

The Carnegie Planet Finder Spectrograph: integration and commissioning

Jeffrey D. Crane^a, Stephen A. Sheckman^a, R. Paul Butler^b,
Ian B. Thompson^a, Christoph Birk^a, Patricio Jones^c and Gregory S. Burley^a

^aObservatories of the Carnegie Institution for Science,
813 Santa Barbara Street, Pasadena, CA, USA 91101

^bDepartment of Terrestrial Magnetism, Carnegie Institution for Science,
5241 Broad Branch Road, NW, Washington, DC, USA 20015

^cLas Campanas Observatory, Colina el Pino, Casilla 601, La Serena, Chile

ABSTRACT

The Carnegie Planet Finder Spectrograph (PFS) has been commissioned for use with the 6.5 meter Magellan Clay telescope at Las Campanas Observatory in Chile. PFS is optimized for high precision measurements of stellar radial velocities to support an ongoing search for extrasolar planets. PFS uses an R4 echelle grating and a prism cross-disperser in a Littrow arrangement to provide complete wavelength coverage between 388 and 668 nm distributed across 64 orders. The spectral resolution is 38,000 with a 1 arcsecond wide slit. An iodine absorption cell is used to superimpose well-defined absorption features on the stellar spectra, providing a fiducial wavelength reference. Several uncommon features have been implemented in the pursuit of increased velocity stability. These include enclosing the echelle grating in a vacuum tank, actively controlling the temperature of the instrument, providing a time delayed integration mode to improve flatfielding, and actively controlling the telescope guiding and focus using an image of the target star on the slit. Data collected in the first five months of scientific operation indicate that velocity precision better than 1 m s⁻¹ RMS is being achieved.

Keywords: spectrograph, spectroscopy, echelle, high resolution, radial velocity, iodine cell, extrasolar planet

1. INTRODUCTION

The Carnegie Planet Finder Spectrograph (PFS) has recently been commissioned at the 6.5 meter Magellan Clay telescope at Las Campanas Observatory (LCO) in Chile. PFS (Figure 1) is a high resolution spectrograph that uses dioptric optics, an R4 echelle grating in a near-Littrow configuration, and a prism cross-disperser to produce spectra completely covering wavelengths from 388 to 668 nm with a resolving power of 38,000 when using a one arcsecond wide slit. PFS is being used to conduct a long-term observing program aimed at discovering and characterizing extrasolar planets around nearby stars.

PFS sits in a gravity invariant location on the telescope's nasmyth platform. However, this platform is shared by multiple instruments, and space limitations dictate that each instrument be removed from the platform when not in use. It was therefore necessary to design PFS to be mobile, an extra challenge for an instrument for which stability is an important goal.

The main optomechanical portion of PFS is mounted on an optical table inside an insulating enclosure in which is embedded a thermally controlled liquid jacket designed to maintain an instrumental temperature near 25°C. A pre-slit assembly that contains calibration lamps, a molecular iodine absorption cell, a slit-viewing guide camera, and control electronics for the main science CCD is positioned outside the thermal enclosure, but

Further author information:

J.D.C.: E-mail: crane@obs.carnegiescience.edu, Telephone: 1 626 304 0217

S.A.S.: E-mail: shec@obs.carnegiescience.edu, Telephone: 1 626 304 0219

R.P.B.: E-mail: paul@dtm.ciw.edu, Telephone: 1 202 478 8866

I.B.T.: E-mail: ian@obs.carnegiescience.edu, Telephone: 1 626 304 0225

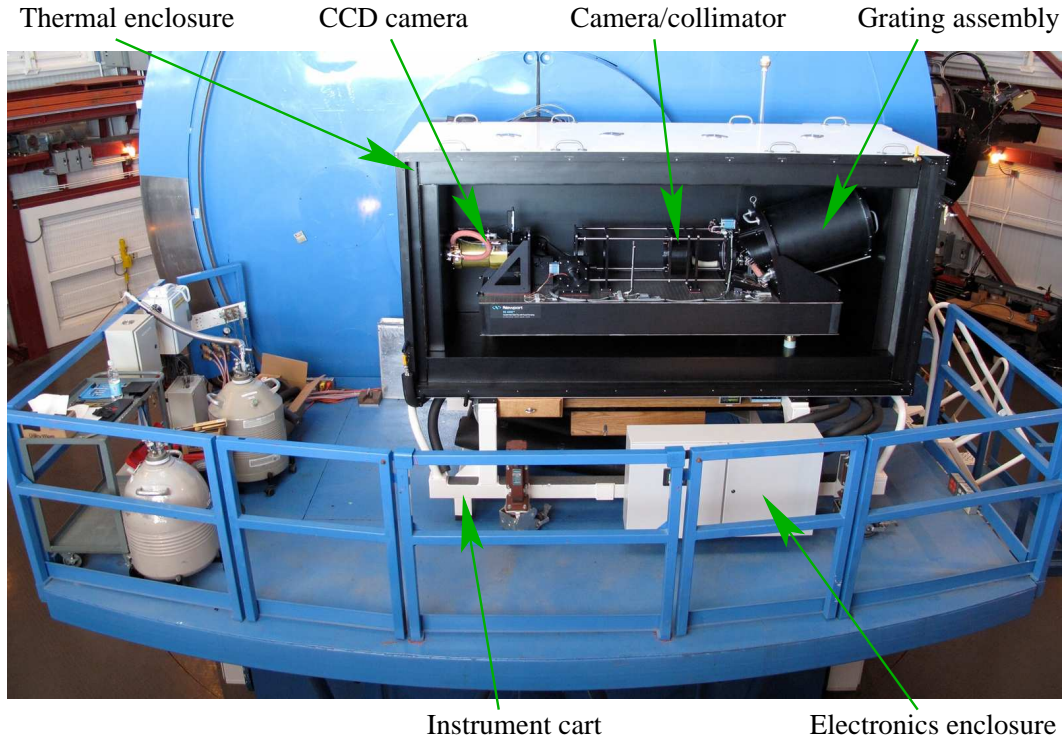


Figure 1. PFS is shown in position on the telescope's nasmyth platform. When this photo was taken, one side of the thermal enclosure had been temporarily removed to allow engineering access to the interior.

mechanically attached to the optical table using thermally insulating tubes that contain focal reducing optics and provide convenient paths for electrical cables.

This is the third in a series of papers describing PFS. Each paper is complementary, requiring that all be referenced to gain a complete perspective on the functionality of the instrument. Crane et al.⁴ (2006; hereafter C06) describes the scientific motivation and goals for the project, an overview of the the main detector and the concept of time delayed integration, and mechanical details of the bulk instrument structure, slit assembly, camera/collimator mount, and disperser assembly. The optical design is also primarily described in C06, but several modifications and a more refined description of the pre-slit optics are given in Crane et al.⁵ (2008; hereafter C08). C08 also describes the mechanical design of the pre-slit assembly, Hartmann mask, photomultiplier assembly, and CCD camera, as well as the assembly and functional design of the thermal control system.

Since the publications of C06 and C08, some minor optical and mechanical modifications have been made, electronic control of various mechanisms has been implemented, and control and interface software has been developed. PFS was disassembled, shipped to LCO, reassembled and tested on the Clay telescope. Scientific use began on 2010 January 1. Since then, three observing runs have been successfully completed, representing a total of 24 nights of telescope time. This paper describes these activities and presents early performance results based on the first few months of operation.

2. OPTICS

The PFS optical design was described in C06 and C08. With the exception of fold mirrors and the reflection grating, all optics are dioptric. All but one of the pre-slit optics consist of small, inexpensive stock singlet and doublet lenses and serve to focus calibration lamp light, image the instrumental field of view on the guide camera, and reduce the telescope's f/11 beam to f/5 at the spectrograph's entrance slit. The primary spectrograph optics are custom made from standard Ohara and Schott glasses and calcium fluoride, and work in double-pass as both collimator and camera.

