

Journal of Astronomical Telescopes, Instruments, and Systems

AstronomicalTelescopes.SPIEDigitalLibrary.org

Capabilities and performance of the Automated Planet Finder telescope with the implementation of a dynamic scheduler

Jennifer Burt
Bradford Holden
Russell Hanson
Greg Laughlin
Steve Vogt
Paul Butler
Sandy Keiser
William Deich

Capabilities and performance of the Automated Planet Finder telescope with the implementation of a dynamic scheduler

Jennifer Burt,^{a,*} Bradford Holden,^a Russell Hanson,^a Greg Laughlin,^a Steve Vogt,^a Paul Butler,^b Sandy Keiser,^b and William Deich^a

^aUCO/Lick Observatory, 1156 High Street, Santa Cruz, California 95060, United States

^bCarnegie Institution of Washington, Department of Terrestrial Magnetism, Washington, DC 20015, United States

Abstract. We report initial performance results emerging from 600 h of observations with the Automated Planet Finder (APF) telescope and Levy spectrometer located at UCO/Lick Observatory. We have obtained multiple spectra of 80 G, K, and M-type stars, which comprise 4954 individual Doppler radial velocity (RV) measurements with a median internal uncertainty of 1.35 ms^{-1} . We find a strong, expected correlation between the number of photons accumulated in the 5000 to 6200 Å iodine region of the spectrum and the resulting internal uncertainty estimates. Additionally, we find an offset between the population of G and K stars and the M stars within the dataset when comparing these parameters. As a consequence of their increased spectral line densities, M-type stars permit the same level of internal uncertainty with $2\times$ fewer photons than G-type and K-type stars. When observing M stars, we show that the APF/Levy has essentially the same speed-on-sky as Keck/high resolution echelle spectrometer (HIRES) for precision RVs. In the interest of using the APF for long-duration RV surveys, we have designed and implemented a dynamic scheduling algorithm. We discuss the operation of the scheduler, which monitors ambient conditions and combines on-sky information with a database of survey targets to make intelligent, real-time targeting decisions. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License.

Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JATIS.1.4.044003](https://doi.org/10.1117/1.JATIS.1.4.044003)]

Keywords: ground-based telescopes and instrumentation; data mining; observatory operations and science operations scheduling. Paper 15009P received Mar. 12, 2015; accepted for publication Oct. 8, 2015; published online Nov. 17, 2015.

1 Introduction

The Doppler velocity technique now has two decades of success in enabling extrasolar planetary detections. It has produced candidates with masses approaching that of Earth, and has been especially successful in detecting long period planets. In recent years, Doppler velocity confirmations have proven vital to gain an understanding of the planetary candidates discovered by photometric space missions, especially Kepler. Doppler velocity campaigns are responsible for 31% of the 1,854 confirmed planetary discoveries in the past three decades, but they account for 87% of the 330 confirmed planets with periods longer than 1 year.¹ Ground-based facilities, furthermore, are amenable to the operation of long-term surveys due to their relatively low construction and operational costs along with their ability to be upgraded as instrumentation improves.

In order to find analogs of our own solar system, we need to extend the catalog of successful radial velocity (RV) planet detections to encompass longer-period planets (particularly true Jupiter analogs) and smaller mass, short-period planets. This means, observing efforts must increase their temporal baselines and cadence of observation to more effectively populate each planet's RV phase curve.

The Automated Planet Finder (APF), located at the Mt. Hamilton station of UCO/Lick Observatory, combines a 2.4-m telescope with a purpose-built, high-resolution echelle spectro-

graph, and is capable of 1 ms^{-1} Doppler velocity precision.² Eighty percent of the telescope's observing time is specifically dedicated to the detection of extrasolar planets. This time is shared evenly between two exoplanet research groups, one at UC Santa Cruz and one at UC Berkeley. Time is allocated in whole night segments, with a schedule developed quarterly by the telescope manager. Target lists and operational software are developed separately as the two exoplanet groups are focused on different types of planet detection/follow up. For a description of the UC Berkeley planet detection efforts, see Ref. 3. The remaining 20% of telescope time is dedicated to at-large use by the University of California community. All users are allowed to request specific nights if it is beneficial to their science goals (e.g., to obtain RV values while a planet transits its star), and such requests are taken into account by the telescope manager when setting the schedule.

The APF leverages a number of inherent advantages to improve efficiency. For example, its Levy spectrometer, a high-resolution prism cross-dispersed echelle spectrograph with a maximum spectral resolving power of $R \sim 150,000$, is optimized for high precision RV planet detection.^{2,4} A full description of the design and the individual components of the APF is available in Ref. 2.

To support long-running surveys, we have developed a dynamic scheduler capable of making real-time observing decisions and running the telescope without human interaction. Through automation and optimization, we increase observing efficiency, decrease operating costs, and minimize the potential

*Address all correspondence to: Jennifer Burt, E-mail: jaburt@ucolick.edu

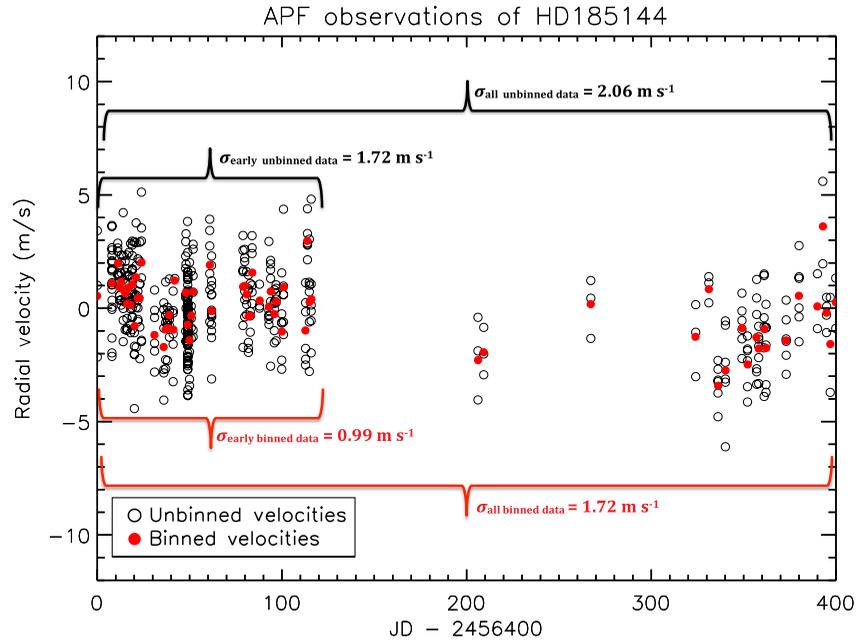


Fig. 1 Our radial velocity (RV) dataset for the RV standard star HD 185144 showing the individual exposures (open circles in black) and the resulting binned velocities (solid circles in red). The binned velocities have a much smaller RMS value (0.99 m s^{-1} for the first 120 days and 1.72 m s^{-1} over the entire 400 day span, compared with the 1.72 and 2.06 m s^{-1} exhibited by the unbinned data over the same respective time spans) due to their increased signal-to-noise. Additionally, the individual exposures are each $\sim 60 \text{ s}$ long while the binned velocities span at least 5 min, so they average over the p-modes of the star. This paper uses the individual exposure data when carrying out all further analyses and calibrations. The change in the RMS values between the first 120 days and the full data set is largely due to a change in observing strategy. After the first 4 months we started observing HD 185,144 less frequently and with fewer exposures in each observation to allow for more targets to be observed each night.

for human error. The scheduler's target selection is driven by balancing scientific goals (what we want to observe based on scientific interest, required data quality, and desired cadences) and engineering constraints (what we can observe based on current atmospheric conditions and physical limitations of the telescope). To address these criteria, we need to know how the velocity precision extracted from a given stellar spectrum depends on inputs that can be monitored before and during each observation. We have assessed the influences of the various inputs by analyzing 16 months of data taken on the APF between June 2013 and October 2014.

The plan of this paper is as follows: in Sec. 2, we describe the current APF RV catalog, paying special attention to the variety of spectral types and the frequency of observations. In Sec. 3, we evaluate the relations between velocity precision and parameters including stellar color and V band magnitude, airmass, seeing, date of observation, and atmospheric transparency and we explain how these relations inform the nightly decision-making process executed by the observing software. In Sec. 4, we outline the parameters that were assumed to be important prior to on-sky observations, but that have since been determined to have little relevance. In Sec. 5, we describe how the relevant relations are integrated into the scheduling software, and we discuss its structure, its dependencies, and its capabilities. In Sec. 6, we discuss other automated and semiautomated observatories and highlight both the similarities and differences between those systems and that employed by the APF. Finally, we conclude in Sec. 7 by reviewing the application of the APF's automated observing strategy to the telescope's current and upcoming observing scientific campaigns.

2 Dataset Description

2.1 Description of Observing Terminology

This paper differs from most publications discussing precision RV work in that all of the plots, equations, and discussions presented are based on individual, unbinned exposures of stars. It is well known that pulsation modes (p-modes) in stars cause oscillations on the stellar surface, adding noise to the RV signal. It has also been well-documented that the noise imparted by these p-modes in late-type stars can be averaged over by requiring total observation times longer than the 5- to 15-min periods typical of the pulsation cycles.^{5,6} Thus, the majority of RV publications, especially those dealing with exoplanet detections, presents binned velocities and error values. That is, they take multiple, individual exposures of a star during the night and then combine (bin) them to create one final observation with its own velocity and internal error estimate.^{3,6,7} The binned observations therefore contain more photons than any individual exposure, but, more importantly, average over the pulsation modes on the star, and therefore exhibit a measurably smaller scatter (Fig. 1).

