression system

The Lightning protection system for the APF telescope conforms to Australian Standard AS/ NZS 1768:2003 which is comparable to the United States standard NFPA 780. Both of these standards specify design and installation guidelines for commercially available components to protect against lightning damage. The number, size and locations of the terminals were optimized using a rolling sphere analysis as described in the "Icestorm II Lightening Protection-09 January 2006" report.

The APF dome has six air terminals (three on each side of the shutter) with supporting tripods that are attached to 0.5" diameter down conductors. The down conductors enter the dome at four locations near the base of the arch girders and are connected to brush assemblies that run on the vertical surface of the azimuth ring beam. The ground path is completed through the ring beam and its steel support structure which is connected to the master electrical ground grid at four locations. The installation details for the system are described in the "Lick APF Lightening Protection Scope of Work RFQ". During assembly on site the "M" head described in the project scope of work was rejected by UCO and replaced with a custom brush assembly provided by EOS to increase the current carrying capacity of the interface at the ring beam.

2.2.2 Dome Shutters Operating Modes

As described in section 2.2.1.1., the dome shutter consists of two panels: a front shutter panel and a rear shutter panel. Each shutter panel has its own servo controller, and each panel can be moved independently of the other; it is also possible to move both panels in unison.

When the dome shutter is closed, the rear edge of the front shutter panel and the front edge of the rear shutter panel meet at a dome shutter angle of 58 degrees. This also corresponds to the "home" position for each shutter panel. Thus, the dome shutter can be commanded to close either by commanding both panels to drive to "home" or to drive to 58 degrees.

When the shutter is opened for observing, it can be operated in two distinct modes. Objects whose position in the sky is between 41 and 81 degrees of elevation can be observed in either mode; objects whose position is above 81 degrees elevation should be observed in split-shutters mode, while those with positions below 41 degrees elevation should be observed in up-and-over mode. Shutter panel positions are expressed as dome shutter angles.

For additional details, see the EOS Dome vignetting report.

2.2.2.1 Split-shutters mode

In the split-shutters mode:

- the front shutter panel is moved forward so that its front edge is at the forward end of travel limit for this panel, which puts its rear edge at a dome shutter angle of 30 degrees.
- the rear shutter panel is moved towards the rear so that its front edge is at a dome shutter angle of 110 degrees.

When the dome shutter is operated in split-shutter mode, the telescope can observe objects positioned between 41 and 90 degrees of elevation. In this mode, observing objects at elevation angles below 41 degrees will cause the object to be vignetted by the front shutter panel.

2.2.2.2 Up-and-over mode

In the up-and-over mode (also referred to as the non-split-shutters mode):

- the rear shutter panel is moved all the way towards the rear so that its rear edge is at the rear end of travel limit for this panel; this puts its front edge at a dome shutter angle of 154 degrees.
- the front shutter panel is moved so that its rear edge meets the front edge of the rear panel at a dome shutter angle of 154 degrees.

When the dome shutter is operated in up-and-over mode, the telescope can observe objects between 15 and 81 degrees of elevation. In this mode, observing objects at elevation angles below 15 degrees will cause the object to be vignetted by the bottom edge of the dome shutter opening while observing objects at elevation angles above 81 degrees will cause the object to be vignetted by the front edge of the front shutter panel.

2.2.2.3 Wind-shielding option

In either the split-shutters or the up-and-over modes, the dome shutter opening is quite long and a large cross-section of the telescope is exposed to the wind. On windy nights, the telescope experiences significant windshake when observing objects whose position in the sky results in the dome shutter pointing into the wind. To reduce the exposure of the telescope to the wind, a wind-shielding option can be enabled.

At present, this wind-shielding option is only available when the EOS Observatory Server process is running. Due to several unresolved problems with the implementation of the server, it is not normally running during observing. As a result, the wind-shielding option is not currently available to observers.

However, it would not be difficult for us to implement a wind-shielding option that would operate without the Observatory Server. We have not yet done so, both due to lack of resources and because of concerns about the safety of this option. In particular, we are concerned about the safety of moving the shutter panels above an unprotected, up-looking primary mirror. Because of that concern, our current practice is to only move the dome shutter panels when the primary mirror cover is closed.

Regardless of whether we decide to implement our own wind-shielding option or use the one provided by the Observatory Server, the result of enabling the wind-shielding option depends on whether the split-shutters or up-and-over mode has been chosen.

2.2.2.3.1. Wind-shielding in split-shutters mode

If wind-shielding is enabled in split-shutters mode, both shutter panels actively track the path of the telescope so that the rear edge of the front panel is kept positioned just below the telescope's incoming beam while the front edge of the rear shutter panel is maintained just above that beam. As a result, the dome shutter opening will be kept as small as possible while ensuring that the telescope's beam is not vignetted.

2.2.2.3.2. Wind-shielding in up-and-over mode

If wind-shielding is enabled in up-and-over mode, both shutter panels actively track the path of the telescope, but they act as if they were a single panel whose front edge remains just above the telescope's incoming beam. In this mode, for objects observed at elevation angles above 15 degrees, the wind-shielding option will be less effect because the portion of the dome shutter area that lies below the telescope's incoming beam will remain exposed to the wind.

2.2.3 Remaining Problems and Concerns

Aside from the problems enumerated below, the APF Dome hardware and software is fully operational and capable of supporting the functions required for local, remote, and robotic operations.

2.2.3.1 Wind can blow rain under back of dome shutter

Since its installation in 2006, the APF Dome has had problems with water leaks. Most of these have been repaired, and the dome has remained dry through several heavy rain storms.

Unfortunately, there is still a problem in which wind-blown rain can blow under the back of the dome shutter slot near the top of the dome; from there, it can flow along the track for the dome shutter panels and eventually drip down into the dome, quite often directly onto the mirror cover for the primary mirror. Since the 4 petals of that cover do not form a water-tight seal, rain water can (and has) found its way onto the primary mirror.

***Figure 2.12**: Shutter in open position showing chunks of ice.*

Worse yet, in cold weather, water that leaks into the dome shutter track via this path can freeze, forming chunks of ice that are not easily seen until the dome shutter panels are opened at least partway (see photo above). Were such ice chunks to break free over an open primary mirror, serious damage could occur. Because of this concern, we have deferred the use of the wind-shielding option until this leak is resolved.

2.2.3.2 Noise on linear potentiometer signal

As discussed in section 2.2.1.4.1.2, the dome azimuth rotation servo system uses a linear potentiometer attached to the "trailer hitch" assembly to measure the residual rotational offset between the telescope and the dome. The analog signal for this potentiometer runs from the trailer hitch on level 2 of the dome down to the dome control cabinet on level 1. Once inside that cabinet, this signal encounters an extremely noisy electrical environment. Inside the cabinet, this signal is processed by a noise filter board (the small green circuit board just to the left of the bank of relays) which is supposed to remove the high frequency noise from what should be a very low frequency signal.

Unfortunately, this noise filter board does not appear to work correctly, and the signal from the potentiometer (as received at the input to the Baldor Vector drive controller for the master dome azimuth drive motor) is extremely noisy. As a result of the noise on this signal, when the azimuth drive power is enabled and the dome azimuth drive motors are servoing, they make a horrible chattering sound as they oscillate about their commanded position. This can't be good either for the motors, the gear boxes, or the Baldor controllers. Because it greatly increases the background noise level in the dome, it makes it hard to hear other noises that might indicate problems in other systems.

2.2.3.3 Cold Dome azimuth drive can't keep up with telescope

After working reliably for over two years, starting with the onset of cold weather in late 2011, the dome azimuth drive started having problems keeping up with the telescope drive when the telescope was slewed in azimuth at its full rated speed of 4 degs/sec. During such slews, as the dome fell behind the telescope, the trailer hitch mechanism would become sufficiently offset to trip one its limits, thereby stopping the motion of the telescope and dome and aborting the slew operation. As a work-around to this problem, during periods of cold weather we temporarily reduced the telescope's top azimuth slew speed from 4 degs/sec to 3 degs/sec. Resolving this problem is still a pending warranty item with EOS.

During his service visit in late July, Jak Gray tried to address this problem. Unfortunately, since cold weather is not common on Mt. Hamilton in July, he could not reproduce the problem. However, the replacement of the Baldor Vector drive controller that failed during Gray's visit might have resolved this problem. Alternatively, this problem might have been solved by the BEI

anti-dither device which now filters the telescope azimuth encoder signals before they are input to that Baldor controller. Now that cold weather is once again at hand, we should soon know whether or not this problem has been resolved.

2.2.3.4 Repeated failure of dome rear shutter home switch

Each dome shutter panel has an associated “home” switch which is activated by a mechanical detente whenever the panel reaches the home position, i.e., a dome shutter angle of 58 degrees. While these switches are exposed to the weather and can get quite wet, they are supposed to be water tight and rated for outdoor use.

Unfortunately, the home switch for the rear shutter panel has failed twice due to problems with corrosion. When it fails in this way, the switch reports that the shutter panel is at the home position, regardless of its actual position. Because of the lack of exceptional-condition handling in the EOS software the controls the dome shutter panels, the consequences of these home switch failures have been quite serious.

As noted in section 2.2.1.1 above, the standard procedure for closing the dome shutter is to command both dome shutter panels to their home positions. The procedure that performs that function uses the following logic:

- Step 1: If the home switch for the rear shutter panel indicates that the panel is already at home, then that panel should be driven in the reverse direction until the switch no longer reports that the panel is at home.
- Step 2: If the home switch for the rear shutter panel indicates that this panel is not at home, then this panel should be driven forward until that home switch trips.

Unfortunately, when the home switch for the rear shutter panel fails such that it always reports that this panel is at home (regardless of its actual position), then step 1 of the homing procedure will cause this panel to be driven in the reverse direction until it hits the final end-of-travel limit (that limit can only be cleared by manually driving this shutter panel using the local control push-buttons, and thus requires on-site intervention). Hence, requests to close the dome shutter result in the rear shutter panel being driven fully open.

2.2.3.5 Repeated failure of dome azimuth Baldor vector drives

As noted in section 2.2.1.4.1.3 above, each of the two Baldor Vector drive controllers for the dome azimuth drive system has failed. In both cases, we believe the failures occurred in the power supplies for these units, although we are certain of that only with respect to the failed master drive. We are not sure what is causing these failures.

2.2.3.6 Unreliable operation of vent door #3

One of the four vent doors (door #3) still fails to operate reliably, despite repeated attempts to re-align it. At present, this vent door is tagged out.

2.2.3.7 Excessive light leaking into dome during daytime

The dome is not very light tight, so quite a bit of light leaks in during the daytime. Some of this light finds it way inside the spectrometer, whose enclosure is also not completely light tight.

2.2.3.8 External chiller for dome cooling trips off periodically

The external chiller that provides chilled water to the fan cooling units for the APF Dome will trip off periodically; when it does, it requires manual intervention by on-site personnel to restart it. The APF monitor task does not reliably detect when the chiller trips off, so alarms do not get issued until level 1 of the dome starts to overheat.

2.2.3.9 Fan cooling units (FCUs) do not restart after power loss

2.2.4 Next Steps

2.2.4.1 Add or improve weather seals on dome shutter

2.2.4.2 Improve filtering of linear potentiometer signal

2.2.4.3 Continue to pursue 2.2.3.3 with EOS (warranty item)

2.2.4.4 Inspect rear shutter limit switch every six months

2.2.4.5 Find funds to upgrade Baldor vector drives to new type

2.2.4.6 Improve mechanical alignment of vent door #3

2.2.4.7 Plug gaps and blacken non-opaque areas of dome wall

2.2.4.8 Add hardware/software to detect chiller tripping off

2.2.4.9 Update monitor software to restart FCUs after outage

2.3 Overall Concerns

The following concerns apply to both the APF Telescope and Dome:

2.3.1 Lack of Critical Spare Parts

We lack spare parts for most of the critical components of the telescope and dome. A failure of any of these components could render the facility inoperative for weeks or months. Since many of these parts are, or will soon be, out of production, the longer we wait, the more difficult it will become to obtain the needed spares. In December of 2011 we gathered costing and availability data for all telescope and dome spares. The detailed list of all the critical spares as well as a complete spares list and pricing for the spectrometer can be found here. The pricing summarized in the spread sheet does not include sales tax and should be inflated to 2012 dollars.

2.3.2 Difficulty of Maintaining/Replacing Obsolete Hardware

Much of the hardware in the APF control systems is obsolete, and support for these systems will become more difficult to obtain over time. For example, all of the PC-based systems delivered by EOS and EOST run Windows/XP. Several of the EOST-designed circuit boards use components that are no longer manufactured.

2.3.3 Inadequate Documentation from Dome/Telescope Vendors

The set of mechanical drawings and electronics schematics that we received from these vendors is incomplete and/or not in the proper format; in many cases, the drawings and schematics we received do not appear to correspond to our hardware. Repeated efforts to remedy this situation have proven unsuccessful.

2.3.4 Lack of Access to Source Code (locked in Escrow)

All of the source code for the EOS and EOST application software and all of the firmware for the field programmable gate arrays on the EOS-built circuit boards is currently locked in escrow and is not accessible to us. Unless these vendors go out of business altogether, we will not have access to this escrowed software and firmware until 2019. Because we lack access to this information, many aspects of the EOS and EOST systems remain a black box; that strongly limits our ability to debug or correct problems with these systems.

The only controller firmware to which we have access is that for the Delta Tau PMAC board in the telescope control computer.

2.3.5 Inadequate Staffing (and Diminishing Availability of Staff)

An ongoing problem with the APF Project over the past several years has been inadequate staffing to complete the various tasks involved in making this facility fully

operational. Partly, that is due to UCO having to take on numerous tasks which should have been completed (under the terms of the contract) by either EOST or EOS. But it is also due to more widespread budget problems and the competing demands of other projects.

APF has had no full time staff assigned to it for the last several years, and that includes its Project Manager. As a result, pending problems have typically been addressed slowly and in a piecemeal fashion. That is not the most effective strategy for bringing a complex project like APF to a successful conclusion.

These staffing problems are likely to become more acute as key, long -term members of the APF project team either have, or will soon be, retired, including:

Jeff Lewis, UCO Instrument Lab Supervisor, retired July 1, 2012

Bob Kibrick, Research Astronomer, mostly-retired July 1, 2012

Barry Alcott, Lick Electronics Lab Supervisor, retiring December 10, 2012

There are no plans to refill any of these 3 positions. In addition to these retirements, there have been other, budget-driven reductions in staffing both at the UCO Technical facilities on the UCSC campus and at the facilities on Mt. Hamilton; additional retirements that will occur next year.

These staffing issues represent a significant challenge for the continuity of the project and its overall completion

3 Spectrometer

The heart of the APF is the Levy Spectrometer (hereafter referred to as the Levy). The Levy is a high-resolution prism cross-dispersed echelle specifically optimized for high precision radial velocity research. The basic optical design concept of the Levy follows that used for MIKE, the Magellan Inamori Kyocera Echelle (Bernstein et al. 2003 SPIE 4841, 1694) on the Magellan Clay telescope of The Carnegie Institution of Washington. An RV-optimized version of MIKE called PFS (Planet Finding Spectrometer) was later commissioned at Magellan (Crane et al. 2010 SPIE 7735, 773553) and also uses an optical scheme quite similar to APF's Levy.

Unlike with precision radial velocity spectrometers at Keck and at Magellan, where midnight mountaintop temperature variations are typically < 20 ^{0}F throughout the year, the thermal environment atop Mt. Hamilton presents quite challenging large seasonal temperature variations. The following plot shows midnight temperatures as measured from the top of the 3-m dome from 2007-2011. Diurnal and seasonal variations of 60 ^{0}F peak-to-peak in the midnight temperature are the norm atop Mt. Hamilton. It is not uncommon on Mt. Hamilton (especially during the summer months) to see large temperature variations (of order 10 ^{0}F) over the course of a single night and quite often over time scales of 1 hour or less. Thus, a principle concern in the design of a highly stable spectrometer in the face of such temperature swings is holding the focus, scale, and spectrum position constant.

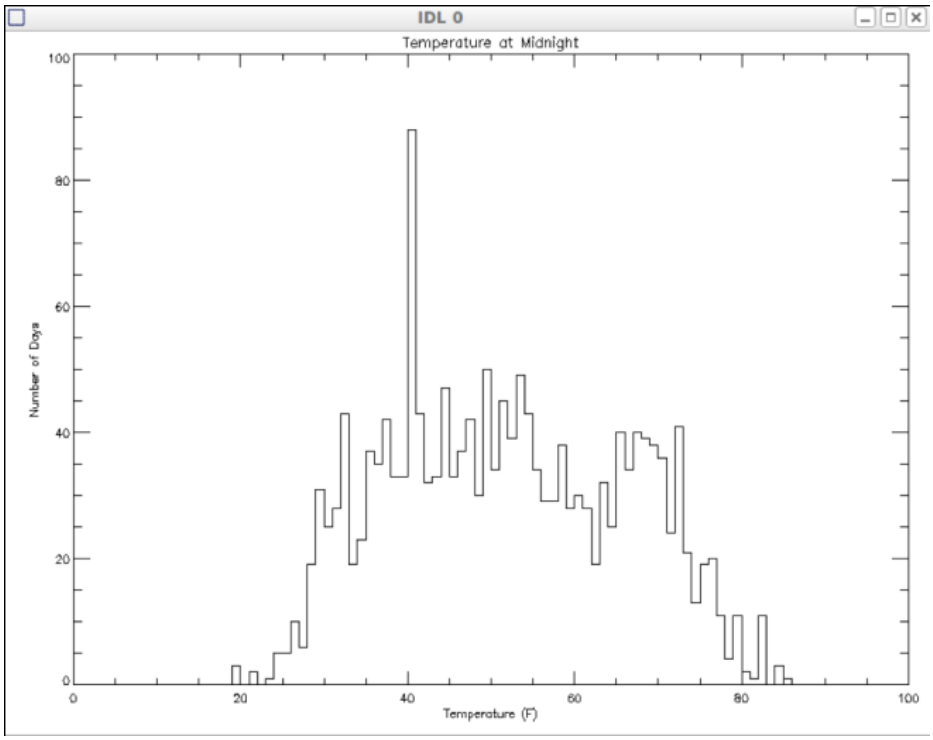

Figure 3.1*: Midnight temperatures as measured from top of 3-m dome from 2007-2011.*

With the Levy, this challenge was addressed on five different levels. The first line of defense against dome interior temperature swings was to incorporate a high degree of passive athermalization explicitly into the optical design. Every optical component in a spectrometer has a coefficient of thermal expansion (CTE) that makes its dimensions and sometimes even its optical properties vary with temperature. For glass lenses with curved surfaces, the radius of curvature of each surface varies with temperature, giving rise to optical power and focal length variations as a function of temperature. Also, both the thickness of any lens and its refractive index will vary with temperature. With curved mirrors, again, the curvature of its surface changes with temperature, giving rise to power and focal length changes. Spacings of optical components will also change due to the non-zero CTE of lens mounts. With replica gratings, where the rulings are typically a plastic structure supported by a low-CTE ceramic substrate, even though the substrate is chosen for its low CTE, that CTE is not identically zero at all temperatures. Expansion or contraction of that substrate changes the grating pitch, producing subtle shifts and dispersion changes in the spectrum.

Most all of these effects were considered in the optical design of the Levy using the thermal modeling tools available in Zemax. Certain optical components were allowed to piston with temperature by using strategic choices of lens mounting materials such that element spacing changes largely counteracted changes due to their CTE. The merit functions in Zemax

were set up to attempt to hold image focus, image scale, and spectrum position as constant as possible in the face of up to +/- 25 C changes in the temperature. In this way, a high degree of passive athermalization of the optical train was built into the optical system from the outset. The echelle grating and cross-dispersing prism were allowed to slightly tip/tilt/rotate to remove components of spectrum shift due to their CTE sensitivity. The spacings of other key lens elements or groups were also allowed to vary as required. These spacing changes were either designed into lens cells by proper choice of cell material, or were built into the determinate structure that supports the optical elements and/or cells. Most of the struts of the determinate structure that supports all the optics are made of Invar-36, a low-CTE metal. But in cases were movement of an element was required to stabilize against temperature, the strut either incorporated a link of higher CTE material, or was made completely of a high-CTE material to accomplish the requisite element motion. For example, the prism has two of its struts made of 17-4 Stainless rather than Invar to accomplish a very slight rotation of the prism with temperature. The most dramatic example of using the determinate structure for athermalization can be seen in the next figure that shows the 6 camera barrel support struts.

Here, the main body of the camera was required to piston some 1.3 mm with respect to other elements in the optical train. This pistoning was accomplished by fabricating these struts as a telescoping combination of Invar and Magnesium struts. Each camera body support strut is actually an outer Magnesium body, with an inner Invar return body, and an innermost Magnesium rod. The telescoping combination effectively acts as a very thermally active strut (with the high CTE of Magnesium) and of twice its overall length. All six of these telescoping struts were closely matched and perform together to provide the requisite 1-2 mm of thermal piston to stabilize focus and scale.

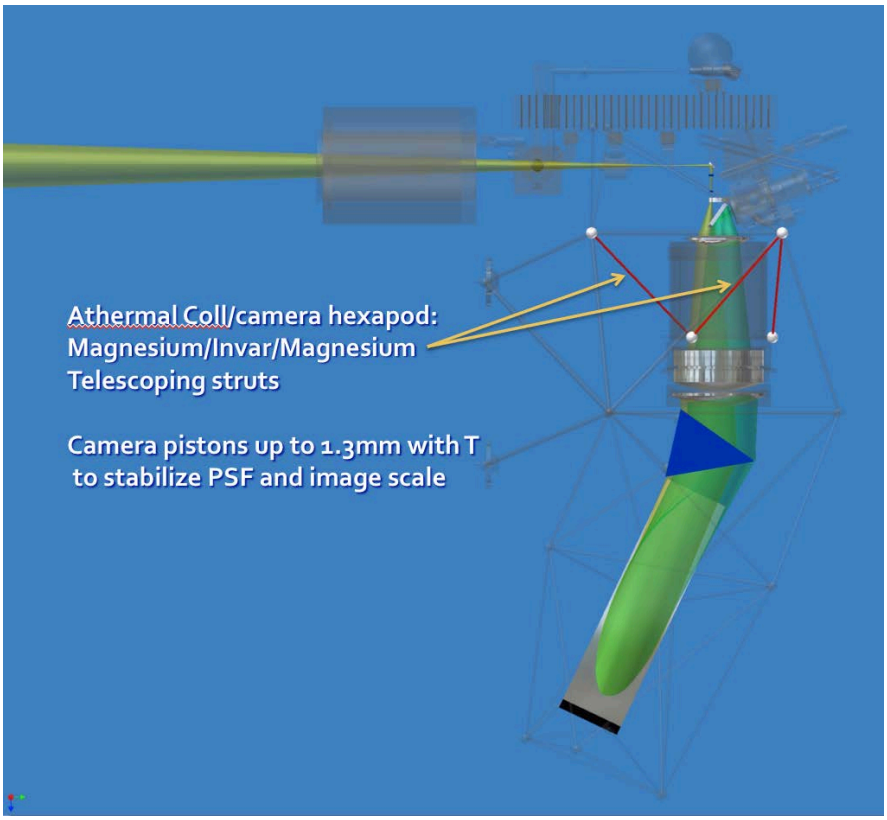

***Figure 3.2**: Piston athermalization of the main camera body.*

The next figure shows a schematic of some of the athermalization process. Here, the main body of the collimator/camera has to piston with respect to the field flattener. Various lenses within the main body's group also have to piston with respect to each other.

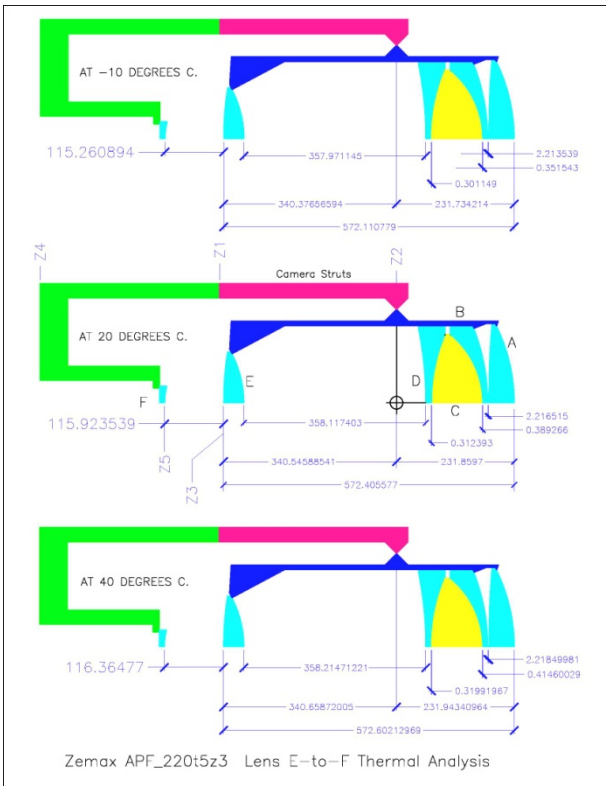

***Figure 3.3**: Representative schematic of some of the lens group athermalization process.*

The second line of defense against temperature-induced instability was to surround the spectrometer with a thermally insulated enclosure. The spectrometer housing is fabricated of insulating material with an R-value of 12. As much as is practical, it has been sealed to minimize air-intrusion and to minimize conductive heat shorts through the enclosure. This is a work in progress, and the electrical feed-thru panel was just recently remounted to reduce mechanical connections with the determinate structure.

The third line of defense was to actively control the temperature inside the spectrometer with heating/cooling loops, and with good air circulation to reduce internal temperature gradients. This is also still a work in progress. We have active cooling loops for the 50-watt flat field lamps housing, the UCAM CCD controller unit, and a heat exchanger. We also have an active heater loop that controls a 90-watt heater panel. We are currently in the process of tuning these cooling/heating loops and adjusting their power capacity through tuning of flow valves. We are also planning to increase the air circulation within the spectrometer. The following figure shows a recent example of some of this tuning activity from Oct 17-19, 2012.

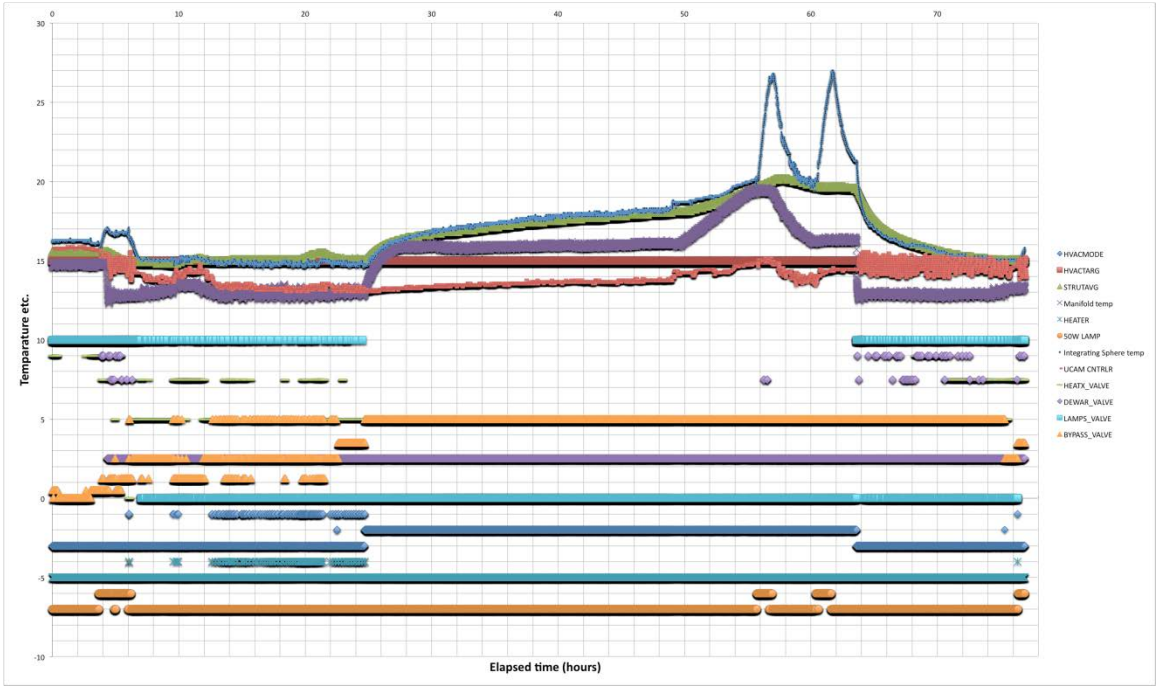

***Figure 3.4**: Example of thermal control tuning activity from Oct 17-19, 2012.*

The most thermally active members in the spectrometer are the 6 telescoping Magnesium struts that support the main body of the collimator/camera lens group. The 12-bit RTC temperature sensors on all 6 struts each have a temperature resolution of about 0.13C. We use the average of all six struts as our control metric, which should allow a factor of $\sqrt{6}$ increase in sensitivity, allowing ~ 0.05 C control.

Even though a very high degree of passive athermalization was achieved, it was not perfect. And because a commercial stage was being used for the dewar focus, that component would have been quite difficult to accurately model thermally. So the 4th line of defense was to be able to do a final tweak of the dewar focus using an extremely fine pitch focus stage that could be calibrated vs. temperature to correct any final uncompensated focus term. Accordingly, the stage pitch was set at 0.2 microns per step and this stage motion was used to do a final focus vs. temperature calibration. The calibration is nearly complete, but is still underway. The next plot shows our most recent calibration. Black points are the focus measurements, with the solid line showing the least squares fit. The data look noisy, but that is only because the focus resolution is so high. 100 steps of dewar focus stage travel corresponds to only 20 microns. The red squares show the focus behavior predicted by the Zemax model and indicate that our theoretical focus athermalization goal was achieved quite closely. Worst-case focus changes are only +/-20 microns over the 8-22 C operating regime, and are easily compensated with the focus stage and a temperature lookup table. In practice, we will control the spectrometer

temperature at some intermediate temperature, probably around 15-18 C, which can be held constant throughout the year.

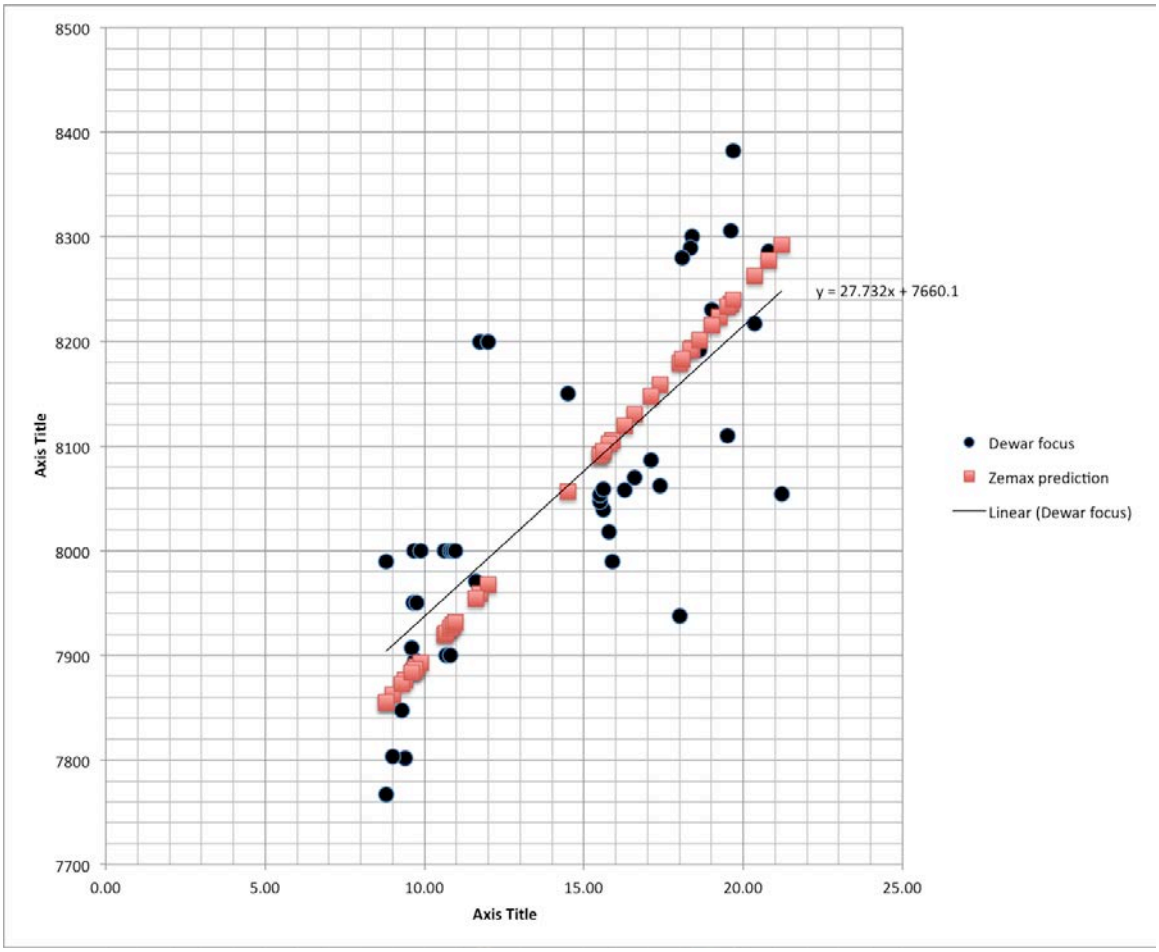

***Figure 3.5**: Spectrometer focus vs. temperature calibration.*

The 5th and final line of defense was to actively air condition the dome interior, to keep the interior dome temperature close to the midnight temperature throughout the day. The APF dome features powerful chiller units that air-condition the dome interior against daytime insulation heating, and large vent doors that facilitate rapid temperature equalization after sunset. The present AC system does fairly well, but currently lacks adequate capacity to keep the dome fully at midnight temperatures during warm summer days, or during sunny winter days that have very cold nights.

3.1 Basic Characteristics

The Levy spectrometer is a prism cross-dispersed echelle. It features a 200mm diameter collimated beam and was designed to achieve a theoretical “throughput” (resolution x slit width product) of about 110,000 arc-seconds. The optical train is supported within a determinate

structure Invar “space-frame” that has been optimized to provide a high degree of passive athermalization to hold constant focus and image scale. A 310-mm clear aperture all-dioptric f/3.17 lens group, used in double-pass, serves as both the collimator and camera. Cross-dispersion is done by a prism, also used in double pass.

- Collimator/camera- EFL 483.4 mm; f/3.17
- Grating- 41 gr/mm R-4 Echelle (214x840mm)
- Detector- E2V CCD42-90 (2048x4608)
- Pixel size- 13.5 microns
- Dispersion- 1.46 A/mm or 0.0197 A/pixel at 5500A
- Spectral coverage- 3743A-9800A with small gaps above 7500A
- Order separation- 8 arc-sec minimum at 7800A, increasing to 13 arc-sec at 3700A
- Max spectral resolution- ~150,000 at 5500A (0.5 arcsec slit)
- Overall system efficiency- ~17-35% at 550nm (preliminary)
- Thermal control- passive athermalization; insulated housing and active control
- Slit lengths and widths- 2x7.5, 1x3, 0.75x3, 0.5x3, 2x3, 1x2, 2x2, 1x4, 2x4 arcsecs
- Maximum slit length- 7.5 arcsecs (without order overlap)
- Filters- none
- CCD readout time- ~30s
- CCD readout noise- 3-4 electrons
- Arc lamps/wavelength calibrations- Thorium-Argon, 50W incandescent, Iodine absorption cell
- Doppler accuracy goal- <1 m/s
- Some fringing visible redward of ~6500A

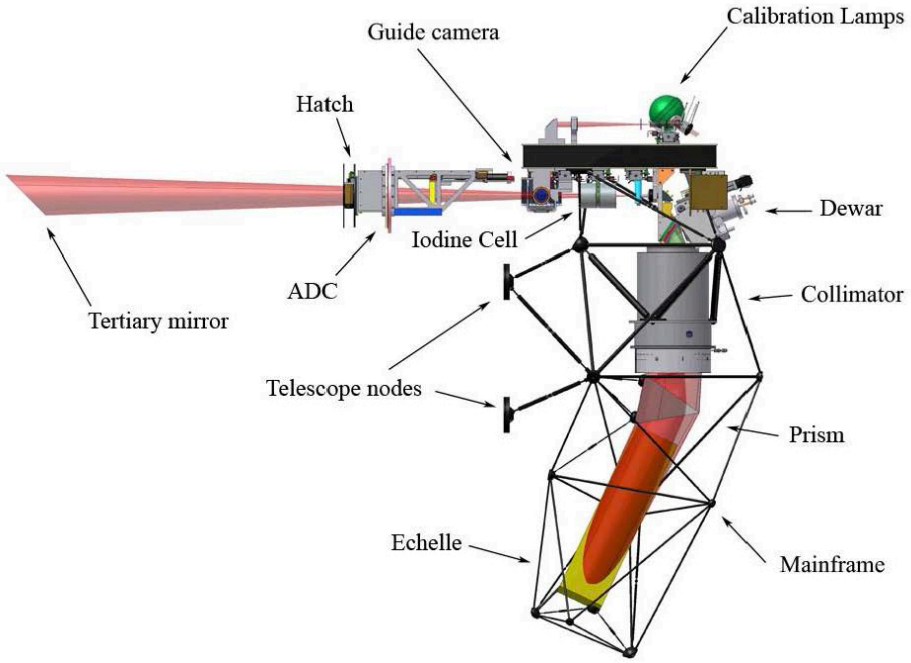

***Figure 3.6**: Schematic side-view of Levy Spectrometer.*

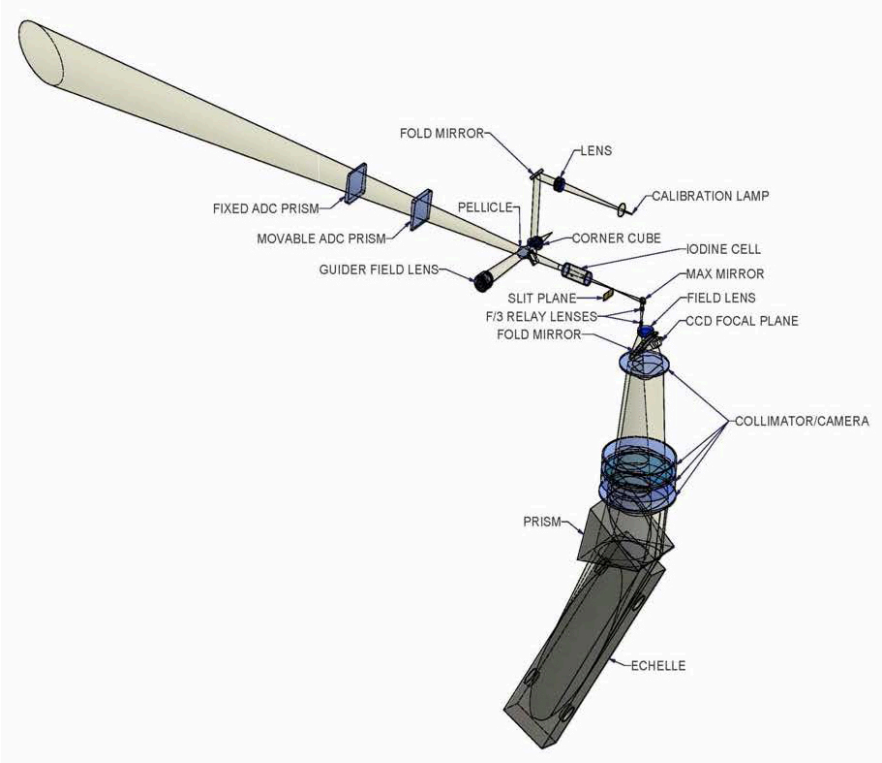

***Figure 3.7**: Schematic view of Levy Spectrometer.*

3.2 Description of the Optical Train

The Levy is mounted at a Nasmyth focus to avoid flexure due to forces that arise from changes in gravity loading. The optical train is supported by a determinate structure (colored rods in the next figure) of six hexapods. All of the pre-slit optics, calibration optics, and stage mechanisms are mounted to a commercial Invar optical table that also defines the determinate structure.