2.1 Optical Modifications

Several minor modifications were made to the PFS optical system relative to the descriptions given in C06 and C08.

Just in front of and below the CCD field flattener, there is a baffle over the injection fold mirrors that incorporates a flat surface angled to match those of the expected light rays at that position. This baffle also happens to be located precisely where light redward of the PFS spectral coverage passes on its way from the diffraction grating to the CCD. Early test spectra of thorium-argon lamp light showed scattered light patterns with several bright, vertical stripes on the images. These originated from bright, red argon emission lines being scattered at glancing incidence from machine tooling marks on the baffle. The only reasonable way to combat this scattered light was to filter it out. BG-38 filters were added to the quartz and thorium argon lamp assemblies, and a separate, thinner, BG-38 filter was added to the front aperture of the instrument. The addition of these filters reduced the scattered light remarkably and it is now very low — typically a few tens of photons per pixel between orders in the iodine absorption region.

There was an unintended consequence of adding the filter to the instrument's entrance aperture. Extracted stellar spectra taken during an engineering run showed interference patterns caused by multiple reflections between the surfaces of the filter. These fringes had unacceptably large peak to valley amplitudes of 2% of the signal level. To alleviate the problem, a new filter mount was built that incorporates a 10 degree tilt. An anti-reflection coating was also applied to the filter. These changes eliminated the visibility of the fringes.

The lengths of the original PFS slits were chosen to maximize the amount of sky flux adjacent to the stellar spectra while maintaining a distinct separation between the reddest orders of the iodine absorption spectral region. After initial science observing began, however, it was decided that data reduction was suffering because the order separation in the red was too small to allow sufficiently accurate modeling and removal of scattered light. The primary planet-hunting program targets are bright enough that sky subtraction is not important, but an accurate accounting of scattered light does matter. For this reason, the slit plate was replaced. The new plate contains sets of slits of different widths with lengths of both 2.5 and 3.7 arcsec. The former are used for the planet program, and the latter are available for guest observer programs that target fainter objects for which sky subtraction is more important.

C06 described the concept for a novel optical system that facilitates active telescope guiding and focusing using an image of the target star that is aligned on the slit. This optical system produces an image on the slit-viewing guide camera that shows both the instrumental field of view with the target star on the slit and two offset images of that same star produced by each half of the telescope pupil *before* the light reaches the slit. Initial on-sky use of PFS demonstrated that this combination of images on one guide camera was very confusing — particularly when multiple stars are in the field of view (as with a double star or cluster). The problem was compounded by the fact that the beamsplitter that enabled the system was a piece of flat glass with an anti-reflection coating on just one side — a coating that, while of high quality, still produced a faint ghost image. To alleviate the annoyance of this confusing slit viewer image, the plate beamsplitter was replaced with an uncoated pellicle beamsplitter and a shutter was placed in front of the doublet lens and dihedral mirrors that produce the offset half-pupil images (refer to C06). The shutter remains closed during target acquisition but can then be opened to enable on-slit guiding.

2.2 Optical Alignment

The PFS optics were initially aligned in the lab, and later re-aligned after shipping the instrument to LCO.

The entrance window of the science CCD camera acts as a field flattener, but its position relative to its mechanical mount is not adjustable. A conjugate lens to this field flattener is located on the output side of the spectrograph entrance slit. This lens (actually a wedge-shaped cutout from a full lens) is also not easily adjustable relative to its mechanical mount. It was important to ensure that the separation between the slit and field flattener conjugate and the separation between the CCD and field flattener matched as closely as possible. The former was established by design, trusting to manufacturing and assembly tolerances. The latter was adjusted by hand, rotating three threaded positioning posts that hold the CCD package and carefully checking the depth of the CCD within the camera head. The spectrograph optics were then aligned relative to

the common CCD camera and slit mount assembly. This assembly was first positioned against three reference pins permanently attached to the optical table.

The collimator assembly (refer to C06 for a full description) contains three widely separated lens groups, each of which is mounted with no adjustability inside an inner cell that itself sits in an outer cell. The inner cells can be adjusted radially relative to the outer cells, but not angularly. The outer cells can be adjusted axially relative to each other by loosening clamps that attach them to their common space frame. Initially using large calipers, and later using a specially designed jig, the separations of the outer cells relative to each other and to the CCD camera mount were established according to the original optical design, taking into account the known, as-built dimensions of the various mechanical parts.

The CCD camera was removed from its mount and replaced with a rotatable assembly holding an alignment laser pointed through a small aperture in the direction of the diffraction grating. The collimator lenses were removed from their outer cells. Graph paper was placed at the location of the cross-dispersing prism, normal to the spectrograph's optical axis, and used as a reference for the initial laser alignment. The laser assembly was rotated and its spot monitored while making lateral and angular adjustments of the laser mount to establish its centering and normal alignment relative to the CCD camera mount. With this established, the collimator lens farthest away from the CCD camera was installed in its outer cell.

Despite the high quality of the anti-reflection coatings on the lenses, laser reflections from both the front and back surfaces of each lens could be seen in a darkened room. Using set screws in the outer cell, the inner cell of the collimator lens was adjusted until the reflected laser spots were well aligned with the exit aperture of the laser itself. After a satisfactory alignment was achieved, the next inner cell-mounted collimator lens was placed in its outer cell and adjusted in the same way, but with the reflected laser light from the first lens blocked to avoid confusion. The final lens closest to the CCD camera was similarly aligned after locking the second lens into place.

Apart from the alignment provided by standard bolt thru-hole tolerances, little special attention was given to the exact axial and lateral positions of the disperser assembly containing the cross-dispersing prism and the echelle grating as these have no significant effect on the focus of the spectrograph. However, the angular position of the entire assembly was carefully adjusted to remove from the CCD the partially cross-dispersed reflection of the slit off of the interior interface of the cemented prism. In addition, careful angular adjustments of the diffraction grating in both azimuthal and altitudinal directions were made both to approximately center the grating blaze function on the CCD and to select the diffraction orders that would be imaged. Fine tuning of this latter angle eventually resulted in a spectral format that spans 388 to 668 nm, including the calcium H and K lines toward the blue end and hydrogen α toward the red end.

Very little adjustability was built into the pre-slit optical mounts since doing so, in most cases, was judged to be both too onerous and non-critical relative to the expected mechanical tolerances of static mounting. However, the alignment of the pre-slit focal reducing optics *relative to* the post-slit optics needed to be checked and adjusted to ensure that the collimated beam would be well aligned with the echelle grating. To facilitate this, a second laser alignment jig was made to fit the entrance aperture of the instrument on the pre-slit assembly (refer to Figures 1 and 4 in C08). Graph paper was taped to the rear, inner wall of the pre-slit assembly and used as a reference to adjust the alignment laser laterally and angularly until it was well centered. The plan was then to ensure that the instrument as a whole would be properly aligned relative to the telescope so that the front wall of the pre-slit assembly could be assumed normal to and centered on the nasmyth optical axis.

Past the spectrograph slit plate, field flattener conjugate, and shutter, two fold mirrors act to derotate the slit's parallactic angle, remove a small angular offset necessitated by the 35mm offset of the slit relative to the spectrograph's optical axis, and orient the spectrograph perpendicular to the nasmyth optical axis (see C06 for a more complete explanation). These fold mirrors also offer the opportunity to adjust the angular alignment of the pre- and post-slit optics relative to each other. With both the graph paper in the pre-slit assembly and the slit plate removed, the position of the laser spot relative to the centers of each collimator lens was observed and compared to the predicted x and y positions for the chief ray in the *ZEMAX* model. The angle of one of the fold mirrors was adjusted until this chief ray was aligned to its as-designed position to within a reasonable tolerance of a few millimeters. The effectiveness of this adjustment was checked by taking an out-of-focus arc

lamp spectrum through a pinhole in the slit plate to examine the shape of the pupil image. The R4 grating in PFS is well matched in width to the collimated beam size, but overfilled in length. The pupil images therefore show a characteristic shape when the optics are properly aligned.