These terms, exposure and observation, will be used specifically throughout this paper to help make clear to the reader whether we are talking about an exposure—a single instance of the shutter being open and collecting photons from the star, or an observation—the combination of all exposures taken of a star over the course of a night. Similarly, exposure time will refer to the open shutter time during one exposure and observation time will refer to the total open shutter time spent on a target

between all exposures. In the case of the star’s observation consisting of only one exposure, then the observation time is equivalent to the exposure time.

2.2 Automated Planet Finder Dataset

The APF’s star list is made up of legacy targets, first observed as part of the Lick–Carnegie exoplanet campaign using the high resolution echelle spectrometer (HIRES) spectrograph on the Keck I telescope. Stars for the Lick–Carnegie survey were selected based on three main criteria:

- Spectral type between G and M: stars with spectral types earlier than F5 experience oscillations that can produce pseudo-Keplerian signals, and are fast rotators with few spectral lines so they are avoided.
- Single stars: no stars with companions closer than $5''$, as this can lead to scattered light from the nearby companion making its way into the spectrograph slit and contaminating the target star’s signal.
- Quiet stars: only stars without large amounts of emission in the Ca II H & K line cores are permitted, as core emission in these lines is an indicator of stellar chromospheric activity⁸ which complicates the data reduction process and can produce pseudo-Keplerian signals.

The resulting Lick–Carnegie target list is comprised of ~ 1800 stars, which have been monitored using Keck/HIRES over the past two decades. When creating the initial target list for the APF, the Lick–Carnegie star list was culled for targets with V magnitudes brighter than 12 and declinations above -20 deg. In order to efficiently prove the APF’s capabilities, we selected stars with suspected short-period planets ($P < 100$ days) that required only 1 to 2 more rounds of phase coverage to verify. The presence and the false-alarm probabilities (FAPs) of these short-period, Keplerian signals were determined by analyzing the existing Keck/HIRES RV datasets using the publicly available Systemic Console.⁹ Systemic allows users to fit planetary signals to RV data and derive the orbital properties, while also providing tools to handle error estimation and assess orbital stability. This selection process resulted in a list of 127 stars.

The calibrations described in this paper are based on data taken with the APF between June 2013 and October 2014. The dataset includes precision Doppler observations of 80 of the 127 stars selected from the Lick–Carnegie survey and this is before the development/inclusion of the dynamic scheduler, thus all data in this paper were obtained using fixed star lists. Our precision Doppler observations encompass spectra of 80 stars and incorporate 600 h of open shutter time. The stars span spectral types from early G to mid M, have $3.5 < V < 12$, and are all located within 160 pc (Fig. 2).

Every APF star has a set of observations containing between one and seven hundred exposures. Individual exposures are restricted to a maximum length of 15 min to avoid cosmic ray accumulation and to minimize uncertainty when calculating the photon-weighted midpoint times. Additionally, we enforce a total observing time limit of 1 h per target per night to ensure that telescope time is not wasted on observing faint objects when conditions are poor.

For each individual exposure, the $\overline{\text{FWHM}}$ [the average full width at half maximum (FWHM) of the star over the integration time] of the guide camera’s seeing disk is logged in the flexible

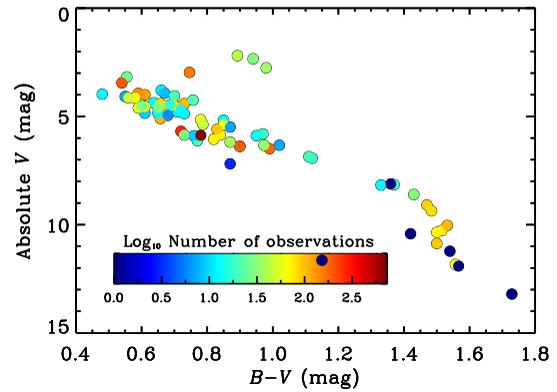


Fig. 2 Color-magnitude diagram for stars observed with Automated Planet Finder (APF) and used in the analysis described herein. Color coding represents the number of exposures that have been obtained for each star. Our dataset spans a wide range in color and magnitude and contains stars with spectral types from F6 to M4.

image transport systems (FITS) image header, along with the total number of photons from the exposure meter and the total exposure time. Colors, magnitudes, and distances for the stars are obtained from set of identifications, measurements, and bibliography for astronomical data (SIMBAD).¹⁰

All spectra are analyzed using the data reduction techniques described in detail by Butler et al.,¹¹ which produce a measurement of the stellar RVs and associated internal uncertainties. The reduction pipeline analyzes each exposure’s extracted spectrum in 2-Å chunks and determines the RV shift for each chunk individually. The final reported velocity is the average over all the chunks, while the internal uncertainty is defined as the RMS of the individual chunk velocity values about the mean, divided by the square root of the number of chunks. Thus, the internal uncertainty represents errors in the fitting process, which are dominated by the photon noise. The internal uncertainty does not include potential systematic errors associated with the instrument, nor does it account for astrophysical noise (or “jitter”) associated with the star and, therefore represents only a lower limit to the accuracy of the data for finding companions.

2.3 Determining Additional Systematic Errors

We use the APF’s 737 individual exposures of HD 185144 (Sigma Draconis, HR 7462), a bright RV standard star, to estimate the precision with which we can measure radial velocities. It should be noted that the RV values produced by our analysis pipeline are all relative velocities and thus have a mean value of zero. We examine the unbinned, or individual, exposures, finding that star’s mean internal uncertainty is $\mu_{\text{int}} = 1.0 \text{ ms}^{-1}$ and that the mean of the absolute values of its velocity measurements is $\mu_{\text{abs}} = 1.8 \text{ ms}^{-1}$. Because μ_{abs} includes effects from the internal uncertainty in addition to other sources of error such as the stellar jitter and the instrument systematics, its value is always higher than the internal uncertainty value for a given exposure.

The difference between these two parameters implies an additional quadrature offset of 1.5 ms^{-1} , which we then generalize as $\sigma_s = 1.5 \text{ ms}^{-1}$ (Fig. 3). We take this offset value to represent the additional error contributions from all other systematics, including the known 5 to 15 min pulsation modes of the star that the binned data averages over and the systematics from the instrument. If, instead, we compared the mean internal uncertainty to the standard deviation of the

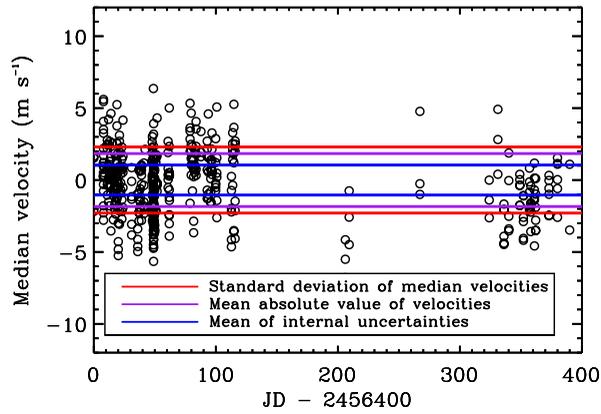


Fig. 3 Radial velocities from APF exposures on the star HD 185144 [G9V, $M_v = 4.68$], 737 points in total. The internal uncertainty estimates produced by the data reduction pipeline are noticeably smaller than the actual spread in the data. The internal error does not account for telescope systematic errors or sources of astrophysical noise in the star. Our average internal uncertainty is $\mu_{\text{int}} = 1.0 \text{ ms}^{-1}$ (blue line) but we find the mean absolute value of the HD185144 velocity measurements to be $\mu_{\text{abs}} = 1.8 \text{ ms}^{-1}$ (purple line). We adopt an estimate that the additional systematic uncertainty in our velocity precision is $\sigma = 1.5 \text{ ms}^{-1}$.

velocity measurements ($\sigma_{\text{vel}} = 2.3 \text{ ms}^{-1}$), we would require a larger additional offset term ($\sigma_s = 2.0 \text{ ms}^{-1}$). We note, however, that the data are strongly affected by the velocities of exposures that fall in the outlying, non-Gaussian tail. We thus choose to determine our error estimate using the mean of the absolute values of the velocity measurements (instead of the standard deviation of the velocities), as doing so mitigates the influence of the outliers.

As stated, the internal uncertainties and velocities used in this analysis are extracted from the individual spectra obtained by the telescope. Our normal operational mode determines the internal uncertainties from data binned on a 2-h time scale, which is the approach used by Vogt et al.² and yields a standard deviation of $\sigma_{\text{bin}} = 1.05 \text{ ms}^{-1}$ for HD 185144. This value is notably smaller than the standard deviation of the individual velocities, $\sigma_{\text{vel}} = 2.3 \text{ ms}^{-1}$, because we deliberately acquired six observations of HD 185144 in order to both average over the pulsation modes and achieve a high precision for the final binned observations.

We use the relations of Wright,¹² which present RV jitter estimates at the 20th percentile, median, and 80th percentile levels to assess the expected stellar activity for HD 185144. We find an estimated median jitter of $\sigma'_{\text{rv}} = 3.5 \text{ ms}^{-1}$, and a 20th percentile value of 2.3 ms^{-1} . Noting that the 20th percentile value matches the σ_{vel} for our exposures, obtained over a timespan of 400 days, suggests that the star is in fact intrinsically quiet. That is to say, we expect only 20% of stars with the same evolution metric, activity metric (both described by Wright¹²), and $B - V$ color as HD 185144 to have activity levels less than 2.3 ms^{-1} . Given that we measure $\mu_{\text{abs}} = 1.8 \text{ ms}^{-1}$, HD 185144 is, at a minimum, much quieter than expected based on its color, activity, and evolution. Even so, it is reasonable to expect that some fraction of the observed uncertainty is due to the star itself (astrophysical noise) or currently unknown planets rather than instrumental effects, thus indicating that the instrument is performing very well.