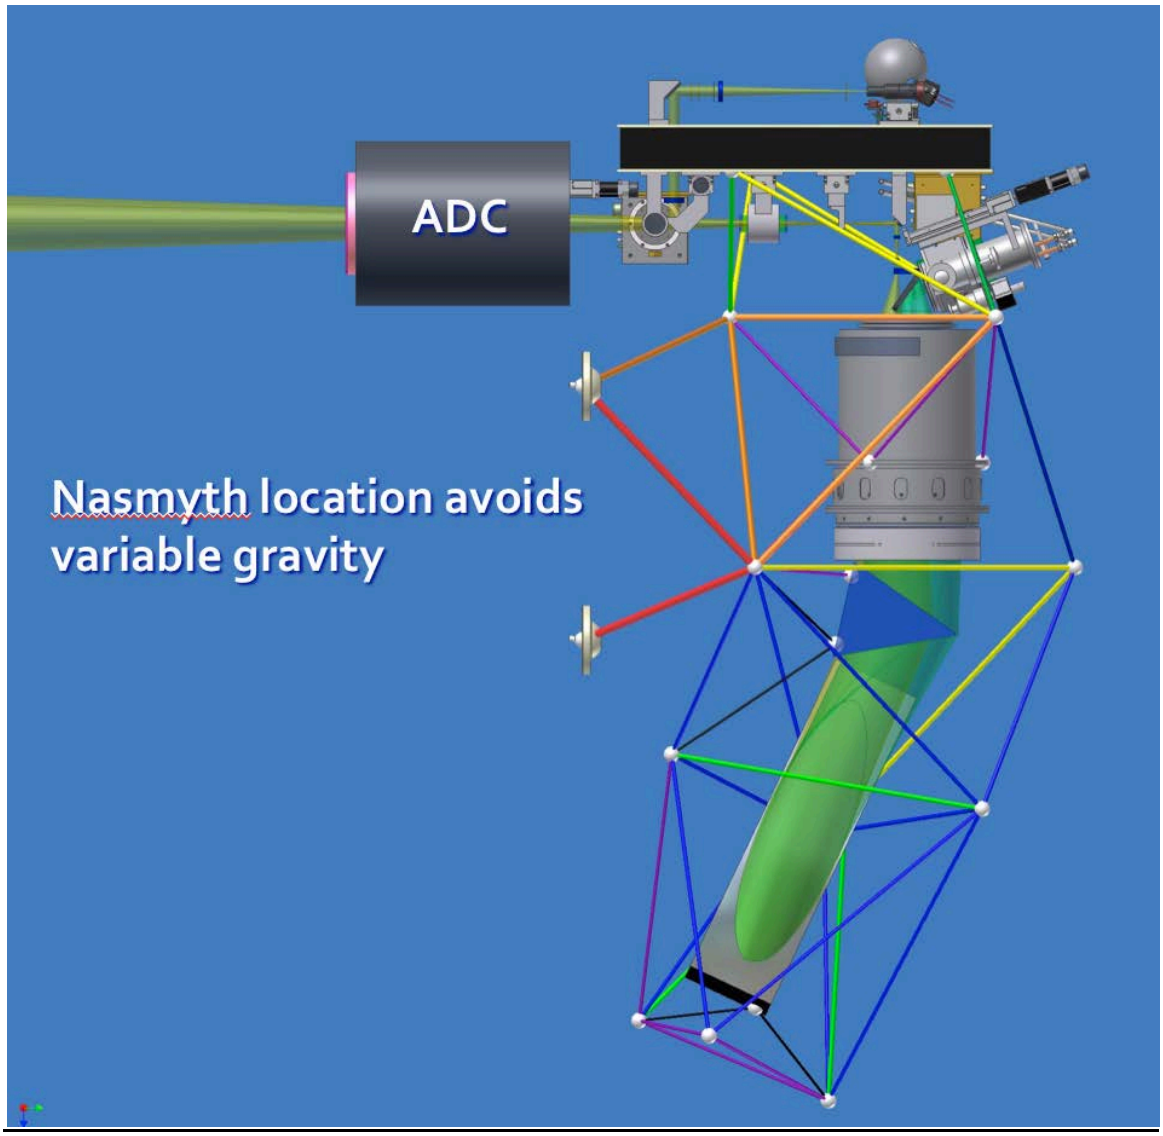

***Figure 3.8**: Spectrometer mounted at Nasmyth focus.*

3.2.1 Atmospheric Dispersion Compensator

Light from the telescope first encounters the front element of the ADC or atmospheric dispersion compensator. The ADC is a two-prism trombone-style compensator with fixed front prism and pistoning rear prism. Both prisms are carried on a mount attached to the elevation bearing and rotate with it. The front ADC prism is also used as the entrance window to the spectrometer and forms an effective seal to minimize dust and moisture intrusion.

Uncorrected atmospheric dispersion produces changes in the order separation on the echelle format, further complicating data reduction by introducing complexity into flat fielding and order extraction tasks. Uncorrected atmospheric dispersion also complicates guiding. The following figure shows the effects of uncorrected atmospheric dispersion. The typical science slit is 1 x 3 arcsecs, much smaller than the rectangle shown.

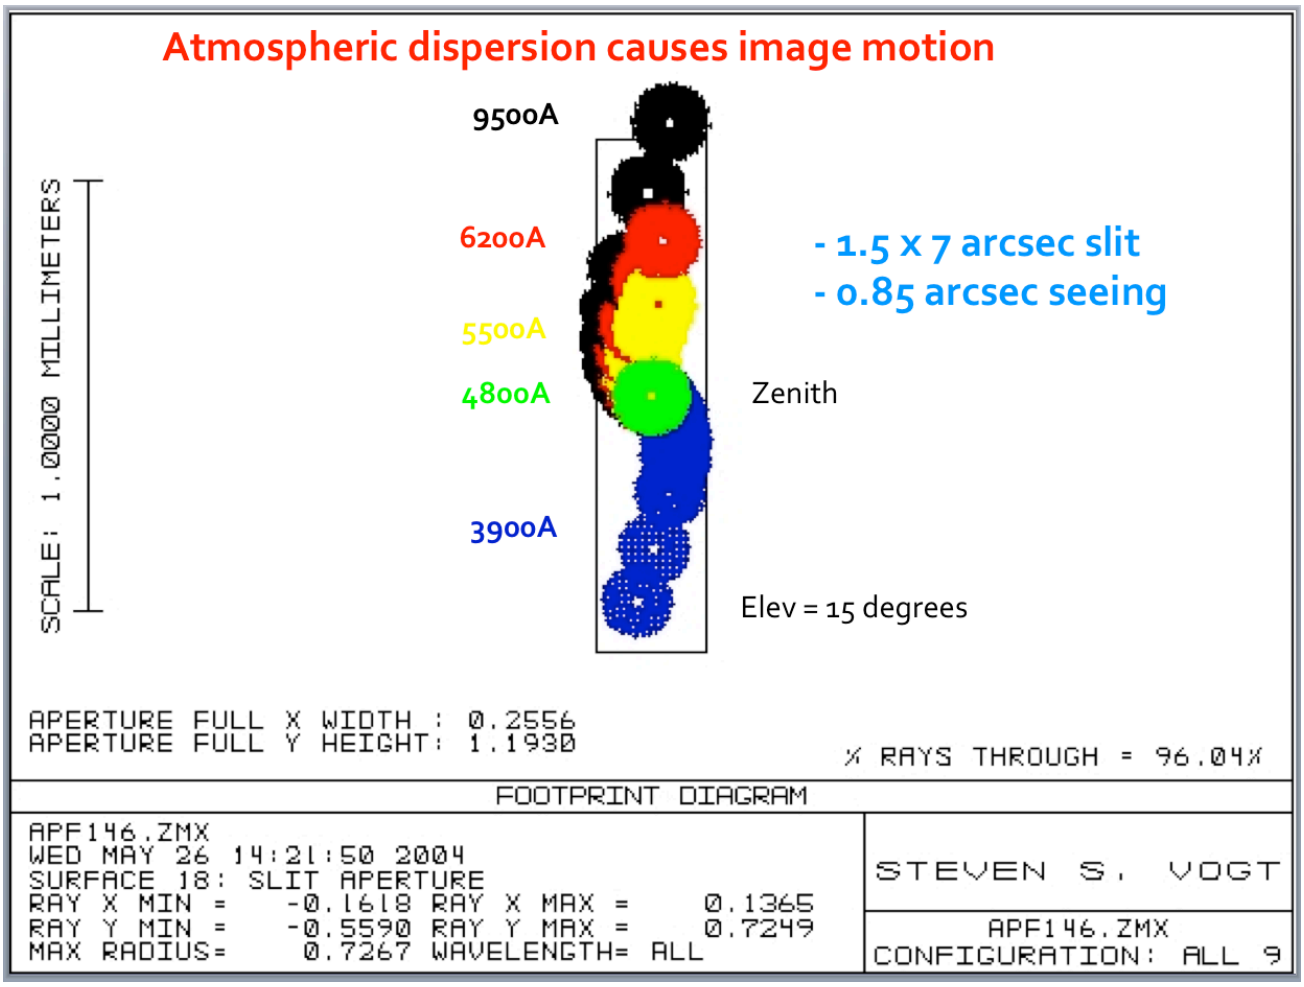

***Figure 3.9**: Effects of uncorrected atmospheric dispersion.*

The next figure shows the two ADC prisms on the trombone slide.

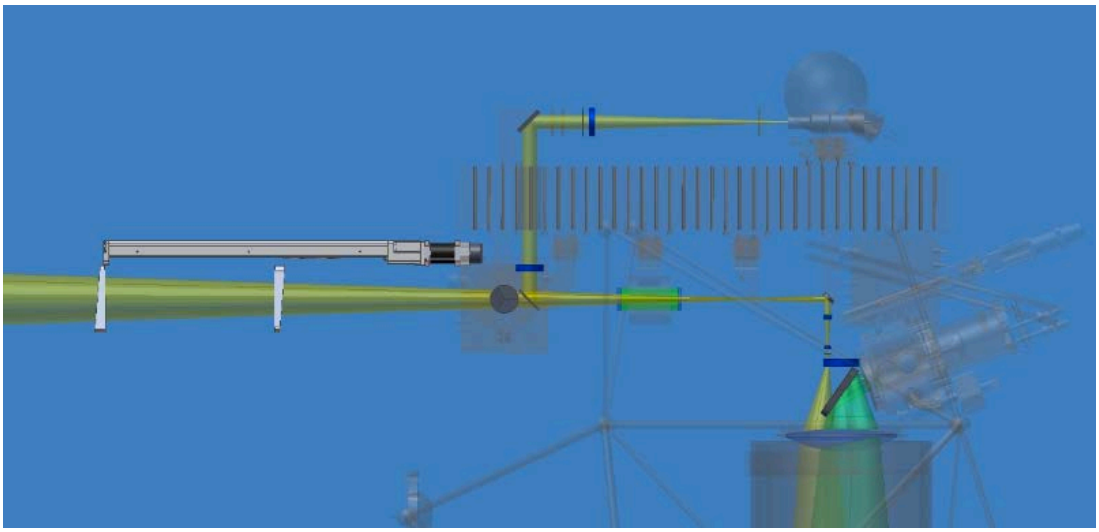

***Figure 3.10**: ADC prisms on trombone slide.*

The ADC provides excellent atmospheric dispersion compensation, effectively eliminating all dispersion down to a working elevation of 15 degrees, as shown in the next figure. For reference, the slit boxes here are only 1.5 arcsecs square.

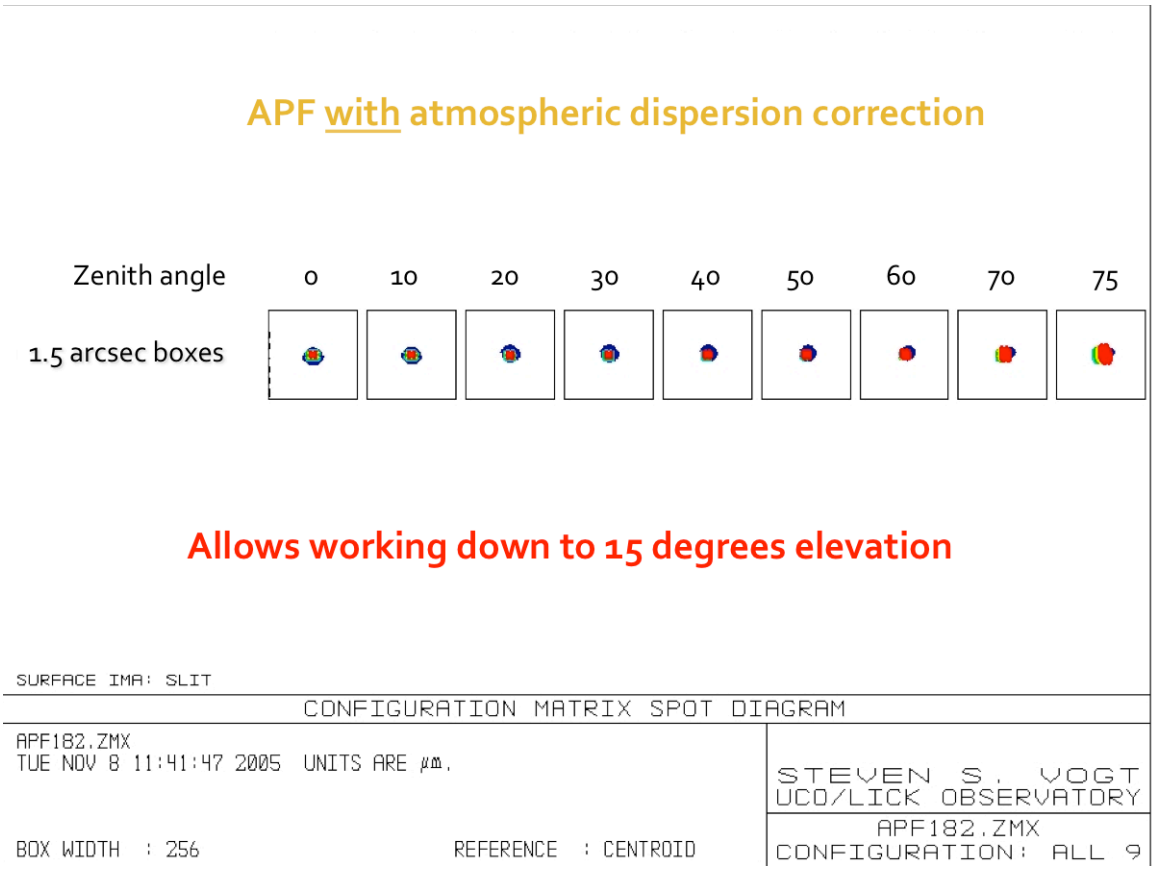

***Figure 3.11**: APF with full ADC correction.*

3.2.2 Acquisition/Guider System

The next optical element encountered is a beam splitter that directs ~4% of the target light to a PhotonMAX-512B CCDTV guider camera made by Princeton Instruments. The beam splitter is a 3.175 mm thick fused silica plate, uncoated on the first side, and with a high performance anti-reflection coating on the second side. Its uncoated front side reflects 4% of the seeing disk to the CCDTV guider, allowing for much higher precision guiding than is usually possible when guiding off of light spilled from reflective slit jaws. Having a full seeing disk to guide on also allows quantitative measures of the seeing width, and accurate photometry to be done with the CCDTV. The symmetric and unvignetted seeing disk image for guiding, coupled with tight specifications on the closed-loop guiding performance of the telescope is expected to further increase the radial velocity precision of the APF facility.

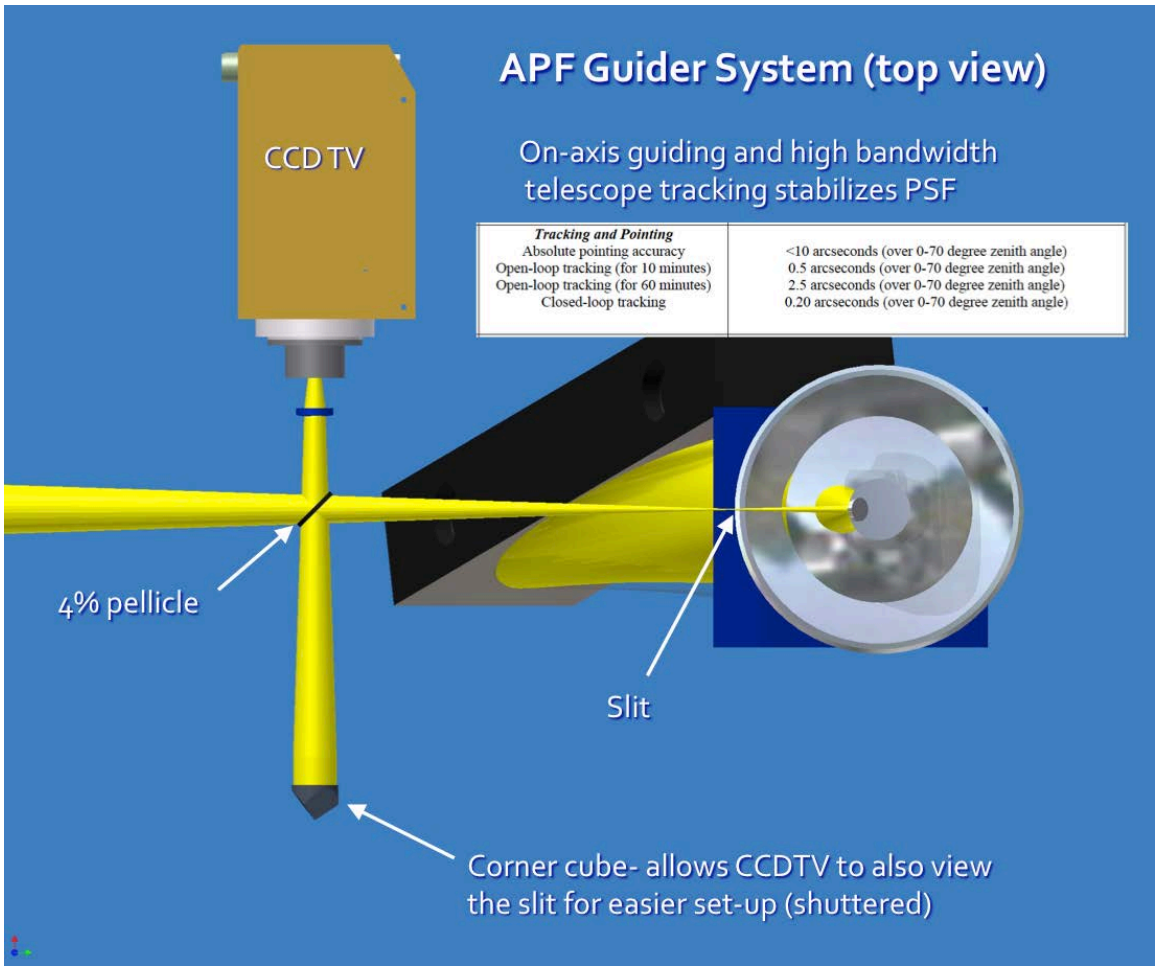

***Figure 3.12**: APF Guider System.*

With the beam splitter arrangement of the guider, the slit itself is not normally viewable. For viewing the slit aperture with the CCDTV, a corner cube and shutter are also provided. When the corner cube shutter is open, and an LED light behind the slit is turned on, the slit aperture can be seen in the CCDTV and a guide box that maps out the slit boundaries is marked and stored in software. When guiding, one only sees the stored guide box for that slit, overlaid on the guider image. The autoguider keeps the seeing disk centered in that slit boundary box. An exposure meter function also keeps track of the time history of the photon flux that falls within the slit boundary box, and computes a photon-weighted time centroid of spectrum exposure, for use in accurate calculation of the barycentric correction. The guider also provides real-time measures of the seeing disk size, shape, and integrated intensity.

A major potential source of error in precision radial velocity work is movement of the star on the slit. Telescope focus errors, and offset errors in guiding both affect the instrumental point spread function, and this translates directly into radial velocity errors, typically at the hundreds of m/s level. It is difficult to precisely quantify this effect, as it depends, among other

things, on the particulars of the spectrometer design and on variations in vignetting in the spectrometer. Basically, anything that produces a variation in the pupil illumination will cause such PSF variations and introduce radial velocity errors. Experience with HIRES on Keck offers one benchmark. If one focuses the telescope as well as possible, and takes a spectrum, then throws the telescope out of focus and refocuses as best as possible, radial velocity errors of some 300 m/s can occur, just from telescope refocus. Similarly, each pixel in a spectrum typically corresponds to ~1000 m/s of velocity, and a typical projected slit width is usually 2-3 pixels. So moving the stellar seeing disk from one edge of the slit to the other effectively shifts the spectrum by some 2000-3000 m/s. Fortunately, much of this is averaged out over long exposures, as long as guiding is reasonable.

The Iodine cell is used to track all these changes, producing a PSF model that becomes part of the data reduction and analysis process. Any subtle changes in the PSF caused by telescope focus error, or misguiding, or vignetting changes, will produce exactly the same PSF changes in the iodine lines. These iodine lines can then be used to map the PSF changes, and remove that error source from the analysis. The Iodine technique does quite a good job at this, but is not perfect. And as one presses ever harder on obtainable precision, at some point, the inability of iodine to adequately track these PSF changes becomes a major contributor to the overall error budget. Other terms in this error budget are photon statistics and stellar jitter.

Another precision RV approach that doesn't use Iodine is the HARPS approach, where the spectrometer is highly stabilized in temperature and pressure, by encasing the spectrometer in a vacuum chamber, and controlling temperature tightly. The light gathered by the telescope is sent down an optical fiber to the spectrometer, and this fiber provides a high degree of scrambling, effectively stabilizing the PSF to very high degree. With this approach, the HARPS group has reported achieving sub m/s RV precision. However, even with such stabilization and fiber scrambling, they found that even small centering errors in guiding on the fiber end caused excess RV errors. So they had to rebuild their guider systems to hold centering of the seeing disk on the fiber much more tightly. And their highest precision work has been done with a dual-fiber scrambler, wherein the output of the first fiber scrambler is then fed optimally into a second fiber. One fiber scrambles predominantly in the azimuthal direction, while the other adds scrambling in the radial direction.

The above discussion provides some context as to how well the guiding needs to work to reach the 1 m/s precision goal with APF. For the past 17 years, our group has used a number of different telescopes for precision RV work. These include the Lick 3-m, the AAT, the Keck 10-m, and the Magellan 6.5-m. All are "conventional" large telescopes, with no extravagant provisions for closed-loop guiding. With such facilities, it is not unusual to see that star wander by as much as ½ of a slit width in any direction randomly throughout an exposure, either from wind gusts,

telescope operator sloppiness, or just normal closed-loop tracking errors. Despite these guiding errors, in general, achieving 2-3 m/s has been pretty routine for the Iodine technique. So it would seem that 0.5 arcsec guiding errors on APF could be tolerated and still reach 2-3 m/s precision. However, reaching 1 m/s is much harder. Despite intense efforts, we have never reached that milestone with either AAT or Lick. AAT bottoms out at more like 2 m/s, and has actually gotten worse in recent years. The precision floor with the Lick-3m is probably more like 5 m/s (depending on who is doing the data reduction). With Keck/HIRES, we have achieved 1-2 m/s precision on sufficiently quiet stars, and where photon statistics is not a limiting factor, but in general it is very hard to beat 2 m/s with HIRES.

With PFS on Magellan, Paul Butler and colleagues have been achieving sub m/s precision over the past 2-3 years. Their results for HD 5499 indicate an overall achieved precision of 76 cm/s held over 2.5 years as shown in the next figure.

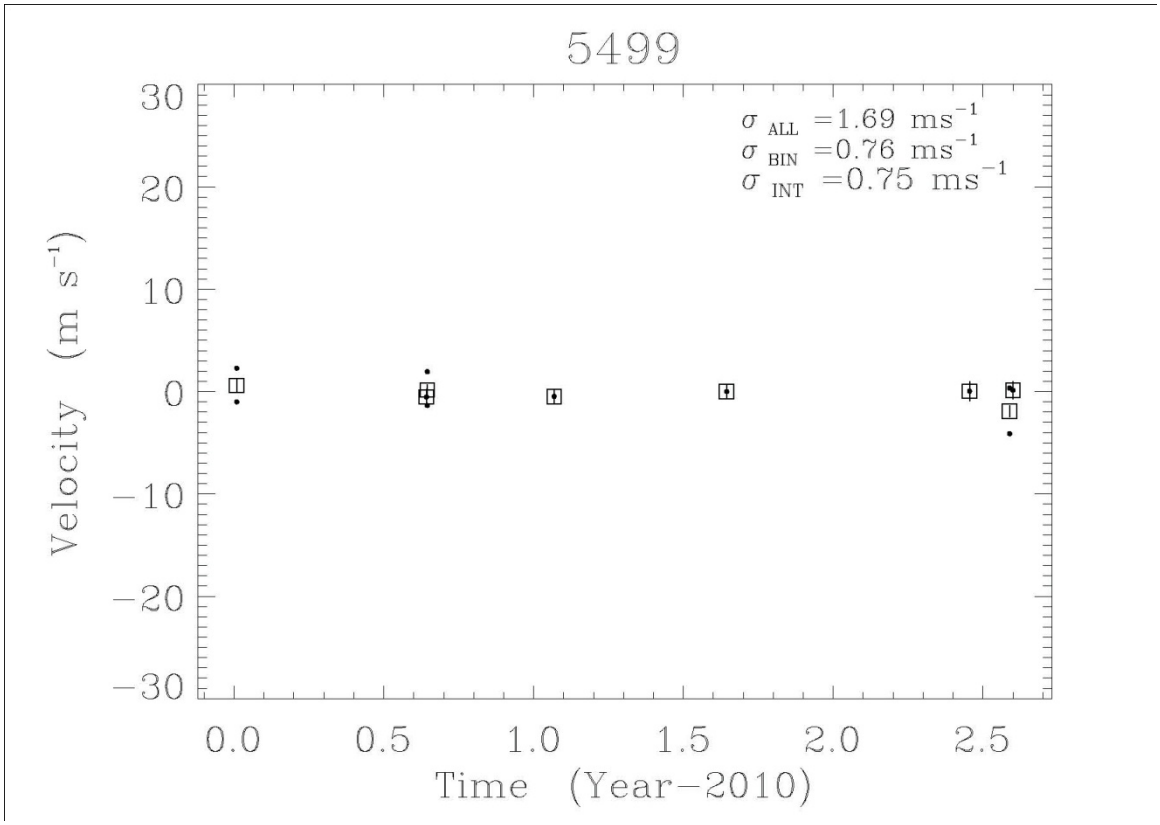

***Figure 3.13**: Precision velocities of HD 5499 from Magellan's PFS spectrometer.*

The PFS precision on HD 5499 is a particularly interesting benchmark since it proves that, with a properly designed instrument, instrumental errors can be reduced to sub-m/s levels, where stellar jitter and photon statistics generally dominate the error budget, and without the

need for (or losses from) a fiber scrambler. And since that 76 cm/s overall PFS precision is the quadrature sum of errors from photon statistics, stellar jitter, and instrumental errors, the instrumental error component itself must already be under 50 cm/s, without scrambling. Since the Levy is optically quite similar to PFS, it seems likely that the Levy should be able to achieve similar sub-m/s performance without a fiber scrambler. At the same time, it is also known from direct HARPS experience that even small guiding errors with a fiber scrambling give rise to excess RV errors. So there is really no lower limit as to how well one would like to guide in an open-ended quest to improve RV precision into the cm/s regime.

Based on extensive experience, it seemed reasonable to conclude that, with a closed-loop tracking RMS spec of 0.2 arcsec, we should expect to be able to reach 1 m/s per exposure with APF. We would have liked to set an even tighter spec on the closed-loop guiding for APF. But demanding too tight a spec here would have greatly inflated the risk/cost undertaken by the telescope vendor, resulting either in no-bids or in unacceptably high bids. So, we pushed hard on the vendors to give us the tightest spec they felt comfortable with within our budget and schedule. We ended up with a closed-loop tracking accuracy specification of <0.2 arcsec RMS, with a desired goal of 0.1 arcsec RMS.

Another common source of guider error occurs when guiding on a seeing disk that is reflected off of slit jaws. Most of the useful guide signal disappears down the slit in good seeing, allowing the star to wander excessively within the slit. Uneven illumination also occurs as starlight from the seeing disk splashes alternately off one slit jaw or the other. That unevenness of illumination, coupled with the finite response and usual sub-optimal tuning of the guider response, leads to excessive motion of the seeing disk within the slit. To mitigate such errors, the APF uses a 4% beam splitter that provides a symmetrical unvignetted seeing disk image to the guide camera. This image is much more constant in time and has a fully illuminated disk that provides a much more precise centroid from frame-to-frame than a heavily vignetted seeing image falling mostly down the slit. APF's CCDTV guider sensitivity was spec'd to be able to guide effectively off of only 4% of the light, and this small amount of light lost to the science beam is a very small price to pay for more accurate guiding.

So the basic philosophy behind the APF guider design was having a telescope with good closed-loop guiding specs, and a sensitive CCDTV guider that is guiding on an unvignetted symmetrical seeing disk. This approach also offers the advantage of having a full seeing disk with which to do quantitative photometry and to make quantitative measurements of seeing width and image shape. Unfortunately, we have learned since completing our rushed Final Site-Acceptance Test, that the telescope has serious persistent low-level oscillations that are variable across the sky and variable over time. At the time of this writing, the telescope does not yet

come close to meeting its <0.2 arcsec RMS closed-loop tracking stability spec. This will be discussed more fully elsewhere in the report.

The idea of incorporating a fiber scrambler into APF was also considered at length a few years back. By then, it was clear that the HARPS team was greatly benefitting from the scrambling advantages of a fiber. And though, from the outset, the Levy was never intended to provide m/s intrinsic long term stability (because of the availability of the Iodine technique), having a fiber scrambler would certainly help to further stabilize the PSF, requiring fewer parameters in the PSF modeling, and thereby probably improving RV precision. The only downside would be light lost to the fiber. Stuffing a seeing disk down a circular fiber aperture is much more lossy than down a similar width long slit. Decentering errors also are more costly with round fiber apertures than with long slits. Traditional fiber scramblers often lose factors of 2 or more due to these factors, not to mention to practical difficulties encountered in mounting the fiber, in surface losses, and in the practical details of aligning and maintaining the optics needed for focal ratio conversion at both ends of the fiber.

APF is however quite well suited for use with a fiber scrambler, and a concept for adding one in the future was developed. Unfortunately, that development work was halted due to budget and schedule constraints. APF has a large “throughput” (resolving power times slit width), of order 110,000 arcsecs. That means, even in typical 1 arcsec FWHM seeing, not much light is lost when entering a fiber end working at resolutions of 60,000-80,000, and there is no need to involve image slicing. To our knowledge, no one has yet demonstrated that image slicing with a fiber can reach 1 m/s precision levels, though various groups have been trying. We will return to a discussion of how to add a fiber scrambler further on down in this report.

3.2.3 Calibration Lamp System

Immediately behind the guider beamsplitter is a small injection mirror that brings in light from calibration lamps mounted above the optical table. These lamps include both Thorium-Argon emission spectrum lamps, and an incandescent quart-iodine tungsten “flat field” source with appropriate filtration to look as flat or white as was practical. The calibration lamp system produces an exit pupil that closely mimics the exit pupil of the telescope, providing an f/15 beam to the spectrometer that appears to come from an image of the primary mirror (the system stop) as formed by M2. The calibration lamp pupil also mimics the central obstruction in the pupil due to M2.

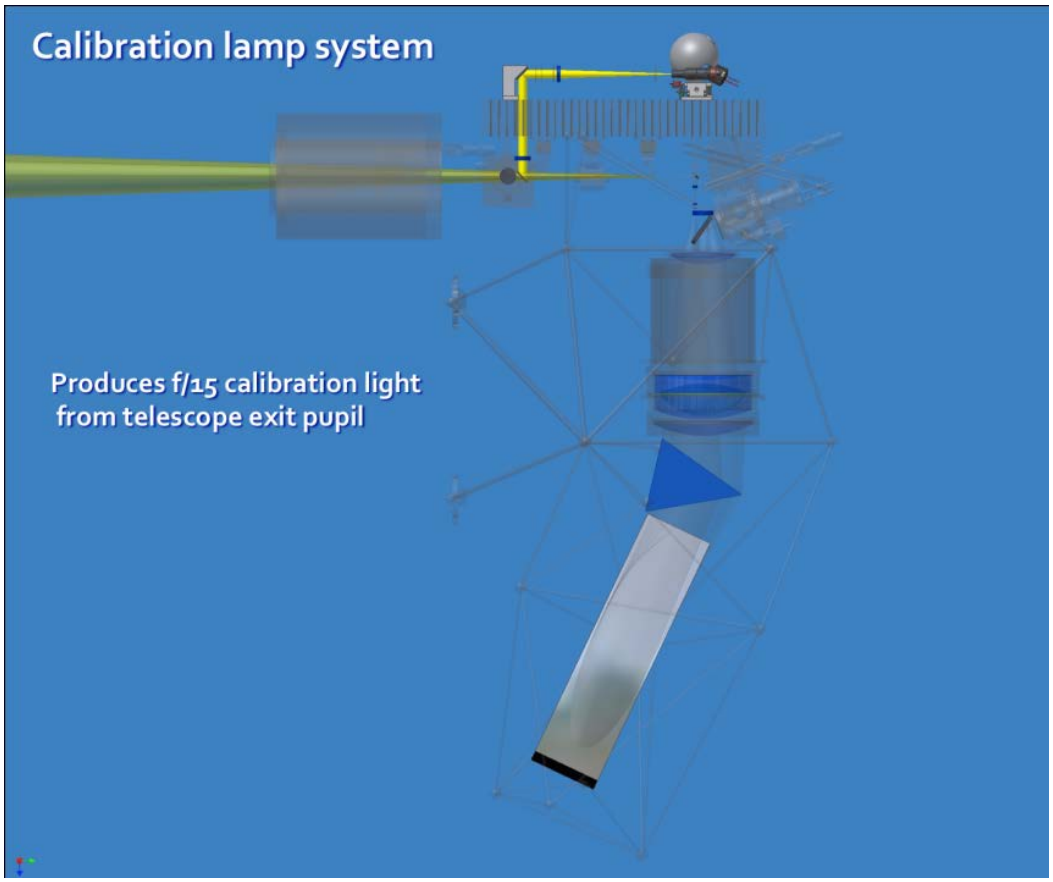

***Figure 3.14**: Calibration lamp system.*

3.2.4 Iodine Cell

Next in the optical train is an Iodine absorption cell. This is a Pyrex cell filled with a small amount of isotopically pure elemental iodine, and heated to 50C, a temperature where the Iodine sublimes into I_2 gas. This I_2 gas has a rich absorption spectrum from 490-600 nm that gets impressed on the stellar spectrum, forming a very stable reference grid for radial velocity determination. The cell is AR-coated on both external surfaces, and temperature controlled to within 0.1C. The cell can be remotely taken into or out of the beam, depending on whether normal radial velocity work or template gathering work is desired. When the cell is out of the beam, a focus-compensating fused silica plate switches in, to keep the telescope properly focused on the slit.