The final step of optical alignment was to determine the optimal separation of the CCD camera from the rest of the optical system. The CCD camera, containing the CCD and the field flattener, attaches to its mount using a shim interface. The best focus was determined for each of a series of multiple shims of different thickness. The resulting relationship between shim thickness and best focus was fit to determine the shim thickness that would produce a minimum focus, and a final shim of that thickness was manufactured.

3. ELECTRONICS

Most instrumental electrical devices are controlled by a programmable logic controller (PLC) manufactured by AutomationDirect (see Figure 2). The PLC, various power supplies, and most device-specific controllers are mounted inside a metal electronics enclosure attached to the side of the instrument cart (see Figure 1). Control electronics that are not mounted in the electronics enclosure include boards for the science CCD (separated from the CCD camera and mounted outside the thermal enclosure for heat management reasons), boards integrated into the guide camera package,² CCD power supplies, a temperature controller for the iodine cell, and an ion pump controller.

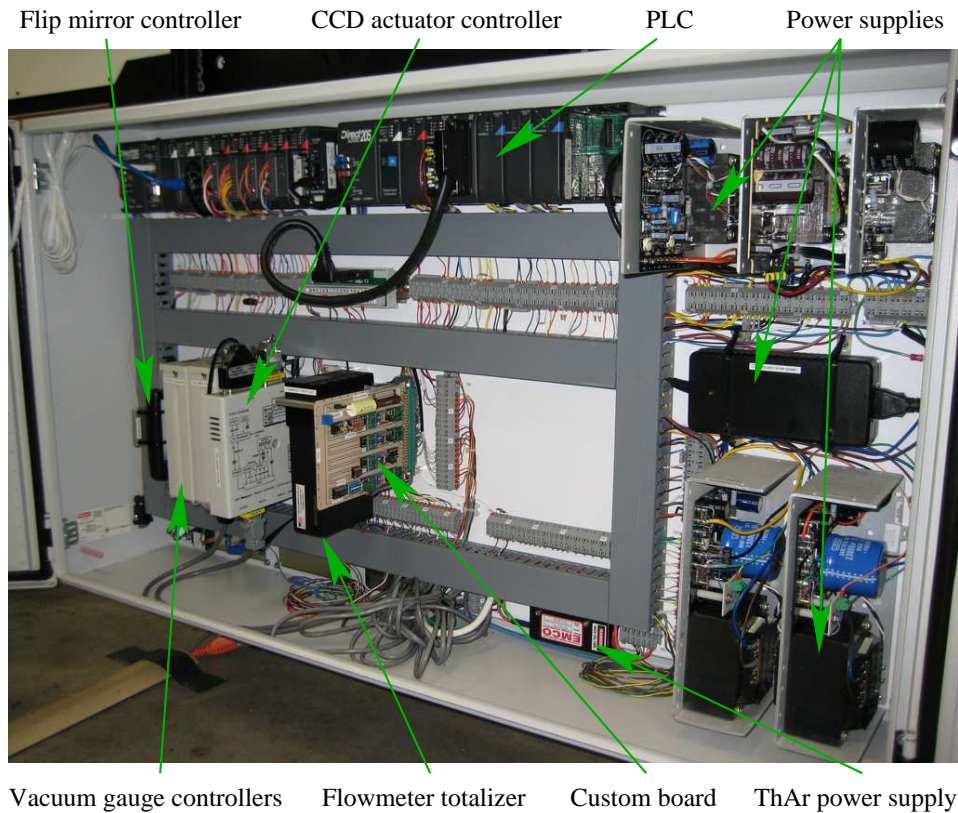


Figure 2. An electronics enclosure attached to the PFS cart contains most of the control electronics for the instrument, including a PLC, the low-level brain for most electronic functions.

A variety of different types of input and output modules are available for the PLC. Analog input modules read voltage signals from positional feedback transducers, thermistor amplifiers, etc. Analog output modules produce variable DC voltages that are used to control customized variable output DC power supplies for the thermal control system. AC output modules power the primary thermal control system heaters. Relay modules

are used to trigger DC and synchronous motor moves, among other things. Pulse outputs from a counter module command a stepper driver that moves the Hartmann mask. The system is versatile and relatively easy to modify and update as necessary.

There are currently two devices in PFS that rely on serial communication. These are the CCD motion actuator for the time delayed integration (TDI) system (see C06, C08 and Section 6.3) and the photomultiplier tube (PMT). PLCs are not designed to easily control these types of devices, although it can be done. A coprocessor module is available that contains its own memory and CPU. Using a version of the BASIC language, it can be programmed to communicate over RS-232 or RS-485 through several ports to serial devices. The CCD actuator and PMT are controlled this way. Setting up the system to work properly has been challenging. Both of these serial devices are critical to the functionality of the TDI system, and timing is important. The coprocessor is capable in many ways, but its deficiencies limit its usefulness for time-critical functions operating beyond a few Hertz. Experience with this system so far makes it clear that when a PLC is to be used for instrument control, serial devices should be avoided when possible.

Several integrated circuits and small, passive electronics are mounted on two custom wire wrap boards mounted in the electronics enclosure. These include support electronics for the synchronous motor that runs the focus actuator and the shutter in the slit viewer/guider system, voltage dividers to condition analog feedback signals from several devices, voltage regulators and thermistor amplifier circuits.

A touch panel in the electronics enclosure door allows control of most basic instrument functions not including imaging with either science or guider CCD cameras. The touch panel is intended mostly for engineering use. The astronomer uses a far more friendly graphical user interface (GUI) that runs on an Apple Mac mini in the telescope control room.

4. SOFTWARE

The PLC control flow is governed by a ladder logic program written in RLL+ (Relay Ladder Logic Plus), an AutomationDirect enhancement of the original RLL. Encoded in this software are all details related to open and closed loop control of various motion devices, operation of calibration lamps, interpretation of feedback from transducers and environmental sensors, and operation of the PFS thermal control system. Within the PLC, a coprocessor module that communicates serially with the CCD actuator and PMT runs on code written in BASIC that keeps track of exposure levels and handles most of the timing for TDI integrations.

CCD data acquisition is controlled by server software running on a Linux-based computer running separately from the rest of the instrument.

These different systems are tied together and made more user-friendly by high level software with a sophisticated GUI that runs on a Mac mini in the telescope control room. The primary windows of the GUI are shown in Figure 3.

The primary instrument functions are controlled from the GUI's Main Window. Among other things, the user may choose a slit, select a diffuser for flat-fielding, control the Hartmann mask, turn on calibration lamps, enable the on-slit guide system, and move the iodine cell into or out of the beam.

A variety of exposure types may be chosen, including the common "Object", "Flat", "Bias", and "Dark", the effects of which are partly to write the corresponding "EXPTYPE" keyword value to the FITS headers when exposures are completed and partly to enable or inhibit the opening of the shutter. Two additional modes, "Quartz" and "ThAr", ease the process of taking calibration lamp exposures. When these modes are selected and an exposure initiated, a flip mirror is moved into position to enable the calibration lamp system, the relevant lamp is turned on, and the exposure is executed. When the shutter closes, the lamp is turned off. In this way, lamp life is extended by eliminating the possibility that the user will forget to turn it off.