To verify that the size of the offset between μ_{int} and μ_{abs} is not specific to the HD 185144 dataset, we compare the mean values

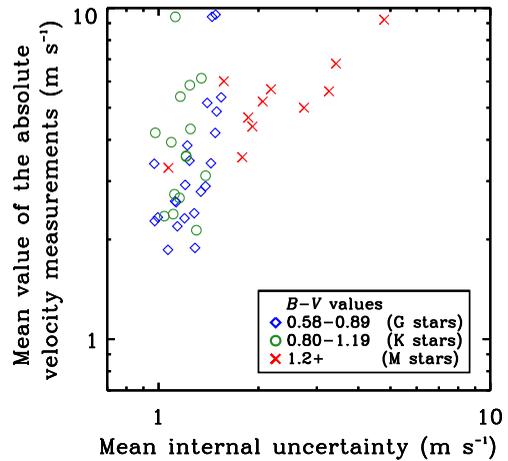


Fig. 4 Mean values of μ_{int} and μ_{abs} for each star observed during the same time span as the HD 185144 analysis presented in Figure 3. The μ_{abs} values are always larger because they include the internal uncertainty in addition to other effects such as stellar jitter and instrument systematics. For bright, quiet stars it is possible to reach μ_{abs} values of 1.8 ms^{-1} .

of these parameters for every star observed by the APF during the time span of Fig. 3. That is, we apply the same procedure detailed above for HD 185144 to all stars observed during the same date range, and then compute and compare the mean values of μ_{int} and μ_{abs} for each star's set of exposures (Fig. 4). As expected, the values for the mean of the absolute RV values are always higher than the mean internal uncertainty values because μ_{abs} includes the effects of the internal uncertainty combined with additional sources of error such as the stellar jitter and instrument systematics. Additionally, some of these stars are planet hosts, and thus display even higher μ_{abs} values because of the planet's influence. However, the quietest, nonplanet-hosting stars are able to reach μ_{abs} values of slightly less than 2 ms^{-1} , suggesting that the additional error offset value determined using HD 185144 ($\sigma_s = 1.5 \text{ ms}^{-1}$) is appropriate.

3 Observing Inputs

To create an efficient and scientifically informative exoplanet survey, we must balance scientific interest in a range of target stars with the observing limitations presented by the weather and the telescope's physical constraints.

3.1 Parameters of Scientific Interest

There are a number of criteria of astronomical importance. For each star, the following parameters are utilized by the scheduling software to determine whether the star is given a high rating for observation:

- Observing priority—a numerical rating that reflects the observers' assessment of potential Keplerian signals in the RV data.
- Desired cadence—dictates the desired wait time between observations of a specific target.
- Desired precision—the average allowable internal velocity uncertainty for measurements of a specific target.

Both observing priority and cadence are determined via the observers' examination of the star's existing data set. As mentioned in Sec. 2, we use the publicly available Systemic Console

to analyze our RV datasets. The console produces a Lomb–Scargle periodogram from the selected RV dataset, which displays peaks corresponding to periodic signals in the data. When a peak surpasses the analytic FAP threshold defined by this method, the observers note the period and the half-amplitude of the signal and use those values to decide upon an observing cadence and desired level of precision that will fill out the signal’s phase curve quickly and with data points of the appropriate SNR. The desired precision is further refined if the observers have some knowledge of the star’s stellar activity, as this sets a lower limit to the attainable precision. Details on the fitting procedures and the statistical capabilities of the Systemic Console are described in detail in Ref. 9.

For each star, this information (along with other characteristics such as right ascension, declination, the V magnitude and the $B-V$ color) is stored in an online Google spreadsheet accessible to team members and the telescope software. This database of target stars drives the survey design and target selection while also being easily accessed, understood, and updated by observers.

3.2 Relating Iodine Region Photons to the Internal Uncertainty of Radial Velocity Values

In perfect conditions (no clouds, no loss of light due to seeing), all stars that are physically available and deemed in need of a new observation (based on their cadence) are simply ranked by observing priority and position in the night sky and then observed one after another, until dawn. Cloud cover and atmospheric turbulence, however, make such conditions rare. Furthermore, stars with low declinations spend only short periods at low-air mass. Consequently, there are a substantial number of constraints that affect the quality of an exposure. In order to maximize the scientific impact of each night’s exposures, the conditions must be evaluated dynamically, and data taken only when the desired precision listed in the database is likely to be attainable.

To identify the targets that can be expected to attain their desired precision within the 1-h observing time limit, the observing program must link the night-time conditions and the physical characteristics of each star with the resulting internal uncertainty over a given exposure time. To this end, we first relate the internal uncertainty of a given exposure to the number of photons that fall in the iodine (I_2) line-dense region of the spectrum (i.e., the $\lambda \sim 5000$ to 6200 \AA bandpass) where our RV analysis is performed.

We fit the relation between internal uncertainty and photons in the iodine region separately for the G and K star dataset (comprised of 2790 G star exposures and 957 K star exposures) and the M star dataset (comprised of 837 M star exposures). Because information on a star’s RV value comes from the location of its spectral lines, we expect the M stars—which contain many more spectral lines—to achieve better precision for the same number of photons. This expectation is validated by Fig. 5, where a vertical offset between the G/K star best-fit line and the M star best-fit line is evident. Comparing the zero points for the best-fit lines (Table 1), we find a factor of 2 difference between the stellar groups. Thus, the M stars reach the same level of internal uncertainty with half the number of photons required by the G and K star population. Additionally, the slopes of both best-fit lines are close to the $m = -2$ expected for shot-noise limited observations with $m = -1.58$ and -1.73 for the G/K and M stars, respectively. The fact that the slopes are shallower than $m = -2$, indicating performance better than shot-noise limited,

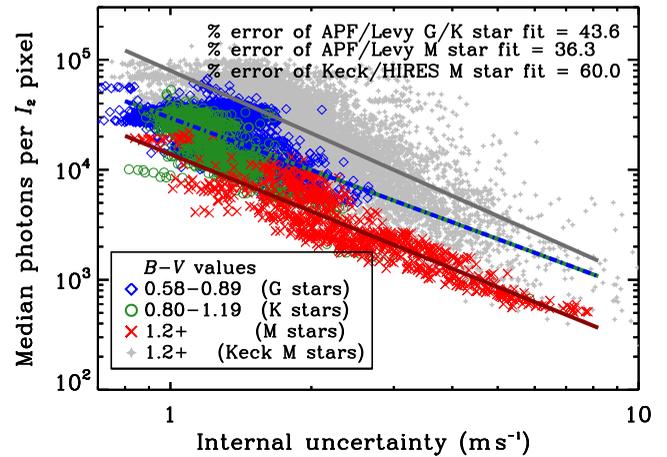


Fig. 5 Observations of G (blue), K (green), and M (red) type stars during the 1.5 years of APF/Levy observations. We plot the individual 2,790 G star exposures, 957 K star exposures and 837 M star exposures. The gray diamonds represent M star data obtained using Keck/high resolution echelle spectrometer (HIRES) as part of the Lick–Carnegie planet survey where the Keck/HIRES pixels have been scaled to match the same $\text{\AA}/\text{pixel}$ scale as the APF/Levy ($0.0183 \text{ \AA}/\text{pixel}$). We find that the G and K stars have the same zero point and the same slope, so we combine these two data sets for this analysis. The green and blue dashed line represents the best-fit to the APF/Levy’s combined G and K star dataset, while the red line shows the best-fit to the APF/Levy M star data set which is fit separately due to the increase of spectral lines in later spectral types. As expected, the APF/Levy M stars show higher data precision for the same number of photons in the I_2 region of the spectrum. The percent errors quoted on the figure are calculated using the scatter in the difference between the observed I_2 photons and the I_2 photons from the best-fit lines, so they represent the scatter of the sample and not the error on the mean. The dark gray line is the best fit to the Keck/HIRES M star dataset. Comparing this to the red line reveals that the APF requires $5.75\times$ fewer photons in the I_2 region to achieve velocity precision comparable to Keck/HIRES on M stars down to at least $M_v = 10$.

is reasonable as the internal uncertainty only accounts for the errors resulting from the extraction of the spectrum from the original FITS files. We emphasize that this is just a piece of the error budget, and does not include other random or systematic errors.

The gray points in Fig. 5 correspond to spectra of M stars obtained with Keck/HIRES since November 2002 as part of the Lick–Carnegie planet search. There are 168 stars represented, all with $B - V > 1.2$, resulting in 8872 individual exposures. In order to compare these individual velocities to those obtained on the APF in a meaningful way, we rescale the Keck/HIRES pixels so that they represent the same range of \AA per pixel as those on the the APF/Levy. This involves two different scaling factors as the HIRES instrument underwent a detector upgrade in 2004 that changed its pixel size from 24 to $15 \mu\text{m}$, resulting in different sampling values. Applying these factors means that all of the data shown in Fig. 5 represent the median number of e^- per pixel, where each pixel covers 0.0183\AA —the native value for the APF/Levy in the I_2 region.

The figure shows that for M stars, the APF requires $\sim 5.75\times$ fewer photons in the I_2 region to achieve velocity precision comparable to Keck. [Based on the work of Bouchy et al.,¹³ we expect that for K dwarfs, the relative speed should scale as the “information content” Q ,¹⁴ which is proportional to the ratio of the resolutions, $Q \propto (R_{\text{APF}}/R_{\text{HIRES}})$. For HIRES, the

Table 1 Values for fit variables in Sec. 3.