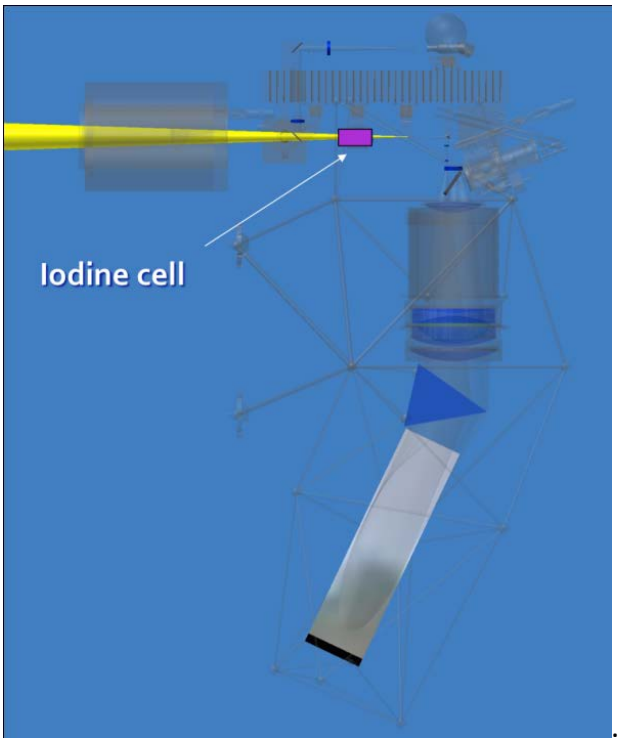

***Figure 3.15**: Iodine absorption cell.*

3.2.5 Slit

The light then encounters the entrance slit of the spectrometer. The slit is actually a series of 9 fixed rectangular apertures, plus a round pinhole. The rectangular apertures provide effective slit dimensions from 0.5 arcsecs in width up to 2 arcsecs, and lengths from 1 arcsec to 7.5 arcsecs. The pinhole projects to a diameter of about 13 um at the final spectrometer focus, producing (when used with either a monochromatic or emission line lamp) a high quality point source image with which to characterize instrumental PSF and/or image quality. The projected image from the pinhole is matched to the size of a single CCD pixel. All slit apertures are tilted by 19 degrees to better align the projected slit image with the row/column axes of the CCD. Behind the slit is an electronic shutter that controls the exposure length.

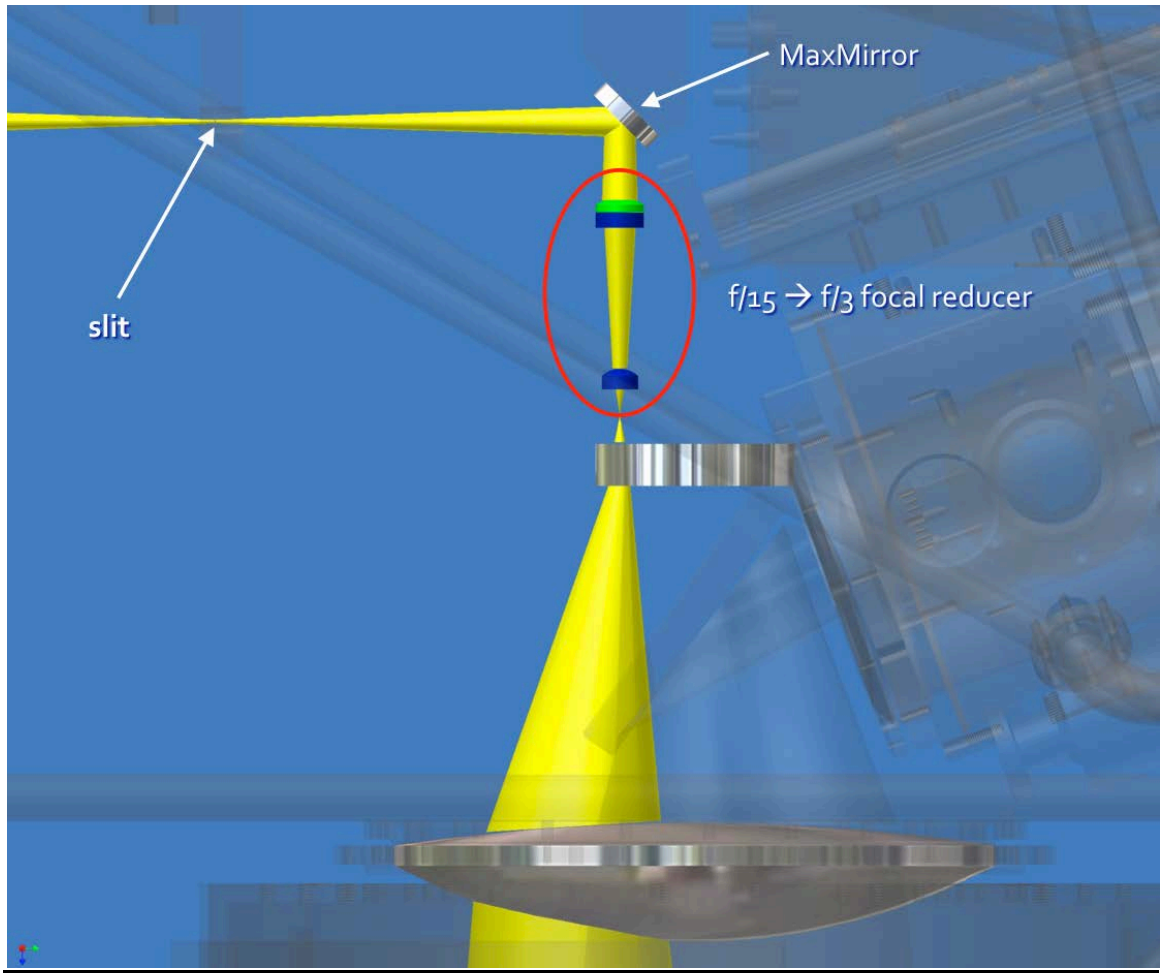

***Figure 3.16**: MaxMirror and focal reducer details.*

3.2.6 Fold Flat

The light is then turned 90 degrees by a very high reflectivity MaxMirror as shown in Figure 3.16 above. The MaxMirror is a commercially available off-the-shelf item. It incorporates a 150-layer multi-layer dielectric coating that achieves a reflectivity above 99% over the entire 350-1300 nm range as shown in the next figure.

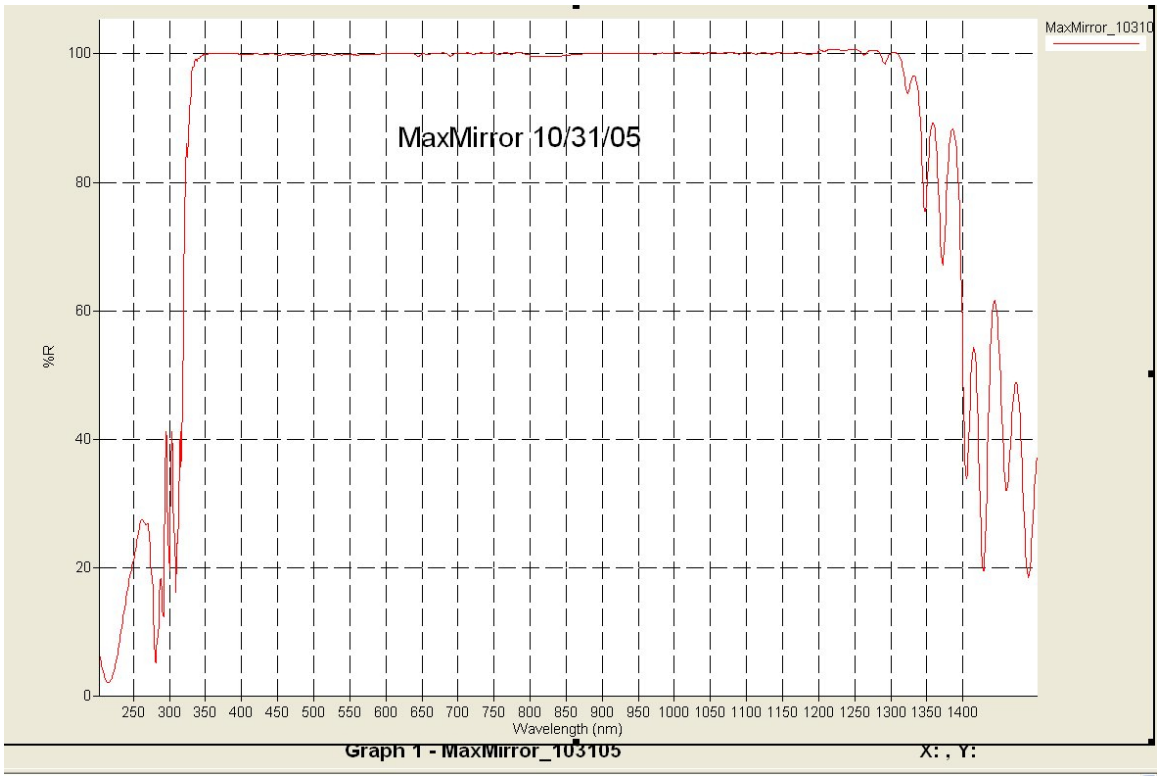

***Figure 3.17**: Reflectivity of the MaxMirror.*

3.2.7 Focal Reducer

Following that is a precision f/15 --> f/3 focal reducer (FRD). The FRD is an optic consisting of a precision doublet of CaF2/S-LAL7 followed by a singlet CaF2 element. Its function is to convert the f/15 beam from the telescope into an f/3 beam, suitable for injection into the collimator/camera unit. Speeding the beam up to f/3 greatly reduces the overall package size of the spectrometer, allowing it to be bolted directly to a Nasmyth focus. Since the entire pupil passes thru this small FRD element, its optical quality needs to be extremely high. The FRD was manufactured in the Lick Optical Labs and we used the highest quality grade glass for these elements, with tight scratch-dig specs. They then received commercial multi-layer AR coatings.

3.2.8 Possible Future Fiber Scrambler

A fiber scrambler could be implemented in a fairly straightforward manner as follows. The output of the FRD produces an f/3 beam that is almost ideal for launch into a fiber. One could mount the FRD and MaxMirror on a rotating fixture such that it rotates the FRD (in the view two figures back) out of the plane towards the observer by some amount, say 90 degrees and locks kinematically into position. The f/3 beam would then fall into a fixed fiber end at the locked position. At the same time, the other end of the fiber would swing into the place where the original FRD output landed, entering the collimator/camera at the original location of the FRD's focal point. An f/3 beam entering the fiber would exit also at about f/3, suffering

negligible focal ratio degradation. By making the fiber long enough and looping around a spindle or equivalent, a high degree of image scrambling would be obtainable. This switch from normal to fiber-scrambled could be done most simply as a manual switch, but it could also be remotely operated if deemed necessary. If, after full commissioning, the Levy fails to perform at sub- m/s level, and the cause is determined to be inadequate stability of the seeing disk at the slit, this fiber scrambler could be added after the fact as fall back insurance.

3.2.9 Collimator/Camera

The expanding f/3 beam then travels vertically downward into the all-dioptric collimator/camera unit. This collimator/camera unit functions as both the collimator in forward pass, and as the camera in return pass (as the beam returns from the echelle). The collimator/camera is a 31-mm clear aperture f/3.17 all-dioptric lens unit with its first and final elements being effectively conventional “field flattener” lenses. The input field flattener lens is the element immediately following the FRD in the above figure. It is a plano-concave lens made of Ohara BAL15Y glass. The output field flattener is identical and forms the CCD dewar window. The expanding f/3 beam from the FRD enters the collimator/camera 25 mm off-axis to provide additional separation between the input and returning beams. The expanding beam then gets collimated into a 200 mm diameter beam and passes on to a cross-dispersing prism.

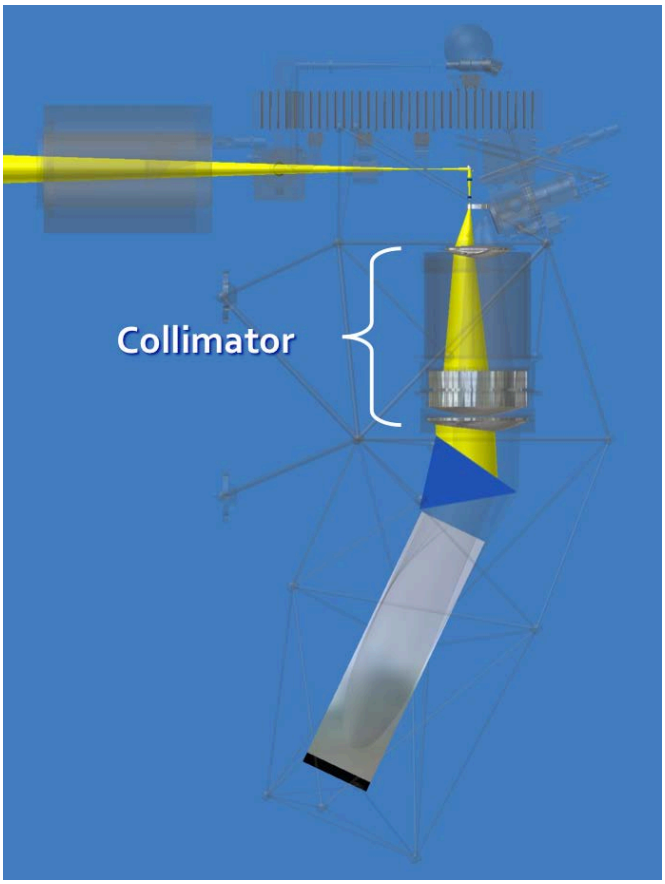

***Figure 3.18**: Collimator/Camera unit.*

3.2.10 Cross-Disperser

The cross-dispersing prism in made of Ohara BSL7Y glass, with a 45-degree apex angle. This prism is used in double-pass and provides an order spacing of about 8 arc-sec minimum at 7800A, increasing to 13 arc-sec at 3700A. The prism was AR-coated using custom dipped solgel coatings done in the Lick Optical Labs. The prism is supported by struts in the determinate structure that attach directly to the glass through bonded titanium pucks. Short aluminum links in the Invar struts provide a very small (few arcsecond) rotation of the prism to eliminate a slight spectral shift term due to thermal changes.

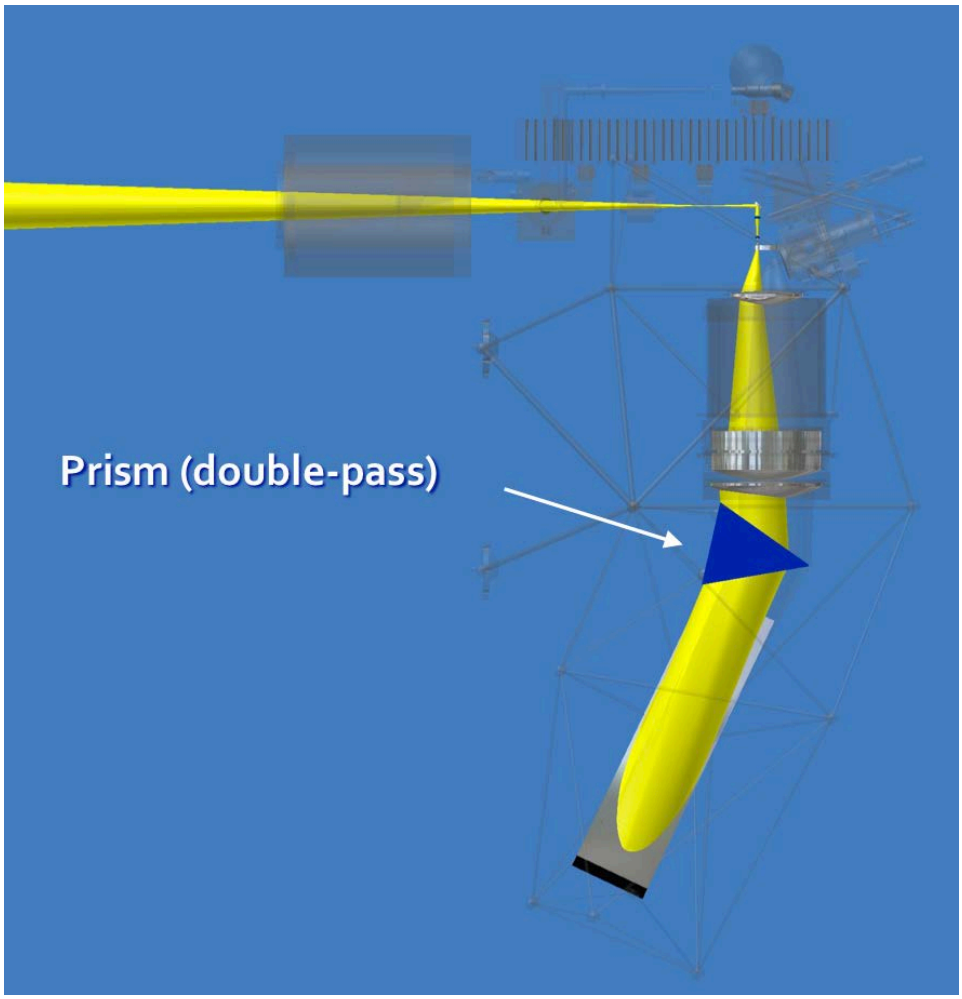

***Figure 3.19**: Cross-dispersing prism.*

3.2.11 Echelle Grating

The beam is then dispersed by a 41.59 gr/mm echelle grating fabricated by the Richardson Grating Laboratory in Spring, 2004. This echelle is a mosaic of two matched segments pulled from the same master that was ruled for the UVES and HARPS spectrometers and replicated onto a Zerodur substrate of size 214 mm x 840 mm x 125 mm thick. The final product has a 16 mm dead gap between the two segments. The mosaic consists of replica grating serial numbers MR166-4-2-1 and MR166-4-3-1 precisely aligned onto a single substrate. With its overall 214 mm by 840 mm format, there is no vignetting of the collimated beam's footprint on the echelle. The echelle is used in pure littrow, with a 0.15-degree out-of-plane "gamma" angle. The gamma angle allows the diffracted beam (green) to clear the input beam (yellow) as the former passes back through the collimator/camera unit.

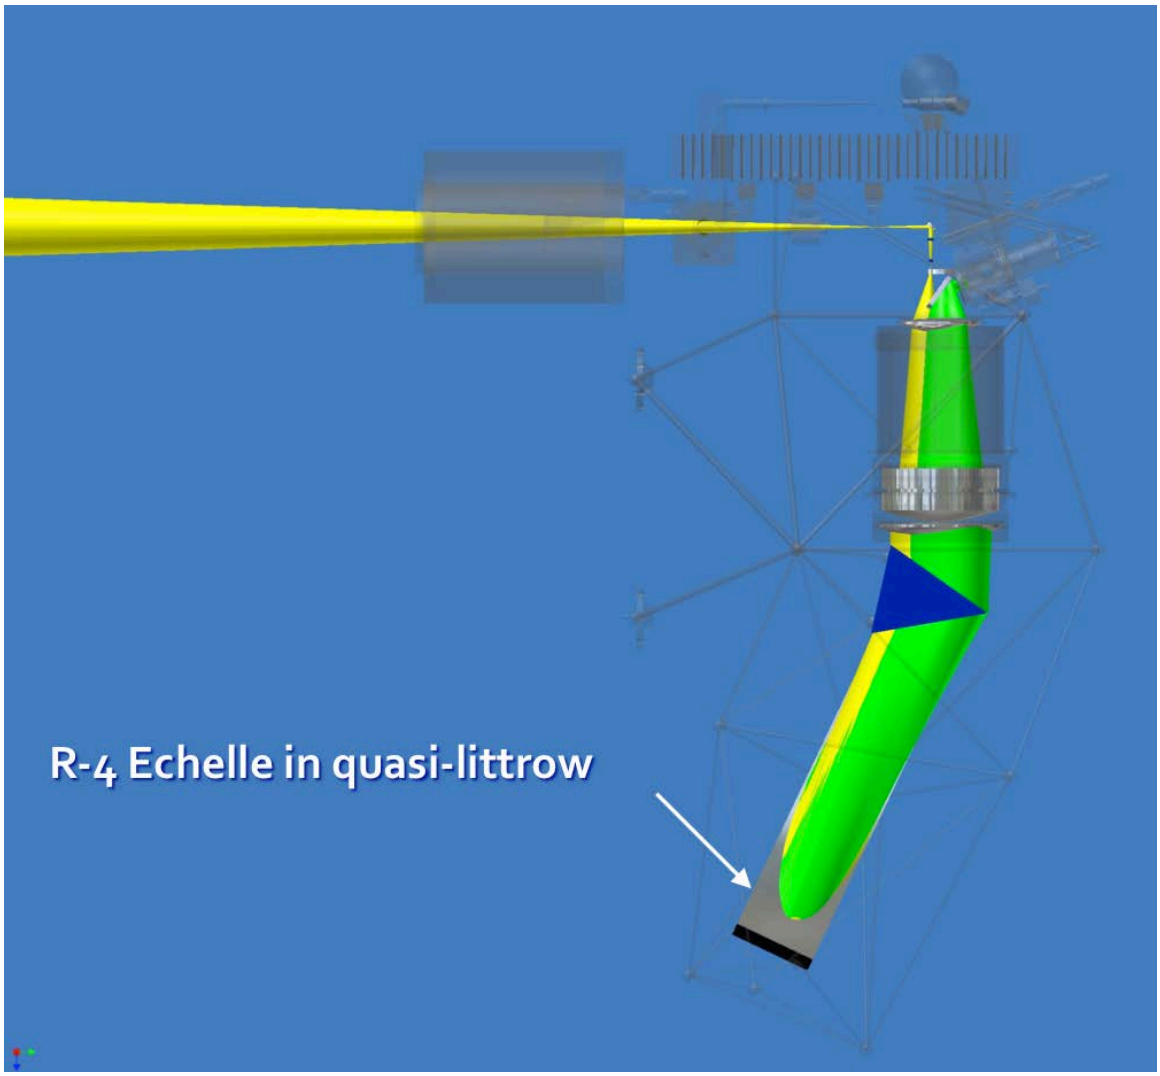

***Figure 3.20**: R-4 Echelle in quasi-littrow mode.*

3.2.12 Fold Flat

The return beam from the echelle travels back up through the prism, doubling the cross-dispersion, and then back through the collimator (which is now functioning as the camera). The beam narrowly clears the incoming beam by virtue of the 25mm decentered input and the echelle gamma angle. The beam is then intercepted by a high reflectance flat mirror and directed to the CCD dewar as shown in the next figure.

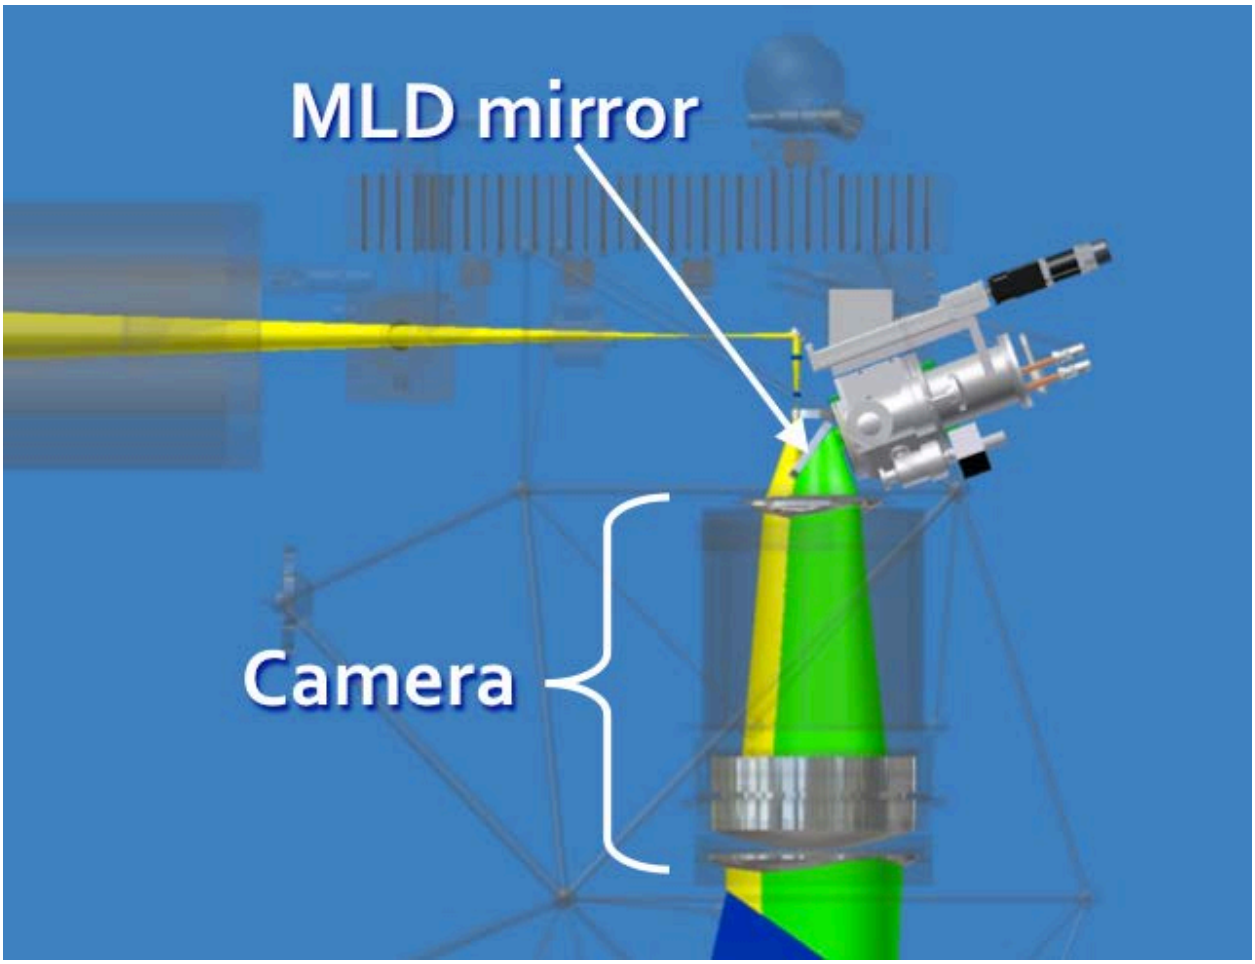

***Figure 3.21**: Camera and MLD fold flat mirror detail.*

This flat mirror is similar in design to the MaxMirror, but much larger. We had it custom-made by MLD Coatings, Inc. in Sunnyvale, CA. Like the MaxMirror, it uses a multi-layer stack of over 100 layers to achieve high reflectivity from 390-1100 nm as shown in the next figure. Thermal stresses from the very thick multi-layer stack, which would otherwise distort the optically flat substrate, were balanced by similar stresses from a dummy multi-layer coating on the back side of the substrate that was painted over black.

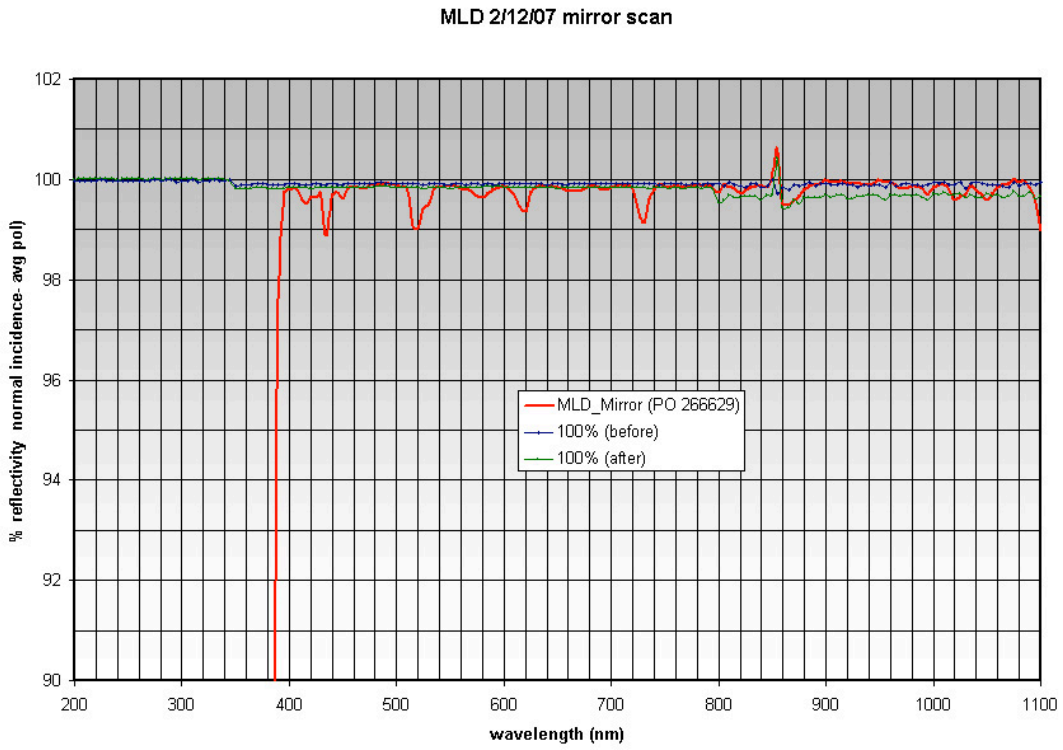

***Figure 3.22**: Reflectivity of the MLD fold flat mirror.*

3.2.13 CCD and Dewar

The beam is deflected by the MLD mirror into the CCD dewar. The window of this dewar is a plano-concave "field flattener" lens of Ohara BAL15Y, essentially equivalent to the field flattener lens immediately following the focal reducer. The CCD is a CCD42-90 back-illuminated high performance CCD manufactured by E2V Technologies, Inc. This CCD has a format of 2048 x 4608 pixels, with a 13.5-micron pixel pitch. Its broadband AR coating delivers quite high Quantum Efficiency over a broad spectral range as shown below.

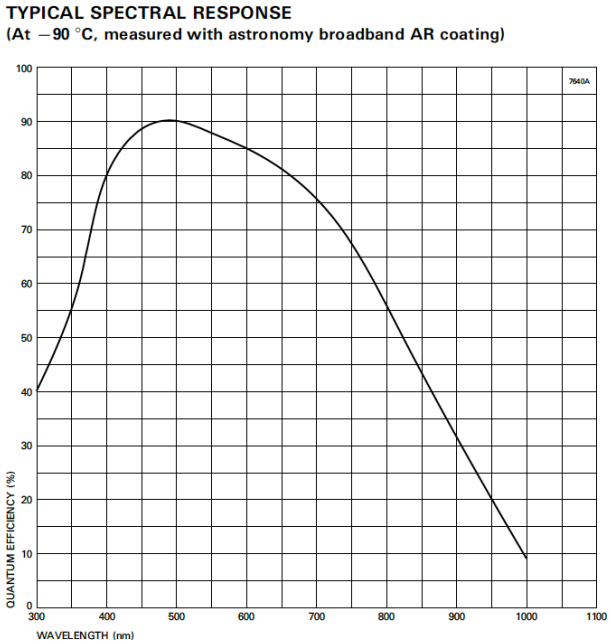

***Figure 3.23**: Quantum efficiency of the E2V CCD42-90 detector.*

The dewar for the CCD was custom-built in the Lick Labs. In place of using LN_2 as a coolant, we opted to use a Cryotiger Polycold closed-cycle cooler. To facilitate fine focusing, the CCD dewar is mounted on a fine-pitch mechanical stage. This stage moves the focal plane 0.2 microns per stage step, with essentially no run-out.

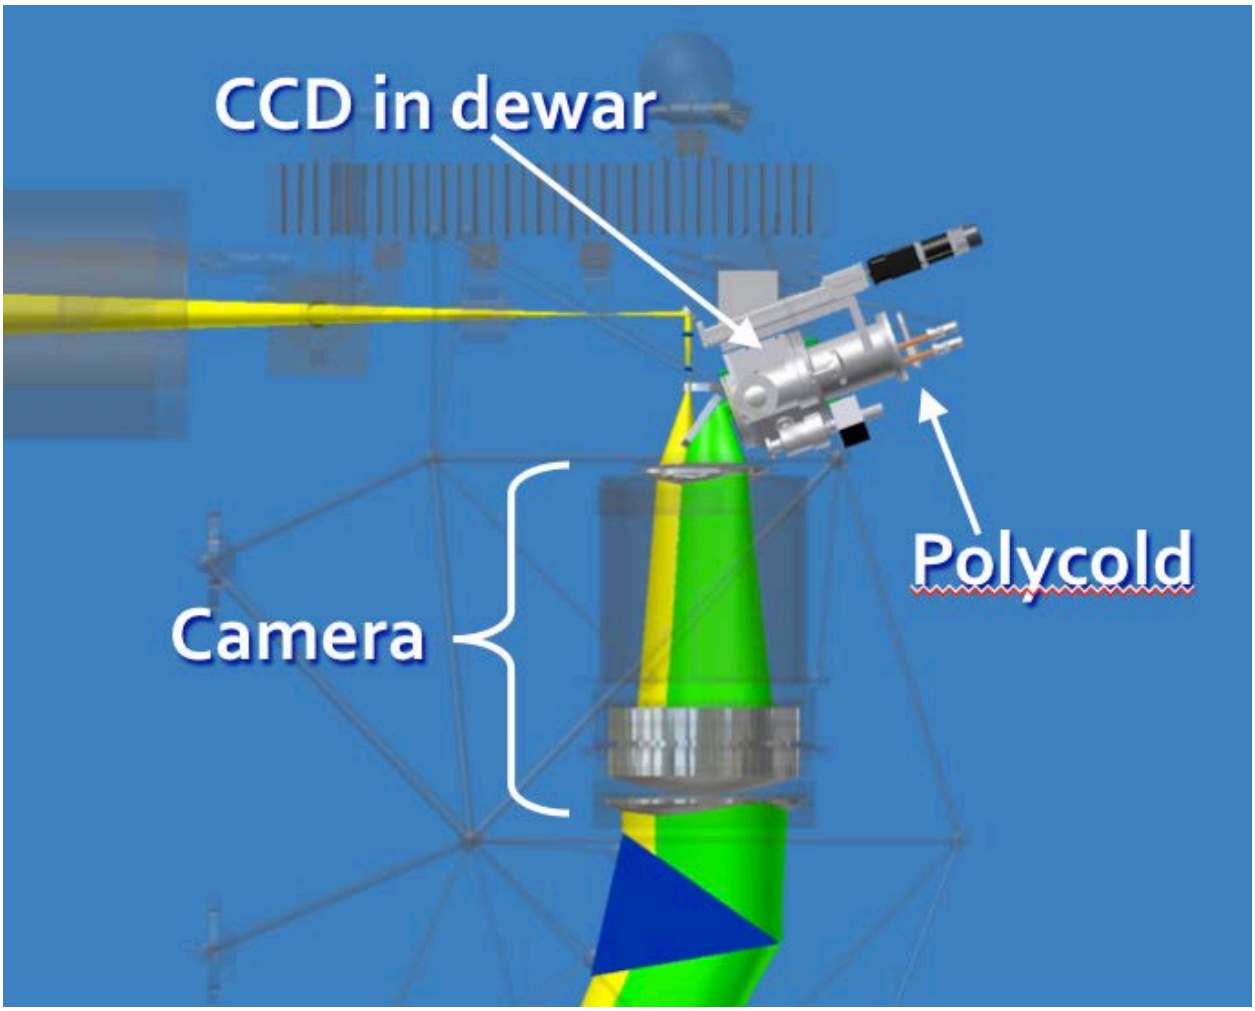

***Figure 3.24**: Dewar and Polycold chiller detail.*

3.2.14 Determinate Structure Support Frame

As described above, all optics are supported via a determinate structure or 'space-frame'. The structure is a collection of six hexapod forms as shown below.