Several exposure modes are selectable. For the "Time" mode, the user enters a target exposure time and the shutter remains open for that period. "Counts" mode substitutes PMT exposure level as the controlling factor in the exposure time. The user enters a count target and the exposure proceeds until that target is reached. This mode is useful for ensuring that all spectra achieve the same signal to noise ratio. An estimate of the total exposure time remaining is provided in a dialog window at the bottom of the GUI and changes based on

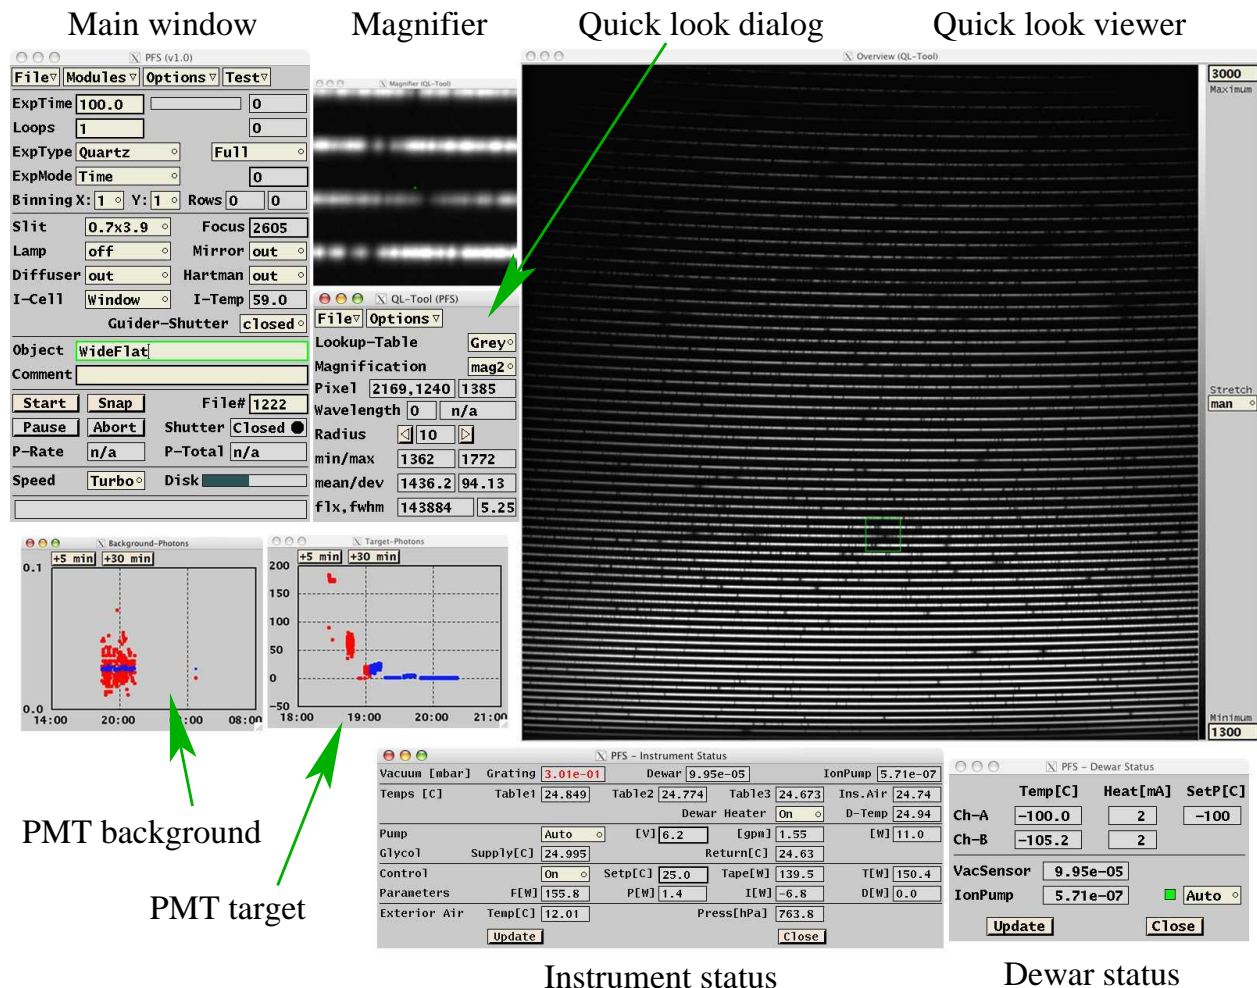


Figure 3. The primary windows of the PFS GUI are shown. Most user interaction with the instrument is via the Main Window where exposures are controlled. A Quick Look tool shows the spectra as they are read out, with optional magnification provided in a separate window. Photomultiplier count levels are monitored both numerically in the Main Window and graphically in separate, task-specific windows. Various instrument and CCD camera status data, such as temperatures, thermal control parameters, and vacuum pressures, are shown in two additional windows.

seeing, guiding variations, and atmospheric transparency. The remaining exposure modes are “TDI/Time” and “TDI/Counts”, which operate similarly to “Time” and “Counts” modes in their use of seconds and counts as exposure limits, but do so in TDI mode where the CCD is translated across a total of 100 rows while the charge is shifted in the opposite direction at the same time. Limitations in the speed at which this process can be properly executed set the minimum exposure time to 80 seconds in either TDI mode. In practice, this means that TDI cannot be used for very bright targets.

When each exposure finishes, a FITS file is written. Before or during the exposure, the user may enter an Object name and Comment in the Main window that are then written to the FITS header. The GUI software also writes the instrumental status data to keywords in the FITS headers, including exposure mode, PMT counts, slit size, the presence of the iodine cell, instrument temperatures, etc.

When the shutter closes and the pixel data readout begins, the image is displayed line by line in a Quick-Look window. This gives the user instant data quality feedback.

During normal operation, the PMT is used all night. When the PMT is enabled, the GUI keeps a running

average of the PMT dark count rate while the shutter is closed. When the shutter opens, target star counts are accumulated while accounting for the dark count rate measured before the shutter opened. Based on the number of counts that are recorded each second during the exposure, an *effective*, flux-weighted exposure mid-point is calculated and written to the FITS header. Properly accounting for this is important for accurately calculating the Earth's barycentric velocity to correct the measured stellar radial velocity.

In addition to the GUI's Main Window and Quick-Look tool, informational windows show instrumental temperatures, thermal control parameters, vacuum pressures, and PMT count rates, all of which are useful as indicators of healthy instrument functionality.

5. THERMAL CONTROL

The physical configuration and principle of operation of the PFS thermal control system is described in C08. It employs a custom liquid jacket embedded in a thick insulating enclosure that surrounds the optical table and its contents. The goal is to maintain the spectrograph at a constant temperature near 25°C throughout the year so that the internal optical focus will remain constant without need for adjustment and the refractive index of the air will be stable.

Temperature sensing is accomplished via Omega Engineering series ON-900 thermistors conditioned by INA330 thermistor amplifiers made by Texas Instruments. The amplifiers produce an analog voltage output that is read by the PLC with 12-bit resolution. The range and resolution of temperature sensitivity is established by the choice of supporting passive electronic components within each circuit. Each PFS thermistor circuit is set for one of two possible ranges. Several sensors are tuned for a coarse range of -5 to 35°C with a (PLC-converted) digital resolution of 0.01°. Other sensors are tuned for a fine range of 23 to 27°C with a digital resolution of 0.001°.

Temperatures are evaluated at several locations. Submersible thermistors monitor the temperature of the glycol solution (made from 85% deionized water and 15% Dowtherm SR-1, an inhibited ethylene glycol) on both supply and return sides of the pump and heating center. The supply side uses a fine sensor while the return side uses a coarse sensor. There are coarse range air thermistors both inside and outside the thermal enclosure. There are three fine range surface mount sensors attached to the optical table; one is near the CCD camera mount, one is near the disperser assembly, and one is near the center of the table. Finally, there is a coarse range sensor on the rear face of the CCD camera's liquid nitrogen dewar.