Variable	Equation	Description	Value(s)
$A_{G/K}, A_M$	Eq. (1)	Constant terms for APF G/K and M stars	4.47, 4.14
$B_{G/K}, B_M$	Eq. (1)	Linear terms for APF G/K and M stars	-1.58, -1.73
A_{Keck}	Eq. (1)	Constant term for Keck M stars	5.03
B_{Keck}	Eq. (1)	Linear terms for Keck M stars	-1.83
$\sigma_{G/K}$	Eq. (1)	RMS for fit to APF G/K stars	0.182 dex, 43.5%
σ_M	Eq. (1)	RMS for fit to APF M stars	0.151 dex, 36.2%
σ_{Keck}	Eq. (1)	RMS for fit to Keck M stars	0.252 dex, 60.0%
α_I	Eq. (2)	$B-V$ term in I_2 photon rate linear regression	-0.0311
β_I	Eq. (2)	Airmass term in I_2 photon rate linear regression	0.158
γ_I	Eq. (2)	Date term in I_2 photon rate linear regression	0.00136
C_I	Eq. 2	Constant term in I_2 photon rate linear regression	-11.2
σ_{I_2}	Eq. (2)	RMS value of $r_{I_2} - r_{I_2\text{obs}}$	0.0499 dex, 11.5%
α_E	Eq. (4)	$B-V$ term in the exposure meter photon rate linear regression	-0.908
β_E	Eq. (4)	Airmass term in exposure meter photon rate linear regression	0.0852
γ_E	Eq. (4)	Date term in exposure meter photon rate linear regression	0.00118
C_E	Eq. (4)	Constant term in exposure meter photon rate linear regression	-21.8
σ_E	Eq. (4)	RMS value of $r_E - r_{E\text{obs}}$	0.0742 dex, 17.8%
δ	Eq. (5)	Constant $B-V$ term in the I_2 photons versus exposure meter photons	4.52
ϵ	Eq. (5)	Linear $B-V$ term in the I_2 photons versus exposure meter photons	-0.196
ζ	Eq. (5)	Quadratic $B-V$ term in the I_2 photons versus exposure meter photons	0.262
$\sigma_{I_2\text{Exp}}$	Eq. (5)	RMS for fit to I_2 photons versus exposure meter photons	0.0463 dex, 11.1%

“throughput” (the resolving power times the angular size of the slit) is $39,000''^{15}$ and for the APF it is $114,000''^{2}$. Normally, HIRES was used with the $0.861''$ slit giving a filled aperture resolution of 45k while for the Levy a $1''$ slit is used, so the ratio of the resolutions is 2.5. The Levy demonstrates a larger than expected improvement over HIRES which could be explained by the increased number of lines in the iodine region for M dwarfs. Further investigation is beyond the scope of the current paper, but we plan to include such analysis in a future publication. We note that the excellent seeing at MK means that some data were observed with a much higher effective R, up to 90 k, which may explain the large scatter we see in Fig. 5.] Speed estimates for the APF/Levy, carried out last year,² show that the telescope and instrument together are approximately $6\times$ slower than Keck/HIRES. Combining these two effects indicates that the APF has essentially the same speed-on-sky as Keck/HIRES for precision RVs of M stars. This is not altogether unexpected, as HIRES was never specifically optimized for precision RV work. The APF’s Levy spectrograph was purpose-built for high precision, RV science and therefore features much higher spectral resolution and finer wavelength sampling than HIRES. Both of these factors, as well as the

significantly higher system efficiency of the APF/Levy optical train over that of Keck/HIRES,² combine to make APF as fast as Keck/HIRES for precision RV work on M dwarfs, at least down to $M_v = 10$.

The functional form of these best-fit lines is given by Eq. (1). Note that the numeric values for each variable [in the case of Eq. (1), A and B] are listed in Table 1. This format will be used for all relations presented in Sec. 3.

Functional form of the fits applied in Fig. 5

$$\log_{10}(N_{\text{med}}) = A + B \cdot \log_{10}(\sigma_{\text{int}}), \quad (1)$$

where N_{med} is the median number of photons per pixel in the I_2 region for a given exposure, and σ_{int} is the estimated internal uncertainty for the resulting RV value.

To assess whether the exposures of these stars represent a Gaussian distribution, we compare the standard RMS and the scatter calculated using a Tukey’s biweight method around each fit. The Tukey’s biweight scatter provides a more robust statistic for data drawn from a non-Gaussian distribution as it less heavily weights the outliers, which are assumed not to be part of a normal distribution.¹⁶ For the G and K star fit, the

standard RMS is 0.182, while the biweight scatter is 0.186. Similarly, for the M stars, the standard RMS is 0.151, while the biweight scatter comes out to 0.155.

In the limit of a true Gaussian distribution, these two metrics would produce the same result. Employing a bootstrap analysis of each method, we find the standard deviation of both the RMS and the biweight scatter of the G and K star sample to be 0.002. Similarly, the standard deviation of both the RMS and the biweight scatter of the M star sample is 0.003. Noting the similarity of these standard deviations with the actual offsets found between the RMS and biweight scatter, we determine that the observations for both sets of stars are drawn from a mostly Gaussian distribution.

3.3 Real-Time Effects

3.3.1 Data selection

Knowing the median number of photons per pixel in the iodine region required to achieve a given level of RV precision enables us to determine the expected exposure time for a star, if we know how quickly those photons accumulate. To determine this rate and the relation between the final exposure meter value and the number of photons in the I_2 region (used to set upper bounds on observing time), we study a subset of the year of APF data described in Sec. 2.

Cuts are applied to the main dataset to select only those observations taken on clear nights and in photometric conditions, as nonphotometric data will induce skew in the results. First, we select only exposures with seeing $\overline{\text{FWHM}} < 2''$, which results in 935 individual observations. We then perform separate, multivariate linear regressions on the photon accumulation rate in the I_2 region of the spectrum and the photon accumulation rate for the exposure meter [Eqs. (2) and (4)] on all remaining data points. In each case, we calculate the variance of deviations from the best-fitting relation for all of the data. We also calculate the variance for all of the points on a given night. We then reject nights using the F -ratio test. Namely, if the standard deviation of a given night is more than twice the standard deviation of the population as measured by the Tukey's biweight, then all observations from that night are rejected. This results in datasets of 865 exposures used in the I_2 photon accumulation rate regression and 816 exposures in the exposure meter photon accumulation rate regression. Points that fall significantly below the regression line are most likely due to clouds, while those falling above the line are likely due to erroneous readings from the exposure meter.

Once the nights with large variance have been removed, we repeat the regressions on the remaining points to determine the actual fits described in Sec. 3.3.2. In Figs. 6 and 7, the points in color are those used to perform the linear regressions, while the points in gray are those we rejected after they were deemed nonphotometric. Figure 8 uses the same set of data as Fig. 7, in order to keep only the normally distributed exposure meter readings.

3.3.2 Linear regressions

To determine the predicted observation time of a star, we perform a multivariate linear regression using the trimmed dataset resulting from the procedure described in Sec. 3.3.1. The regression estimates the rate at which photons accumulate in the pixels of the iodine region of the spectrum, and accounts for: (1) the star's V magnitude, (2) its $B-V$ color, (3) slit loss due to the

current seeing conditions, (4) the airmass based on the star's location, and (5) the modified date of the observation (Fig. 6). The modified date is calculated by subtracting the maximum date from each observation following the selection process described above. This makes the zero point of the relationship the value at the time of the last photometric observation. We use the modified date parameter to address the degradation of the telescope's mirror coatings over time. When the mirrors are recoated, it will introduce a discontinuity in this parameter, and we will then adjust the zero point of all regression fits based on the new throughput estimates and watch for any changes that develop in the slope of the regressions.

The multiparameter fit over these five variables results in a best-fit plane, of which we present a projection in Fig. 6. To help visualize the goodness of fit, we plot the data on one axis, the linear regression combination on the other, and place a 1:1 line on top. This approach is also used when plotting the linear regression in Fig. 7.

The regression gives

$$\log_{10}(r_{I_2}) = -\frac{1}{2.5} \{V_c + \alpha_I(B - V) + \beta_I[\sec(z)] + \gamma_I(\text{MJD}) + C_I\}, \quad (2)$$

with

$$V_c = V_m - 2.5 \cdot \log_{10}(f_t), \quad (3)$$

where r_{I_2} is the photon accumulation rate in the iodine region, z is the angle of the star relative to zenith, MJD is a modified Julian date, and f_t is the fraction of the starlight that traverses the spectrograph slit.

By dividing the number of I_2 region photons necessary to meet our desired RV precision [derived from Eq. (1)] with

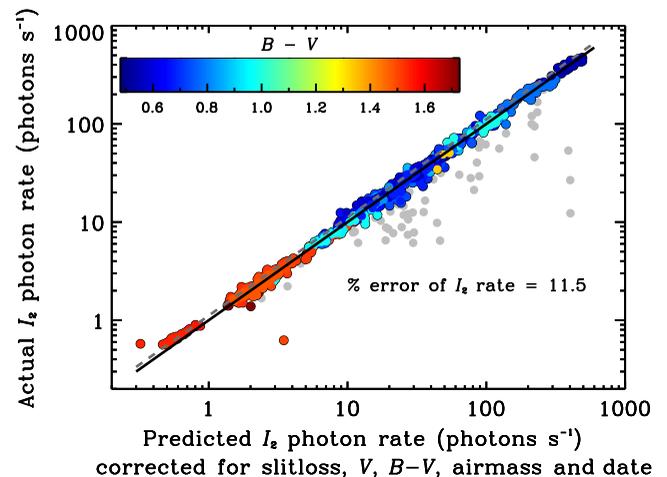


Fig. 6 Multivariate linear regression of the iodine pixel photon accumulation rate that incorporates stellar color, stellar magnitude, atmospheric seeing, and airmass. Colored points are used in calculating the regression, while gray points have been rejected as nonphotometric data as described in Sec. 3.3.1. The black line is a 1:1 relationship, and the gray dashed line shows the relation offset by one standard deviation, which is the limit we use operationally. The strong correlation between the data and the regression line enables prediction of the rate of photon accumulation in the spectrum's iodine region (a value not calculated until the data reduction process) using the stellar properties and ambient conditions. We can thus estimate the observation time required to meet a specific median I_2 photon value, and, in conjunction with Fig. 5, a RV precision.