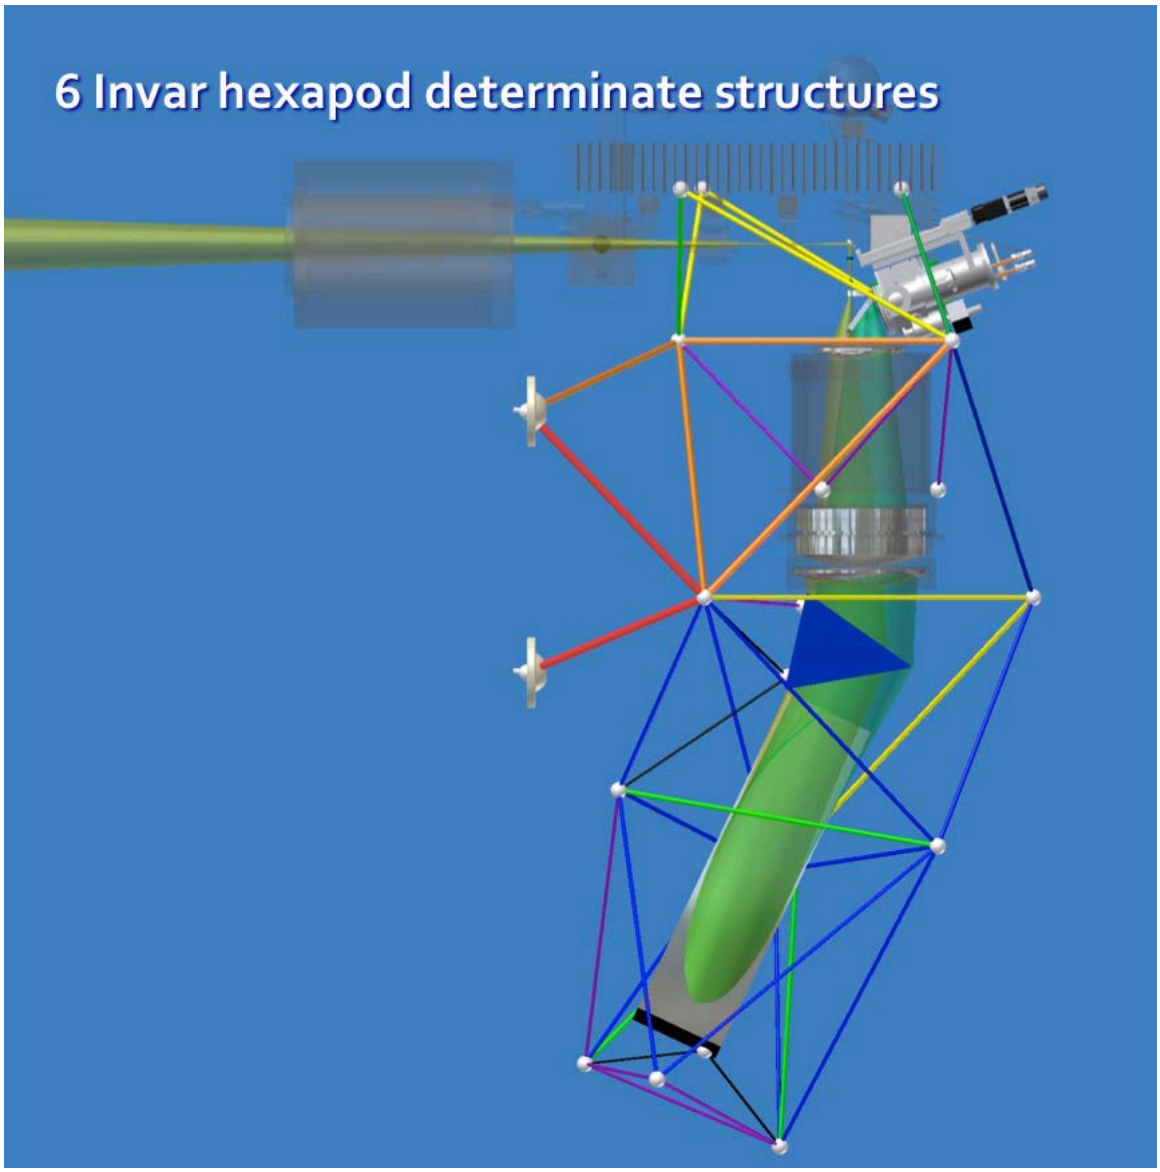

***Figure 3.25**: Space-frame assembly of six hexapod structures.*

The red hexapod struts attach the assembly to 3 mounting points on the telescope fork arm. The purple structure supports the main body of the collimator/camera unit, and consist of temperature-compensating Magnesium pistons that telescope with temperature to provide most of the passive athermalization of the optical train.

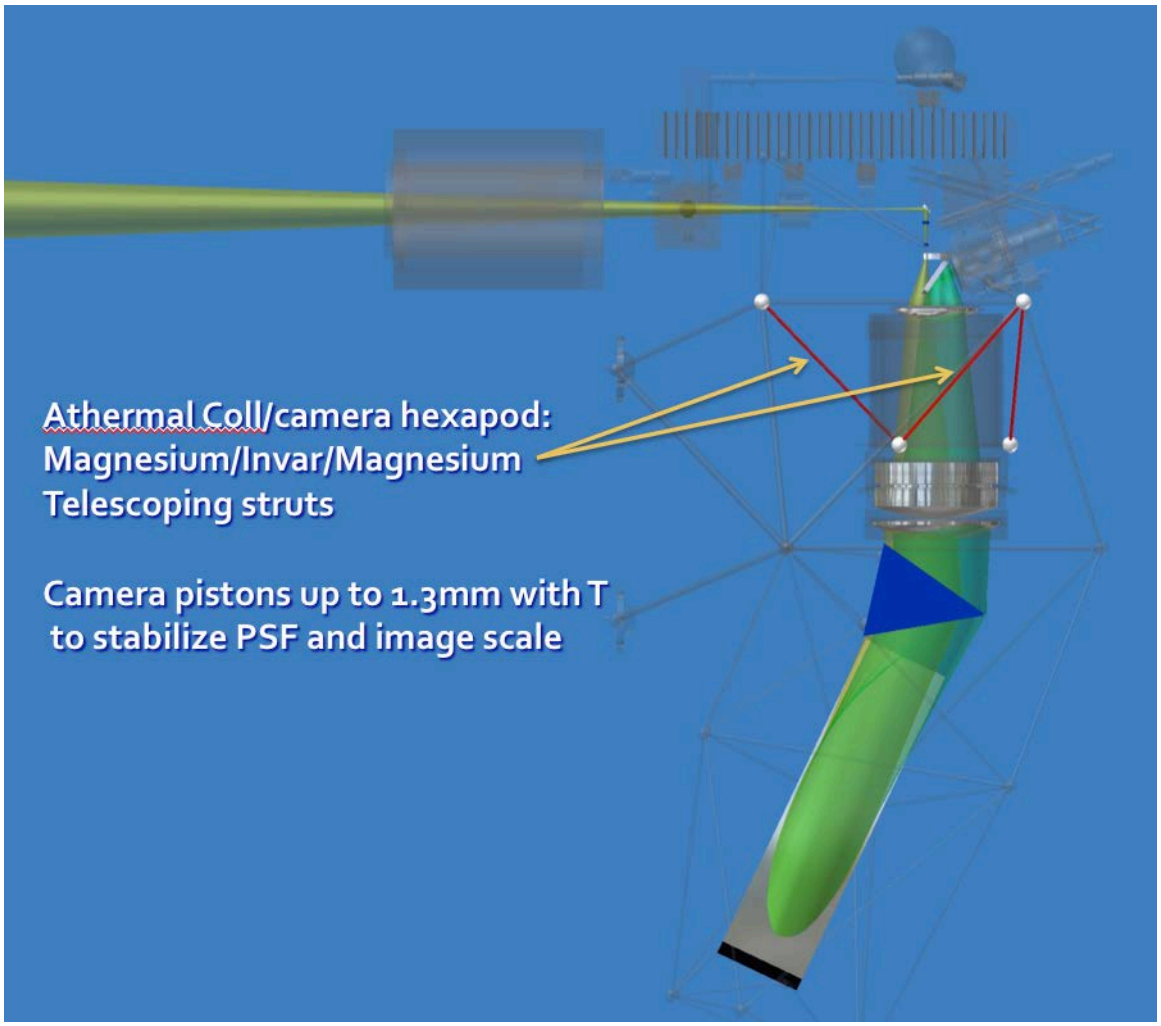

***Figure 3.26**: Hexapod that thermally pistons the Collimator/Camera lens group.*

3.2.15 Enclosure

The entire spectrometer is encased is a split-shell insulating fiberglass enclosure. The enclosure provides R-12 insulation and forms a light-tight barrier around the spectrometer.

3.2.16 Spectrometer Thermal Control System

A NESLAB chiller unit delivers chilled de-ionized water to 3 separate cooling loops within the spectrometer to maintain temperature. A heater pad is also provided for temperature control. There are 3 cooling loops- UCAM (for the CCD control electronics chassis), LAMPS (for the 50-watt quartz-iodide incandescent bulb), and HEATX (for the spectrometer interior as a whole). Valves control the flow rate of each loop, with a bypass valve also provided in case all valves end up closed. RTC temperature indicators are distributed throughout the interior of the spectrometer, providing thermal control data from key areas. The primary thermal control points are the 6 camera struts, and the average of all 6 is used as the control metric- holding this average to some desired set point. The goal is to control “strutavg” to some fixed value

throughout the year. We expect that this set point will be 15-18C, but have not yet finalized that value. A thermal control algorithm can apply either heating (via the heating pad), or coolant (via control of the chiller flow rates) to maintain the desired temperature. The relatively long thermal time constant of the spectrometer (some 4-6 hours) results in only minimal heating and cooling input changes from day to day and season to season.

3.2.17 Coatings

There are a total of 47 surfaces encountered by a photon along the APF optical train. Careful attention to minimizing losses at each surface was necessary to maintain overall efficiency. The two flat mirrors described above were highly optimized by fabricating them as very high performance multi-layer dichroic stacks. The echelle grating efficiency is under complete dominion of the manufacturer Richardson Grating Labs. We chose aluminum as the grating over-coating for best overall longevity. A silver over-coating would have offered perhaps 10-15% efficiency improvement, but with uncertain longevity and minimal prospects to refresh or clean over the long term.

Aside from a few glass-glass interfaces (that are coupled either with index-matching grease or oil), the rest of the surfaces are all air-glass. Unless treated with high-quality anti-reflection (AR) coatings, could easily cost a loss of 3-4% per surface. There are 25 such air-glass surfaces. So we developed in the Lick Labs the capability of applying Sol-gel AR coatings. Solgel is a colloidal suspension of silica nano-particles in an alcohol-based liquid suspension that can be applied either by dip-coat or by spin-coat. Once applied, the liquid quickly evaporates, leaving a stalagmite-like microscopic structure of silica nano-particles. Basically, entering photons can't tell exactly when they encounter the actual surface. This AR coating layer has a surface area effectively 40 times larger than the part it coats, so contamination by dust and particulates is a concern. The coatings are soft, but when properly treated with ammonia soaking for 24 hours, solgel coatings are quite robust. A further treatment with an anti-wetting agent also makes them hydrophobic.

The next figure shows the expected performance of a single 110 nm layer of solgel on the various glass lenses used in the Levy. As can be seen, when properly applied, solgel offers a quite broad AR performance that can also be specifically tailored for the critical 490-600 nm Iodine region.

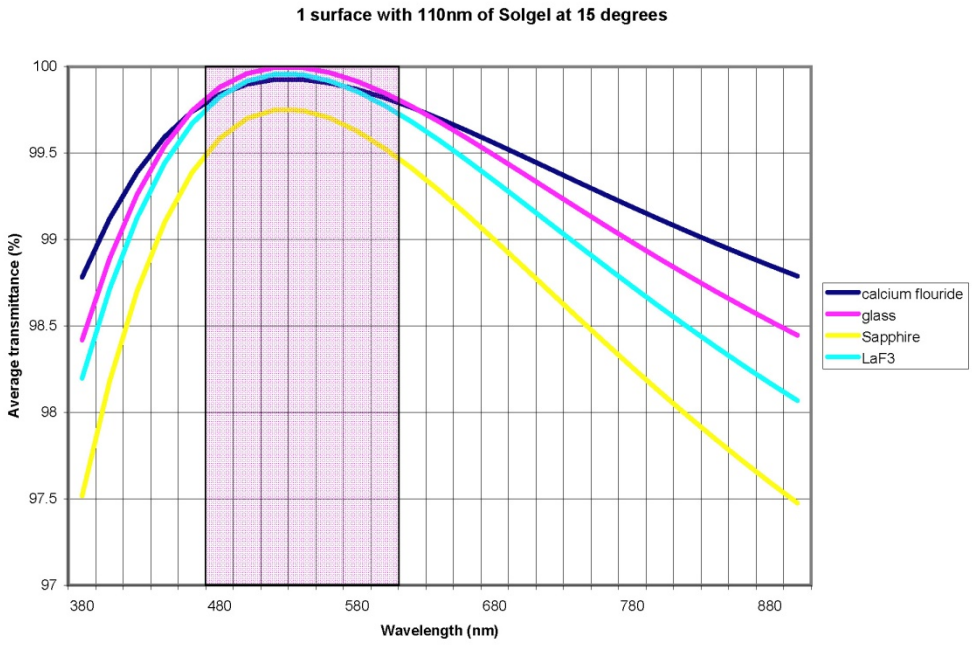

***Figure 3.27**: Transmission of a single 110-nm layer of solgel on various glasses.*

By controlling the spin rate and /or dip rate, the thickness of the solgel can be tuned for optimal performance over any desired band. The solgel also shrinks over the first few hours, and the ammonia soak also shifts the peak of the reflectance curve. All these variables were calibrated and accounted for in optimizing the Levy's AR coatings for the 490-600 nm Iodine region. The next figure shows the theoretical effect of solgel on 28 air-glass surfaces. It illustrates that, even with 25-30 air-glass surfaces, losses due to surface reflections could, in principle, be held to less than 10% over the iodine region with properly applied and maintained solgel coatings for all the lenses.

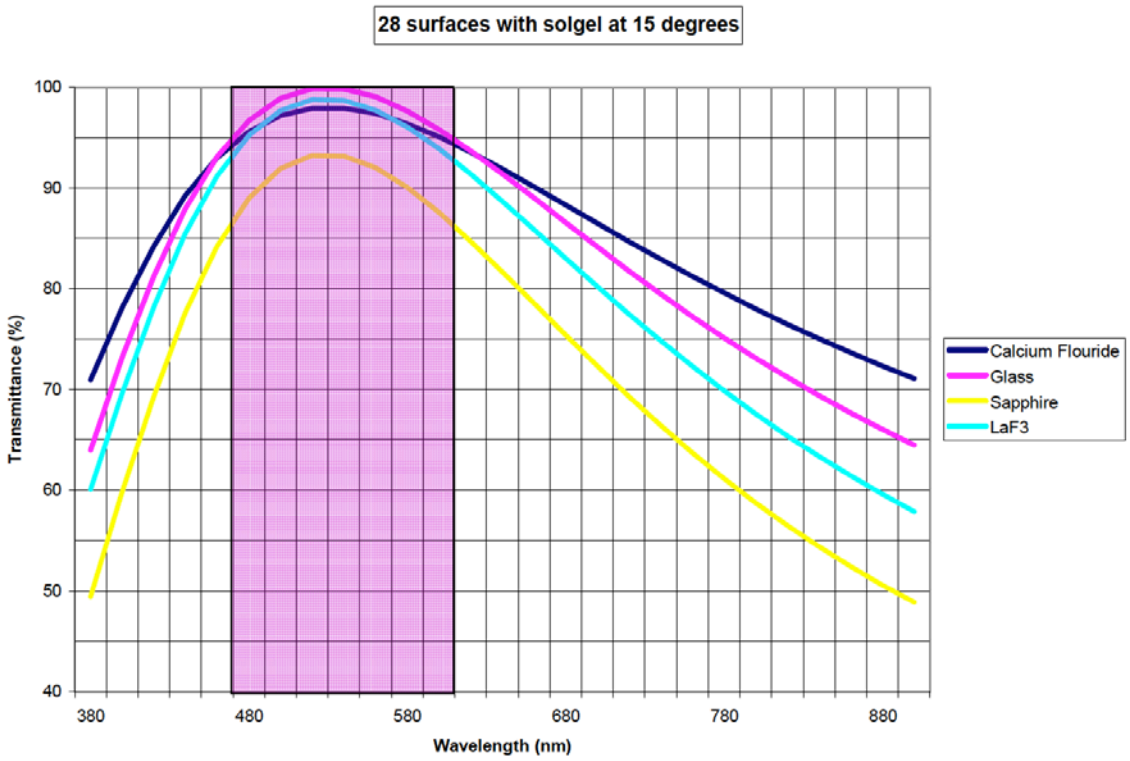

***Figure 3.28**: Theoretical transmission of 28 solgelled air-glass surfaces.*

Most of the lenses of the Levy were solgelled using a spin process as shown in the next figure. Here, solgel is being spun onto the large front element of the collimator/camera.

***Figure 3.29**: Solgel spin-coating process.*

The large flat surfaces for the two ADC prisms and for the cross-dispersing prism were better coated using a solgel dip process. Basically, the optics was immersed in solgel and then withdrawn under precise computer control at a constant rate. Surface tension and wetting of the solgel at the liquid's surface coats the optic with a thickness that is proportional to the draw rate. By controlling the draw rate, any desired thickness can be achieved and thereby tuned for maximum performance over any desired spectral region. The following picture shows the cross-dispersing prism being dip-coated in the Lick Optical Labs.

***Figure 3.30**: Solgel dip-coating process.*

Once coated, each sol-gelled optic was treated with an ammonia soak for 12-24 hours to increase hardness. The sol gel coating was then thoroughly dried by baking in an oven for 48 hours at 90 degrees C. Then a small amount of Hexamethyldisilazane in a petri dish was introduced into the oven and the coated optic soaked in that atmosphere for another 24 hours. The next figure shows one such test where the beading properties of a water droplet on a treated sol-gelled surface are used to assess the hydrophobic quality of the coating.

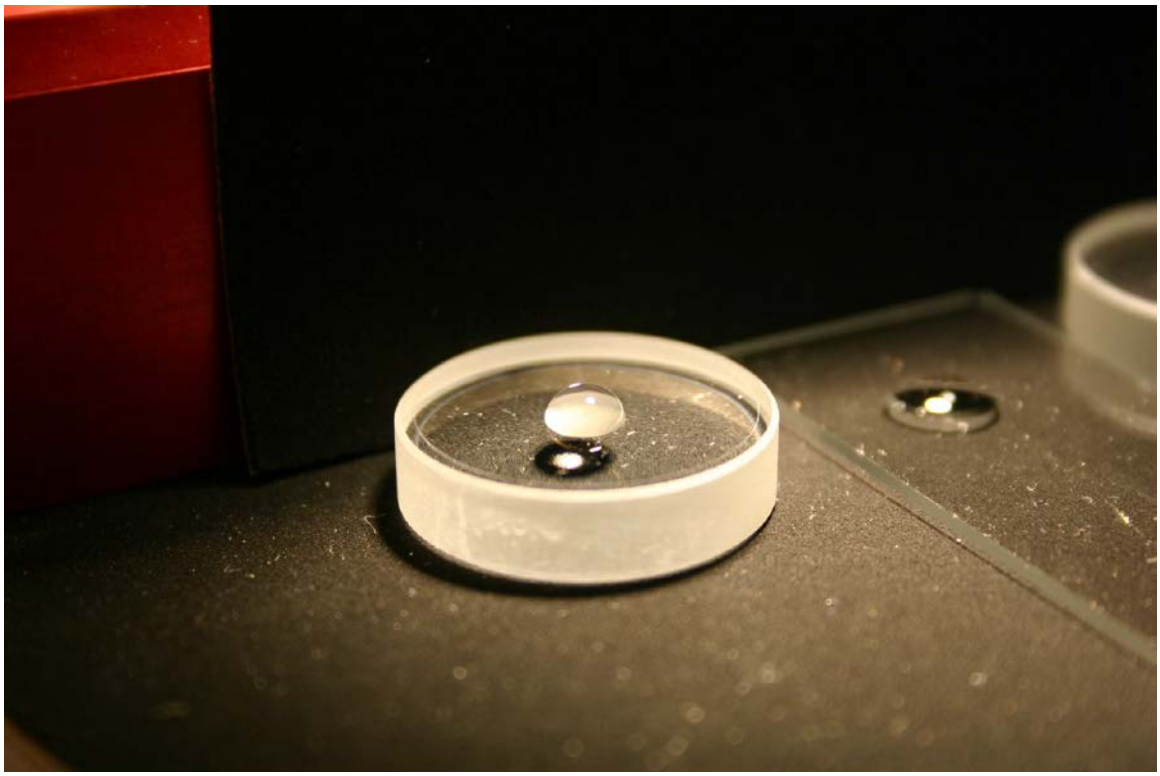

***Figure 3.31**: Water droplet on a treated sol-gelled surface.*

Once AR-coated, each surface was measured with a custom-built reflectance probe as shown in the next figure.

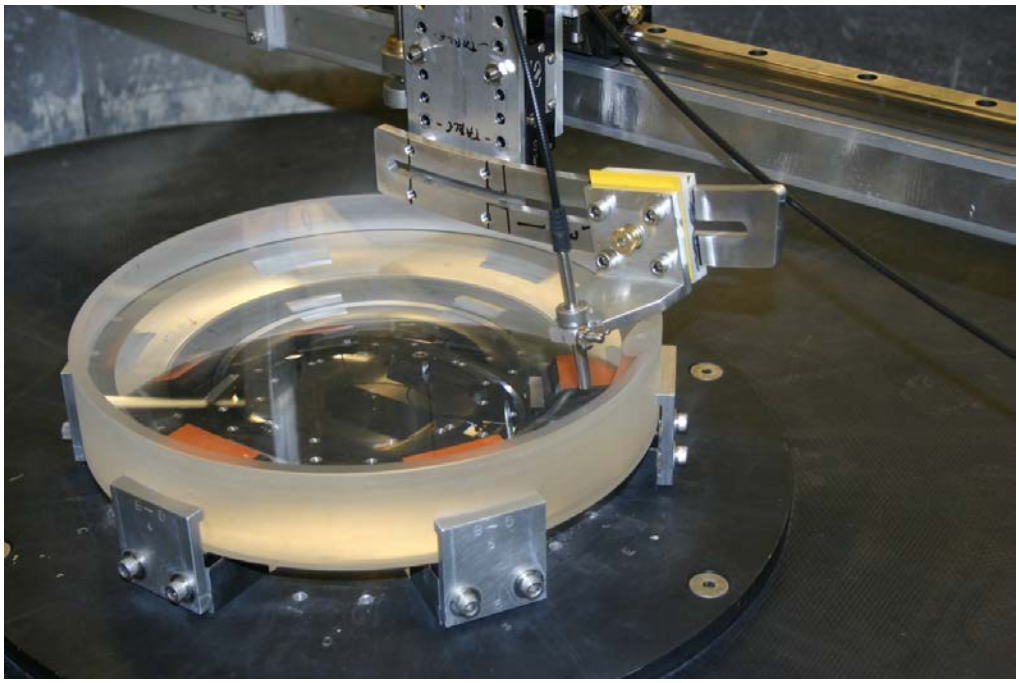

***Figure 3.32**: Measuring reflectance using custom-built probe.*

The next figure shows a representative plot of reflectivity vs. wavelength for both surfaces of the large cross-dispersing prism. Similar plots are available for all the sol-gelled lens surfaces. The coating performance in all cases was tuned for highest performance over the 490-600 nm Iodine region. The red reference line indicates an uncoated fused silica air-glass surface. Blueward of ~380nm, solgel coatings rise rapidly, becoming essentially equal to an uncoated fused silica surface by 300 nm.

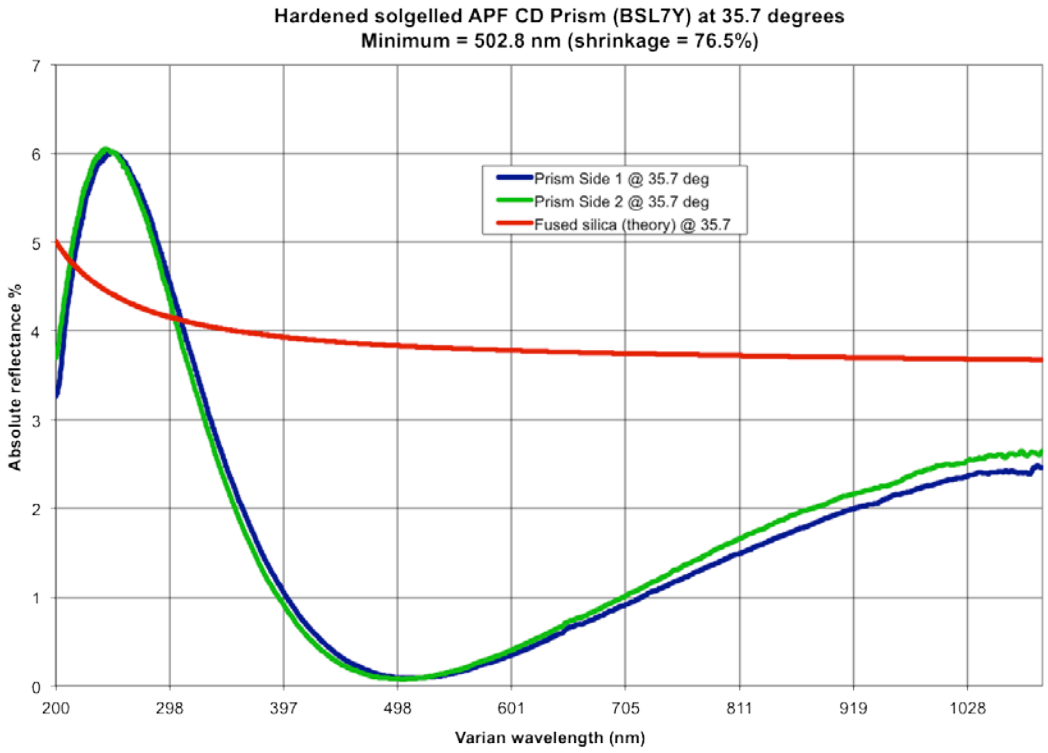

Figure 3.33*: Reflectivity vs. wavelength for both surfaces of the large cross-dispersing prism.*

The next plot shows a typical curve for a solgelled lens- in this case Lens D in the collimator/camera unit. Coating curved surfaces is not as optimal as dip-coating flat surfaces, so we made measurements across each lens, from center to edge. In this case, the surface was concave, and despite our best efforts, ended up with a slight difference between center and edge performance. And though performance across the 490-600 nm iodine region is good, the central region shows higher reflectance, especially redward of 650nm. That extra reflectance gives rise to a ~1% optical ghost in the system that we call the Red Rascal, as will be discussed later.

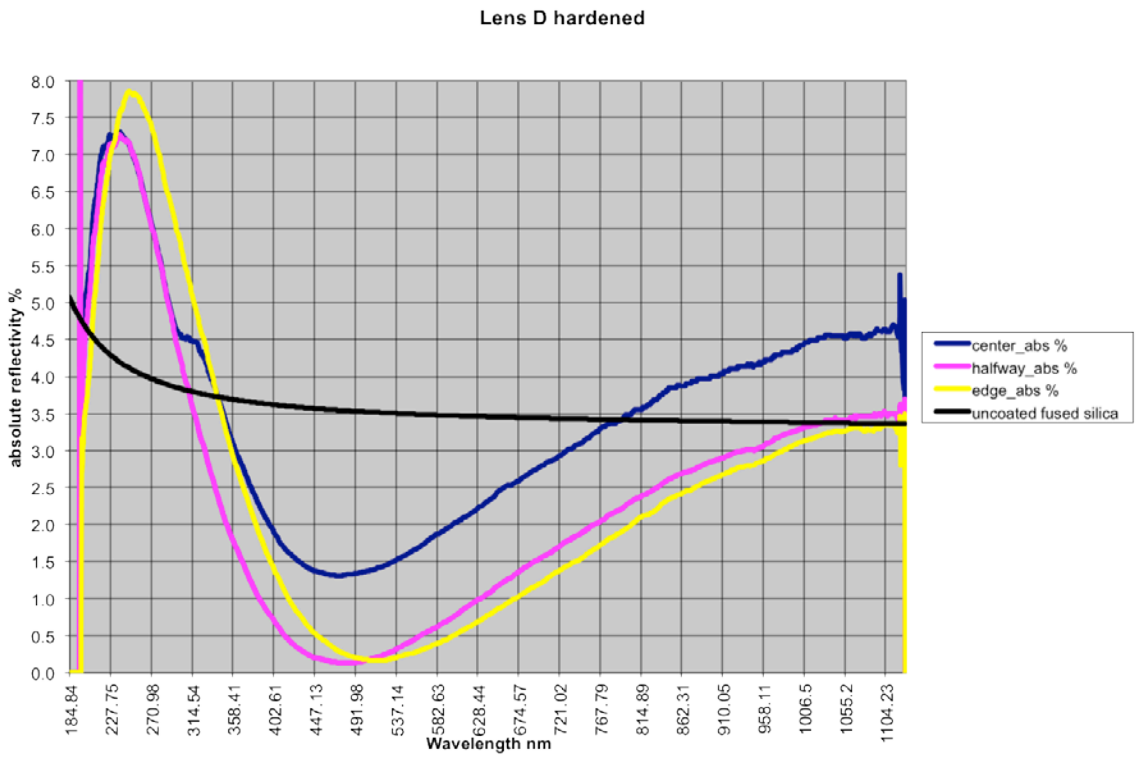

***Figure 3.34**: Measured reflectance for solgelled concave surface of Lens D.*

3.3 Sample Spectra

3.3.1 Thorium Argon

A Thorium-Argon hollow cathode lamp is provided for both wavelength calibration and monitoring of the instrumental point spread function (PSF). The particular lamps used were made with a special admixture of Neon and Argon such that the intensity of the very strong Argon lines redward of 700nm was significantly reduced. The following figure shows one such representative Th-Ar spectrum showing the full frame coverage from 370nm to 950 nm. Here, the figure has been stretched vertically by about a factor of two to make the 2048 x 4608 format squarer for presentation. The very saturated bright lines are the Argon lines that sit redward of 700nm. Here, the top of the echelle format is at about 990nm, while the bottom is at 375 nm.

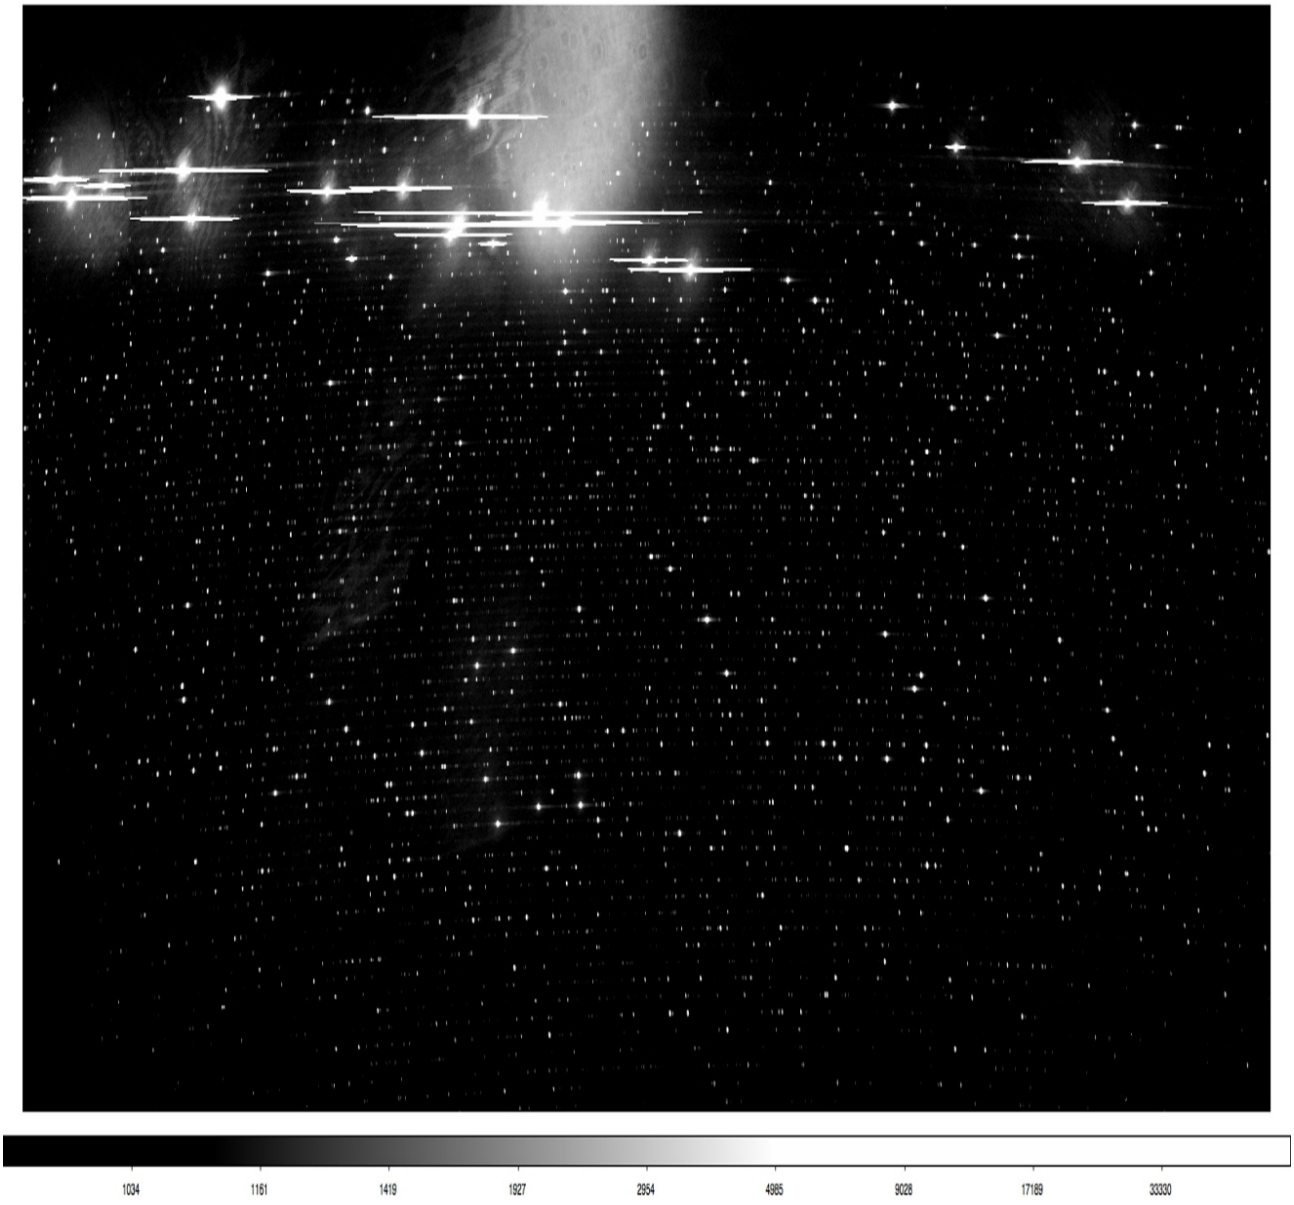

***Figure 3.35**: Thorium-Argon spectrum showing the full frame coverage from 370nm to 950 nm.*

The next figure shows a zoomed-in close-up of the center of the format at the same 2x-vertically stretched aspect ratio.

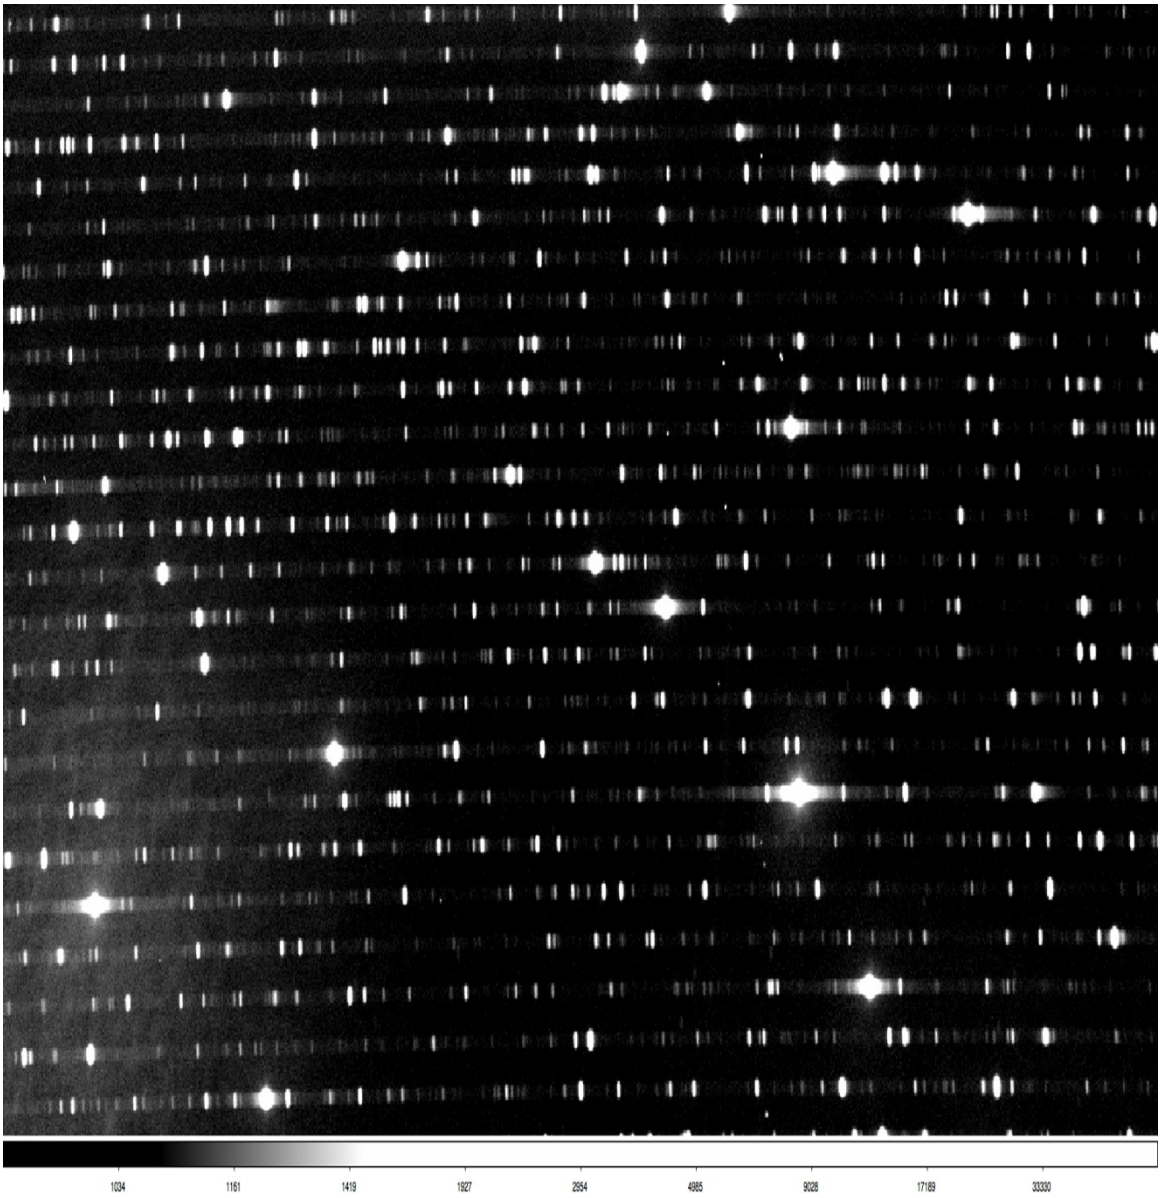

***Figure 3.36**: Close-up of Th-Ar spectrum near the center of the format.*

3.3.2 HD 192281: An O5V Star

The next figure shows the spectrum of the V=7.55 O5V spectrophotometric standard star HD 192281. The telluric A-band and B-band of atmospheric oxygen are prominent in absorption against the star's almost featureless continuum. Also obvious are deep interstellar lines of sodium and singly-ionized calcium. Less obvious at this display magnification is the extensive forest of Iodine lines across the center of the format. Finally, extensive systems of

telluric atmospheric water absorption are seen near the top of the format, as well as pronounced fringing from the CCD redward of 800 nm.