A subset of the thermal control data spanning one week is presented in Figure 4. The exterior air temperature, shown in the top panel, varied by close to 12°C during this period. At this time, the instrument was not in use but being stored in the telescope dome room. Nightly drops in the air temperature caused by opening the dome slit and ventilation louvers are apparent. Note that during this period, the CCD dewar was warm and the heater on the dewar shell (discussed below) was turned off. When operating, this heater responds to internal temperature changes by varying its output and has a slight dampening effect on the instrument temperature variations shown in the figure.

The instrument temperature is maintained almost entirely using a PID loop that controls the temperature of the glycol solution. Attempting to tune the system based on feedback from the actual instrument temperatures would be more difficult because the thick layer of surrounding insulation gives rise (by design) to a very long thermal time constant relative to the shorter term changes in exterior air temperature and especially modulations in the glycol solution temperature. The middle panel of Figure 4 shows that the glycol supply temperature is maintained fairly well at 25°C despite large swings in the exterior air temperature. The glycol return temperature will change according to how much heat is lost to the exterior air and to keeping the instrument warm. The typical differential between the glycol supply and return temperatures is 0.4°.

The flow rate of the glycol solution is controlled by adjusting the supply voltage to the DC pump. That voltage is set proportional to the difference between the glycol setpoint temperature and the exterior air temperature, with limits placed to ensure a minimum flow rate at all times and a maximum power demand from the pump. The flow rate, shown in the lower panel of the figure, clearly mirrors the changes in exterior air temperature.

The amount of heat required to maintain the glycol solution temperature is 11-13 Watts per degree difference between the 25°C setpoint and the exterior air temperature. The variation is partly a function of wind speed.

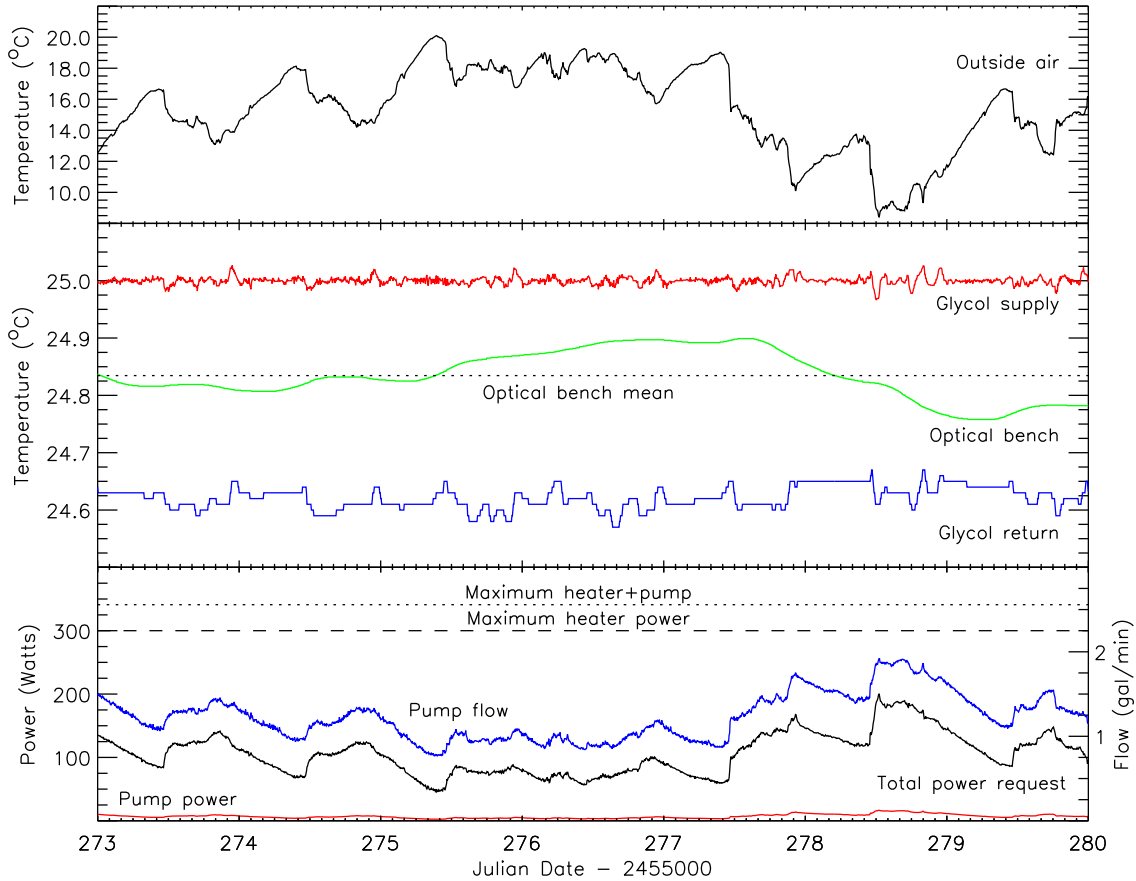


Figure 4. A subset of the available instrument temperatures and thermal control system parameters are shown for a one week span between observing runs when the CCD dewar was warm and the dewar heater was turned off. In addition to the internal instrumental temperature, the middle panel shows the supply temperature of the recirculating glycol solution immediately after it passes through its heating conduit and immediately upon return from the thermal enclosure. The glycol return thermistor is tuned to a digital sensitivity of about 0.01° , which explains the step-like appearance of its trace in the plot. The pump flow, shown in the lower panel, is set proportional to the exterior temperature. During the time period shown, the exterior air temperature varied over close to 12°C , but the internal spectrograph temperature was maintained steady at a level of 0.03° RMS. Slightly better performance ($\sim 0.02^\circ$ RMS) is achieved when the CCD dewar heater is engaged.

The maximum amount of heat that can be added to the glycol solution from heat tape is 300 Watts. In addition, the majority of the heat generated by the pump must end up in the fluid. The instrument temperature can therefore be maintained through exterior air temperature swings to as low as a few degrees below zero. The full heat Wattage is also used when driving the instrument to its setpoint from a low temperature (necessary after opening the thermal enclosure for engineering work, for example).

Not long after the thermal control system was first activated, it became apparent that the instrument temperature was being pulled down unacceptably when the CCD camera was full of liquid nitrogen. To combat this, heat tape was added to the exterior of the CCD dewar shell and is turned on only when the camera is cold (the camera is allowed to warm between observing runs). The dewar heat tape supplies 6.8 ± 1.0 Watts, with the precise level set in proportion to the interior air temperature. This mostly corrects the problem, but it is still evident that the interior instrument temperature is taking small negative hits every time the dewar nitrogen level is refreshed. Current plans are to set the dewar heat tape to a constant level and add a low level variable heater with a low power fan mounted high in the enclosure. This variable heater will respond to changes in

the interior air temperature by more directly supplying an extra Watt or two of heat in an attempt to further improve thermal stability.

The temperature of the optical bench, measured by one of the surface mount thermistors, is shown in the middle panel of Figure 4. The system currently holds a mean temperature with an RMS variation of 0.02-0.03°. While this is reasonably good, it is hoped that this can be improved to perhaps 0.01° in the future as the PID system is further tuned and the interior heating scheme is modified.

6. PERFORMANCE

A stellar spectral image taken with PFS using a 0.5 arcsec wide slit is shown in Figure 5. Several prominent absorption features are labeled.