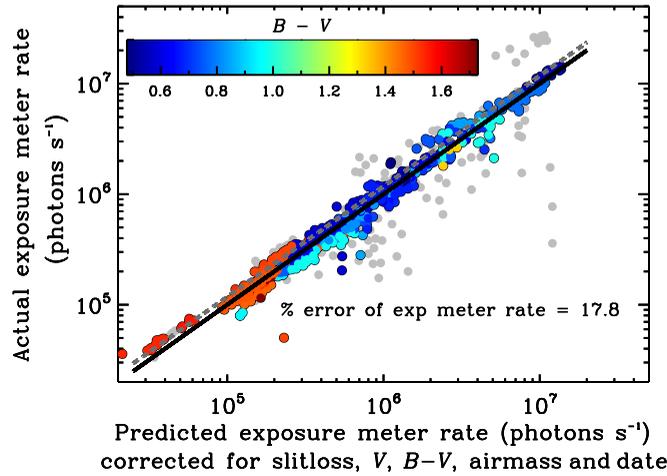


Fig. 7 Multivariate linear regression to the exposure meter photon accumulation rate as measured on the APF guider, which incorporates stellar color, stellar magnitude, atmospheric seeing, and airmass. Colored points are used in calculating the regression, while gray points have been rejected as described in Sec. 3.3.1. The black line is a 1:1 relation, and the grey dashed line shows the relation offset by one standard deviation, which is the limit we use operationally. The strong correlation permits prediction of the expected exposure photon accumulation rate for a given star in photometric conditions, and thus provides a measure of the transparency. Any decrease in exposure photon accumulation rate from what is predicted is presumed to arise from decreases in atmospheric transparency brought about by cloud cover.

the photon accumulation rate in the I_2 region calculated using Eq. (2), we can determine the predicted observation time for a star for a specified internal uncertainty in a given set of conditions. These predicted observation times account for atmospheric conditions such as seeing and airmass, but do not address the issue of atmospheric transparency. The APF lacks an all-sky camera with sufficient sensitivity to assess the brightness of individual stars, meaning that we cannot evaluate the relative instantaneous transparency of different regions of the sky. Instead, we must determine the transparency during each individual observation by comparing the rate of photon accumulation we observe with what is expected for ideal transparency. Although the I_2 region photons provide a straightforward way to determine the predicted exposure times, the number of iodine region photons is available only after the final FITS file for an observation has been reduced to yield a RV measurement. Thus, we cannot monitor in real time the rate at which they are registered by the detector to assess the cloud cover.

Instead, we compute a transparency estimate during each exposure using the telescope’s exposure meter. The exposure meter is created by using series of two-dimensional (2-D) images from the guider camera that are updated every 1 to 30 s depending on the brightness of the target. Rather than guiding on light reflected off a mirrored slit aperture, as is traditionally done, the APF uses a beamsplitter to provide 4% of incoming light to the guider camera as a fully symmetric, unvignetted seeing disk. This allows a straightforward way to monitor how well the telescope is tracking its target and provide real-time corrections to both under and over guiding—both of which smear out the telescope’s point spread function on the CCD and result in broader FWHM values for spectral lines. Guide cameras that utilize the reflected light off of a mirrored slit aperture are significantly more sensitive to these problems in good seeing, as the majority of the light falls through the spectrograph slit. In our experience, the loss of 4% of the star’s light is acceptable if it ensures that the telescope’s guiding

is steady throughout the night and across the different regions of the sky.

After each guiding exposure is completed, the guide camera then passes the 2-D FITS images it creates to the SourceExtractor software¹⁷ which analyzes each image and provides statistics on parameters such as the flux and FWHM, which are in turn used to evaluate the current atmospheric seeing. These guide camera images are also used to meter the exposures. Each image is integrated over the rectangular aperture corresponding to the utilized spectrograph slit, with background photons (determined using adjacent, background-estimating rectangles) subtracted off to determine the number of star-generated photons accumulated by the guide camera.

Analysis of the existing APF data suggests that the exposure meter rate (much like the iodine photon accumulation rate) depends on the star’s color, its V magnitude, the atmospheric seeing (in the form of slit losses), the airmass, and the date of observation. A multivariate linear regression to the exposure meter rate over these five terms results in the correlation displayed in Fig. 7.

Because the exposure meter is rapidly updated, we can monitor photon accumulation in real-time during an observation. Atmospheric seeing is already incorporated into the V_c term in the linear regression [Eqs. (3) and (4)], so any decrease from the expected exposure meter rate likely stems from an increase in clouds and corresponding decrease in atmospheric transparency. The ratio of expected exposure meter rate to observed exposure meter rate provides a “slowdown” factor that the scheduler tracks throughout the night and multiplies by the predicted “clear night” observation times calculated using Eq. (2) to determine a best guess exposure duration.

The regression in Fig. 7 gives

$$\log_{10}(r_E) = -\frac{1}{2.5} \{V_c + \alpha_E(B - V) + \beta_E[\sec(z)] + \gamma_E(\text{MJD}) + C_E\}, \quad (4)$$

where r_E is the photon accumulation rate on the exposure meter, α_E , β_E , γ_E , and C_E have the same meanings as in Eq. (2) but for the exposure meter photon accumulation rate instead of the Iodine photon accumulation rate, z is the star's zenith angle, and V_c is defined in Eq. (3).

3.4 Setting Upper Bounds on Exposure Time

Finally, because of the scatter seen in Figs. 6 and 7, we must ensure that exposures end when photons sufficient to achieve the desired RV precision have accumulated instead of continuing on for extra seconds or minutes. Photons in the I_2 region cannot be monitored during the observation. However, the APF dataset shows a strong relationship between the number of photons obtained in the I_2 region of the spectrum and the photons registered by the exposure meter, which is monitored in real time.

The telescope's guide camera, which is used for the exposure meter, has a broad bandpass and is unfiltered. This generates a strong color-dependent bias when comparing the guider photons to those that fall in the much narrower I_2 region. We apply a quadratic $B - V$ color correction term to produce the relation shown in Fig. 8. Combining this with the equations identified in Fig. 5, we obtain relations that allow us to relate the desired internal precision to the corresponding number of photons in the iodine region of the spectrum, and then to the number of photons required on the exposure meter. The resulting exposure meter threshold is used to place an upper limit on the exposure. This is particularly useful on nights with patchy clouds, where the

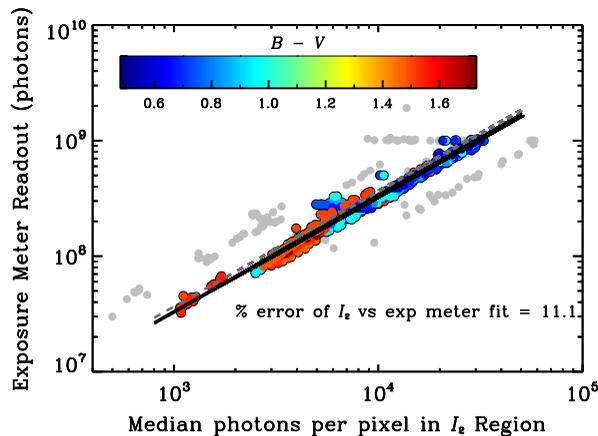


Fig. 8 Color-corrected relationship between the photons in the I_2 region of the spectrum and the photons registered by the exposure meter. The black line shows the best fit, and the gray dashed line shows the relation offset by one standard deviation, which is the limit we use operationally. To increase telescope efficiency, we require a way to ensure observations don't continue when the number of iodine region photons necessary to achieve the desired internal uncertainty has already been achieved. As described in Sec. 3.3.2 there is no way to measure the I_2 region photon accumulation in real time. However, the tight correlation between I_2 photons and photons on the exposure meter (which does update in real time) displayed here allows us to set a maximum exposure meter value based on our desired precision level. Thus, the observation software will end the exposure when the specified exposure meter value is met, even if the open shutter time falls short of the predicted observation time. This is particularly useful for cases where the cloud cover used in calculating the predicted observation time is actually more than the cloud cover on the target, which would result in an erroneously long observation.

cloud cover estimate calculated during the previous observation can be significantly higher than the cloud cover in other parts of the sky—resulting in artificially high-predicted observation times. In this case, the exposure meter can be used to stop an observation if the desired photon count is reached early, improving efficiency.

Fit applied in Fig. 8

$$\log_{10}(R) = \delta + \epsilon(B - V) + \zeta(B - V)^2, \quad (5)$$

where R is the ratio of photons on the exposure meter to photons in the iodine region of the spectrum.

3.5 Combining the Fits

The above relations are combined to enable the scheduling algorithm to select scientifically optimal targets. The following list summarizes the combination scheme.