***Figure 3.37**: Spectrum of the O5V star HD 192281.*

Very faintly visible (at about the 1% intensity level) in the center top 3rd of the above figure, is the Red Rascal ghost. The next figure shows an intensity-stretched version that brings out this feature more clearly.

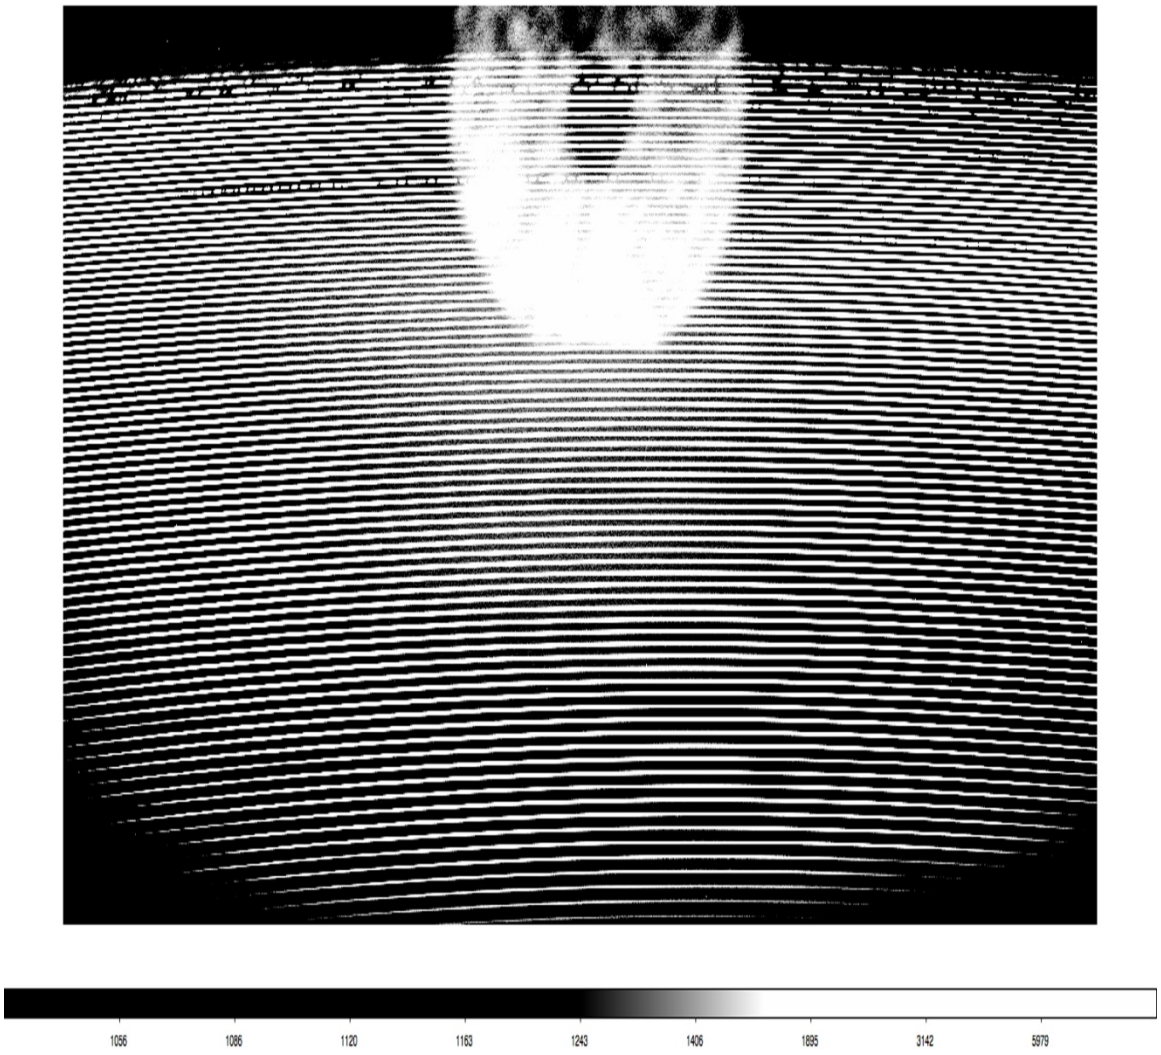

Figure 3.38*: Intensity-stretched version of HD 192281 showing Red Rascal ghost.*

The feature is mostly from light redward of 600 nm reflected off the central region of the concave surface of lens D due to the less-than-perfect AR coating on that surface (as discussed above). The next figure shows a ray trace of this ghost. With a double-pass all-dioptric system of the type used here, tendency for ghosting and narcissus is strong. In the design phase, ghost reflections were explored for all surfaces. Where possible, lenses were bent to deflect their strongest ghosts sufficiently far enough off the optical axis to straddle or otherwise miss the CCD completely. In other cases, curvatures of surfaces were adjusted to spread out a ghost as much as practical to reduce its surface brightness. But despite numerous attempts, it was not possible to completely eliminate all ghosts in such a system. The Red Rascal was the one that remained.

Though annoying, the Red Rascal ghost has actually served as a very useful map of the pupil, allowing one to check for vignetting in the system. Indeed, the central obstruction due to M2 and its spider are visible in the ghost. We will probably also use the Red Rascal as an integral part of a future automated telescope-focusing scheme. Tracking a stellar image slowly across one edge of the slit performs effectively a Foucault knife-edge test. The direction from which the corresponding shadow crosses the Red Rascal determines which side of focus one is on.

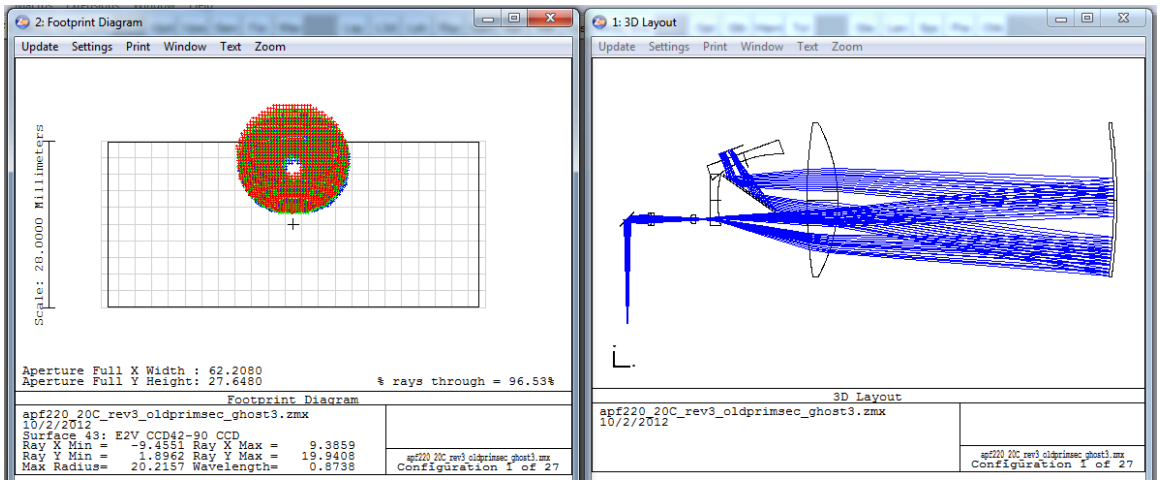

***Figure 3.39**: Ray trace of the Red Rascal ghost.*

Light from the Red Rascal ghost, and other less intense ghosts comprise a background signal that is measurable in the inter-order spacing and easily removed in the spectral extraction along with the sky signal. The next plot shows the extracted spectral order covering the A-band, an order that runs right through the Red Rascal. This absorption band is particularly useful for assessing scattered light in a spectrometer as many of the saturated lines reach true zero intensity at their centers. As this plot and the log version in the next two plots show, the Red Rascal has been removed down to the ~0.1% level in the background-corrected order extracted spectrum.

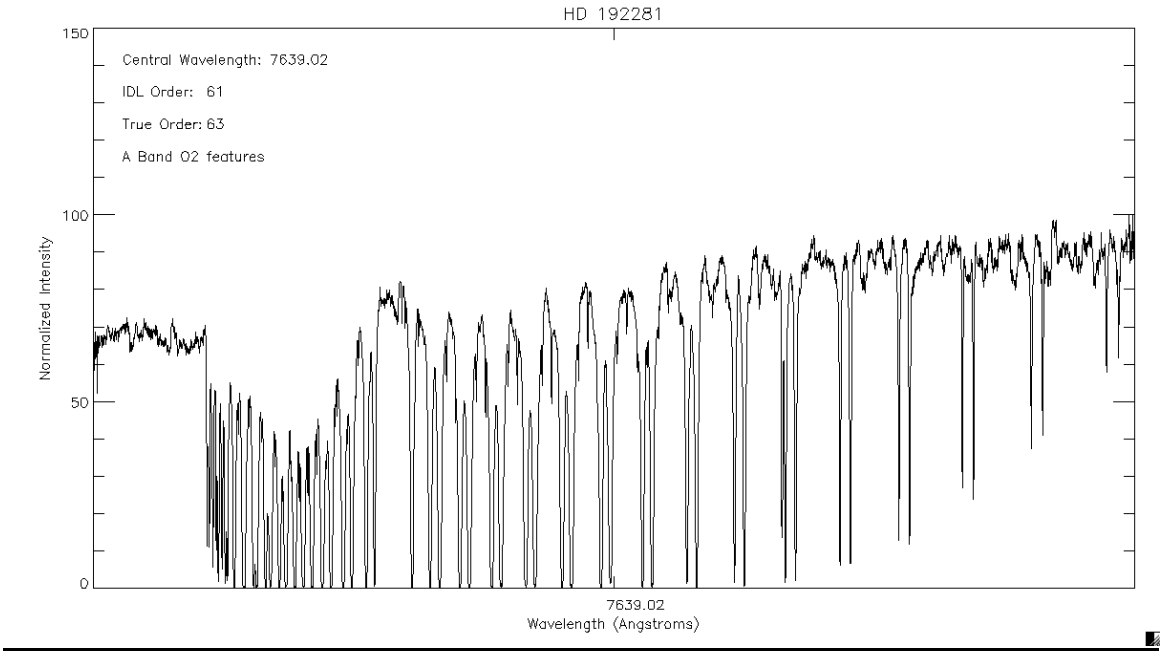

***Figure 3.40**: Extracted spectral order covering the A-band.*

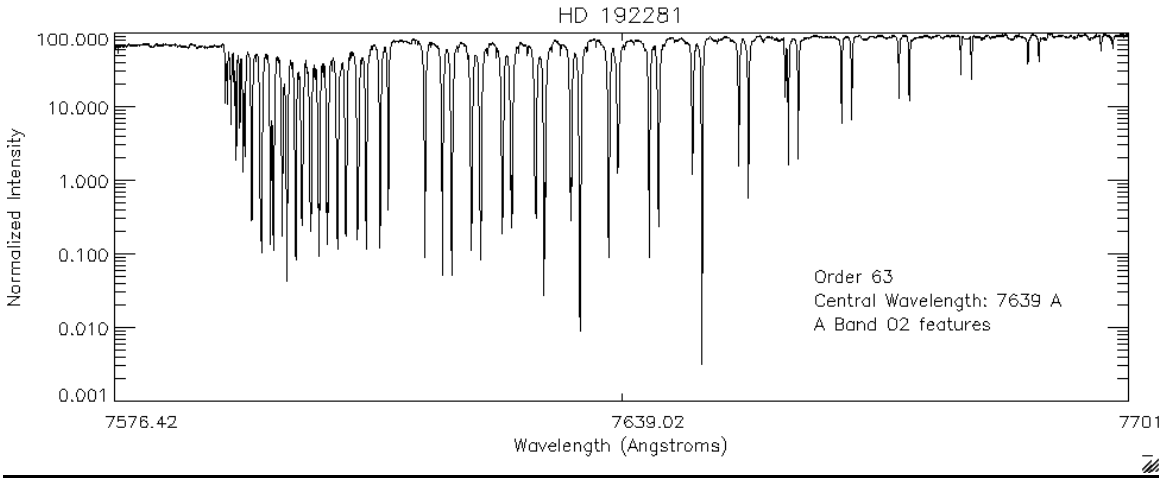

***Figure 3.41**: Extracted spectral order covering the A-band on a log intensity scale.*

The next plot shows the Na D interstellar lines in HD 192281, which also are essentially black in their centers. This spectrum was obtained through the Iodine cell, so numerous Iodine lines are also visible.

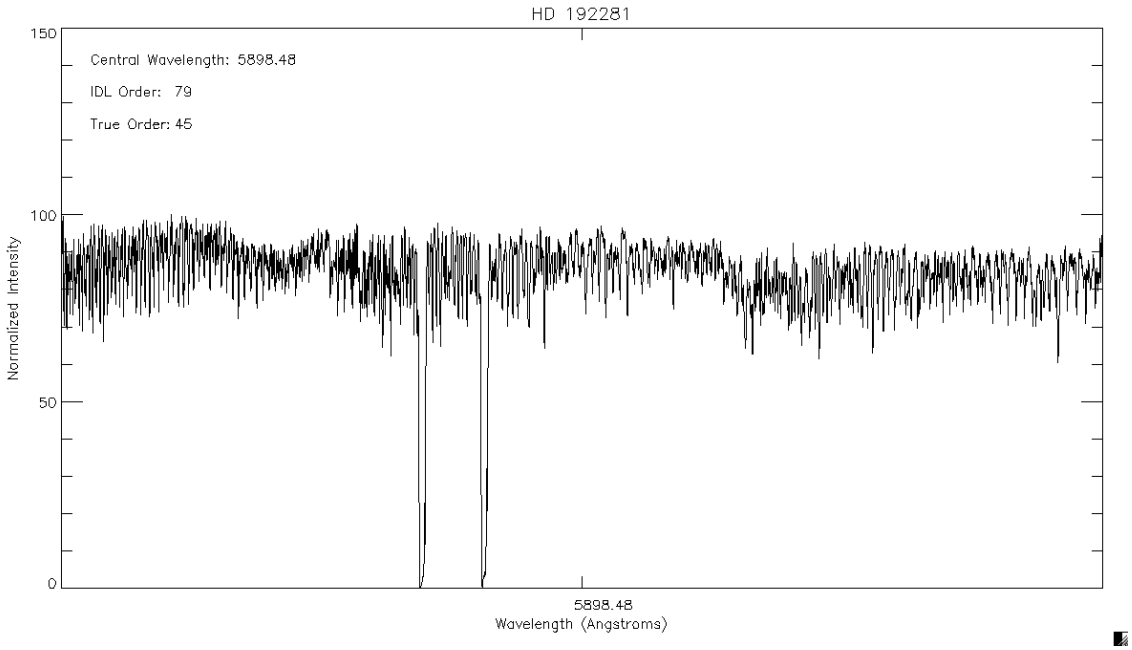

***Figure 3.42**: Interstellar Sodium D lines in the HD 192281 spectrum.*

A better look at the Iodine lines is shown in the next plot, of echelle order 88 near 530 nm.

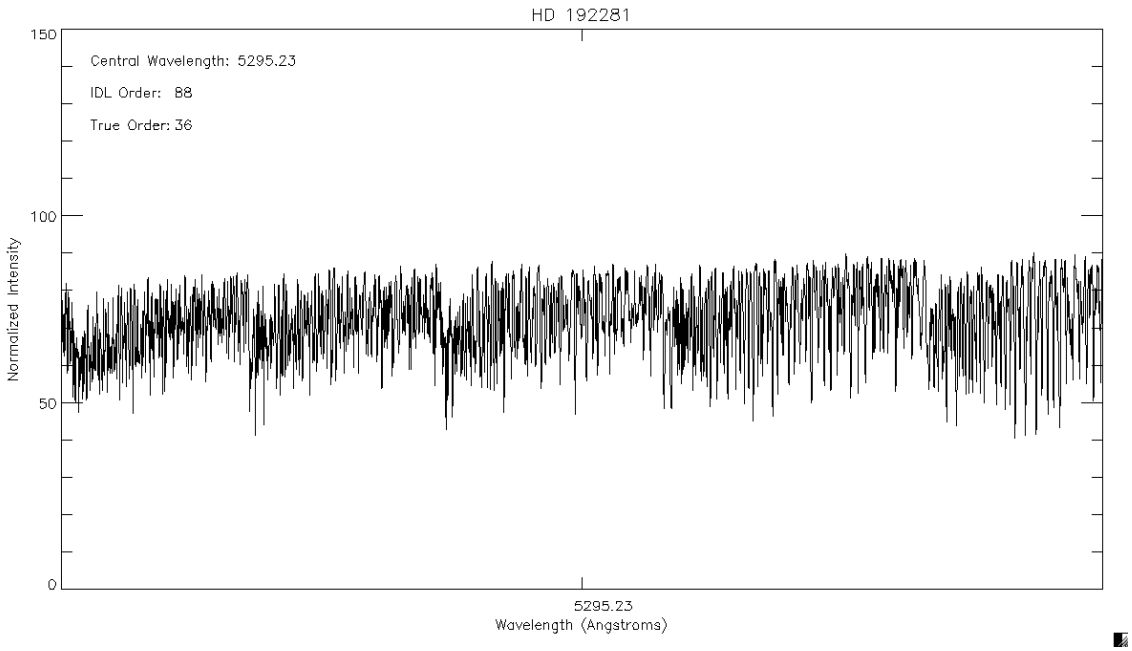

***Figure 3.43**: Iodine lines in the HD 192281 spectrum.*

3.4 Focusing the Spectrometer

Final alignment and focusing of the spectrometer was done primarily by using the Iodine lines. A focus cube of quartz lamp spectra through the Iodine cell was obtained in intervals of

100 steps (20 um) with the dewar focus stage. The Iodine orders were extracted and fully reduced. The PSF of the spectrometer at each wavelength was then derived from the Iodine lines. The next plot shows the output from a typical focus cube. Here, each sub plot shows the sigma of a single-Gaussian fit to the PSF as a function of position along the spectral order (basically the x-pixel coordinate of the CCD). Red points represent orders near the top of the Iodine region, near 600 nm. Green represents orders near the middle of the Iodine region (near 550nm) and blue represent orders near the blue end of the Iodine region, near 510 nm.

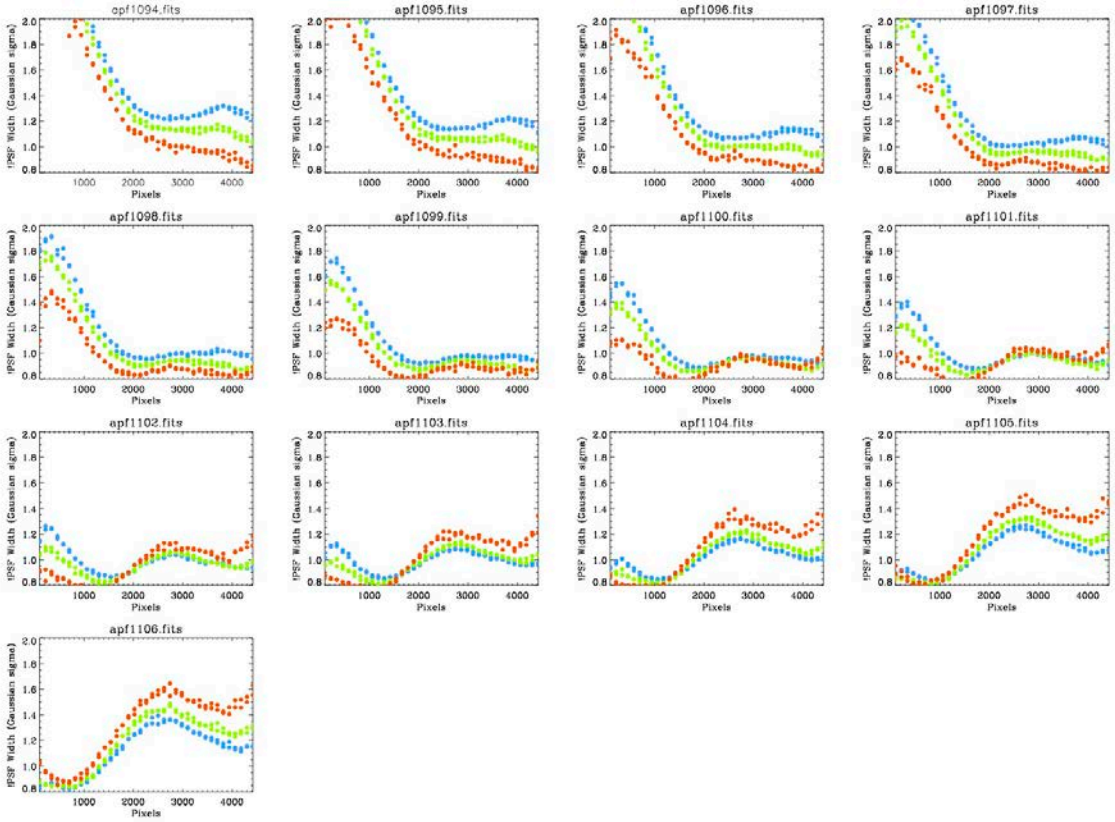

Figure 3.44*: A typical Iodine focus cube.*

The CCD dewar was pitched, yawed, and rotated as necessary to both flatten the above focus curves, and minimize the sigma of the PSF across the entire Iodine region as much as possible. In the above plots, best focus might appear to be at observation apf1102.fits, however best focus formally occurs at observation apf1101.fits. The rising points over at the left edge of apf1101.fits are actually ignored in the formal Iodine analysis as they are within 400 pixels of the CCD edge, not used in the precision RV analyses, and therefore receive zero weight. Ignoring points within 400 pixels of either end, the sigma of a Gaussian fit to the PSF at best focus is ~1 pixel. That translates to a FWHM of 2.35 pixels across most of the spectrum.

From all these points, a quantitative focus metric was then derived, calculated as a weighted mean across the CCD, weighted by the echelle blaze intensity. The next plot shows this weighted mean. Best focus was taken as the minimum of a parabolic fit to each curve. The black points show the best overall focus curve as averaged over the entire Iodine region.

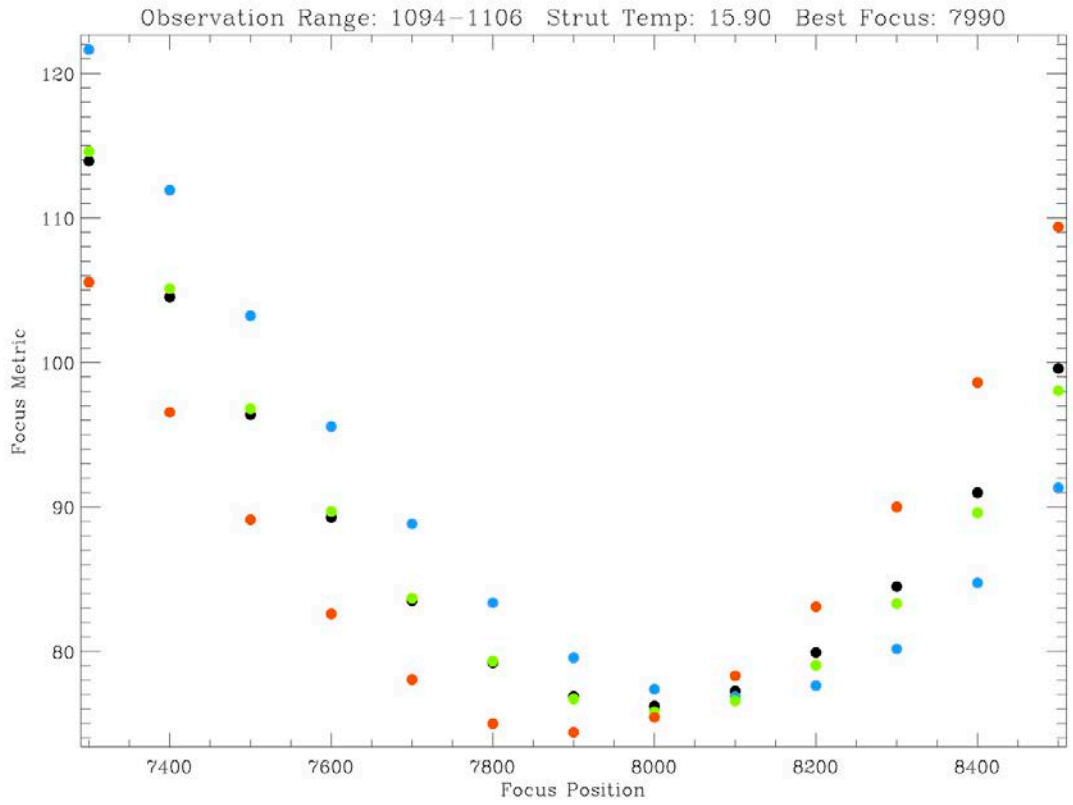

***Figure 3.45**: Weighted mean focus metric.*

3.5 Spectral Resolving Power

The next three plots show the spectral resolving power at best focus across the top, middle, and bottom of the Iodine region, as derived directly from shots of a B star through the Iodine cell. Here, slit widths of 0.5”, 0.75”, 1”, and 2” are shown.

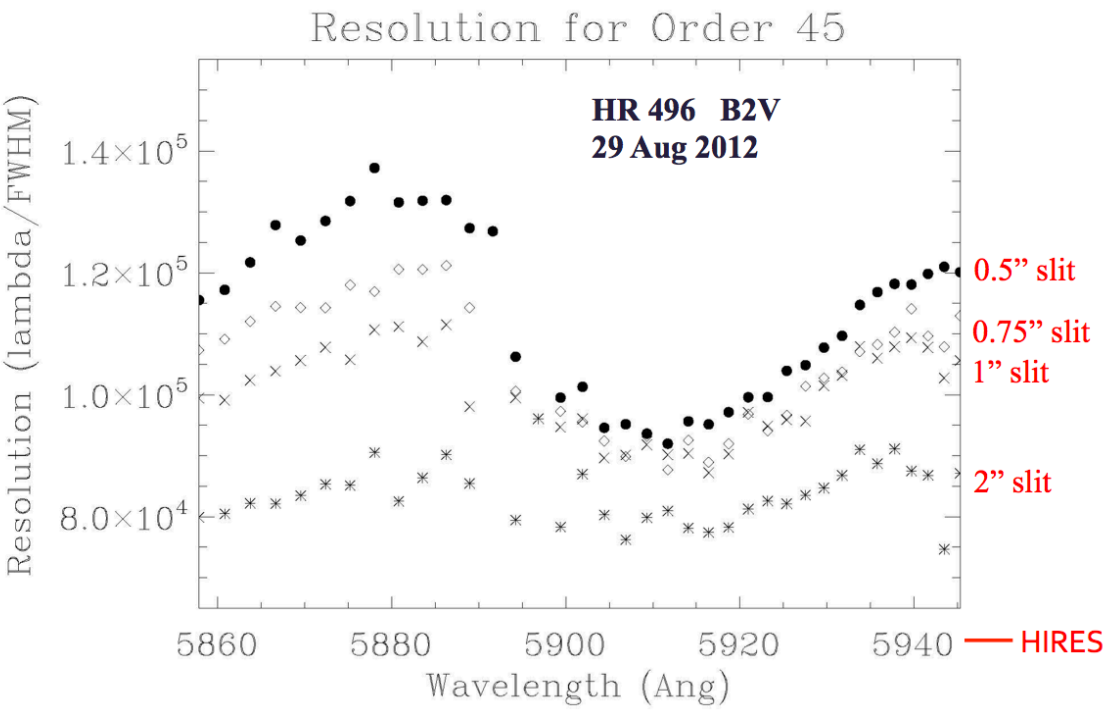
Resolution for Order 45
HR 496 B2V
29 Aug 2012
Resolution (lambda/FWHM)
1.4×10^5
1.2×10^5
1.0×10^5
8.0×10^4
5860
5880
5900
5920
5940
Wavelength (Ang)
0.5" slit
0.75" slit
1" slit
2" slit
HIRES

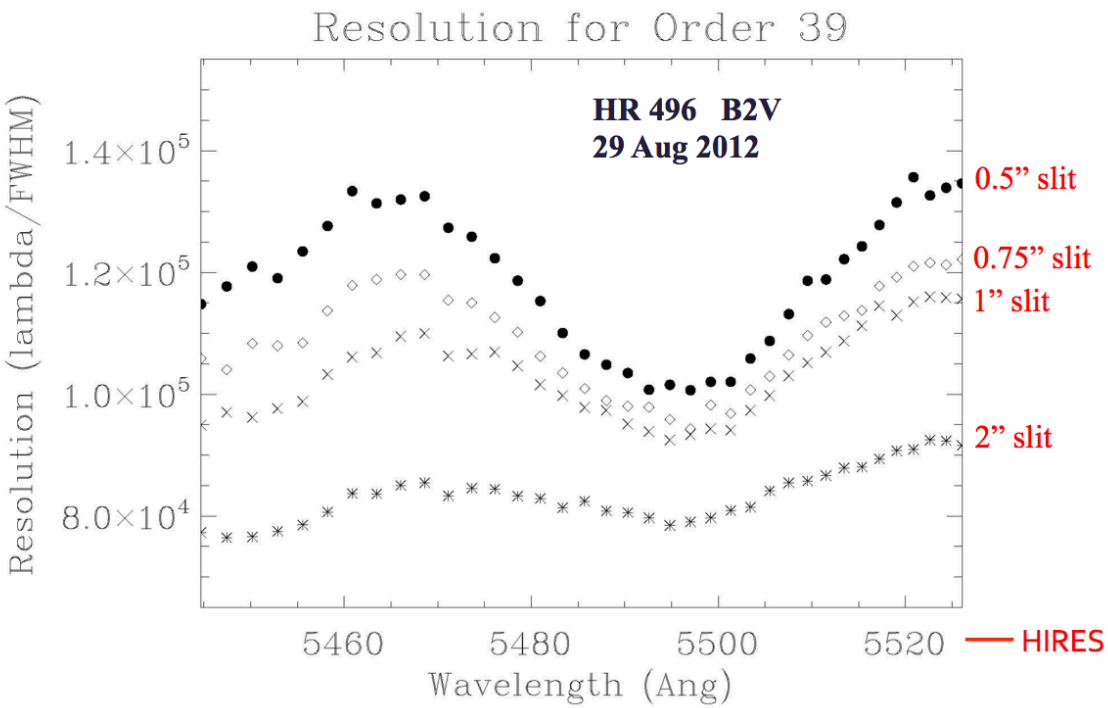
Resolution for Order 39
HR 496 B2V
29 Aug 2012
Resolution (lambda/FWHM)
1.4×10^5
1.2×10^5
1.0×10^5
8.0×10^4
5460
5480
5500
5520
Wavelength (Ang)
0.5" slit
0.75" slit
1" slit
2" slit
HIRES

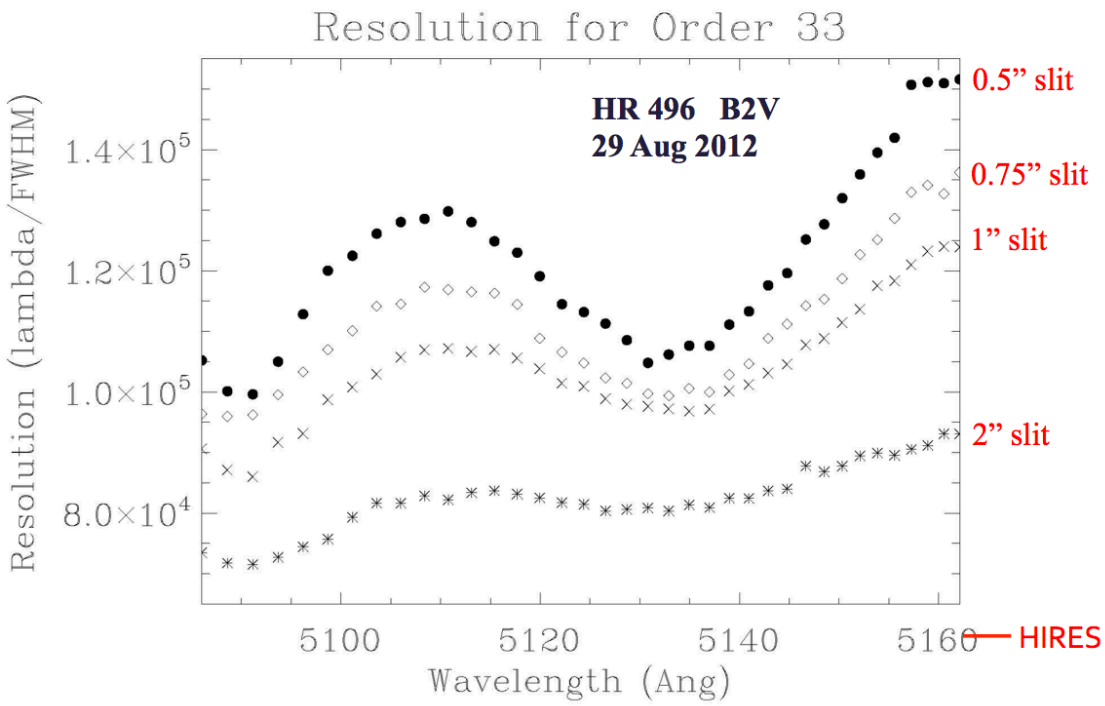

***Figure 3.46**: Spectral resolving power at best focus across the top, middle, and bottom of the Iodine region.*

With the narrow (0.5") slit, resolutions in excess of 100,000 are easily achieved across the entire Iodine region, in some cases reaching as high as 150,000. Since these curves were all obtained at a single best-focus position, further improvements in resolution can be won by slight focus adjustments for any particular region of the echelle format. The standard science decker is intended to be the 1" slit, which produces resolutions of ~100,000 to 120,000 over most of the Iodine region. Even the worst-case 2" slit still produces a resolution of at least 80,000, which is higher than the ~60,000 we currently get with our usual science decker B5 on HIRES. It is instructive to compare the present APF resolution with benchmarks from other workhorse precision RV facilities. A comparison of PSF profiles near the center of the Iodine region for Keck/HIRES, Magellan/PFS, and APF/Levy is shown in the next plot.

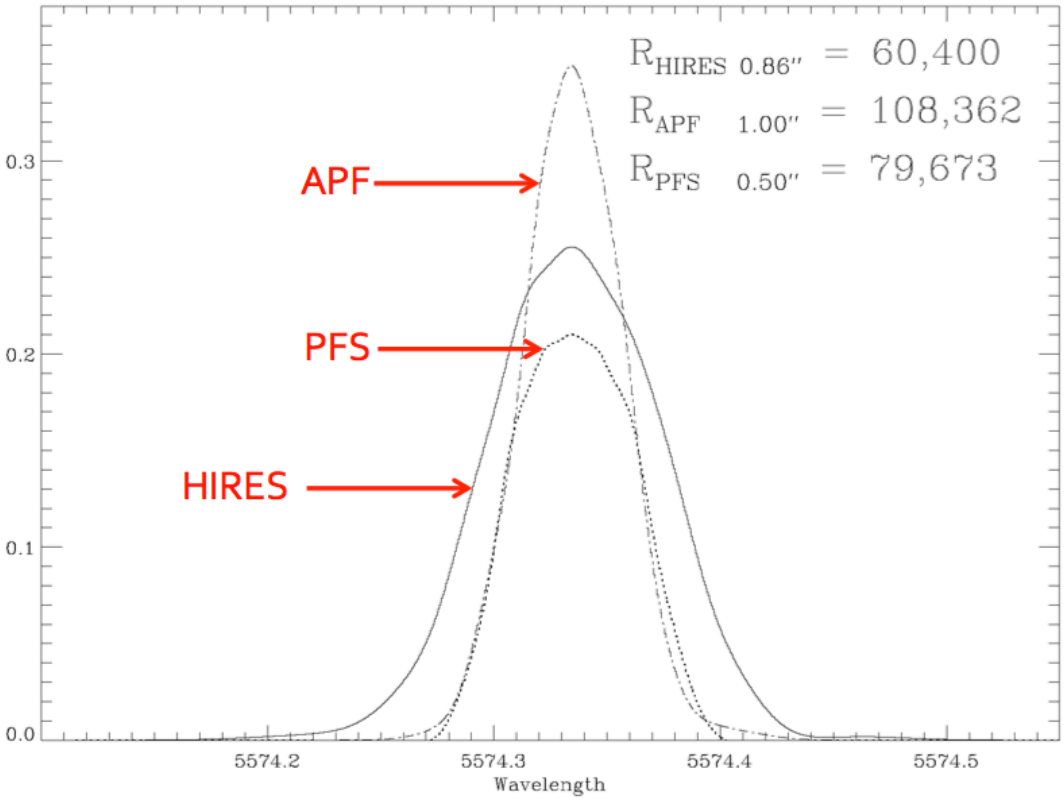

***Figure 3.47**: Comparison of PSF profiles near the center of the Iodine region for Keck/HIRES, Magellan/PFS, and APF/Levy.*

To map the PSF over the entire CCD format, we also used spectra of the Th-Ar lamp obtained through a pinhole which projects to 1 pixel at the CCD. Each emission line then appears as basically a point source, whose width and shape approximates the PSF at that position in the spectral format. The next figure shows the FWHM of Th-Ar lines across the entire echelle format.