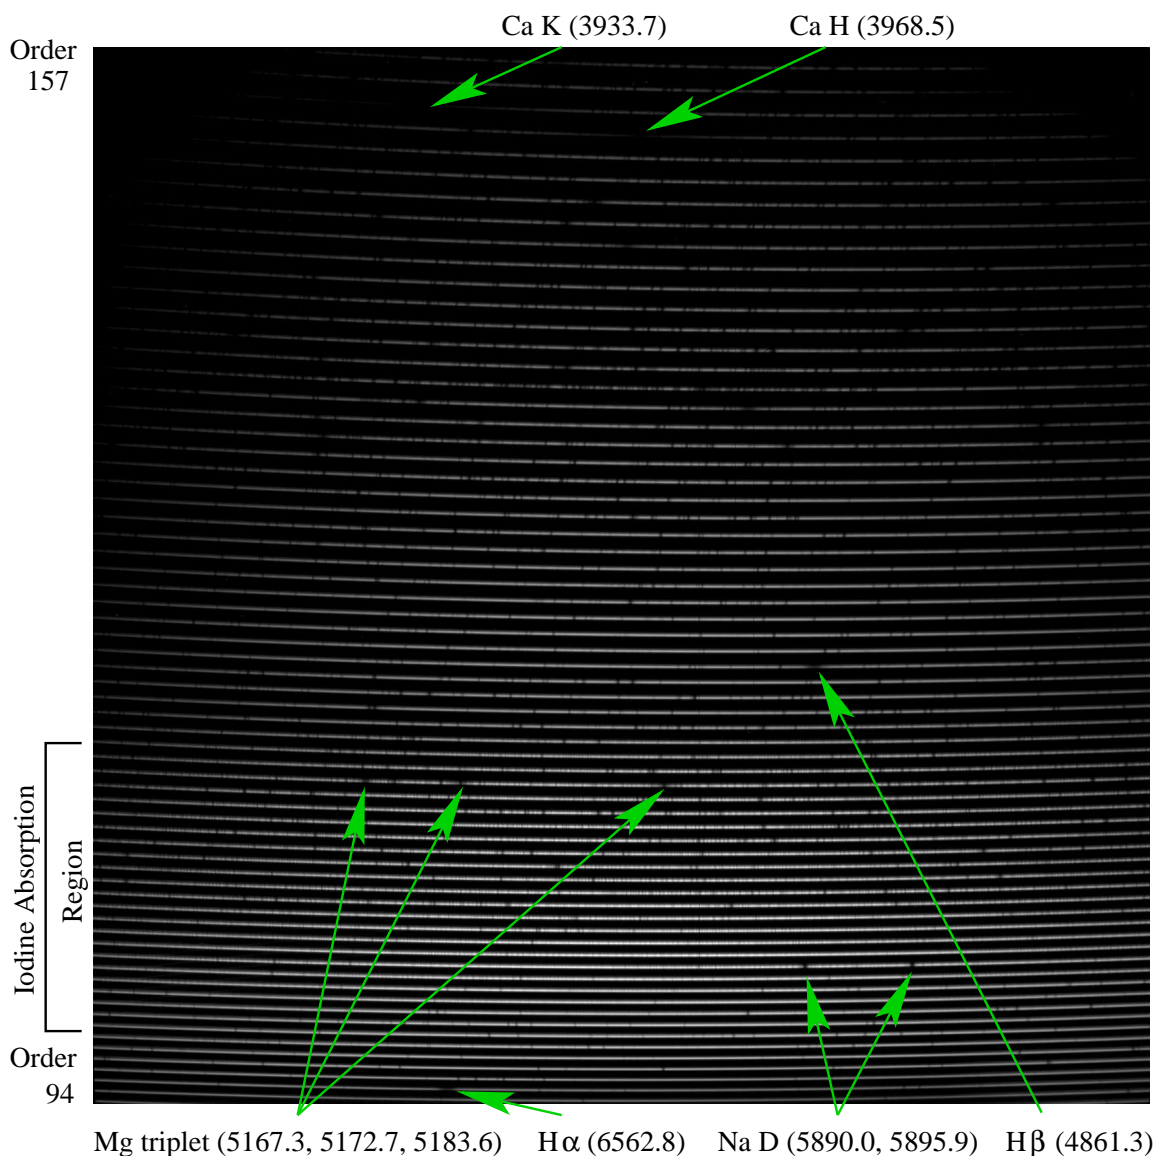


Figure 5. A spectral image taken with PFS is shown. This spectrum was taken at $R \sim 76,000$. Several prominent, well-known Fraunhofer absorption lines are identified. Although difficult to see in this image, at full resolution one can easily see the dense structure of iodine lines that blanket the spectrum between about 500 and 620 nm.

Figure 6 shows sections of two PFS spectra of the same star. One spectrum was taken using a 0.3 arcsec wide slit ($R \sim 126,000$) and the other was taken using a 0.5 arcsec wide slit ($R \sim 76,000$) with the iodine absorption cell in the beam. In these spectral regions, the iodine absorption spectrum is quite dense. The challenge of fitting the spectra to extract precise velocity information is evident.

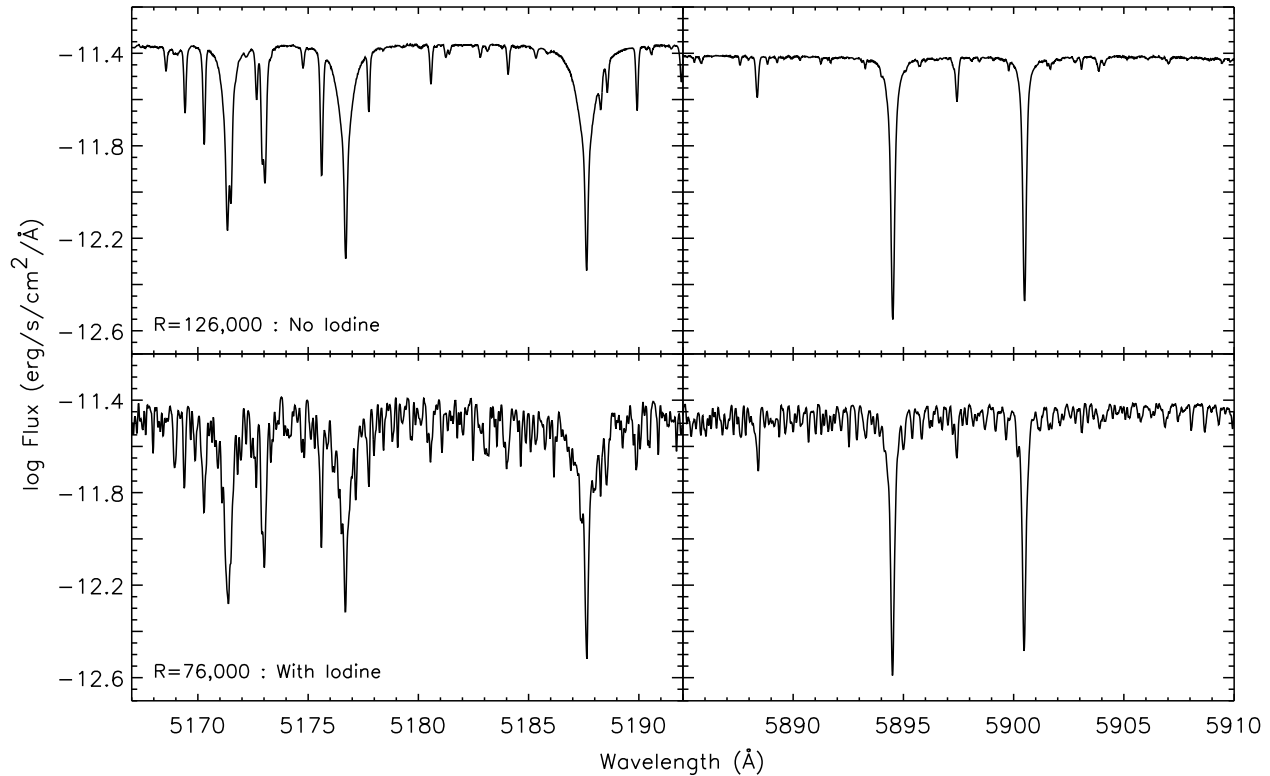


Figure 6. Two small sections of extracted, one dimensional spectra of the same star observed with two different slit widths, with and without the iodine absorption cell in the beam. The spectra on the left show the magnesium triplet while the spectra on the right show the sodium doublet.