Steps to determining an object's predicted observation time and exposure meter cut-off

1. Query spreadsheet for stellar attributes (V , $B - V$, required precision, RA, Dec, observing priority).
2. Use Eq. (1) to determine the desired number of photons in the I_2 region of the spectrum.
3. Use Eq. (2) to calculate the observation time required to obtain the desired total number of I_2 region photons in ideal transparency conditions.
4. Multiply the slowdown factor, calculated during the previous observation using Eq. (4), with the ideal transparency observation time estimates to get the predicted observation time in current conditions.
5. Calculate the exposure meter threshold based on the required number of I_2 region photons using Eq. (5).

We find scatters of 11.5% and 11.1% in Figs. 6 and 8, respectively (Table 1). This means that, even on a photometric night, we may not accumulate the number of photons in the I_2 region necessary for the desired precision as the photon arrival rate could be too low. To increase the likelihood of getting enough photons to reach the desired number of photons in the I_2 region, we increase the observation time estimate and the exposure meter threshold by 11.5% and 11.1%, or one standard deviation. By implementing this padding factor, we ensure that 84% of the time we observe a target, we will obtain the desired number of I_2 photons. However, this does not necessarily guarantee that we will achieve the desired internal uncertainty, due to the scatter in the relation between the I_2 region photons and the uncertainty estimates seen in Fig. 5 and quantified in Table 1.

Using these predicted observation times, the scheduler can evaluate whether any potential target can be observed at its desired precision within the 1-h observation time limit. Combining the predicted observation times with the targets' coordinates determines whether it will remain within the allowed 20-deg to 85-deg elevation range during the exposure. Stars that satisfy all these criteria are then ranked based on their observing priority, time past cadence requirement, and distance from the moon, with the highest scoring star being selected for observation. The scheduler then transmits the necessary information for the selected star, including its expected observation time and exposure meter threshold, to the observing software,

breaking up the total observation time into a number of individual exposures if necessary.

4 Dismissed Factors

Exposures obtained during 2013 and 2014 indicate that some factors initially suspected to be important need not impact target selection considerations. For example, the original observing protocol avoided targets within 45 deg of the direction of any wind above 5 mph to avoid wind shake in the telescope. We find, however, no discernible increase in the internal uncertainty (indicated by the color scale in Fig. 9) as a function of wind speed or direction. This resilience likely stems from an effective wind-shielding mode for the dome shutters, which opens them just enough to ensure that there is no vignetting of the target star.² In addition, substantial effort has been put into tuning the telescope's servo motors in order to "stiffen" the telescope and thus mitigate the effect of wind gusts that do manage to enter the dome.¹⁸

We also previously assigned higher priority to targets with elevation in the 60-deg to 70-deg range, removing scientifically interesting targets that were closer to the horizon from

consideration. As shown in Fig. 10, however, there is no significant loss in velocity precision as a function of elevation from 90 deg down to 20 deg, thanks to the telescope's atmospheric dispersion corrector (ADC), which works down to 15 deg. The telescope has a hard observing limit of 15 deg because of the ADC's range and because working at lower elevations leads to vignetting by the dome shutters. We still enforce an elevation range of 20 deg to 85 deg to avoid mechanical problems in telescope tracking at the high elevations. Working at elevations approaching our lower limit does result in longer predicted observation times (due to the airmass term), which can result in stars being skipped over in favor of other, higher-elevation targets.

Finally, we no longer assign a weighting value to the slew time necessary to move between targets. The APF is capable of moving at 3 deg s⁻¹ in azimuth and 2 deg s⁻¹ in elevation, which means that a direct slew to a target across the sky would take only 1 min. Because the CCD takes approximately 40 s to read out each observation, this slew time factor is small enough to be considered unimportant. Furthermore, with the introduction of the wind-shielding mode, the telescope's movement was

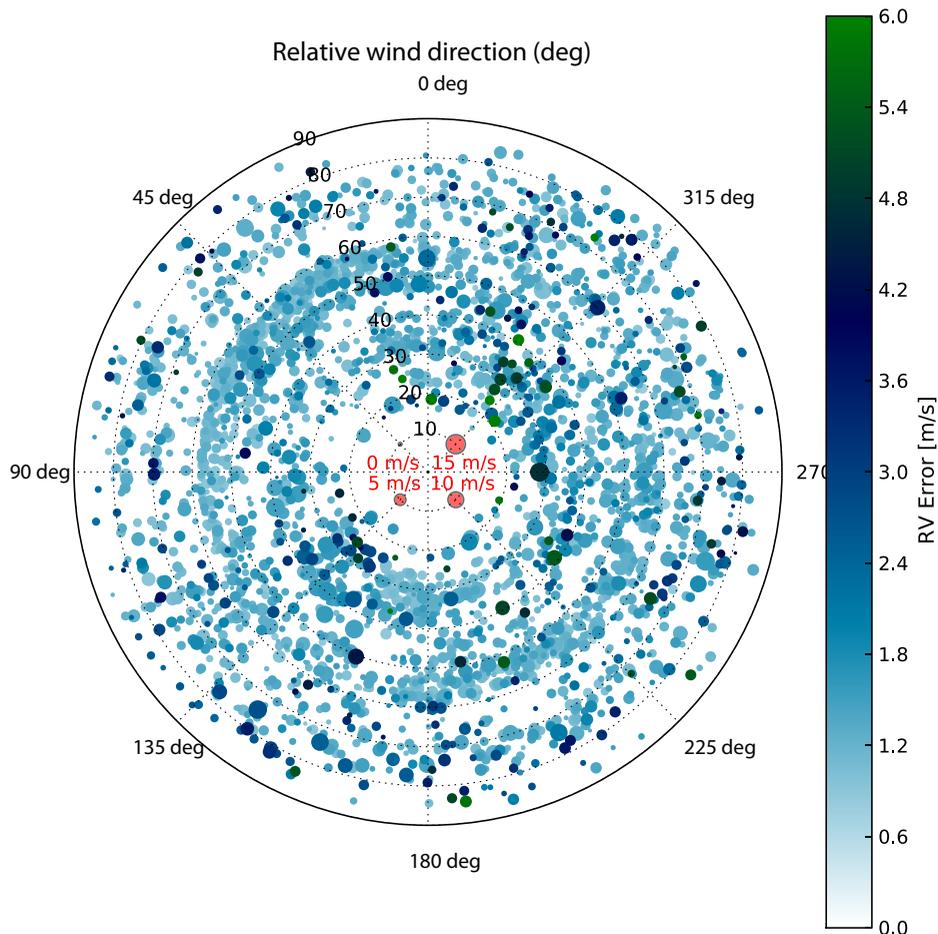


Fig. 9 Wind speed (point size) and direction (azimuthal position) plotted for 3155 individual exposures reveals no strong correlation between pointing near/into the wind and the estimated internal uncertainty displayed in the color bar. The exposures represented in this figure were obtained before we had the means to determine condition-based exposure times. Thus, all exposures were run until they reached their exposure meter threshold, or up to a static maximum exposure time of 900 s and then terminated, regardless of the number of photons collected by the exposure meter. This means the wind-based effects are not being mitigated by longer exposure times, and wind direction can thus be ignored when deciding which stars are considered viable targets for the next observation.

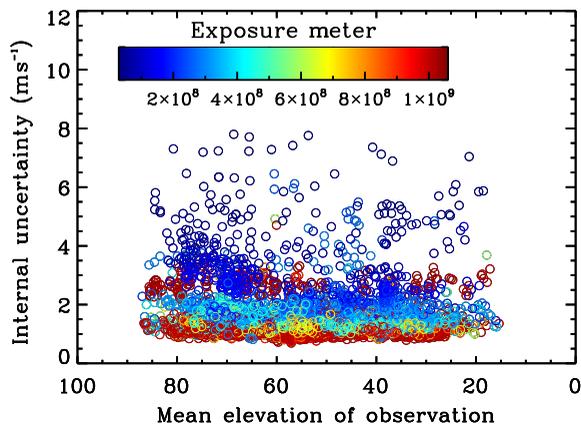


Fig. 10 The RV precision as a function of the elevation shows no strong correlation, once we compare observations with a fixed exposure meter value. As our linear regressions in Figs. 4 and 5 account for the decline in the photon accumulation rate with decreasing elevation, we do not need to add an additional term to account for other elevation effects such as increased seeing. Similar to Fig. 9, the observations presented here were taken before the adaptive exposure time software was in use. Thus every exposure has a static maximum exposure time of 900 s and the low elevation effects are not being mitigated by allowing for longer exposures. The telescope's ADC only functions down to 15 deg, the same elevation at which the telescope begins to vignette on the dome shutter, thus providing the lower limit for our observations.

altered so that it must first drop to a “safe elevation” of 15 deg before rotating to the target azimuth and then moving upward to the appropriate elevation. This is done to protect the primary mirror from falling debris while the dome shutters are reconfigured to minimize wind effects for the next observation during the slew. An additional result is that all slews take approximately the same amount of time, which provides justification for discounting slew time as an input to target selection.

5 Dynamic Scheduler Overview

In Sec. 3, we described a method to predict observation times for targets given their precision requirements and the current atmospheric conditions. To automate the determination of these observation times and the selection of the optimal target at any time throughout the night, we have implemented a dynamic scheduler (written in Python) called Heimdallr.

Heimdallr runs all of the APF's target selection efforts and interfaces with pre-existing telescope control software so that, once it submits an observation request, it waits until that set of exposures is completed before reactivating. Telescope safety, system integrity, and alerts about current weather conditions are monitored by two services called apfmon and checkapf. Each of these software sets has the ability to override Heimdallr if they detect conditions that pose a threat to, or represent a problem with, the telescope. This ensures that the facility's safety is always given priority.