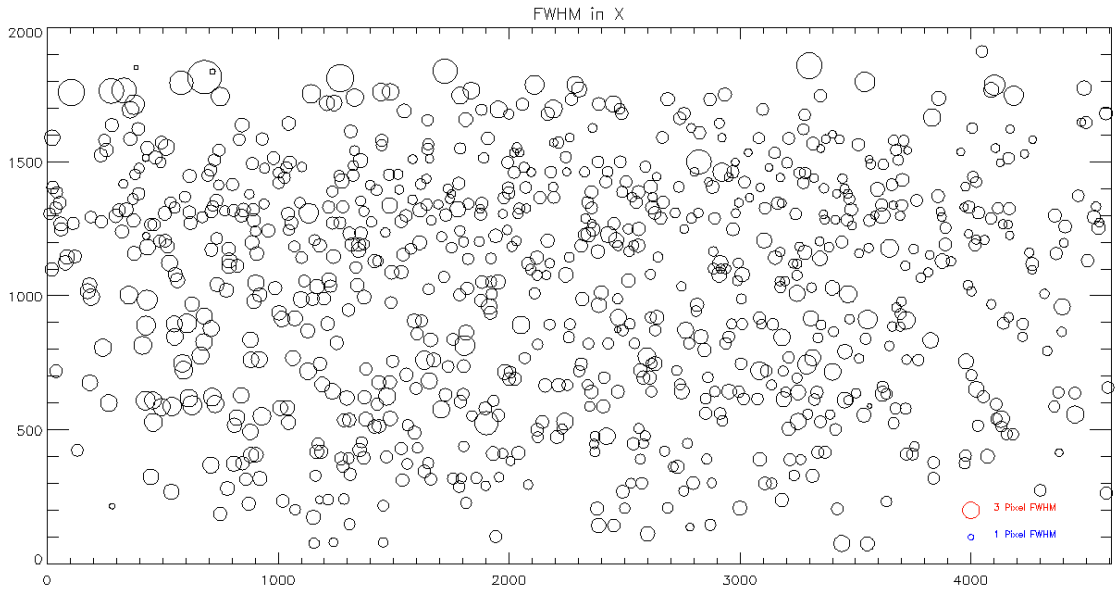

***Figure 3.48**: FWHM of Th-Ar lines across the entire echelle format.*

Even though the spectrometer PSF has been optimized explicitly for the Iodine region, the PSF is of fairly uniform high quality all the way out to the corners of the CCD.

3.6 Stability

We are still in the process of optimizing the stability of the spectrometer. Our primary tool for gauging stability is to take long series of Th-Ar spectra through the pinhole slit plate. An example of one such spectrum is shown below.

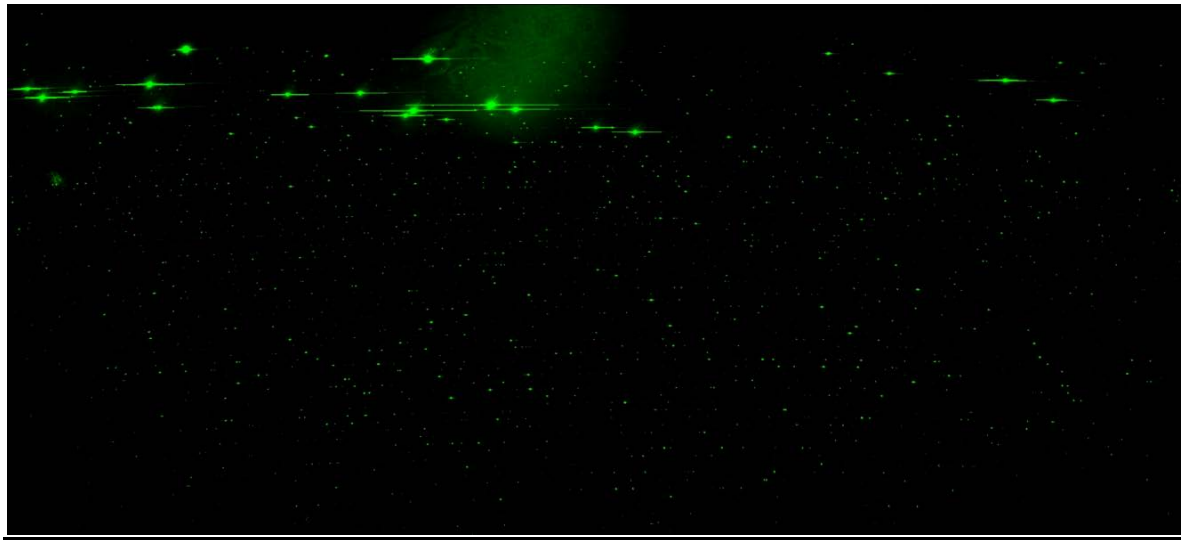

***Figure 3.49**: Thorium-Argon spectra through pinhole slit plate.*

Each pinhole Th-Ar spectrum creates essentially "a thousand points of light" across the CCD. An idl routine written by Jenny Burt locates and measures the centroids of all points,

filtering out saturated lines, cosmic rays, and CCD defects. The mean centroid in x and y is then computed over all lines and used as a proxy for any global shifts in the image over x and y.

The next figure illustrates one such recent stability run. Here, Th-Ar spectra were taken sequentially over a roughly 5-day period. Red points represent the shifts of the spectrum in Y (cross-disperser direction) relative to the first spectrum of the set. Blue points show the shifts in X (echelle dispersion direction). Green points show the temperature of the dome interior air. We run similar plots against many of the environmental variables being monitored and captured in the FITS keywords. The goal of all of these tests is to establish the quiescent stability level of the spectrum's position and also to ascertain if/how shifts in the spectrum are correlated with variations in the many environmental variables being monitored (temperature, pressure, etc.).

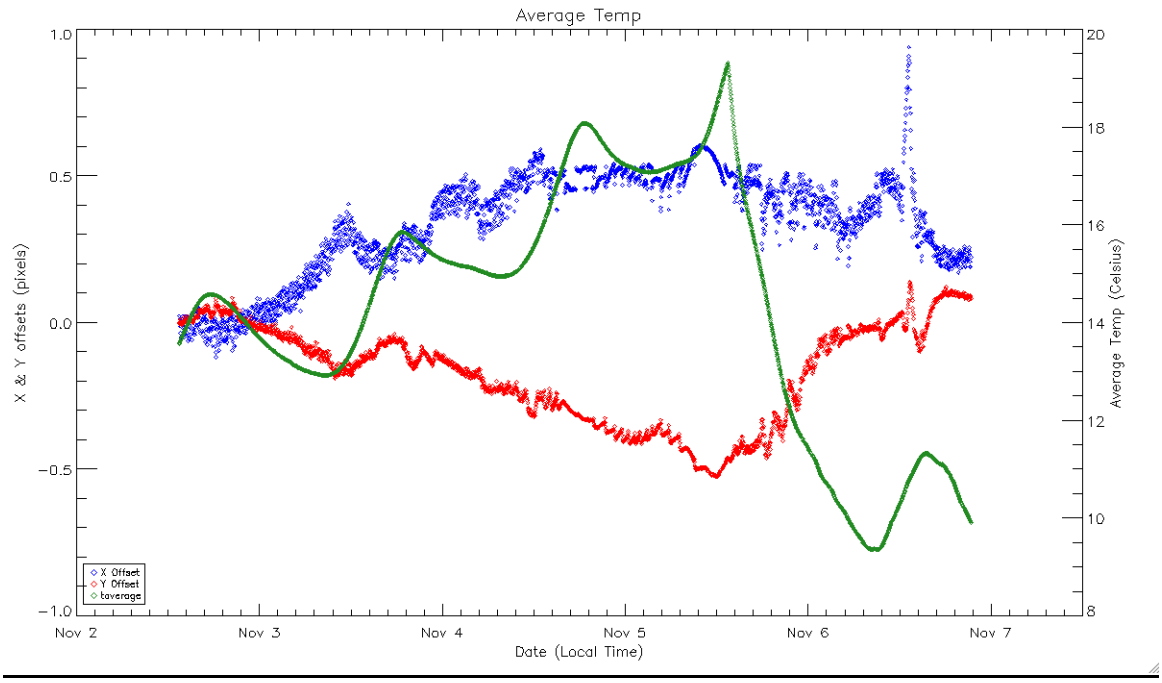

***Figure 3.50**: Five-day stability run of pinhole Thorium-Argon spectra.*

The next plot shows a more detailed view of the above plot, centered on Nov 5.0. This expanded view reveals a clear correlation between spectrum X-shifts and the cycling of the cooling flow through the heat exchanger and/or the heater turning on/off. Here, green crosses are Y-shifts of the spectrum, while red plus signs are X-shifts. Blue vertical lines mark the heater turning on/off. The thick purple line is the heat exchanger valve, cycling between 0 and 25% open. We use these temperature history plots to make improvements to the thermal control system to reduce temperature-driven spectral movements. In this case, it is likely that the heat exchanger fan may be blowing directly on a strut that needs shielding from such intense direct air flow. This may indicate that we need to cover all struts with some insulating sleeves such that they can only sense temperature through highly diffused/mixed air flow.

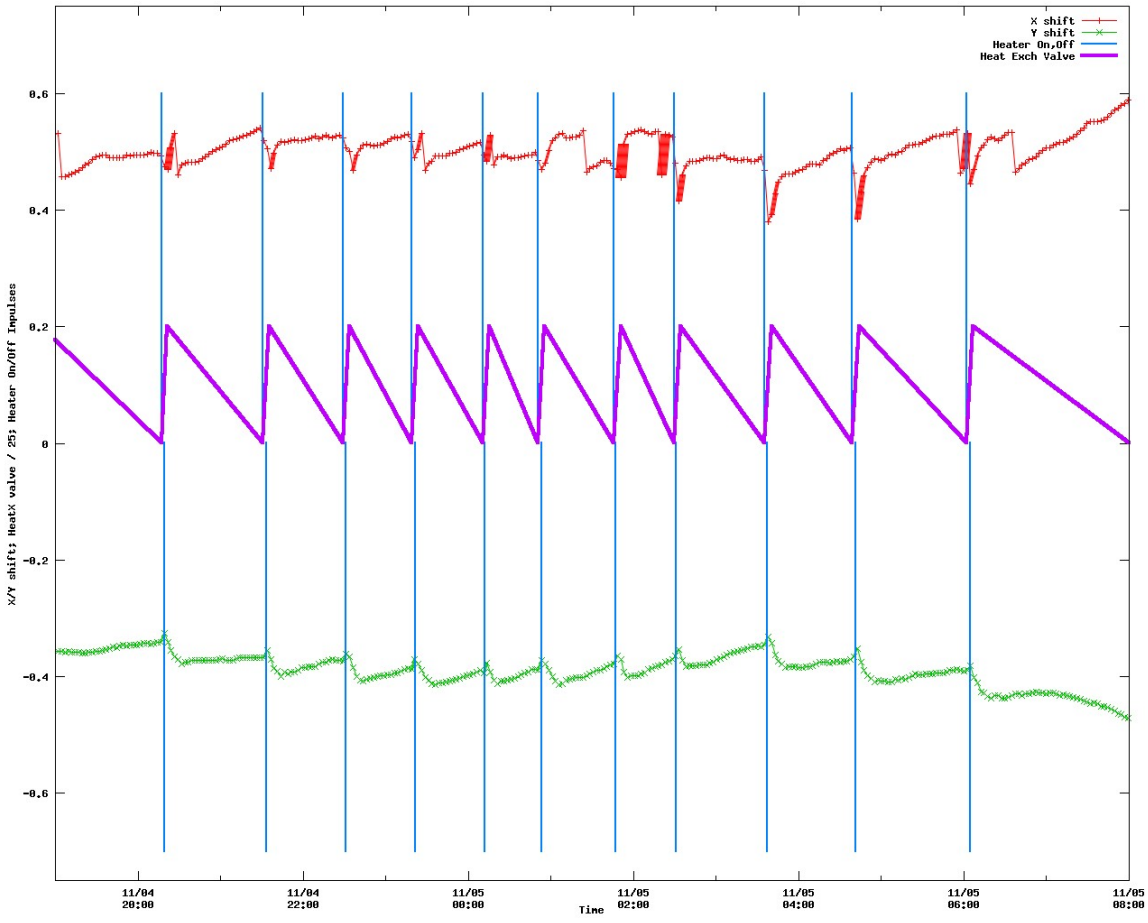

Figure 3.51*: Expanded detail of Figure 3.50.*

From other similar stability testing, we learned that changes in the external spectrometer air temperature were causing dimensional changes in the insulated enclosure, and that was somehow getting transmitted into either the dewar mount, or into the determinate structure of the spectrometer, perhaps causing temperature-related spectrum shifts. We then identified several places where the upper lid of the enclosure was making contact with the breadboard when the shell's latches were tightened down. The electrical feed-through panel was also producing a moment on the optical table when the shell was clamped. So we remounted the feed-through panel to disconnect it completely from the table. Further stability experiments are underway to verify that we have solved, or at least greatly reduced these problems. This is still very much work in progress.

In seeking to minimize spectrum motion on the CCD, it is implicitly assumed that RV errors scale with spectrum stability. This may or may not be true, but seemed a reasonable and conservative starting point. In fact, spectrum drifts at HIRES are often as large as 1-2 pixels

within a single night. Despite this drift, RV precision of 1-2 m/s has been achieved with HIRES. So clearly, the Iodine technique has been quite effective in dealing with such drifts. With the Levy, we are now at a level of < 0.5 pixel/day drifts, and still making progress on reducing drift. Atmospheric pressure variations also produce spectrum drifts and are not currently being controlled in APF. Zemax modeling predicts that typical atmospheric pressure variations during the course of a night produce shifts of ~0.1 pixels, while the most extreme historic record pressure variations would cause X shifts of up to 0.3 pixels, and Y-shifts of 1.6 pixels. So, at < 0.5 pixels/day, we are already close to the limits set by pressure variations, and to stabilize further would require pressure-controlling (or evacuating) the spectrometer.

4 Science Testing and Observations

4.1 Efficiency

In preparation for assessing the end-to-end efficiency of the optical train, a detailed spreadsheet model was developed. The Levy optical train contains 47 surfaces (including the telescope mirrors and CCD). The next figure shows this spreadsheet efficiency model.

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
		Efficiency	Efficiency		from 10nov2005 estimate	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Efficiency	Col J /Col F
Surface #	Surface type	550 nm	490nm	Comments	500 nm	350 nm	400 nm	450 nm	490nm	550 nm	600 nm	650 nm	700 nm	750 nm	800 nm	850 nm	900 nm	950 nm	(Old/new)
1	Atmosphere	1	1	no atmosphere yet...	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
2	central obstr	0.969	0.969	414mm OD M2 baffle + 2350mm OD M1	0.976	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.969	0.993
3	M1	0.860	0.860	Fresh Al -5% as per 10/3/12 633nm measurement	0.920	0.860	0.860	0.860	0.860	0.860	0.850	0.850	0.836	0.830	0.817	0.827	0.855	0.870	0.935
4	M2	0.932	0.932	Original fresh coated - 5% as per 10/3/12 meas	0.960	0.797	0.915	0.929	0.932	0.932	0.931	0.929	0.927	0.925	0.923	0.920	0.919	0.919	0.971
5	M3	0.917	0.915	Original fresh coated - 6.7% as per 10/3/12 meas	0.960	0.766	0.893	0.911	0.915	0.917	0.916	0.915	0.913	0.912	0.910	0.907	0.905	0.904	0.953
6	ADC1_S1	0.9982	0.9947	freshly hardened solgel measured	0.997	0.955	0.973	0.987	0.995	0.998	0.999	0.997	0.995	0.993	0.990	0.988	0.985	0.982	0.998
7	ADC1_glass	0.998	0.998	13mm of PBM2Y at 490nm for all waves	1.000	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998
8	ADC1_S2	0.9982	0.9947	freshly hardened solgel measured	0.997	0.955	0.973	0.987	0.995	0.998	0.999	0.997	0.995	0.993	0.990	0.988	0.985	0.982	0.998
9	ADC2_S1	0.9953	0.99078	freshly hardened solgel measured	0.997	0.953	0.972	0.984	0.991	0.995	0.997	0.996	0.994	0.991	0.989	0.987	0.983	0.981	0.994
10	ADC2_glass	0.998	0.998	13mm of PBM2Y at 490nm for all waves	1.000	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998
11	ADC2_S2	0.9953	0.99078	freshly hardened solgel measured	0.997	0.953	0.972	0.984	0.991	0.995	0.997	0.996	0.994	0.991	0.989	0.987	0.983	0.981	0.994
12	Guider BS front	0.96	0.96	4% beamsplitter	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	1.000
13	Guider BS rear	0.9993	0.9972	fresh solgel wag- copied from F_1 (Flat side)	1.000	0.964	0.982	0.993	0.997	0.999	0.999	0.997	0.995	0.993	0.990	0.988	0.986	0.984	0.997
14	I2_S1	0.99	0.99	wag guess at ZC&R AR coating	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
15	I2_glass	0.999	0.999	6mm of BK7 at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
16	I2_S2	0.96	0.96	uncoated pyrex	1.000	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960
17	I2_gas	0.7	1	guess at 30% for I2 absorption	1.000	1.000	1.000	1.000	1.000	0.700	0.900	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
18	I2_S3	0.96	0.96	uncoated pyrex	1.000	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960	0.960
19	I2_glass	0.999	0.999	6mm of BK7 at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
20	I2_S4	0.99	0.99	wag guess at ZC&R AR coating	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
21	Slit	0.840	0.840	slit1 (2"x7.5") loss: in 1.7" FWHM seeing	1.000	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840	0.840
22	MaxMirror	0.99	0.99	1% for dust on MaxMirror	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	1.000
23	J_1	0.995	0.995	ECI bid curves for bbar on CaF2	0.997	0.980	0.995	0.994	0.995	0.995	0.995	0.995	0.995	0.996	0.997	0.996	0.992	0.985	0.998
24	J_glass	0.9999	0.9999	CaF2 6.9mm thick at 490nm for all waves	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
25	J_2/H_1	0.99	0.99	wag for grease-coupled interface	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
26	H_glass	0.9999	0.9999	S-LAL7 3 mm thick at 490nm for all waves	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
27	H_2	0.995	0.994	ECI bid curves for bbar on LAL7	0.997	0.980	0.994	0.993	0.994	0.995	0.994	0.994	0.995	0.996	0.996	0.995	0.990	0.985	0.997
28	G_1	0.995	0.995	ECI bid curves for bbar on CaF2	0.997	0.980	0.995	0.994	0.995	0.995	0.995	0.995	0.995	0.996	0.997	0.996	0.992	0.985	0.998
29	G_glass	0.9999	0.9999	CaF2 6.8mm thick at 490nm for all waves	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
30	G_2	0.995	0.995	ECI bid curves for bbar on CaF2	0.997	0.980	0.995	0.994	0.995	0.995	0.995	0.995	0.995	0.996	0.997	0.996	0.992	0.985	0.998
31	F_1 (flat side)	0.9994	0.9972	freshly hardened solgel measured	0.997	0.964	0.982	0.993	0.997	0.999	0.999	0.997	0.995	0.993	0.990	0.988	0.986	0.984	1.001
32	F_glass	0.999	0.999	BAL15Y 10mm at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
33	F_2 (ccv side)	0.9992	0.996	freshly hardened solgel measured	0.997	0.948	0.976	0.990	0.996	0.999	0.999	0.997	0.994	0.992	0.989	0.986	0.983	0.980	0.999
34	E_1 (shallow side)	0.9968	0.99866	freshly hardened solgel measured	0.997	0.993	0.997	0.999	0.999	0.997	0.995	0.992	0.990	0.987	0.985	0.983	0.981	0.980	1.002
35	E_glass	0.996	0.996	I-FPL-51Y 40mm thick at 490nm for all waves	1.000	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996
36	E_2 (steep side)	0.9984	0.99853	freshly hardened solgel measured	0.997	0.987	0.994	0.997	0.999	0.998	0.997	0.995	0.993	0.992	0.990	0.988	0.986	0.984	1.002
37	D_1	0.984	0.9979	freshly hardened solgel measured	0.997	0.966	0.981	0.986	0.998	0.984	0.980	0.976	0.972	0.969	0.965	0.962	0.960	0.958	1.001
38	D_glass	0.999	0.999	11mm of I-BSM51Y at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
39	D_2/C_1	0.99	0.99	wag for oil-coupled interface	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
40	C_glass	0.999	0.999	100mm CaF2 at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
41	C_2/B_1	0.99	0.99	wag for oil-coupled interface	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
42	B_glass	0.999	0.999	10mm of I-BSM51Y at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
43	B_2	0.9821	0.9735	freshly hardened solgel measured	0.997	0.944	0.953	0.965	0.974	0.982	0.986	0.988	0.988	0.987	0.986	0.984	0.982	0.981	0.977
44	A_1 (shallow side)	0.9973	0.9979	freshly hardened solgel measured	0.997	0.994	0.997	0.998	0.998	0.997	0.996	0.995	0.993	0.992	0.990	0.988	0.987	0.985	1.001
45	A_glass	0.99	0.99	50mm of S-FPL51 at 490nm for all waves	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
46	A_2 (steep side)	0.9972	0.9993	freshly hardened solgel measured	0.997	0.993	0.998	1.000	0.999	0.997	0.995	0.992	0.990	0.988	0.985	0.983	0.981	0.979	1.003
47	Prism_1	0.9989	0.9985	freshly hardened solgel at 35.7deg	0.997	0.976	0.988	0.996	0.999	0.999	0.997	0.995	0.992	0.989	0.986	0.984	0.981	0.979	1.002
48	Prism_glass	0.9895	0.9895	105mm of BSL7Y at 490nm for all waves	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
49	Prism_2	0.9987	0.9985	freshly hardened solgel at 35.7deg	0.997	0.979	0.989	0.997	0.999	0.999	0.997	0.994	0.991	0.988	0.985	0.982	0.980	0.978	1.002
50	Echelle	0.600	0.600	Dekker and Hoose (1992) ESO proc.	0.650	0.555	0.570	0.585	0.600	0.600	0.610	0.610	0.610	0.600	0.600	0.600	0.610	0.620	0.923
51	Prism_2	0.9987	0.9985	freshly hardened solgel measured	0.997	0.979	0.989	0.997	0.999	0.999	0.997	0.994	0.991	0.988	0.985	0.982	0.980	0.978	1.002
52	Prism_glass	0.9895	0.9895	105mm of BSL7Y at 490nm for all waves	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
53	Prism_1	0.9989	0.9985	freshly hardened solgel measured	0.997	0.976	0.988	0.996	0.999	0.999	0.997	0.995	0.992	0.989	0.986	0.984	0.981	0.979	1.002
54	A_2 (steep side)	0.9972	0.9993	freshly hardened solgel measured	0.997	0.993	0.998	1.000	0.999	0.997	0.995	0.992	0.990	0.988	0.985	0.983	0.981	0.979	1.003
55	A_glass	0.99	0.99	50mm of S-FPL51 at 490nm for all waves	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
56	A_1 (shallow side)	0.9973	0.9979	freshly hardened solgel measured	0.997	0.994	0.997	0.998	0.998	0.997	0.996	0.995	0.993	0.992	0.990	0.988	0.987	0.985	1.001
57	B_2	0.9821	0.9735	freshly hardened solgel measured	0.997	0.944	0.953	0.965	0.974	0.982	0.986	0.988	0.988	0.987	0.986	0.984	0.982	0.981	0.977
58	B_glass	0.999	0.999	10mm of I-BSM51Y at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
59	C_2/B_1	0.99	0.99	wag for oil-coupled interface	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
60	C_glass	0.999	0.999	100mm CaF2 at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
61	D_2/C_1	0.99	0.99	wag for oil-coupled interface	1.000	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990	0.990
62	D_glass	0.999	0.999	11mm of I-BSM51Y at 490nm for all waves	1.000	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
63	D_1	0.984	0.9866	freshly hardened solgel measured	0.997	0.966	0.981	0.986	0.987	0.984	0.980	0.976	0.972	0.969	0.965	0.962	0.960	0.958	0.990
64	E_2 (steep side)	0.9984	0.99853	freshly hardened solgel measured	0.997	0.987	0.994	0.997	0.999	0.998	0.997	0.995	0.993	0.992	0.990	0.988	0.986	0.984	1.002
65	E_glass	0.996	0.996	I-FPL-51Y 40mm thick at 490nm for all waves	1.000	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996
66	E_1 (shallow side)	0.9968	0.99866	freshly hardened solgel measured	0.997	0.993	0.997	0.999	0.999	0.997	0.995	0.992	0.990	0.987	0.985	0.983	0.981	0.980	1.002
67	MLD mirror	0.96	0.968	MLD meas @ 55 deg (fresh) wrt MaxMirr	0.960	0.107	0.981	0.976	0.969	0.972	0.973	0.972	0.961	0.960	0.955	0.955	0.956	0.953	1.009
68	FF_1 (ccv side)	0.9985	0.993	freshly hardened solgel measured	0.997	0.950	0.970	0.985	0.993	0.998	0.999	0.998	0.995	0.993	0.989	0.986	0.983	0.981	0.996
69	FF_glass	0.998	0.998	BAL15Y 10mm at 490nm for all waves	1.000	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998
70	FF_2 (flat side)	0.998	0.9964	freshly hardened solgel measured	0.997	0.967	0.979	0.991	0.996	0.999	0.998	0.996	0.994	0.991	0.988	0.985	0.982	0.980	1.000
71	CCD	0.87	0.9	E2V 42-90 data sheet at T= -90C	0.825	0.550	0.800	0.880	0.900	0.880	0.850	0.810	0.750	0.670	0.560	0.440	0.310	0.200	1.091
	TOTAL	0.142302366	0.204128954		0.367	0.009	0.158	0.202	0.228	0.164	0.208	0.219	0.202	0.186	0.172	0.162	0.158	0.152	0.620
		Efficiency	Efficiency		from 10nov2005 estimate														
	Surface	550 nm	490nm	Comments	500 nm														
		#surfaces	Avg eff/surface	Combined degradation effect		Degraded total @350nm	400 nm	450 nm	490nm	550 nm	600 nm	650 nm	700 nm	750 nm	800 nm	850 nm	900 nm	950 nm	
	# solgel surfaces	25	0.99000	0.778		0.007	0.123	0.157	0.177	0.127	0.161	0.171	0.157	0.145	0.133	0.126	0.123	0.118	
	# telescope mirrors	3	0.90000	0.729		0.006	0.115	0.147	0.166	0.119	0.151	0.160	0.147	0.136	0.125	0.118	0.115	0.110	
	effect of solgel+tel mirrs			0.567		0.005	0.090	0.114	0.129	0.093	0.118	0.124	0.114	0.106	0.097	0.092	0.090	0.086	
	# of surfaces encountered	46	0.995	0.794		0.007	0.126	0.160	0.181	0.130	0.165	0.174	0.160	0.148	0.136	0.129	0.125	0.120	
						350	400	450	490	550	600	650	700	750	800	850	900	950	

***Figure 4.1**: Spreadsheet efficiency model.*

The model predicts that, with pristine coatings, the APF optical train should be able to achieve a peak end-to-end efficiency of ~23%, as shown in the next plot. However, the APF coatings have already been in place now for ~4 years. While every effort has been made to keep the coatings well protected, clean, and dry, stored solgel coatings might conceivably degrade with time and/or exposure to humidity and air-borne contaminants. The figure below also

shows the expected overall efficiency for all surfaces degraded by 0.5% and by 1%. All APF surfaces uniformly degraded by 1% throughout would drop the overall APF efficiency to just under 15%. Finally the purple line on the plot below shows the measured efficiency of HIRES with no losses from slit or telescope mirrors included. While not generally appreciated, HIRES is running at only about 7% efficiency when used with its B5 decker in 1" seeing.

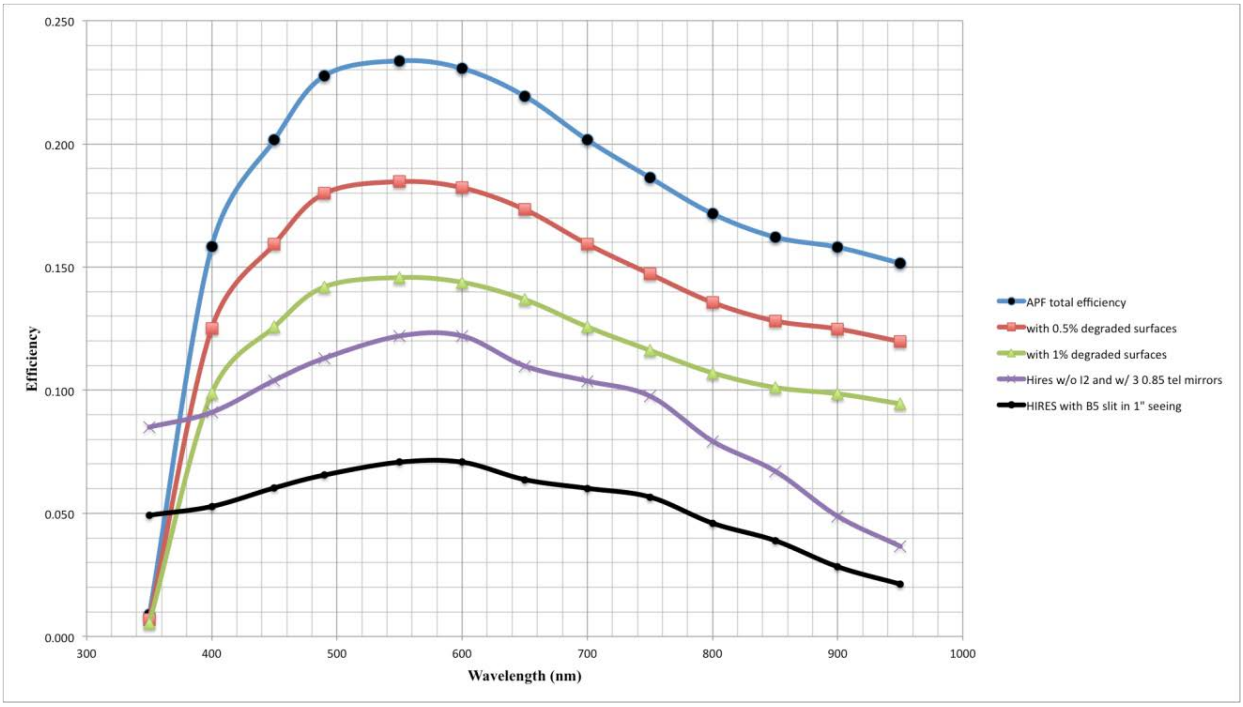

***Figure 4.2**: Expected overall efficiency of the complete APF optical train.*

Measurements of the overall efficiency of the APF system (telescope + spectrometer + CCD) have been greatly complicated by persistent high frequency oscillations of the telescope while tracking. These oscillations, mostly at 33 and 66 hz, cause elongation and/or Lissajous figure distortions of the seeing disk, resulting in unpredictable light loss with even a wide slit. There is no easy provision in APF to remove the slit, so all of our measurements were made with the widest, longest available slit- a 2” x 7.5” slit. For these measurements, we used several well-known photometric standard stars.

The next figure shows the measured efficiency at blaze peaks that was obtained on the night of July 23, 2012 using the V=7.55 spectrophotometric 05V standard star HD 192281.

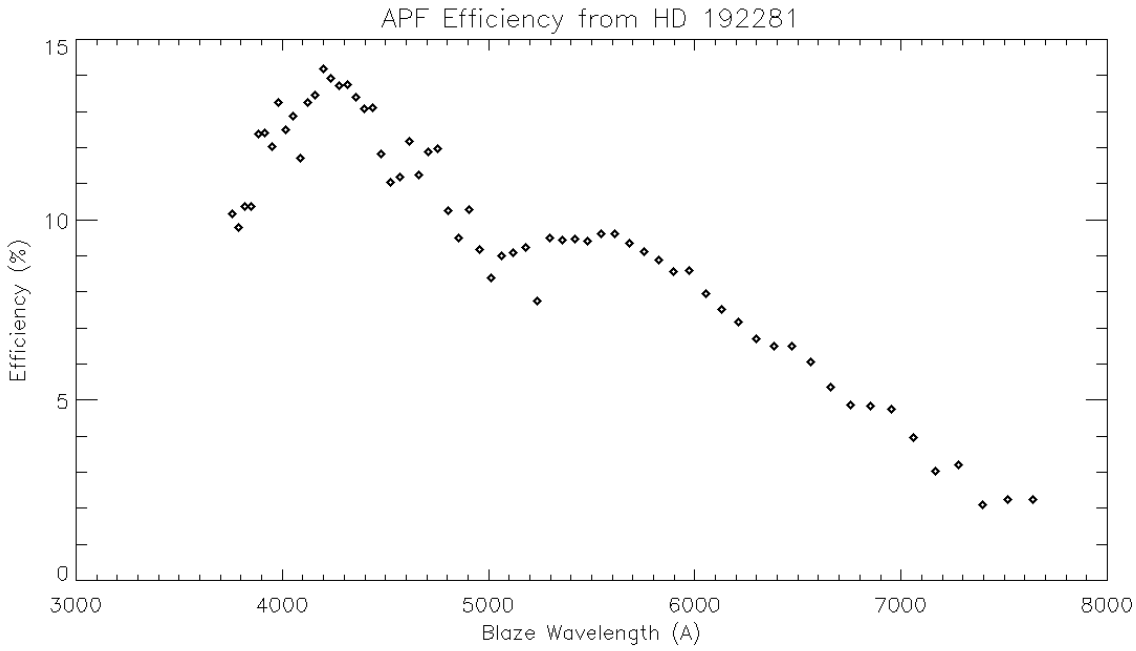

***Figure 4.3**: Measured efficiency of APF optical train at blaze peaks.*

These results may be compromised by an unknown contribution from telescope high frequency oscillations, so should be taken as a lower limit at this early stage. A much better and more meaningful speed measurement will be to see how long it takes to reach equal RV precision to Keck/HIRES on stars with well-known RV behavior as will now be discussed.

4.2 Test Observations of Known RV Standard Stars

A set of stars with well-known RV variations was selected to assess both the relative speed of APF (with respect to the benchmark Keck/HIRES) and also the RV precision of the facility. A short list of moderately bright stars was selected that were either known to have no RV variations (RV null stars), or were known to harbor short period planets in well-confirmed, uncomplicated systems. The list (below) was chosen such that all targets would be available for 1st-half nights, for at least 4-5 months in Fall, 2012.

HD168746	+18:21:49.78244	-11:55:21.6571	G5	7.946
HD185144	+19:32:21.59026	+69:39:40.2354	K0V	4.7
HD185269	+19:37:11.74092	+28:29:59.5055	G0IV	6.669
HD187123	+19:46:58.11299	+34:25:10.2810	G5	7.83
HD209458	+22:03:10.77207	+18:53:03.5430	G0V	7.65
HD217014	+22:57:27.98004	+20:46:07.7912	G5V	5.49
HD217107	+22:58:15.54115	-02:23:43.3840	G8IV	6.18
HD221354	+23:31:22.20763	+59:09:55.8647	K2V	6.74
HD1461	+00:18:41.86724	-08:03:10.8024	G1.5V	6.453
HD7924	+01:21:59.11491	+76:42:37.0335	K0	7.18
HD9826	+01:36:47.84216	+41:24:19.6443	F8V	4.1
HD10700	+01:44:04.08338	-15:56:14.9262	G8V	3.5
HD10780	+01:47:44.83458	+63:51:09.0040	K0V	5.63
HD12846	+02:06:30.24337	+24:20:02.3736	G2V	6.9
HD25665	+04:09:35.03944	+69:32:29.0145	G5	7.7
HD32147	+05:00:48.99977	-05:45:13.2303	K3V	6.22

***Figure 4.4**: Test stars for 1st-half night observing during Fall, 2012.*

Due to the persistent problems with the telescope servo-oscillations, coverage of the list has been less than hoped for. Often, a target had to be skipped to avoid the telescope going into potentially damaging oscillations. Despite these problems, some half-dozen observations and corresponding templates have been obtained so far, on a small sample of stars with known RV behavior. In the discussion that follows, we note that the last shots taken on local date Oct. 20, 2012 were affected by strong oscillations and/or bad observing conditions for HD 168746, HD 185269, HD 187123, HD 209458 and HD 10700. Only HD 185144 was not affected by oscillations on that night.

A caveat- the velocity precision tests discussed below were taken in parallel with significant activity occurring in and around the spectrometer, and with almost complete lack of thermal control of the spectrometer. The next figure shows the temperature history of the spectrometer over the ~1 month period of these observations. The thick black line shows the average of the camera barrel strut temperatures. These struts are the most thermally active components in the spectrometer, and their average is an excellent measure of the temperature

variations experienced by the optics over this period. Red arrows mark the epochs of on-sky observations.