6.1 Efficiency

The PFS instrumental efficiency profile, estimated with reference to the flux standard HD49798,⁶ is presented in Figure 7. These curves show the percentage of photons incident on a 3 arcsec wide slit that are detected by the science CCD. The peak efficiency for the spectrograph is around 23%. Taking into account telescope losses and atmospheric extinction, PFS would detect 1 photon/sec/Å near 5700Å from a target at one airmass with AB magnitude 19.0 when using a wide slit. This is comparable to the throughput of MIKE,¹ another echelle spectrograph in use on the Magellan Clay telescope.

6.2 Velocity Precision

The most direct way to evaluate the velocity precision performance of a spectrograph is to repeatedly visit a known “stable” star over an extended period of time and evaluate the variation in repeat velocity measurements. The difficulty, of course, is in identifying a star that is sufficiently stable. Fortunately, many years spent executing radial velocity planet-hunting programs has yielded as a natural byproduct many stars that are demonstrably stable to the performance limits of the instruments that have observed them.

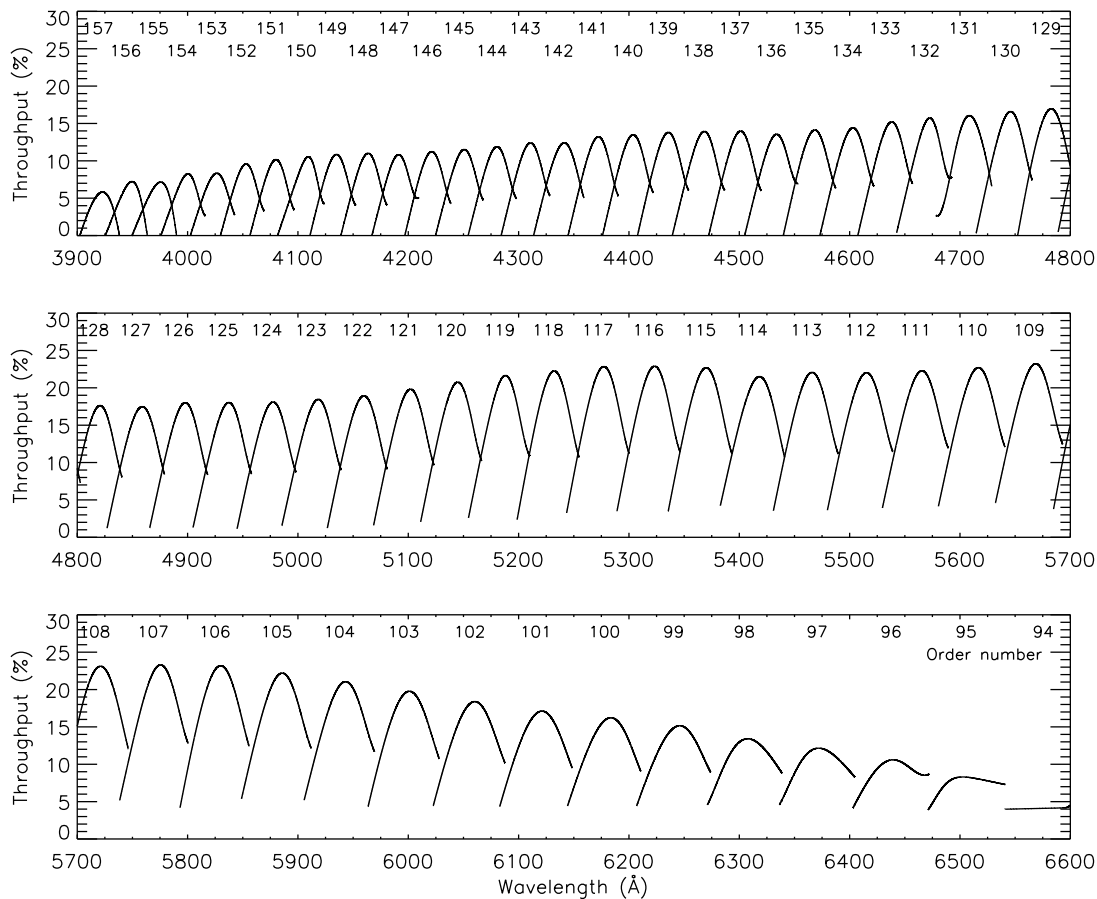


Figure 7. The PFS instrumental efficiency for each spectral order is expressed as the percentage of photons incident on the slit that are detected by the science CCD. Atmospheric extinction and telescope losses have been estimated and removed.

HD53706, a dwarf star of spectral type K0, was chosen as a suitable candidate for repeat observation by PFS during the winter and spring. Long-term monitoring at the Anglo Australian Telescope showed it to be stable to the limits of the UCLES spectrograph — at least 2.2 m s^{-1} .

Between 2010 January 1 and 2010 June 1, HD53706 was observed with PFS on seventeen different nights — ten times using TDI and seven times without using TDI. On most nights, two spectra were taken with the total time elapsed from the start of the first exposure to the end of the second exposure being about five minutes. Integrating over several minutes is important to average over helioseismology timescales. A velocity for each spectrum was determined using the method described by Butler et al (1996).³ The velocities for adjacent spectra collected on a single night were then averaged. The results are presented in Figure 8. The error bars in the figure show the RMS variation in the velocity solutions from all included wavelength “chunks” that contributed to the measurement.

When all velocity measurements are included, the RMS velocity variation is 1.05 m s^{-1} . When velocity measurements from spectra taken only in TDI mode are included, the velocity RMS is 1.25 m s^{-1} . When measurements from spectra taken only in non-TDI mode are included, the velocity RMS is 0.66 m s^{-1} .

6.3 Time Delayed Integration

C06 and C08 described the use of time delayed integration (TDI) during PFS exposures. The assumption was that detector defects at the sub-pixel level that cannot be resolved through traditional flatfielding may be quite deleterious to velocity measurements, particularly when the intended measurement precision is equivalent to one

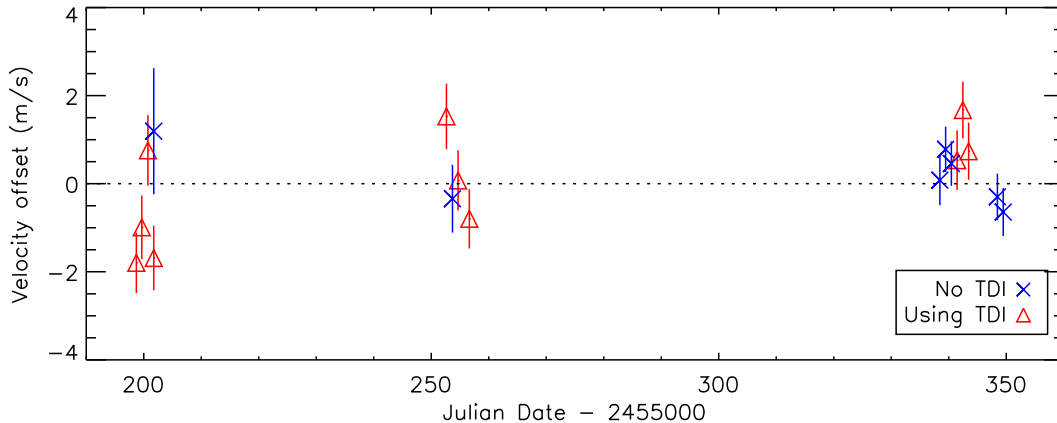


Figure 8. Velocity offsets from the mean for HD53706, a stable star observed multiple times in both TDI and non-TDI modes using PFS over a five month period. TDI and non-TDI data are indicated by different symbols. Error bars represent the RMS variation of the velocity solutions from many different wavelength “chunks” within each observation. The RMS velocity variations for all data, TDI data only, and non-TDI data only are 1.05, 1.25 and 0.66 m s^{-1} , respectively.

one-thousandth of a pixel! The goal of using TDI is to minimize the effects of these nonrecoverable errors by effectively averaging them out. Each pixel of the final image gives spectral information that has been collected in 100 separate detector pixels as the TDI procedure physically moves the CCD in the cross-dispersion direction 100 times per exposure while simultaneously shifting the charge in the other direction.