Additionally, while directing the night's operations, Heimdallr uses other pre-existing utilities including openat-night, prep-obs, and closeup which, as their names suggest, open the facility prior to nightfall (or when night time conditions warrant), prepare the instrument and optical train for observing, and close the facility, securing the telescope when conditions warrant. The setting of the guider camera, the configuration of dome shutters and the control of telescope movement to avoid

interference with the cables wrapped around the telescope base are handled by yet another utility called scriptobs.

5.1 Observing Description

Typically, Heimdallr initiates in the afternoon and prompts the instrument control software to focus the instrument by obtaining a dewar focus cube (series of exposures of the quartz halogen lamp taken through the iodine cell) using the standard observation slit. Once the software determines a satisfactory value for the instrument's dewar position (via a simple linear least squares parabolic fit to the focus values), it proceeds to take all of the calibration exposures that are required by the data reduction pipeline. Upon completing these tasks, Heimdallr then waits until dusk, at which point it consults checkapf to ensure that there are no problematic conditions in the weather and then apfmon to ensure facility readiness. If both systems report safe conditions, it then opens the dome, allowing the telescope to thermalize with the outside air.

Heimdallr runs a main loop that continuously monitors a variety of keywords supplied by the telescope. At 6 deg twilight, the telescope is prepped for observing and begins by choosing a rapidly rotating B star from a predetermined list. The B star has no significant spectral lines in the I_2 region and serves as a focus source for the telescope's secondary mirror while also allowing the software to determine the current atmospheric conditions (as described in Sec. 3.3.2).

At 9 deg twilight, Heimdallr accesses the online database of potential targets and parses it to obtain all the static parameters described in Sec. 3.1. It then checks the current date and time and eliminates from consideration those stars not physically available. The scheduler then employs the stars' coordinates, $B - V$ colors and V magnitudes, combined with the seeing and atmospheric transparency determined during the previous observation, to calculate the predicted observation time for each target given their desired precision levels. Stars unable to reach the desired precision within the 1-h maximum observation time are eliminated from the potential target list.

The remaining stars are ranked based on scientific priority and time past observing cadence. The star with the highest score is passed to scriptobs to initiate the exposure(s). When the exposures finish, Heimdallr updates the star's date of last observation in the database with the photon-weighted midpoint of the last exposure and begins the selection process anew.

When the time to 9 deg morning twilight becomes short enough that no star will achieve its desired level of precision before the telescope must close, Heimdallr shuts the telescope using closeup and initiates a series of postobserving calibration exposures. Once the calibrations finish, Heimdallr exits.

5.2 Other Operational Modes

In addition to making dynamic selections from the target database during the night, Heimdallr can also be initiated in a fixed list or ranked fixed list operational mode. The fixed list mode allows observers to design a traditional star list that Heimdallr will move through, sending one line at a time to scriptobs. Any observations that are not possible (due to elevation constraints) will be skipped and the scheduler will simply move to the next line.

The ranked fixed list option allows users to provide a target list that Heimdallr parses to determine the optimal order of observations. That is, after finishing one observation from the

list, the scheduler will then perform a weighting algorithm similar to what is employed by the dynamic use mode to determine which line of the target list is best observed next. Heimdallr keeps track of all the lines it has already selected, so that they do not get initiated twice, and will re-analyze the remaining lines after each observation to select the optimal target. This option is especially useful for observing programs that have a large number of possible targets but do not place a strong emphasis on which ones are observed during a given night.

6 Comparing with Other Facilities

There are a number of automated and semiautomated facilities that perform similar observations. Examples include HARPS-N, CARMENES, the Robo-AO facility at Palomar, and the Las Cumbres Observatory Global Telescope (LCOGT) network. In order to place the APF's operations in context relative to these other observatories, we will briefly discuss the approaches of these other dedicated RV facilities (CARMENES, HARPS-N, and the NRES addition to LCOGT) and of the more general facility with queue scheduling (Robo-AO). We will then highlight the common approaches along with discussing what is simpler for our facility as it is dedicated to a single-use.

6.1 Radial Velocity Surveys

RV surveys generally require tens or hundreds of observations of the same star to detect planetary companions. Traditionally, high-precision velocities were obtained on shared-use facilities and observing time was limited. More recently, however, purpose built systems (such as HARPS) have presented the opportunity to obtain weeks or months of contiguous nights.

The HARPS-N instrument, installed on the Telescopio Nazionale Galileo at Roque de los Muchachos Observatory in the Canary Islands, is a premier system for generating high-precision velocities. At present, it is primarily devoted to Kepler planet candidate follow up and confirmation. The system allows users to access XML standard format files that define target objects using the Short-Term-Scheduler GUI to guide the process and then assemble the objects into an observing block. When an observer initiates an observation, these blocks and their associated observing preferences are passed to the HARPS-N Sequencer software, which places all of the telescope subsystems into the appropriate states, performs the observation, and then triggers the data reduction process.¹⁹ Thus, while the building of target lists has been streamlined, nightly observation planning still requires attention from an astronomer or telescope technician.

The forthcoming Calar Alto CARMENES instrument is expecting first light in 2015.

CARMENES will also be used to obtain high-precision stellar RV measurements on low-mass stars. Its automated scheduling mechanism relies on a two-pronged approach: the off-line scheduler which plans observations on a weekly to nightly time scale based on target constraints that can be known in advance, and the on-line scheduler which is called during the evening if unexpected weather or mechanical situations arise and require adapting the previously calculated nightly target list.²⁰

Finally, the LCOGT network will soon implement the Network of Robotic Echelle Spectrographs (NRES), six high-resolution optical echelle spectrographs slated for operation in late 2015. Like all instruments deployed on the LCOGT network, the scheduling and observing of NRES will be autonomous. Users submit observing requests via a web interface

which are then passed to an adaptive scheduler which works to balance the requests' observing windows with the hard constraints of day and night, target visibility, and any other user specified constraints (e.g., exposure time, filter, airmass). If the observation is selected by the adaptive scheduler, it creates a "block" observation tied to a specific telescope and time. This schedule is constructed 7 days out, but rescheduling can occur during the night if one or more telescopes become unavailable due to clouds, or if new observing requests arrive or existing requests are canceled. In this event, the schedule is recalculated, and observations are reassigned among the remaining available sites.²¹

6.2 Automated Queue Scheduling

Robo-AO is the first fully automated laser guide star adaptive optics instrument. It employs a fully automated queue scheduling system that selects among thousands of potential targets at a time with an observation rate of ~ 20 objects h^{-1} . Its queue scheduling system employs a set of XML format files that use keywords to determine the required settings and parameters for an observation. When requested, the queue system runs each of the targets through a selection process, which first eliminates those objects that cannot be observed, and then assigns a weight to the remaining targets to determine their priority in the queue at that time. The optimal target is chosen, and the scheduler passes all observation information to the robotic system and waits for a response that the observations were successfully executed. Once the response is received, the observations are marked as completed and the relevant XML files are updated.²²

6.3 Comparing our Approach to Other Efforts

We have designed this system based on our RV observing program carried out over the past 20 years at Keck and other facilities. This experience, coupled with the pre-existing software infrastructure, has guided the development of both the dynamic scheduler software itself and our observing strategy.

Comparison to these other observing facilities and the strategies they employ emphasizes some shared design decisions. For example, the APF has a similar target selection approach to that of the Robo-AO system. Both utilize a variety of user specified criteria to eliminate those targets unable to meet the requirements and then rank the remaining targets, passing the object with the highest score to the observing software. Additionally, as is common with all of the observing efforts mentioned above, our long-term strategy is driven by our science goals and is in the hands of the astronomers involved with the project.

Several differences are also notable. The first is that we lack an explicit long-term scheduling component in our software. Our observing decisions are made in real time in order to address changes in the weather and observing conditions that occur on minute to hour time scales, and to maximize the science output of nights impacted by clouds or bad seeing. However, for a successful survey there must also be a long-term observing strategy for each individual target. We address this need via a desired cadence and required precision for every potential target. By incorporating the knowledge of how often each star needs to be observed and a way to assess whether the evening's conditions are amenable to achieving the desired precision level, Heimdallr adheres to the long-term observing strategy outlined by our observing requirements and doesn't need to generate separate multiweek observing lists.

A second difference is that the final output of the scheduler is a standard star list text file, one line in length. This format has been in use in UCO-supported facilities such as Lick and Keck Observatories for more than 20 years, and thus is familiar to the user community. The file is a simple ASCII text file with key value pairs for parameters and a set of fixed fields for the object name and coordinates. This permits read-by-eye verification of the next observation if desired and allows the user to quickly construct a custom observing line that can be inserted into the night's operations if needed. Observers can furthermore easily make a separate target list in this format for observations when not using the dynamic scheduler (see Sec. 5.2)

Finally, we use a Google spreadsheet for storing target information and observing requirements, as opposed to a more machine-friendly format such as XML. Although Lick Observatory employs a firewall to sharply limit access to the APF hosts, it only operates on the incoming direction. Thus, it is straightforward and uncompromising from a security standpoint for our internal computers to send a request to Google and pull the relevant data back onto the mountain machines. This approach provides team members with an accessible, easy-to-read structure that is familiar and easily exportable to a number of other formats. Google's version allows for careful monitoring of changes made to the spreadsheet and ensures that any accidental alterations can be quickly and easily undone. The use of the Google software also allows interaction with a browser, so no custom GUI development is required. Therefore, we are taking advantage of existing software to both minimize our development effort, and make it as easy as possible to have the scientists update and maintain the core data files that control the observations.