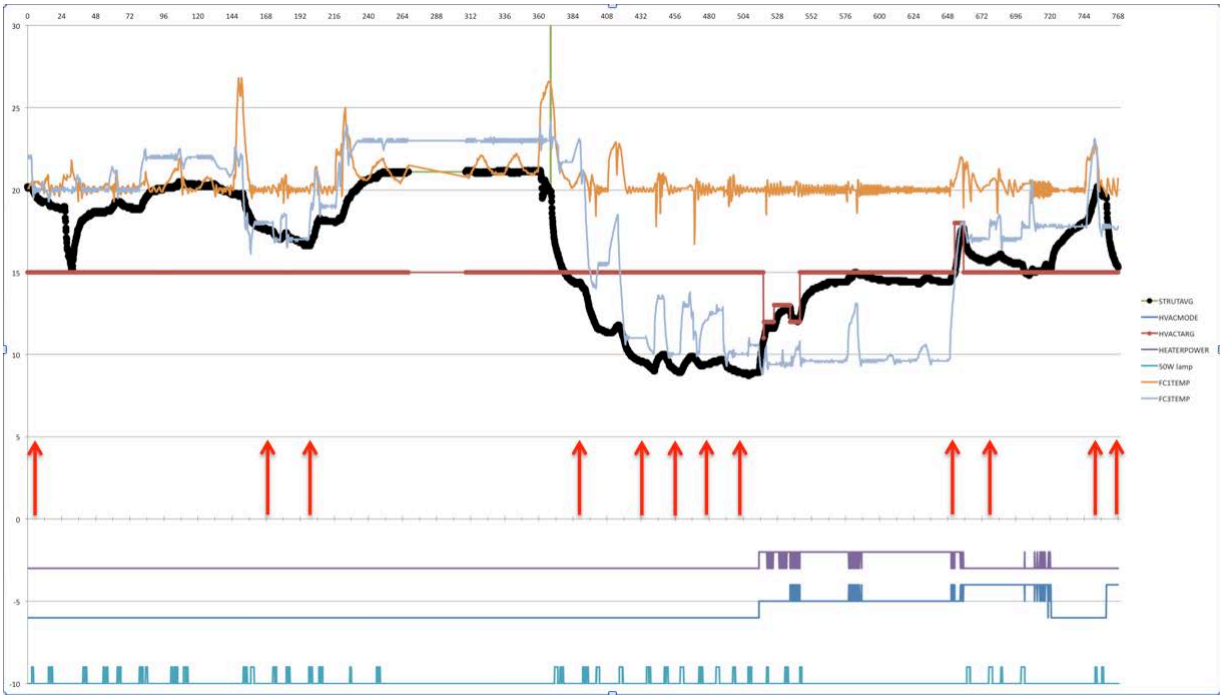

***Figure 4.5**: Temperature history of the spectrometer over ~1 month observation period.*

As is evident, the spectrometer underwent significant thermal variations, up to 14C throughout the on-sky tests. Also, many of these observations were plagued by pronounced high frequency telescope oscillations, whose contribution to the overall error budget is yet unknown. So the velocity precision results shown below are likely to be worst-case values.

4.2.1 HD 168746

This V=7.75 G5V star has a well-known 6.402-d planet with a mass of 0.23 M_{jup}. A screenshot of the SYSTEMIC console is shown in the next figure. We have a total of 180 RV's on this star, the most recent 7 of which are from APF. A global fit to all these data reproduces the known orbit quite closely, and further improves the orbital parameters slightly, by adding another ~33 orbital phases to the observational time base. The RMS of the global fit is dominated by the relatively poor CORALIE data.

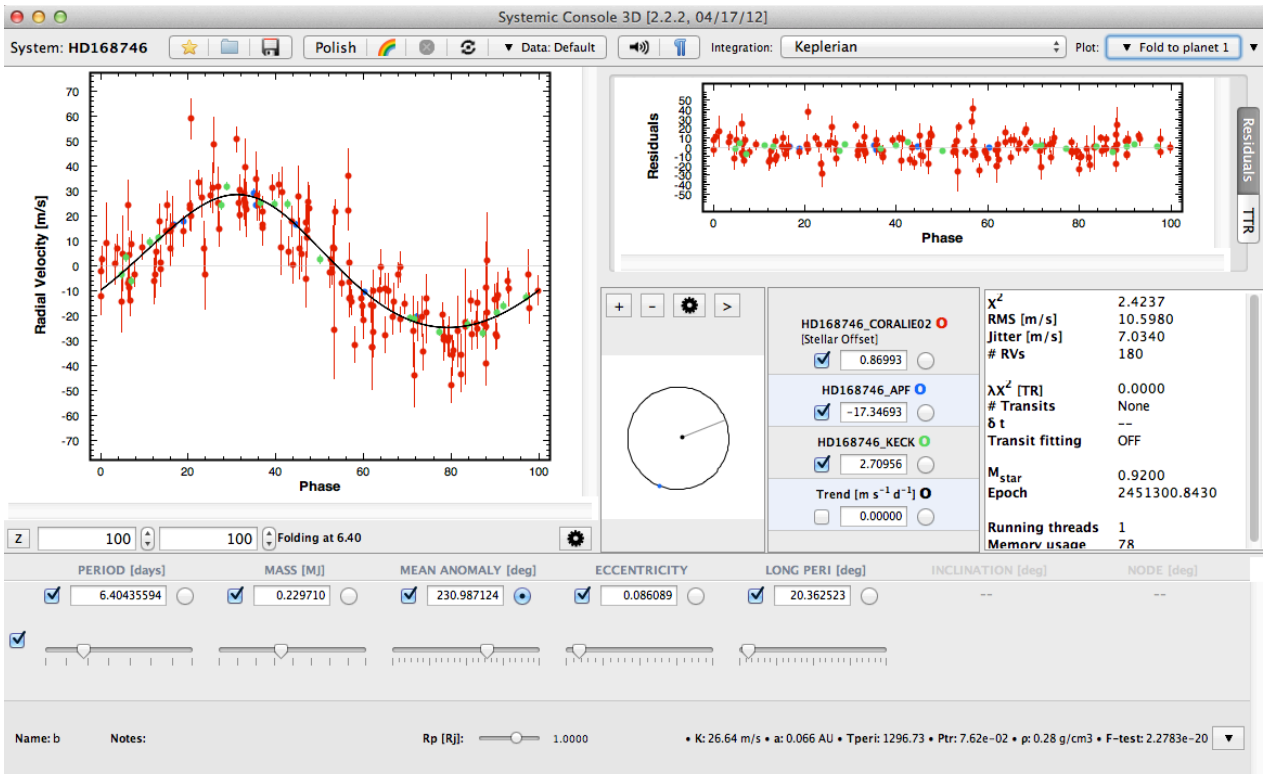

***Figure 4.6**: HD 168746 Keplerian model, all velocities.*

We then removed the CORALIE and Keck data, leaving only the seven APF points to fit. A fit to this orbit allowing only Voffset to adjust yielded an RMS=8.85 m/s.

We then allowed the MA to vary along with the velocity offset in the final fit to APF velocities alone. In principle, one should not need to use the MA as an independent variable, but rather should just phase shift the original MA by 2π (T - T_0) / P, where (T - T_0) is the time difference between the first points of the CORALIE and APF data sets, and P is the orbital period. However, in this case, the additional APF points actually improved the overall orbit, so we allowed MA to adjust, finding a 25.88-degree difference between the expected and re-optimized MA. The fit obtained by allowing both Voffset and MA to adjust achieved RMS=1.553 m/s, in reasonable agreement with the mean of the formal internal errors of 1.24 m/s. The optimized fit is shown in the next figure. Using HIRES velocities alone, the RMS=3.14 m/s for this star. So APF precision seems to be about a factor of two higher than Keck/HIRES at this point. A typical exposure time on this star with Keck/HIRES is 262s whereas it was 900s with APF. So, ignoring APF's better precision on this V=7.75 star, APF is only about a factor of 3.44 times slower than Keck/HIRES.

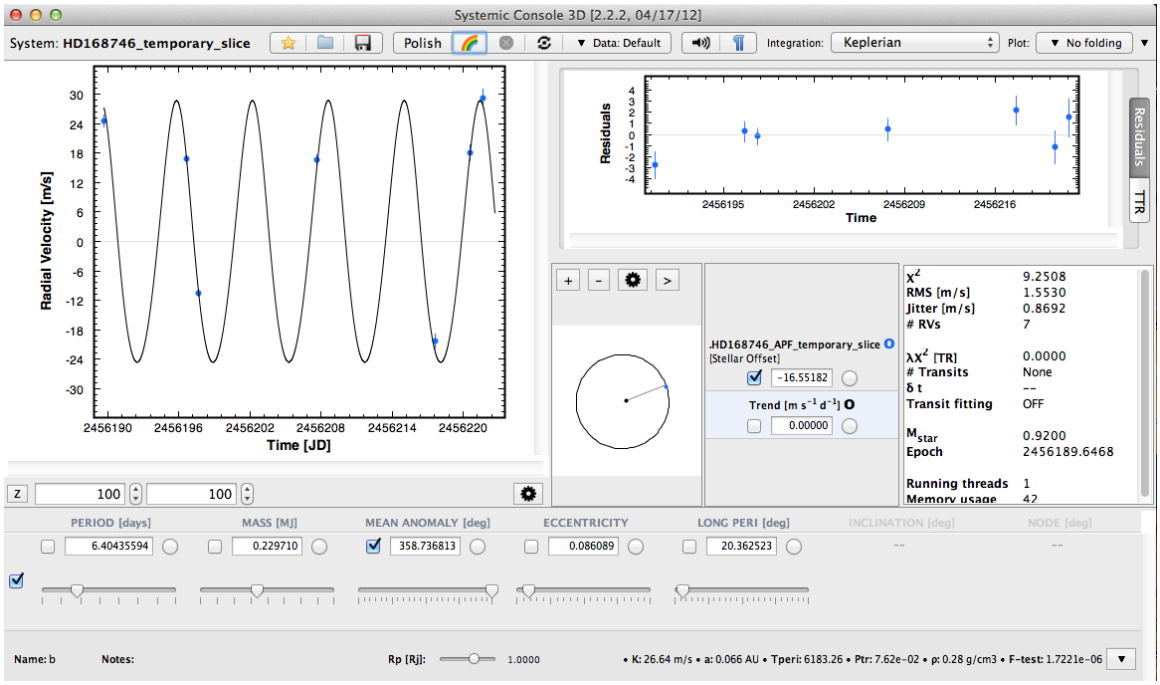

***Figure 4.7**: HD 168746 Keplerian model, APF velocities alone.*

4.2.2 HD 185144

This star is Sigma Dra, a well-known V=4.7 K0V star that has long been used by other teams as an RV "null" star, or constant velocity standard. Our SYSTEMIC console fit to the ~15-year data string (next figure) shows that, in fact, there may well be a long-period sub-Saturn-mass planet in a ~20-year orbit. But for the purposes of assessing stability over time scales of a few years or less, this is effectively a very stable "null" star.

The RMS of the fit is largely dominated by HIRES uncertainties at 1.78 m/s.

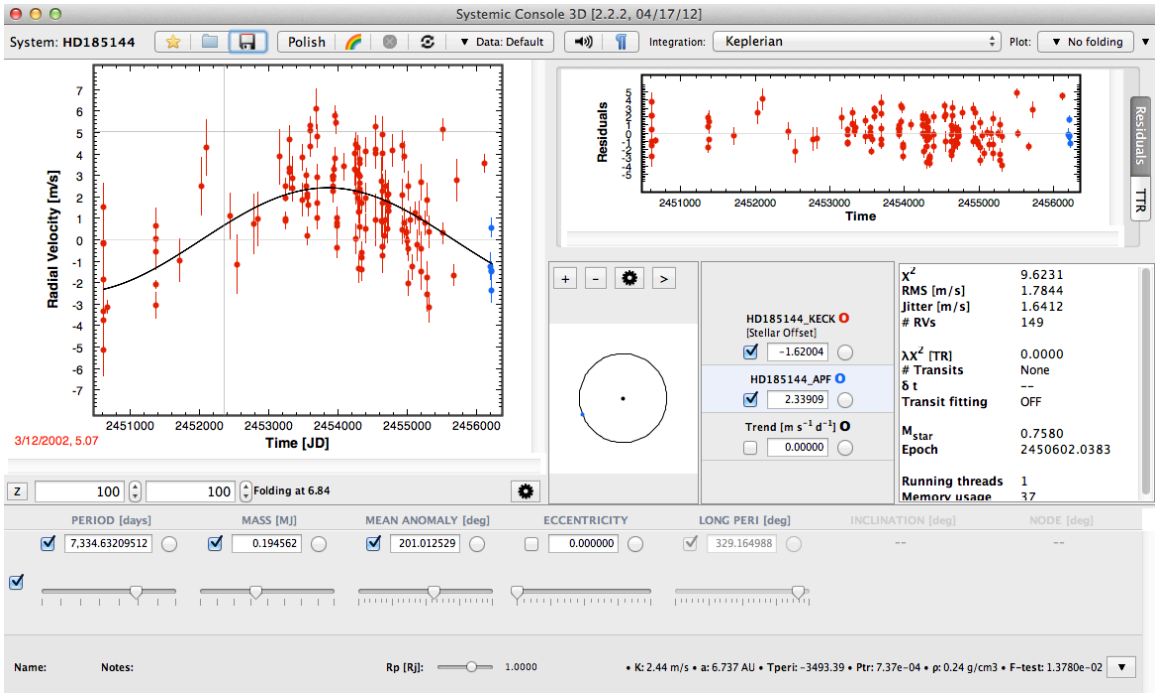

***Figure 4.8**: HD 185144 Keplerian model, all velocities.*

The next plot shows the APF RV's, fit to a simple constant-velocity model. The five APF points show an RMS=97 cm/s. The mean internal error for these points is 51 cm/s. This star did somewhat better than others on the target list due, ironically, to its very northerly declination of 69 degrees. In many cases, when setting up on this star, we encountered serious telescope oscillations and skipped observing this star under such cases. But once the PMAC tuning improved, we were able to get back on this star without the serious oscillations. So the data set for HD 185144 is probably the least affected of the entire suite of test stars in terms of systematic errors due to telescope oscillations.

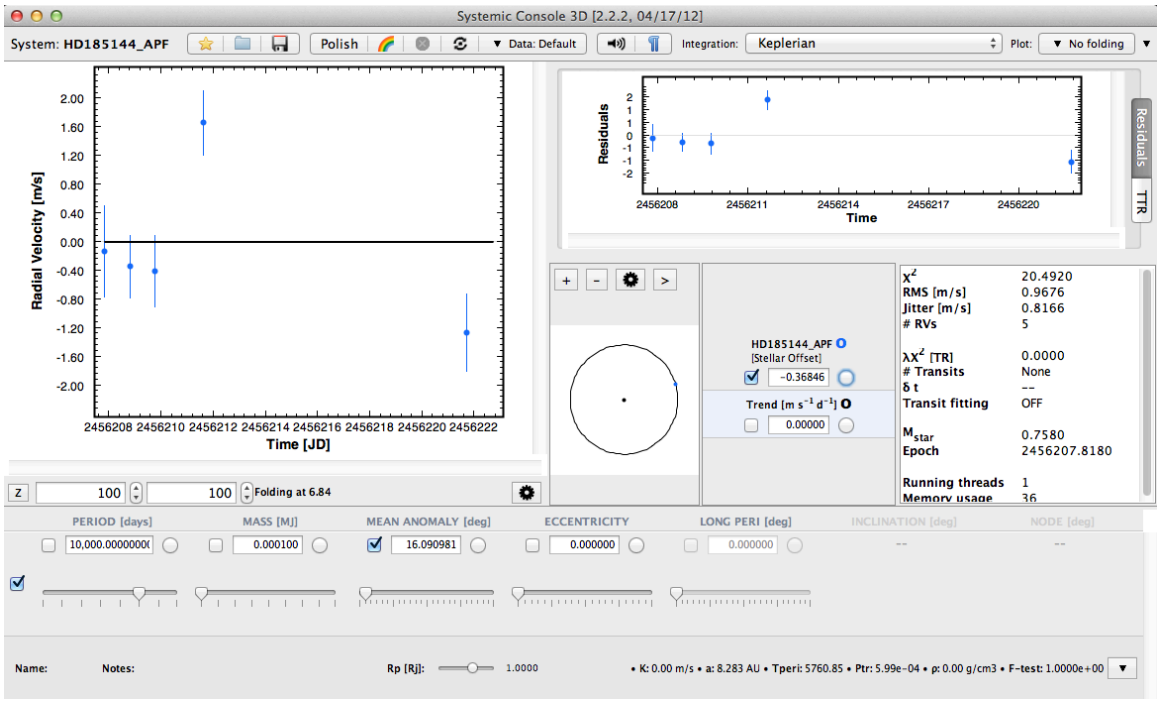

***Figure 4.9**: HD 185144 Keplerian model, APF velocities alone.*

A typical exposure time on this star with Keck/HIRES is 8.2s, whereas it was 79s with APF. So, for this star, APF is a factor of 9.7 times slower than Keck/HIRES. However, the RMS achieved with APF appears to be a good factor of 2 higher.

4.2.3 HD 185269

This V=6.7 G0IV star hosts a well-known 6.8-d planet of ~1.0 M_{jup} that was discovered with the Lick/Hamilton way back in 2006. Our team has not been following this star at Keck. The Lick data is quite noisy by comparison with APF, but the orbit was already well constrained by only the Lick data. Adding our APF data produced a slight improvement to the orbital parameters by virtue of it greatly extending the time base, and by supplying a crucial point that (fortuitously) occurred right at the phase of minimum velocity. The RMS of the global fit was 8.6 m/s and dominated entirely by the Lick velocity uncertainties as shown in the next figure.

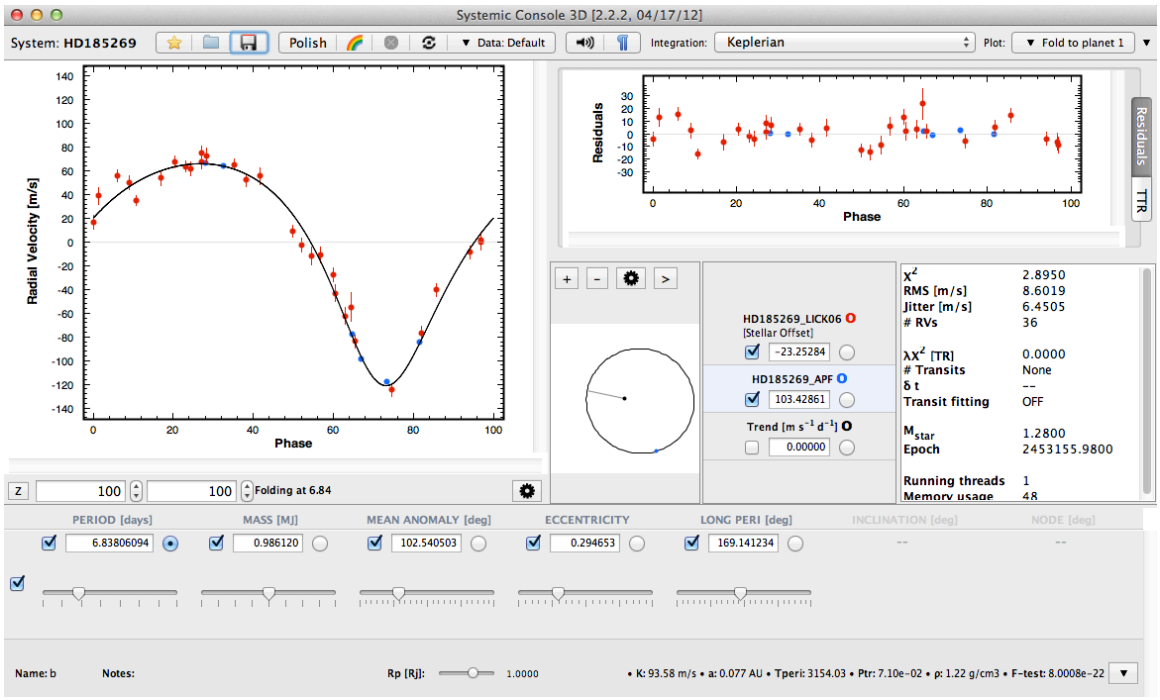

***Figure 4.10**: HD 185269 Keplerian model, all velocities.*

Removing the Lick points, and allowing the velocity offset to re-adjust and a MA polish to shift the epoch to the first APF velocity yielded an RMS=1.57 m/s for the APF points with respect to the 1-planet model. In this case, the MA=335.517985 derived from the optimized fit is very close to the value of 335.534392 expected from the epoch differences. The RMS of the fit agrees quite well with the results above for HD 168746, APF is delivering a precision at least as good as 1-2 m/s. The mean internal uncertainty of the HD 185269 velocities was 1.24 m/s, again quite similar to that for HD 168746.

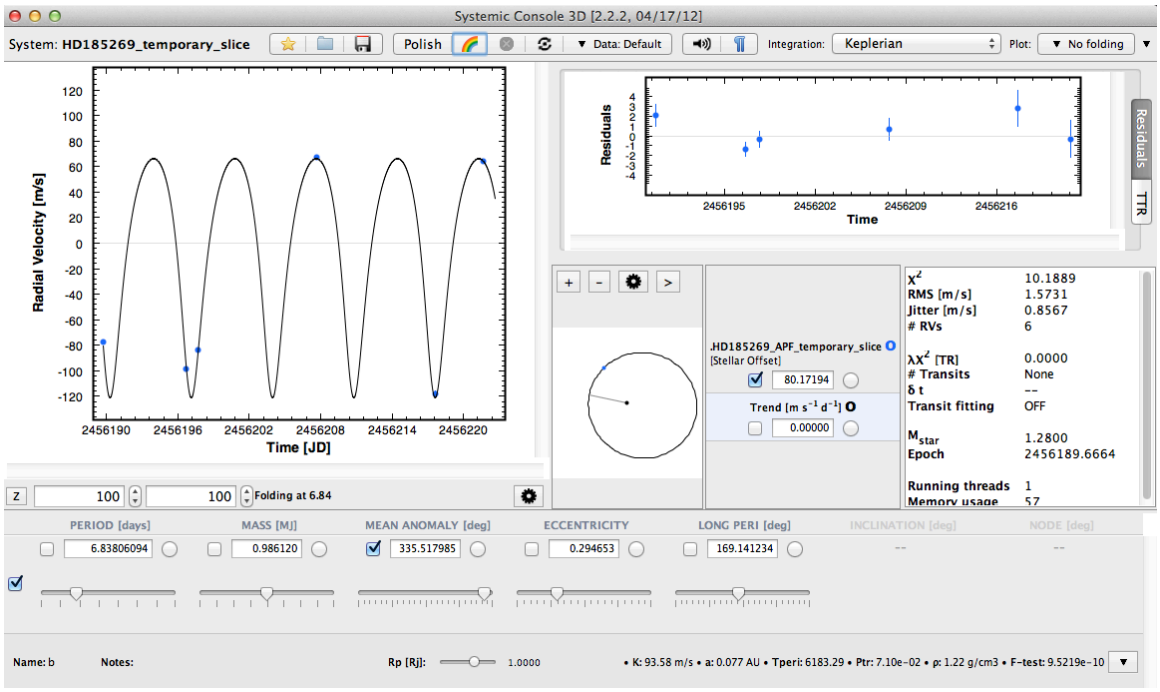

***Figure 4.11**: HD 185269 Keplerian model, APF velocities alone.*

4.2.4 HD 187123

This V=7.83 G5V star has a well-known 3.1-d planet with a mass of 0.52 M_{jup}, plus a1.8 M_{jup} planet in a 9.2-year orbit. Both orbits are well constrained by a total of 142 ELODIE and Keck/HIRES velocities. Our additional 4 new APF velocities produced only minor tweaks to the orbital parameters. Again, the RMS=7.44 m/s of the global orbital solution is dominated by the large ELODIE uncertainties.

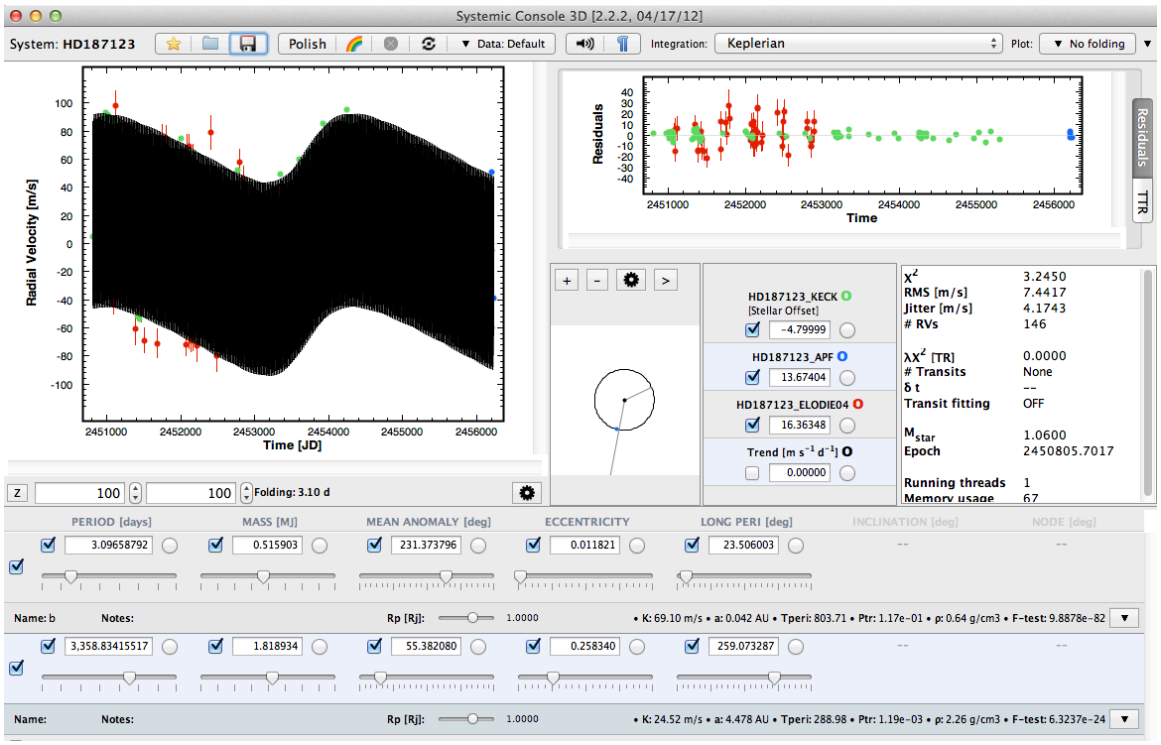

***Figure 4.12**: HD 187123 Keplerian model, all velocities.*

Removing both the ELODIE and Keck/HIRES velocities, and polishing on velocity offset and Mean Anomaly yields the fit below. The derived MA=211.838596 was slightly different than the MA=234.801997 expected from the epoch difference of the first data points, though the RMS when using the latter would have been 17.4 m/s. Clearly adding the APF points has improved the orbit. The RMS of the APF points to this orbit is 2.44 m/s. The mean internal error for these points is 1.12 m/s.

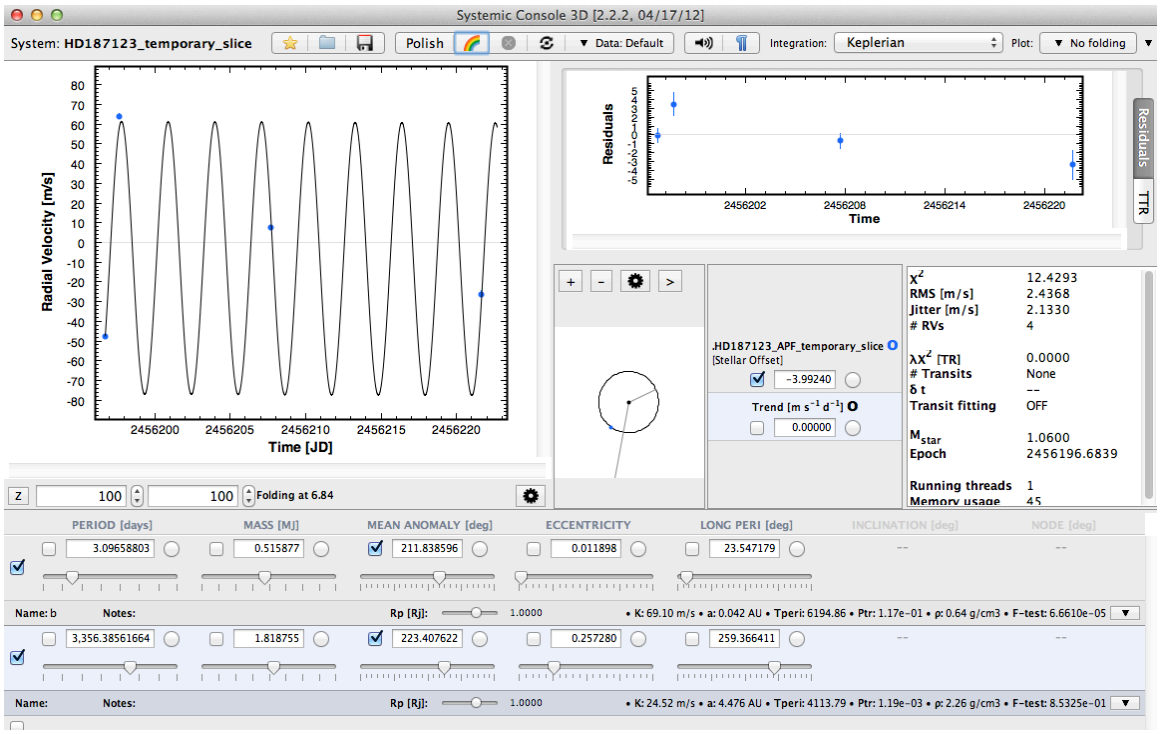

***Figure 4.13**: HD 187123 Keplerian model, APF velocities alone.*

Much of that RMS is due to a clear linear trend in the residuals, suggesting that we may be seeing a systematic trend of -0.208 m/s/day in this star's velocities that is purely instrumental in origin. Alternatively, a slight polish of the orbital period and omega also removes the apparent trend in the residuals as shown below. Here, the RMS=2.05 m/s.

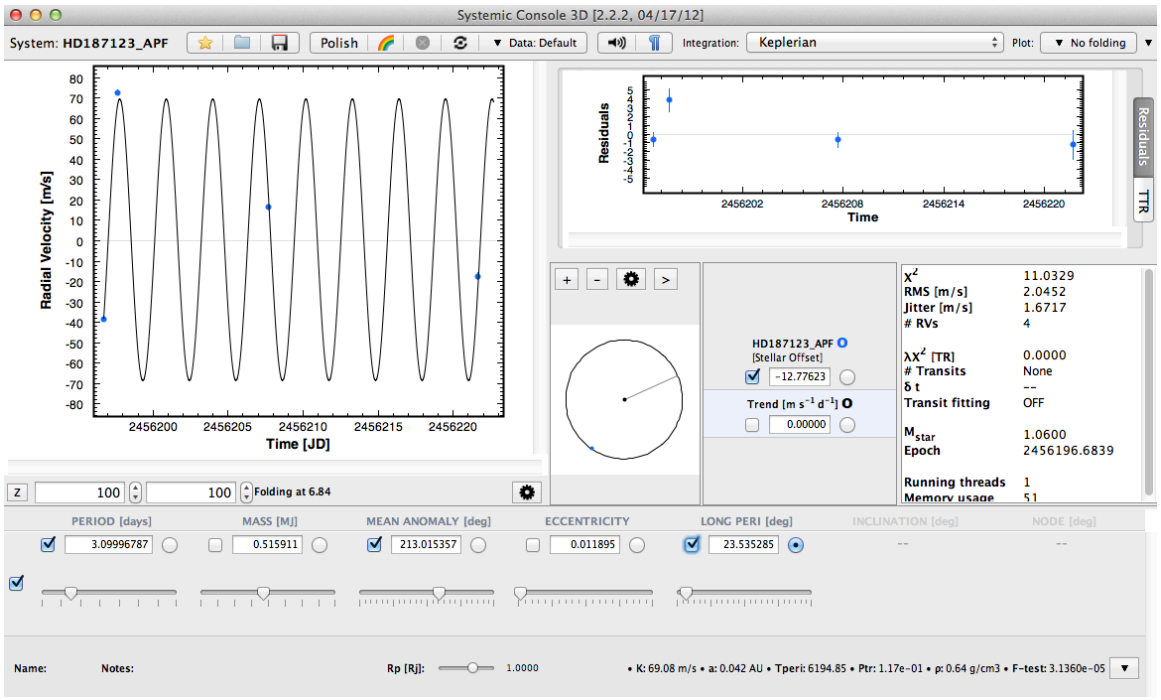

***Figure 4.14**: HD 187123 Keplerian model, APF velocities alone with MA and omega polish.*

A fit using only Keck/HIRES data achieves an RMS=2.56, very similar to the RMS=2.44 for APF data alone. So for this star, RV precision looks roughly similar to HIRES, even with likely instrumental drifts contributing. A typical exposure time on this star with Keck/HIRES is 138s, whereas it was 642s with APF. So, for this star, APF is a factor of about 4.7 times slower than Keck/HIRES.

4.2.5 HD 209458

This V=7.65 G0V star is one of the most well known exoplanet stars, hosting the first known transiting planet. By virtue of its transiting orbit, and extensive study by many groups, the planet's orbit has been extremely well determined. The only complication is with observations that might occur during the actual transit. Such observations would either simply be ignored, or would have to be corrected for the known Rossiter-McLaughin effect. But that happens only about 3-4% of the time, and is easily managed. The next figure shows the SYSTEMIC fit to all the data. There are a total of 269 RV's of which the most recent seven are from APF. Again, the RMS=14.3 m/s of the global fit is completely dominated by the CORALIE and ELODIE data.

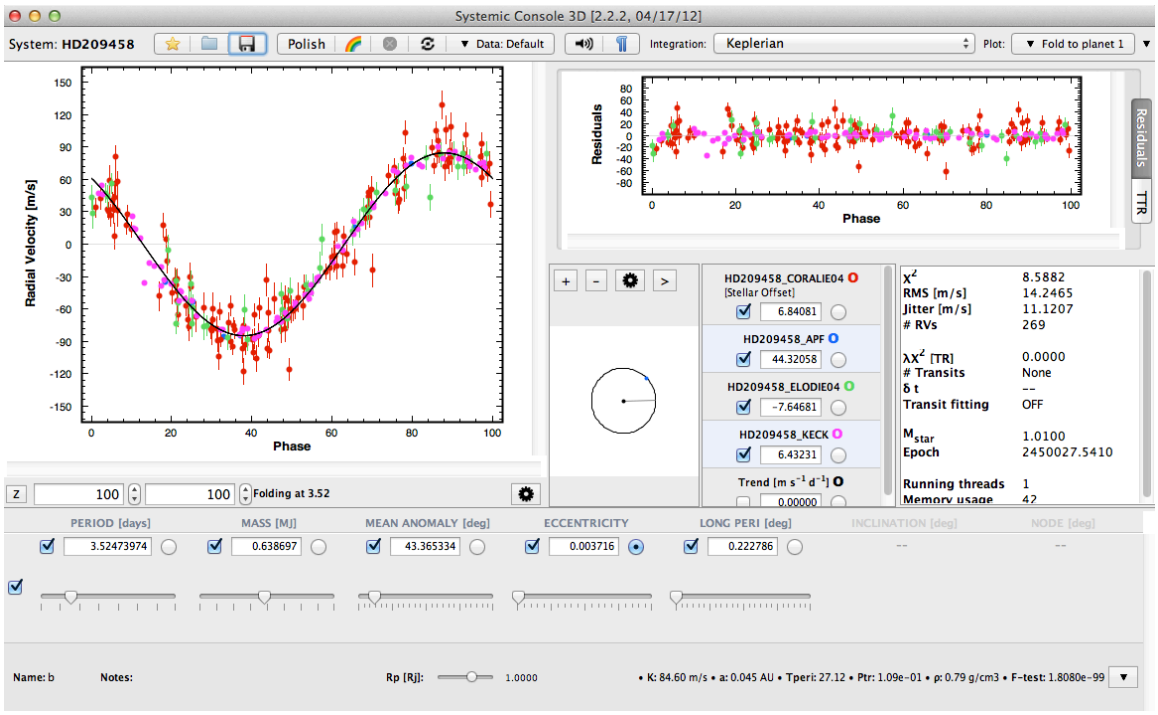

***Figure 4.15**: HD 209458 Keplerian model, all velocities.*

A fit of the APF data alone to this orbit, with velocity offset adjust and MA polish is shown in the next figure and achieves an RMS=4.06 m/s. The mean internal error for these points is 1.59 m/s. The HIRES-only fit achieves an RMS=6.1 m/s. So APF is delivering significantly higher precision than HIRES on this star. Again, as with HD 187123, the RMS is largely due to a clear downward linear trend of -0.403 m/s/day in the residuals, which, like HD 187123, is likely instrumental in origin.