The macroscopic effect of TDI is shown in Figure 9. A region of a quartz lamp spectral image is shown in which at least two obvious spots are visible when the data are collected in a standard, timed exposure mode. The same region of the image is shown on the right when the data are collected in TDI mode. The spots are smeared out. The intention with this method is to combat problems at a much smaller scale than what is shown in the figure, but of course those cannot easily be demonstrated visually. Nevertheless, the easily seen, gross effect of TDI shown in the figure illustrates its functionality. It also raises an important point. While errors that result from defects in single detector pixels may be minimized for those individual image pixels, an unfortunate consequence of using TDI is that those defects are then spread across many output image pixels, some of which might otherwise have been “clean”. It is difficult to predict without evidence whether, on balance, these competing positive and negative effects should produce a net positive or a net negative result. However, one can draw conclusions from the data presented in Figure 8.

The HD53706 stability results lead to the conclusion that TDI does *not* have the intended effect of improving velocity precision. The current trend implies that using it results in slightly poorer performance. At best, we might expect that further testing will show that whatever benefit TDI might have is negligible, particularly considering the extra instrumental effort required to make it work. It is unclear whether this result is due to a flawed implementation of TDI within PFS, or whether TDI itself is simply of no benefit to radial velocity science.

7. DISCUSSION AND FUTURE WORK

The PFS velocity precision, both with and without use of TDI, satisfies the original goals set for the project. Moreover, if TDI is abandoned, as it will be if further testing confirms the current trend, then the PFS velocity precision appears to be well under 1 m s^{-1} . This is a thrilling result by itself, but especially when one considers that there is still work remaining that could improve the final velocity precision.

One of the unusual features of the PFS design is the novel pre-slit optical system that allows active, on-slit telescope guiding and focus. Although the optomechanical structure and electronics for this system are in place, the actual telescope control has not yet been implemented. When this implementation is finished, slit illumination should improve markedly, resulting in a more evenly illuminated pupil and a more symmetric point

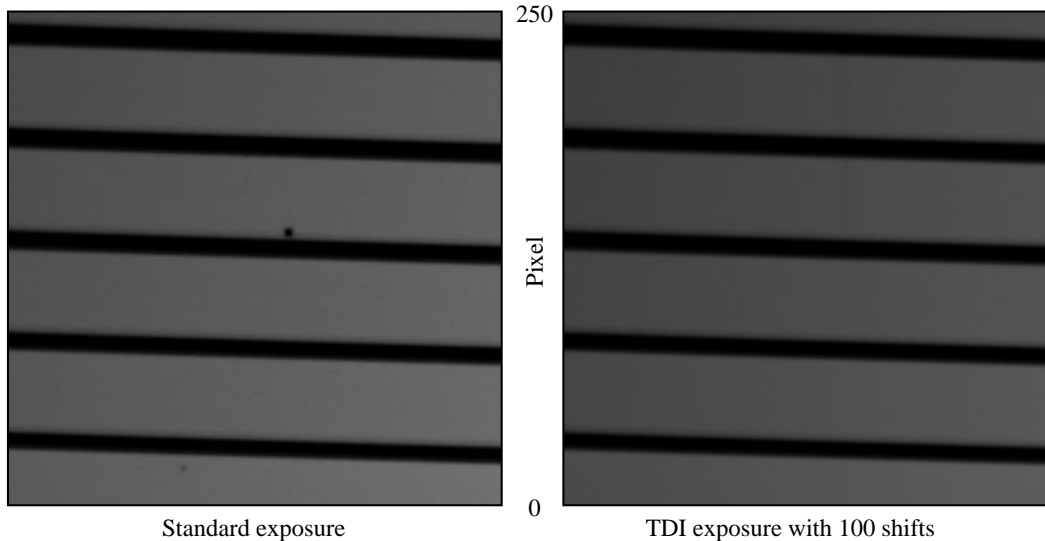


Figure 9. A comparison of images taken both with and without TDI enabled. Defects can be seen in the left panel, but are visually much reduced in significance in the right panel. The smaller, lighter defect toward the lower left of the frame is undetectable by eye in the TDI image, but the much more significant defect near the center can still be seen in the right panel as a vertical smear. Note that TDI mode is less intended for these types of macroscopic defects, and more for much smaller, sub-pixel scale detector defects.

spread function (PSF). The hope is that this will result in an improvement in the PSF modeling fits that are integral to the iodine cell method.

As stated in Section 5, the thermal control system is still being improved. The PID control parameters are not yet optimized. A variable output heater currently attached to the CCD camera dewar shell and modulated proportional to the instrument’s interior air temperature will be replaced by a system including a constant output dewar heater and a low level variable air heater. More careful work can be done to seal “hot spots” around the exterior of the thermal enclosure where heat is escaping. The net effect of these future improvements may lead to further stabilization of the instrument temperature to perhaps 0.01°C RMS or better. Whether or not this small improvement in thermal stability will translate to a noticeable improvement in *velocity* stability remains to be seen. Regardless, the feeling is that active thermal control is probably very beneficial to long-term velocity stability.

Finally, data reduction software is critically important. Improvement of the radial velocity code is a constant, unending labor. Several specific ideas for improvements to the current code are under development and one might very reasonably hope that they could contribute to increases in velocity precision.

ACKNOWLEDGMENTS

We thank Jerson Castillo, Megan Crane, Jorge Estrada, Wendy Freedman, Tyson Hare, Earl Harris, Charlie Hull, Matt Johns, Sharon Kelly, Vincent Kowal, Pat McCarthy, Greg Ortiz, Frank Pérez, Robert Pitts, Vgee Ramiah, Scott Rubel, Jeanette Stone, Robert Storts, Alan Uomoto and Steve Wilson for their contributions to this project. We also thank Rebecca Bernstein, Mario Mateo, Eric Persson, and Andrew Szentgyorgyi for helpful comments in preparation for commissioning. This work was supported by the National Science Foundation under award number AST-0116529, a generous gift from Mr. Frank H. Pearl, and the Carnegie Institution for Science.

REFERENCES

- [1] Bernstein, R., Shectman, S. A., Gunnels, S. M., Mochnacki, S., & Athey, A. E. “MIKE: A Double Echelle Spectrograph for the Magellan Telescopes at Las Campanas Observatory,” *Proc. SPIE* 4841, 1694 (2003)

- [2] Burley, G., Thompson, I., & Hull, C. "Compact CCD Guider Camera for Magellan," in *Scientific Detectors for Astronomy, The Beginning of a New Era* 300, 431 (2004)
- [3] Butler, R. P., Marcy, G. W., Williams, E., McCarthy, C., Dosanjh, P., & Vogt, S. S. "Attaining Doppler Precision of 3 m s^{-1} ," *PASP* 108, 500 (1996)
- [4] Crane, J. D., Sheckman, S. A., & Butler, R. P. "The Carnegie Planet Finder Spectrograph," *Proc. SPIE* 6269, 626931 (2006)
- [5] Crane, J. D., Sheckman, S. A., Butler, R. P., Thompson, I. B., & Burley, S. B. "The Carnegie Planet Finder Spectrograph: a status report," *Proc. SPIE* 7014, 701479 (2008)
- [6] Turnshek, D. A., Bohlin, R. C., Williamson, R. L., II, Lupie, O. L., Koornneef, J., & Morgan, D. H. "An atlas of Hubble Space Telescope photometric, spectrophotometric, and polarimetric calibration objects," *AJ* 99, 1243 (1990)