7 Observing Campaigns on the Automated Planet Finder

The APF has operated at high precision for over a year, and has demonstrated precision levels of $\sigma \sim 1 \text{ ms}^{-1}$ on bright, quiet stars such as Sigma Draconis.² The telescope's slew rate permits readout-limited cycling and 80% to 90% open shutter time, allowing for 50 to 100 Doppler measurements on clear nights.

7.1 Lick–Carnegie Survey

Heimdallr's design dovetails with the need to automate the continued selective monitoring of more than 1000 stars observed at high Doppler precision at Keck over the past 20 years.²³ The APF achieves a superior level of RV precision and much-improved per-photon efficiency in comparison to Keck/HIRES for target stars with $V < 10$. As a primary user, we can employ the APF on 100+ nights per year in the service of an exoplanet detection survey.

At present, 127 stars have been prioritized for survey-mode observation with the APF. This list emphasizes stars that benefit from the telescope's more northern location, and gives preference to stars that display prior evidence of planetary signals. We adopt a default cadence of 0.5 observations per night. When a star is selected, it is observed in a set of 5- to 15-min exposures with the additional constraint that the total of a night's sequential exposures (the observation time) on the star is less than an hour. Additionally, information obtained at Keck has, in some cases, permitted estimates for the stellar activity of specific target stars. In these cases, observation times can be adjusted to conform to a less stringent desired precision.

With its ability to predict observation times, Heimdallr readily achieves efficiencies that surpass the use of fixed lists, and indeed, its performance is comparable to that of a human observer monitoring conditions throughout the night.

7.2 Transiting Exoplanet Survey Satellite Preccovery Survey

Heimdallr is readily adopted to oversee a range of observational programs, and a particularly attractive usage mode for APF arises in connection with NASA's Transiting Exoplanet Survey Satellite (TESS) mission. TESS is scheduled for launch in 2017, and is the next transit photometry planet detection mission in NASA's pipeline.²⁴ Transit photometry observations of potential planet signatures generally require follow-up confirmation, with RV being the most common. Currently, there is a dearth of high-precision RV facilities in the northern hemisphere, and HARPS-N is heavily committed to Kepler planet candidate followup. NASA recently announced plans to develop an instrument for the 3.5 m WIYN telescope at Kitt Peak Observatory capable of extreme precision Doppler spectrography to be used for follow up of TESS planet candidates.²⁵ However such an instrument will require time for development and commissioning and thus is not a viable candidate for pre-launch observations of TESS target stars. The APF's year-round access to the bright stars near the north ecliptic pole, which will obtain the most observation time from TESS, makes it an optimal facility to conduct surveys in support of the satellite's planet detection mission.

At present, comparatively little is known about the majority of TESS's target stars. We have little advance knowledge of which stars host properly inclined, short-period transiting planets observable by the satellite. We will thus start with a default value for the observing cadence and be poised to adapt quickly should hints of planetary signatures start to emerge. Additionally, initial desired precisions (and the corresponding observation times) must rest on the suspected stellar jitter of the targets.

The availability of on-line databases to track all of the science-based criteria will be crucial for moving back and forth quickly between this project and further Lick–Carnegie follow up. Evaluations of the value of APF coverage are governed by three scientific criteria (priority, cadence and required precision), along with supporting physical characteristics (RA, Dec, V and $B - V$) for each target. Thus when beginning RV support for TESS, Heimdallr can easily be instructed to reference the TESS database when determining the next stellar target (instead of the Lick–Carnegie List).

TESS observations also provide an excellent test bed for experimenting with alternate observational strategies. For example, Sinukoff et al.²⁶ stated that obtaining three 5-min exposures of a star spaced approximately 2 h apart from one another during the night results in a 10% increase in precision over taking a single 15-min exposure. The TESS stars that will be monitored by the APF are all located in the north ecliptic pole region, meaning that the slew times will be almost negligible. It is thus likely that an observing mode that subdivides exposures to improve precision could be very valuable. In short, the APF is extremely well matched to the TESS Mission.

Acknowledgments

We are pleased to acknowledge the support from the NASA TESS Mission through MIT sub award #5710003702. We are also grateful for partial support of this work from NSF grant

AST-0908870 to SSV, and NASA grant NNX13AF60G to RPB. This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France. This material is based upon work supported by the National Aeronautics and Space Administration through the NASA Astrobiology Institute under Cooperative Agreement Notice NNH13ZDA017C issued through the Science Mission Directorate.

References

- I. Zolotuhkin, “The extrasolar planets encyclopaedia,” 1995, <http://exoplanet.eu/> (2 September 2015).
- S. S. Vogt et al., “APF-The Lick Observatory Automated Planet Finder,” *Publ. ASP* **126**, 359–379 (2014).
- B. J. Fulton et al., “Three super-earths orbiting HD 7924,” *Astrophys. J.* **805**, 175 (2015).
- M. V. Radovan et al., “A radial velocity spectrometer for the Automated Planet Finder Telescope at Lick Observatory,” *Proc. SPIE* **7735**, 77354K (2010).
- N. C. Santos et al., “The HARPS survey for southern extra-solar planets. II. A 14 Earth-masses exoplanet around μ Arae,” *Astron. Astrophys.* **426**, L19–L23 (2004).
- X. Dumusque et al., “Planetary detection limits taking into account stellar noise. II. Effect of stellar spot groups on radial-velocities,” *Astron. Astrophys.* **527**, A82 (2011).
- J. Burt et al., “The Lick–Carnegie exoplanet survey: Gliese 687—a Neptune-mass planet orbiting a nearby Red Dwarf,” *Astrophys. J.* **789**, 114 (2014).
- R. W. Noyes et al., “Rotation, convection, and magnetic activity in lower main-sequence stars,” *Astrophys. J.* **279**, 763–777 (1984).
- S. Meschiari et al., “Systemic: a testbed for characterizing the detection of extrasolar planets. I. The systemic console package,” *Publ. ASP* **121**, 1016–1027 (2009).
- M. Wenger et al., “The SIMBAD astronomical database. The CDS reference database for astronomical objects,” *Astron. Astrophys.* **143**, 9–22 (2000).
- R. P. Butler et al., “Attaining Doppler precision of 3 ms^{-1} ,” *Publ. ASP* **108**, 500 (1996).
- J. T. Wright, “Radial velocity jitter in stars from the California and Carnegie planet search at Keck Observatory,” *Publ. ASP* **117**, 657–664 (2005).
- F. Bouchy, F. Pepe, and D. Queloz, “Fundamental photon noise limit to radial velocity measurements,” *Astron. Astrophys.* **374**, 733–739 (2001).
- P. Connes, “Absolute astronomical accelerometry,” *Astrophys. Space Sci.* **110**, 211–255 (1985).
- S. S. Vogt et al., “HIRES: the high-resolution echelle spectrometer on the Keck 10-m Telescope,” *Proc. SPIE* **2198**, 362–375 (1994).
- T. C. Beers, K. Flynn, and K. Gebhardt, “Measures of location and scale for velocities in clusters of galaxies—a robust approach,” *Astron. J.* **100**, 32–46 (1990).
- E. Bertin and S. Arnouts, “SExtractor: software for source extraction,” *Astron. Astrophys.* **117**, 393–404 (1996).
- K. Lanclus et al., “Tuning a 2.4-meter telescope. . . blindfolded,” *Proc. SPIE* **9145**, 91454B (2014).
- R. Cosentino et al., “Harps-N: the new planet hunter at TNG,” *Proc. SPIE* **8446**, 84461V (2012).
- A. Garcia-Piquer et al., “CARMENES instrument control system and operational scheduler,” *Proc. SPIE* **9152**, 915221 (2014).
- J. D. Eastman et al., “NRES: the network of robotic echelle spectrographs,” *Proc. SPIE* **9147**, 914716 (2014).
- R. L. Riddle et al., “The Robo-AO automated intelligent queue system,” *Proc. SPIE* **9152**, 91521E (2014).
- R. P. Butler, “The HIRES/Keck precision radial velocity exoplanet survey,” *Astrophys. J.* (2016).
- G. R. Ricker et al., “Transiting exoplanet survey satellite (TESS),” *Proc. SPIE* **9143**, 914320 (2014).
- T. Beal, “NASA needs Kitt Peak telescope for exoplanet duty,” 2014, The Arizona Daily Star, http://tucson.com/news/blogs/scientific-bent/nasa-needs-kitt-peak-telescope-for-exoplanet-duty/article_20f88d08-9c03-5ae4-8eab-d4984f4bdbff.html (2 September 2015).
- E. Sinukoff, “Optimization of planet finder observing strategy,” in *Search for Life Beyond the Solar System. Exoplanets, Biosignatures and Instruments*, D. Apai and P. Gabor, Eds., p. 4P, Cambridge University Press (2014).

Jennifer Burt is a PhD candidate in astronomy and astrophysics at the University of California Santa Cruz. She received her BA degree in astronomy from Cornell University in 2010 and anticipates completion of her PhD in 2016. Her current research interests include exoplanet detection, observational survey design, and understanding the effects of stellar activity on RV data.

Bradford Holden is an associate research astronomer for the University of California Observatories. He received his PhD in astronomy and astrophysics from the University of Chicago in 1999. He is an author on 99 journal papers. His current research interests include detector development, survey planning, and automation of observatory operations.

Sandy Keiser has been with the Carnegie Institution’s Department of Terrestrial Magnetism for over 20 years working in the fields of astronomy, astrophysics, seismology, and cosmochemistry. She graduated from the College of William and Mary in 1978. During the last 7 years, she has been a coauthor on 10 journal papers in astrophysics.

Biographies for the other authors are not available.