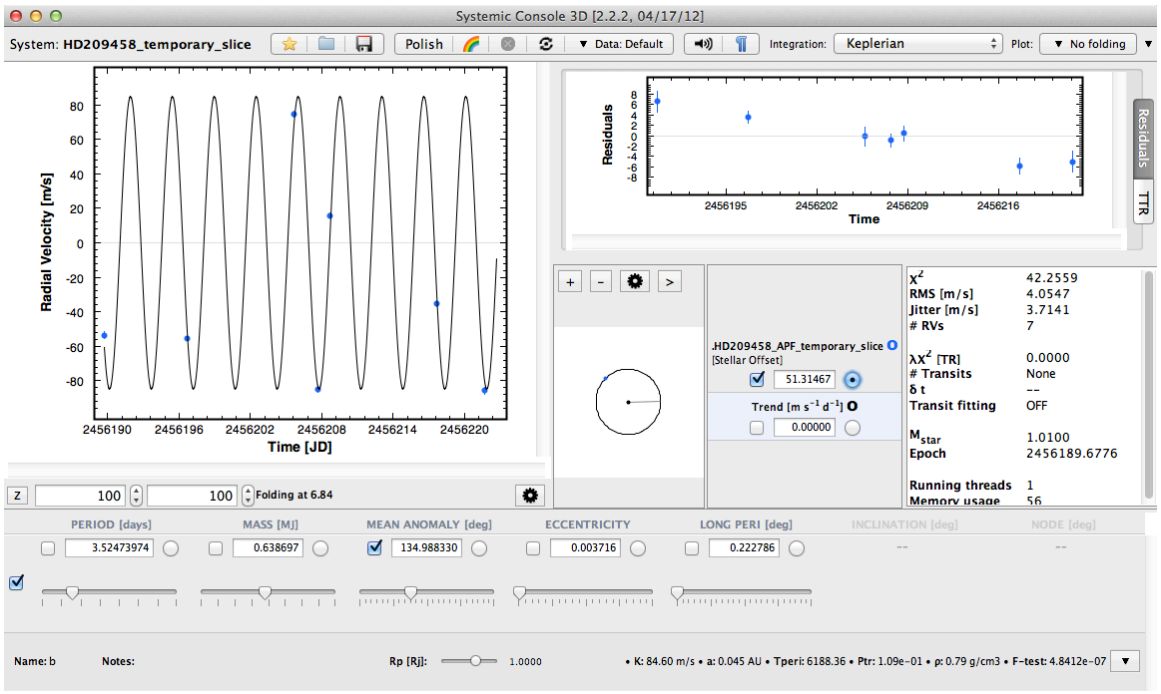

***Figure 4.16**: HD 209458 Keplerian model, APF velocities alone.*

A typical exposure time on this star with Keck/HIRES is 150s, whereas it was 580s with APF. So, for this star, APF is about a factor of 3.9 times slower than Keck/HIRES.

4.2.6 HD 10700

This V=3.5 G5V star is Tau Ceti, the 18th nearest star to the Sun, and quite similar in spectral type. This nearby bright star has received huge attention by the world's precision RV teams, both for planets and for astroseismology, and is probably one of the best known "null" stars in the sky by now. Unfortunately, it is only just rising near midnight, so has been difficult to get as intense coverage on as the other stars in our 1st-half-nights observing schedule. An in-hand total of 557 HIRES, Keck, MIKE, PFS, and APF RV's on this star reveal no hint of planets. The HARPS team has an additional 4864 velocities that also show no evidence of detectable Keplerian signals. The SYSTEMIC console constant-velocity fit to these 557 velocities achieves an RMS=4.17 m/s as shown in the next figure.

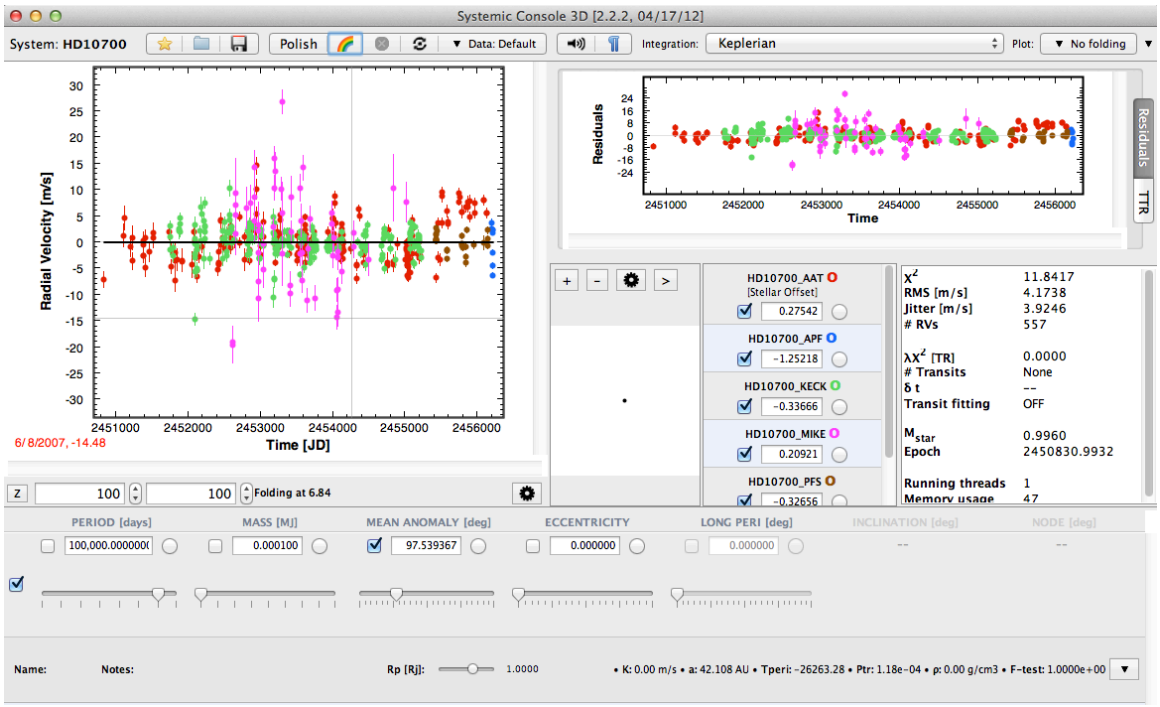

***Figure 4.17**: HD 10700 Keplerian model, all velocities.*

A similar fit of the seven new APF velocities alone to a constant velocity model is shown below and achieves an RMS= 3.51 m/s. The mean internal error for these points is an astonishing 49 cm/s. Most of the RMS is due, again, to a pronounced downward linear trend in the residuals. Furthermore, the last point of these 7, like the last point in all the other observations (excepting for HD 185144) was seriously affected by telescope oscillations, both audible and visible with the accelerometers. But instead of moving on to other stars, we went ahead and obtained these observations as part of the overall commissioning testing of the telescope. If we delete this final velocity for HD 10700, and detrend, the RMS reduces to 89 cm/s. So it seems likely that there may indeed be an instrumental trend in our velocities over this test interval, of some -0.4 m/s/year. That is hardly surprising at this early stage, given the fact that the spectrometer has not been well temperature stabilized and has been opened and worked on repeatedly during this 1-month test period.

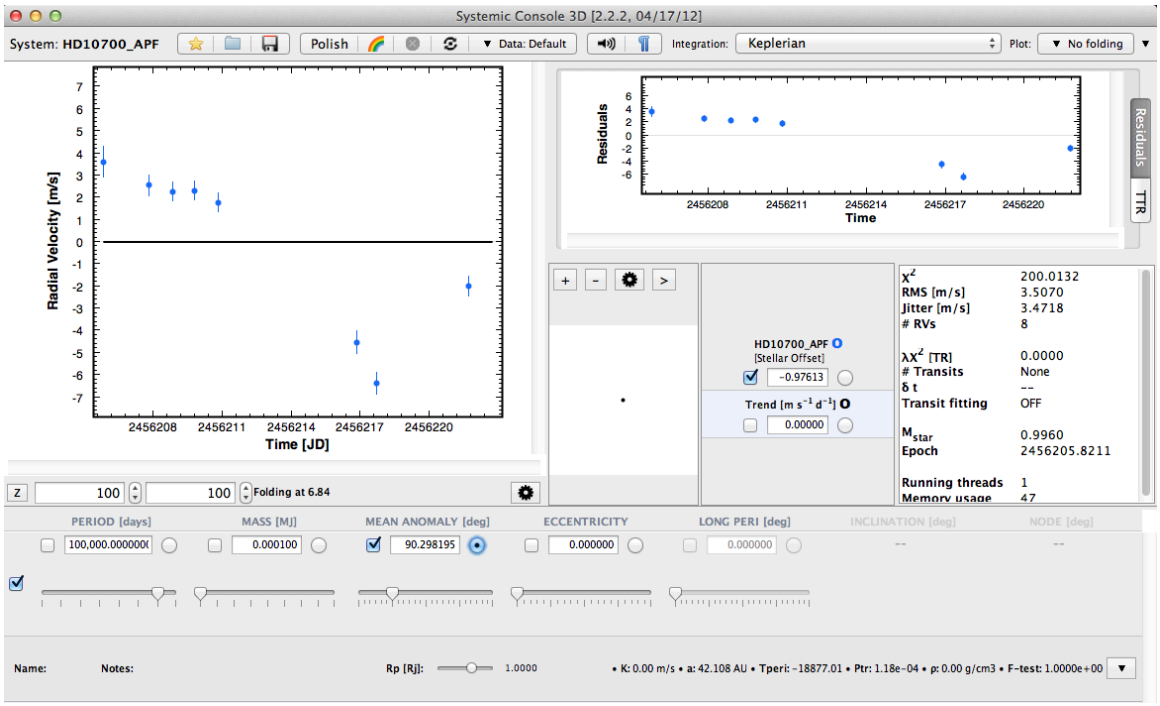

***Figure 4.18**: HD 10700 Keplerian model, APF velocities alone.*

A typical exposure time on this star with Keck/HIRES is 4.1s, whereas it was 26s with APF. So, for this star, APF is about a factor of 6.4 times slower than Keck/HIRES.

Additional data is required to understand whether such an instrumental trend exists, what might be its cause, and how to mitigate. But at this early stage, these RV's look quite encouraging. It seems likely that our precision is at least as good as ~1-2 m/s on moderately bright stars, already better than Keck/HIRES. In cases where we have intentionally avoided taking data during telescope oscillations, such as HD 185144, our bright-star precision may even be sub-m/s. The RMS=96 cm/s for this star is the quadrature sum of stellar jitter, photon statistics, and instrumental errors, implying that the instrumental errors component of that sum must be <1 m/s, though clearly much more data will be required to understand APF precision limits. In parallel with these early test observations, we have identified several sources of mechanical instability in the spectrometer that make it sensitive to temperature changes in the dome. We are also not yet running with good temperature control inside the spectrometer. Finally, we need to eliminate, or at least greatly reduce, the telescope oscillations before we can make any further serious progress on assessing APF velocity precision.

A comparison of exposure times with Keck/HIRES vs. APF on the above RV-standard stars reveals that APF is delivering equal or better precision than HIRES in only about 4-6x the exposure time. The ratio of Keck/APF primary mirror collecting areas results in an expected factor of 17.4, whereas we are seeing only 4-6x, leading to the conclusion that the optical train

of APF (telescope mirrors + spectrometer + slit transmission when used with its nominal 1” slit science slit) is about a factor of 2.9-4.4x more efficient than the optical train of Keck/HIRES (when used with its nominal 0.86” science slit).

The Keck web site lists the measured HIRES efficiency at 18% in the iodine region, not including the 3 telescope mirrors. If it is assumed that each aluminum Keck telescope mirror is 0.85 - 0.91 efficient, that implies an overall end-to-end efficiency for the Keck/HIRES optical train of ~11-13% without including any slit losses. Our nominal HIRES B5 science slit (0.861 x 3.5 arcsecs) transmits about 58% of the light in 1" FWHM seeing. So the operational efficiency of HIRES when used in exoplanet hunting with the B5 slit in 1" seeing is about 6-8%. Since APF is some 2.9-4.4x more efficient (in similar seeing), that implies that the effective end-to-end efficiency of the APF optical train (telescope + spectrometer + CCD) is about 17-35%. These are still very rough numbers, but do suggest that no significant amount of light is missing from the APF optical train. At the same time, under these conditions, the spectral resolution delivered by APF is about 80% higher than that of HIRES.

5 Safety Systems

The APF contains extensive sensors for monitoring its health, and two high-level software daemons to implement the operational rules for ensuring both human and facility safety, regardless of whether observing is carried out in a local, remote, or autonomous mode.

5.1 Safety Sensors

The safety sensors include:

- **Software and Hardware Limits:** the telescope has the usual software and hardware limits on its motion. However, it does not have a brake, and if power is cut (say, due to a hardware or software E-stop condition) the telescope will at best coast to a stop, or at worst can be driven by wind gusts into their hard limits. Both telescope and dome have been demonstrated at Site Acceptance Testing to survive such abuse, but it is a concern.
- **External Monitoring:** Outside the facility, the APF has a Vaisala weather station that monitors wind speed and direction, humidity and dew point, presence of rain or dew. The APF software also monitors the other weather stations on Mt. Hamilton, to increase robustness in case of failure of any single sensor or station.
- **Facility Access Monitoring:** sensors keep track of whether the front door and vents are open, and the position of the dome shutter. Web-enabled security cameras and microphones are installed on all levels.
- **Internal Facility Monitoring:** there are ambient temperature sensors on all levels, at several places on the telescope, and humidity and dew point sensors on all levels. The Multistack facility chiller's fault indicator and other indicators are monitored (but at this writing, we are dissatisfied with the value of the chiller's signals).
- **Dome drive monitoring:** there is a motor over-temperature cutoff switch on the dome azimuth drive.
- **Instrument Monitoring:** There are ambient air temperature and humidity sensors inside the spectrograph, temperature sensors on six struts of the spectrograph frame, spectrograph motors, CCD controller electronics enclosure, iodine cell, the spectrograph electronics rack, and in the science dewar. "Loopback" signals alert if any spectrograph cable is broken or disconnected. There are coolant flow meters to monitor the flow rate from the instrument chiller and from a bypass valve.

- **E-stop signals:** in addition to facility E-stop signals that cut power to the telescope and dome drives, there is a button to toggle an indicator that staff may be working inside the instrument, thereby forbidding commands from being sent to either lamps or stages.
- **Low-level over-temperature protection:** the iodine cell has been outfitted with a thermostatic switch to cut heater power, and similar automated power shutoff is provided for the calibration lamps, CCD controller, and the Omega controller that drives the iodine cell.
- **Acceleration monitoring:** As described in section 2.1.3.2, accelerometers have been installed on both telescope axes. At this writing, the displays of constantly-updating power spectra must be monitored by humans (primarily due to cost concerns, it is unclear whether we will monitor them via software).

Not all electronics are temperature-monitored; the following components are not monitored:

- EOST telescope control rack (telescope control computer and custom interface cards).
- EOS dome and guider computers.
- APF instrument computers.
- Network gear (firewall/router, switches).

5.2 Safety Daemons

It has long been recognized that some observers fail to follow the Observatory-mandated procedures for ensuring facility safety (e.g. obeying weather limits), and therefore Lick Observatory deployed an automated system in 2007 to ensure that the Nickel telescope may only be opened after Lick staff release the facilities for the upcoming night's observing, and automatically close them in poor weather. The same software has been in use at the APF since 2009 under the name "*checkapf*", and is the first of the APF's two high-level software safety daemons. These safety systems have been fully reliable at commanding the systems to close if conditions are not demonstrably good. (Note that this is a statement about the software, but to achieve safety, the facility also depends on the hardware working properly at these critical times.)

5.2.1 The checkapf Dome Closure Daemons

Checkapf keeps track of the current weather and whether the instrument and telescope have been "released" by Observatory staff for operations. Each day at noon *checkapf* automatically sets the state to not-released. The instrument and telescope may be released at different times -- the instrument might be released earlier, so that the observer can start afternoon calibrations, whereas the telescope may not be released until weather conditions are satisfactory, or some maintenance is completed, etc. *Checkapf* will automatically close the dome

if the weather is poor; the telescope has not been released; a remote observer has not supplied contact information; or a telescope technician has otherwise decided to forbid opening. During remote observing, the dome will be closed if the observer fails to press a software "deadman switch" at least once in each twenty minute window. *Checkapf* requires regular updates of critical information such as weather data, even if the data is not changing, as assurance that the other systems are live and functioning. If *checkapf* does not receive updates regularly, it closes the dome.

The Observatory staff use the *checkapf_tt* GUI (telescope technician's GUI) to interact with *checkapf*. *Checkapf_tt* displays the Open/Close conditions; the remote observer's contact information; and whether the mirror cover, dome shutter, and vents are open or closed. The technician can select whether the applied rules are for a local or remote observer; tell *checkapf* to disregard specific weather sensors that are malfunctioning; set special "force-close" conditions (e.g. "ash in sky due to wildfire"); and grant or remove permission for a remote observer to open the dome.

The remote observer uses the *checkapf_obs* GUI to interact with *checkapf*. It displays whether the mirror cover, dome shutter, and vents are open or closed; and summarizes the weather conditions and whether it is ok to open. The observer's GUI lets the observer enter his or her contact information; and when the dome is open, the GUI provides the deadman countdown timer, which starts counting down from 20 minutes each time it is pressed.

5.2.2 The apfmon Monitor Daemon

In addition to *checkapf*'s monitoring of external conditions, the APF has a second high-level safety daemon, "*apfmon*" (APF Monitor), which monitors internal conditions. *Apfmon* monitors a wide range of inside-dome and inside-instrument conditions:

- Monitors if the computer-controlled power outlets are turned on for essential components, including the Omega controller, and the guide camera and cameralink power.
- Monitors the status output bits for spectrograph power supplies.
- Monitors each spectrograph stage for overall readiness, reporting if it is homed, locked out, in a limit, in an appropriate control mode, and so on.
- Monitors the iodine cell temperature, and warns if it is not being held at the setpoint.
- Monitors the instrument chiller flow rate, and warns if it is low.
- Monitors temperature, humidity, and dew point outside the dome, inside the dome, and inside the spectrograph enclosure, and reports if it is unsafe to open the dome or enclosure due to the possibility of condensation.

- Warns if the dome cooling settings for the three levels are consistent with the current state of the shutter and vents (e.g. upper-level cooling should be turned off when the dome is open).
- Monitors ambient temperature, compared to cooling setpoints, and alarms if there is a significant difference (at this writing, this feature is commented-out from the active configuration).
- Verifies that telescope positions are broadcasting as expected.
- Alerts if the PMAC's clock is significantly in error, because that would cause pointing errors.
- The dome shutter limit switches are known to have a failure mode which can prevent the shutters from closing without human intervention, so *apfmon* warns if the limit switch levels are inconsistent with the shutter position.
- Monitors the CCD temperature inside the dewar.
- Monitors the relative temperature at various points in the spectrograph, so as to check on how close the instrument is to isothermal.
- Monitors the spectrograph's circuit board temperatures and instrument rack temperatures.
- Monitors whether autonomous tasks are running (vs failed).
- Monitors whether essential services are alive, by monitoring their heartbeat broadcasts.
- The following three items will be added in the near future:
- Monitor if UPS tripped.
- Monitor whether the front door is opened without a concomitant press of the OK-to-continue-operations button, and if the telescope is in a remote or robotic mode, stop operations.

Note that all the above are only monitors (at this writing). They put status into KTL keywords (keywords are defined in section 2.1.3.2.2) for GUI displays and generate email alerts, but they do not otherwise respond to problems when they are detected.

Apfmon does provide some active controls:

- The heater mat power and coolant flow rates (for the calibration lamps, CCD controller electronics, and instrument heat exchanger) are controlled to hold temperature constant to approximately 0.1C.

- (not yet implemented): For the calibration lamps that have spares, keep track of which lamps have failed (as indicated by the lamp's power being on, but no current is flowing), and therefore which lamps should be used as the 'default' lamps.

5.3 Known Flaws in Safety Systems

- There are no brakes on the telescope drives. Should power be cut due to an E-stop condition, the telescope and dome should coast to a stop -- unless winds are strong, in which case the wind can push the telescope into a hard stop. If the telescope is slewing when power is cut, it will coast in azimuth, and since it coasts at a different speed than the co-rotating dome, the telescope will tend to lose alignment from the dome shutter, and can strike the enclosure structure.
- The dome shutter limit switches have failed twice, but they do not fail-safe: after failure, the standard close-shutter command will result in it driving to the full-open position because of the erroneous limit signals. (*Checkapf* avoids the failure mode by commanding the shutter to the closed position, without issuing the actual "close-shutter" command. (The APF monitor watches for inconsistencies between the limit switch value and shutter position, but does not currently act to stop the shutter from running into a limit. Perhaps it should.)
- At the present time, we do not have a good monitor for Multistack failure -- in particular, we can't reliably tell if the coolant is cold. This will be addressed in the near future by direct temperature monitoring of its output.
- It is essential that the cable bundle supplying power and signals to the spectrograph remain “soft”, rather than become a stiff coupling. For example, if it is so short that it rigidly couples the telescope and dome, then the dome cannot correctly track the telescope. When the dome tries to adjust its position relative to the telescope, the stiff coupling will push or pull the telescope along, causing the dome to drive harder and harder, and eventually into a limit, as it tries to correct.

At this writing, many alarm conditions simply generate warning messages on GUI's and/or email to key Observatory staff. We have not yet made decisions regarding how all conditions should be handled when in robotic mode. It is likely that we will implement a combination of automated recovery attempts, and robust alarms to Observatory staff that can get attention more quickly than an email.

6 Operations

6.1 Afternoon Operational Issues

The APF has a few operational tasks that do not have a counterpart at the other Lick Observatory facilities:

- The APF's co-rotating dome and floors have many potential interference points with the non-rotating part. Therefore, operation of the telescope is generally preceded with a walkthrough inspection to check that there are no obstructions. We have not yet decided if it's acceptable to begin operations without a walkthrough, should the front door have remained closed since the telescope was last released.
- Each afternoon, the dome's upper floors must have their temperature control setpoint set to the predicted temperature for the start of observing. How to accurately predict the temperature is an open question.

Other afternoon tasks are similar to those required at other telescopes:

- Check that the telescope is free for motion, and that any tarps or covers have been removed;
- Clear E-stops and other safety switches, and that any auto/manual switches are set for computer control;
- Check that all software services are running;
- Ensure that all end-user software (VNC sessions and the GUI's that run in them) is up and operational;
- Respond to any alarms being reported by monitor software;
- Home telescope if needed;
- Verify that camera and instrument are operational.
- Release the instrument to the observer. If conditions permit, release the telescope to the observer.

6.2 Nighttime Operational Issues

The key observing-time operational issues are generally the same as for other telescopes (e.g. monitoring weather), but the APF has certain distinct issues of its own. These include:

- poor response to E-stop conditions (discussed above);
- potentially serious consequences for dome shutter limit switch failure (also discussed above);

- ensuring the UPS has sufficient capacity to be able to close the dome in the event of power outage. The UPS could fail either because its batteries become too old without replacement, or because a series of power outages during one night run down the batteries.

6.3 Local Operations

A local operation, in which the observer is on site at the APF or at an observing station in a nearby building, is similar to operations at all other Mt Hamilton telescopes.

The APF-specific setup issues are discussed under the Afternoon Operations section, above.

6.4 Remote Operations

Remote operations are quite similar to local operations. After an Observatory staff member has carried out steps requiring local presence, and has released the telescope to the observer, it doesn't matter whether the observer is physically present at Mt Hamilton or in a remote observing room on a distant campus. This sweeping statement does not indicate an overly-casual approach to remote observing. Rather, it reflects two important features of remote observing at Mt Hamilton:

- The safety features implemented by *checkapf* are well-tested and have a demonstrated track record; and
- Observatory staff is on hand to respond to hardware failures that could otherwise endanger the facility.

6.5 Robotic Operations

The APF facility must be able to execute planet-hunting activities fully autonomously. This task is less daunting than it may sound at first blush, because the autonomous tasks do not require reasoning capabilities. Instead, the basic observing task can be encapsulated in simple scripts (“Turn on this lamp; take a 5-second exposure; ...”), while safety issues are handled by the *checkapf* and APF Monitor systems.

The autonomous facility is made up of the following systems:

- The basic observing subsystems: telescope and dome control, instrument control, and camera control.
- The safety subsystems: the *checkapf* and *apfmon* daemons to ensure the dome is closed in poor weather and that electronics are kept at safe temperatures.
- Automated recovery tools: there are several conditions, such as bad weather, that can temporarily interrupt observations. The automated recovery tools do the work necessary to restore the APF to its state prior to interruption.

- The robot observer: a set of scripts that carry out a sequence of observations, following the same rules that a human would follow.

The first two of the above systems are required for any kind of observing, whether local, remote, or robotic.

The APF monitor system greatly simplifies the task of writing robotic observing software, because it allows us to summarize complex status conditions in one or a few simple keywords. The monitor tracks arbitrary conditions, where each "condition" is based on the value of one or more keywords, and can use hierarchical sets of conditions. For example, each stage's calibration status, control mode, lockout status, and so on are treated as distinct monitored conditions. All the conditions for each stage are gathered together into single overall stage "condition group"; all the stages' condition groups form a combined-stages group; and the combined-stages group is a member of an overall "instrument configuration" group. The status of the instrument configuration group is an enumerated keyword that indicates if the entire instrument is ready to operate, supplemented by a string keyword with an explanatory message.

We will create additional condition groups in the APF monitor, to make it reasonably simple to write robotic control scripts. These include:

- an "ok-to-open" group that will determine whether it is OK to open the APF dome, based on *checkapf*'s information about whether it is OK for the dome to be open at all; whether the telescope is released; the current state of the telescope homing; whether the mirror covers are open or closed; the time of day, and the current azimuth.
- an "ok-to-move" group that will indicate whether it is OK to move the telescope, including homing the telescope, based on whether the telescope has been released, and whether the front door has been opened without someone pressing the ok-to-continue-ops button.
- a group to both monitor the UCAM state, and try to recover from the various UCAM failure conditions for which power-cycling the controller and/or restarting the UCAM software are known to often be successful.

Robotic observations can be interrupted by several recoverable situations. The automated recovery daemons (not yet written) will handle the following situations when executing robotic observations:

- *Checkapf* does not allow the dome to be open, typically for weather reasons. The recovery daemon will wait until the dome is allowed to be open, and then re-open it.
- The UCAM will sometimes hang in ways that can be recovered by power-cycling the controller or restarting the combined UCAM GUI/control daemon, and the UCAM

software uses keywords to indicate when these things should be tried. The recovery daemon will restart UCAM in this situation.

A robotic script simply has to check the status keywords to know if the instrument is in a ready-to-operate state for any kind of exposures, and check the ok-to-open group when the script is ready to begin on-sky observing, and so on. It should also refuse to run unless *checkapf* is configured for a robotic observer, to ensure that the appropriate safety rules are being enforced.

Similar to the deadman software switch for a human remote observer, the robotic script will be required to regularly set a 'deadrobot' keyword to indicate that it is running correctly, else *checkapf* will shut the dome. This and a related set of "task" keywords will be used by the APF monitor to track the status of autonomous operations.

The robot observer consists of the following components:

1. **Calibration script:** the *calibrate* script takes a series of images using Th-Ar or Quartz lamps through several slits and pinholes, and a series of darks. A command-line argument controls whether it reverses the order of images, so that it is suitable for both start and end of night calibrations.
2. **Dewar focus script:** The "*focuscube*" script is used to take a series of spectra while stepping through a range of focus values. At this writing, the script does not do an automatic best-focus analysis. The focus is sufficiently stable that it is probably not necessary to change the focus often.
3. **M2 Focus script:** The "*m2focus*" script will be used to take a series of guide camera images while stepping through a range of M2 focus values. It will perform an automatic best-focus analysis.
4. **Single-object script:** this script commands the telescope to slew to a particular object and track it, and then acquire a spectrum through an observer-specified slit. The APF has an exposure meter and the ability to stop the exposure when the integrated source counts reach a desired level. At this writing, objects are acquired by relying on the telescope's blind pointing accuracy to ensure the star is within a user-set acquisition radius in the guider's field of view, and then using the guide images to let us center the stellar image on the desired slit. By default, the current algorithm assumes the brightest object in the field is the desired star. On the occasions when there are multiple stars in the field, an alternate acquisition algorithm can be specified that will offset to a nearby pointing star, then narrow the acquisition radius to exclude all but the target star, and offsets to the target.

5. **Dynamic scheduler:** A critical component of a robotic observing system is an efficient scheduler for selecting the best object to observe at each point in a night. The scheduler needs to account for variations in observing throughout the night. For example, if poor weather forces the dome to close for an hour, the scheduler might pick up where it left off in a sequence of targets, or it might skip some and move on to others.

   The APF commissioning team has a prototype scheduler that can be used to dynamically select the next object to observe, taking into account not only prior nights' observing, but also any successful observations during the current night. As a result, regardless of whether an observation of an object went well or poorly, the same action can be taken: the scheduler is invoked to select the next object to observe. (This scheduler is currently a work in progress, and is not yet ready for integration into the robotic system.)

6. **Robot master:** this script runs the *calibrate* script and does other afternoon setup tasks. Starting between sunset and 6-degree twilight, it opens the dome, runs the *m2focus* script, checks the seeing (derived from autoguider images), and so on. From 10- or 12-degree twilight until 12-degree dawn, it repeatedly invokes the scheduler to select the best object to observe, and then invokes the single-object script to carry out the observation.

6.6 Concerns/Unsolved Problems

1. Each afternoon, the setpoints for cooling the upper levels of the APF dome need to be set to the predicted outside temperature at dome-opening time.
2. Network accessibility: we have not determined the acceptable degree of autonomy for the APF. In particular, is it acceptable to continue operations if the Observatory staff cannot control *checkapf* through the *checkapf_tt* GUI from their normal stations at the 120" dome?
3. Further testing is required to determine how well the autonomous system can recover from UCAM failures.
4. The Multistack chiller is not on UPS power. We need to improve its ability to restart after a power outage and after the common event of an extended period of unclean power during a transition to generator power.
5. We need to debate the importance of the deadman switch: why does the human matter; or, is a deadrobot switch sufficient to the need?
6. As mentioned earlier, we need to be able to measure the Multistack output temperature.
7. What to do when a dome shutter limit switch fails?

8. How shall we verify that the UPS has sufficient capacity to handle closure?
9. Will Observatory staff always be on hand to address serious hardware failures such as dome shutter failure? Or will the APF be operated robotically, even if staff is not on duty? If staff is not on hand, how does that affect our response to issues such as chiller failure?

6.7 Next Steps

In addition to tasks implied in the Concerns/Unsolved Problems section, the following software tasks remain to make the APF fully ready for autonomous operations:

1. Add the handful of new items to the APF monitor, as described in the Safety Systems section.
2. Add the ok-to-open, ok-to-move, and ucam-ok groups to the APF monitor.
3. Create the automated recovery tools for dome re-opening and UCAM failure.
4. Make the calibration and focuscube scripts robust enough for use in the robot observer, and modify them to set their state in keywords, and use the APF monitor status keywords.
5. Create the m2focus script.
6. Create the single-object script.
7. Make the dynamic scheduler into a callable script for the robot observer.
8. Create the robot master.
9. Create a “global” startup script.

Beyond the above list, there remain various tasks that may be done after general autonomous operations have begun, including such things as:

1. Standard web pages for observers who will use the APF.
2. Documentation of UCO-created software.
3. Guide camera image acquisition system does not build at present.
4. Strong observer data privacy (optional).

6.8 General Operations

The dome and telescope have been design for maximum reliability with a manageable amount of required regular preventive maintenance. Detailed maintenance manuals exist for the dome (IceStorm2 Mechanical Maintenance Manual) and telescope (Lick Automated Planet Finder 2.4 meter Maintenance and Installation Manual) in accordance to the contract requirements. The telescope manual includes a regular maintenance schedule on page 18. UCO has developed a separate maintenance schedule for the dome and spectrograph systems which can be found here.

The Levy spectrometer requires almost zero preventative maintenance. The only notable exception is bi-annual service (change filter, check for algae) of the instrument chiller. The stage and motor assemblies are a proven design that has been in service at the Keck Observatory for 12 plus years without a single failure. The CCD dewar is cooled by a Brooks Automation cryotiger that is also maintenance free. The dewar vacuum has been steady for the 8 or more years the CCD has been cooled. The only consumable items in the spectrometer are the calibration lamps (Quartz and Thorium/Argon). Both of these lamps have a current monitor system to detect a lamp failure and an onboard spare which is automatically inserted if and when a lamp burns out.

The AC power at Mt. Hamilton is notoriously unreliable, which has created an ongoing issue with the building chiller. Faults on the 3-phase power to the chiller often cause the system to shut down and require a manual reset. We currently are working to increase our immunity to this problem and we now have an automated alarm system that sends emails to the mountain techs when the system needs resetting. All critical items in the dome are powered by a 16KVA APC Symmetra UPS which has sufficient reserve power to continue an observation for at least five minutes before activating an automatic shutdown procedure. This is more than sufficient time to allow the backup diesel generator to start and stabilize mountain power. The UPS recently had all it batteries replaced, which will be done on a seven/7 cycle; the UPS also has two new power modules. Power module failure has been a problem with this unit, which is surprising considering our experience with the same model UPS in our campus laboratory. According to the manufacturer, the power at Mt. Hamilton is within the normal operating range expected by the UPS and should not be contributing to the premature failures of the power modules. This remains an item we plan to carefully monitor and purchase spares to support.

Mirror cleaning and coating are always major tasks at any observatory. We plan to wash the primary mirror on a quarterly basis and probably recoat the aluminum every 2 to 3 years as required in the Shane vacuum tank. This was done successfully in April of 2009 using the lifting fixture and mirror handling cart designed by EOST. The mirror removal process is documented in the "Primary Mirror Assembly Installation /Removal Procedures". What is not obvious in this document is that mirror removal from the telescope requires a section of the dome on level 3 to

be removed. This is necessary because the dome diameter at this plane is less than 3x the mirror diameter. Once the dome panel is removed the mirror and cart can be rolled clear of the telescope to a location where a mobile crane can lift the mirror out of its cell and into a shipping box for the short drive to the Shane 3m telescope dome.

M2 an M3 will be recoated as required at the UCO coating labs on campus. We currently can reliably produce protective silver coatings that show little to no degradation in accelerated aging tests that are comparable to 4 to 5 years of normal service.

Appendix 1 APF Operation Costs

APF Operation costs				material	hrs/yr		
	1. MATERIALS AND SUPPLIES			$10,000			
	2.TELESCOPE CLEANING				64		
	3. TELESCOPE MAINTENANCE				200		
	4. COMPUTERS SERVICES			$4,918			
	5. MIRROR ALUMINIZING		every 3 yr	$1,000	300		
	6. FACILITY MAINTENANCE				100		
	7. UTILITIES						
	8. OTHER	(nightly visit)			50		
	ME from SC				52		
	Elect tech SC				104		
	Software SC				385		
					1255		
				$15,918	0.883803	FTE	$104,298.28

Appendix 2 Post-Review Budget

			Post Acceptance Review				
Mechanical/Elec		PRIOITY					
		1	re-align telescope AZ encoder		40		$1,000.00
		1	de glitch and fix periodic error in AZ encoder		80		$1,000.00
		1	re-tune elevation and Az telescope axis		100		$200.00
		2	lantronix server for UCAM UDS1100 ($130)		2		$270.00
		2	Insulate struts on spectrograph		16		$200.00
		3	Dome Chiller Leak Repair		16		$300.00
		3	Sight gauge for NesLabs chiller		8		$500.00
		3	service building chiller		20		$4,000.00
		3	fix water hammer problem with NesLabs Chiller		32		$500.00
		4	New servo amps		8		$600.00
		5	Fix light leaks.		16		$140.00
		5	Shield back of shutter and fix rain leaks		40		$2,000.00
		6	Improve slack in Az cable wrap		16		$140.00
		6	install spares cabinet		4		$500.00
		6	Brakes on Elevation (and Az?) Axis		20		$0.00
		7	Replace Baldor Dome motor controllers		40		$4,000.00
		8	Electronic documentation				$0.00
		8	Install exterior indicator sign near door		2		$50.00
		8	75 volt Keepco DC power supply failure light is on				$0.00
		8	One of the AC power outlet controllers is unmarked (no label to indicated what is controlled by each channel)		1		$0.00
		8	State of all indicator lights should be recorded for normal operation		1		$0.00
		8	CCD controller needs a note instructing staff to wait XX? seconds on power down for cap. to discharge before restating. Failure to do so will prevent HV for coming on.		0.5		$0.00
		8	typical operating current of ion pump should be recorded		0.5		$0.00
		8	Cryotiger operating pressure should be recorded.		0.5		$0.00
		8	Right El fork wiring needs to be cleaned up MVR 10/9/12		8		$70.00
		9	Install overhead crane		40		$3,000.00
				sub total		511.5	

software							
	Infrastructure SW						
			check apf bug fix		4		
		2	APF web site		40		
		2	SW documentation		80		
	Facility Safety						
			Multi Stack chiller monitor		4		
			Ground FL temp senor alarm		8		
			Monitor UPS		4		
			Determine if UPS can close the Dome		16		
			change apfmon to take-action items		8		
	User Safety						
			Monitor Door open flag to determine if OK to continue operations		4		
			Keep track of Cal lamp spares		16		
		2	Improve thermal control		24		
			guide Camera image acquisition		40		
			convert extant tools from engineering to operational		16		
			Modify UCAM's camread		8		
			Global startup Script		8		
			convert helper scripts to daemons		8		
			pointing SW suit into CVS		4		
			Automate downloading of data from IERS		16		
			Auto notify host reboot		8		
			Implement Glue SW		40		
		2	implement observer privacy		80		
	Autonomous Operations						
		2	Monitor and resp to Accelerometers		40		
		2	Implement apfmon condition groups OK to …		32		
			Implement UCAM auto recovery		8		
			robustify calibration scripts		8		
			implement M2 focus script		12		
			implement single object script		12		
			convert dynamic scheduler		16		
			implement robot master		40		
				sub total		604	
				Totals		1115.5	$18,470.00